A SEMI-QUANTITATIVE CATASTROPHIC RISK LIKELIHOOD PRIORITIZATION FRAMEWORK FOR THE METALLURGICAL INDUSTRY

S. Hlouschko* and M. Cramer
Hatch
2800 Speakman Dr, Mississauga, Ontario, Canada L5K 2R7
(*Corresponding author: stefan.hlouschko@hatch.com)

M. Pergler
Balanced Risk Strategies
195 Glebe Ave, Ottawa, Ontario, Canada K1S 2C6

ABSTRACT

The paper discusses a semi-quantitative risk reduction decision framework for catastrophic risks at a single asset or among multiple assets in the metallurgical industry. Catastrophic risks can be identified through a comprehensive audit program, but the implementation of critical controls can be costly and face resistance due to the belief there is a very remote likelihood the consequences of these risks will materialize. The framework combines elements of the risk bowtie and layer of protection analysis and is based on documented reference statistics. These are conditioned by structured expert judgment using order-of-magnitude heuristics to capture how the specific context differs from the quantified reference classes. Combined with order-of-magnitude consequence assessment, the framework allows for more robust risk prioritization than typical, purely qualitative risk heat maps, and leads naturally to implementation of effective mitigation actions.

KEYWORDS

Arsine, Bowtie, Catastrophic risk, Decision making, Delphi, Likelihood, LOPA, Risk management

INTRODUCTION

Catastrophic risks can be defined as the risks with highest impact and yet low likelihood. Consequences of such risks in ongoing operations include multiple fatalities, significant environmental damage, asset destruction, long term operations stoppage, reputational damage, and legal or regulatory penalties. Financial and management resources are typically easier to assign to high impact risks with perceived high likelihood, since the impression of urgency exists. Conversely, more stringent justification is often required to assign resources to high impact risks with perceived low likelihood, since as the saying goes: it will never happen here.

Industry wide and within individual companies, a variety of approaches are used to assess low likelihood risks, evaluate potential controls, and audit implementation of these controls. These methods include qualitative and quantitative heat maps, the risk bowtie, layer of protection analysis, quantitative risk assessment including fault tree analysis, event tree analysis etc. Their applicability depends on impact severity, failure mode complexity, available data and expertise, and organizational capacity, among other factors. Most risk assessment is done under considerable uncertainty and this can lead to analysis-paralysis rather than a timely move to action.

Decision making under uncertainty can be supported through effective consensus building between risk assessment stakeholders, which includes site personnel, industry experts, and management. Each group understands problems in different ways and at different levels of resolution. Careful framing of the risk problem to be solved coupled with workshop facilitation techniques such as the Delphi method have been used successfully by practitioners to progress to meaningful risk reduction and follow-up verification. Such techniques also move the perception of the risk assessment process from that of a check-the-box theoretical exercise to effective decision support and help address human biases in dealing with risk through the amplification of reasoning (System 2) versus purely intuitive (System 1) thinking. (Kahneman, 2002).

Finally, evaluating risks between multiple assets is challenging where the assessment framework lacks transparency and assets are located in different cultures and are at different levels of operational maturity. It is, however, often crucial in supporting the allocation of risk reduction resources across an asset portfolio. A comprehensive visual-quantitative framework which explicitly indicates how subjectivity is applied can permit effective cross-cultural and hierarchical communication, allow for risk comparison between assets, and support more effective dialogue on the sharing of mitigation best practices.

METHOD OVERVIEW

The proposed risk reduction decision framework is an adapted layer of protection analysis (LOPA) and risk bowtie applied using the Delphi method and half order of magnitude heuristics in a workshop setting. Its key aim is to arrive at an estimate of catastrophic event frequency and motivate follow up mitigating actions. A detailed explanation follows in the next section.

An absolute estimate of low likelihood event frequency is possible using LOPA, but depends on data quality, failure mode complexity, and the ability to calibrate the result to reality (engineering standards, benchmarks, equipment data etc.). However, in many instances the event in question has not occurred within living memory or has no reliable analogues, and investigations of any past events may have failed to identify all root causes. As well, crucial barriers often lack good effectiveness data, expert opinion can be hard to incorporate, and assessments can over index on illusory analogues where data is available (UK Health and Safety Executive, 2009). These complicating factors combined with LOPA's inherent conservatism make convincing absolute estimates of event frequency impractical in many cases.

A relative risk comparison between multiple scenarios grounded in the same assumptions offers another way to act with imperfect information and can support some scenarios with complex failure mechanisms or poor data. Comparing several scenarios and targeting improvement in those with the highest relative likelihood can lead to unquantified absolute risk reduction in the short term, during which there may be an opportunity to collect site specific data to support an absolute estimate or decide on more detailed quantitative risk analysis.

ANALYTIC BASIS

The risk bowtie and LOPA form the analytic basis for the proposed framework. These methods are applied widely in risk assessment generally and have begun to appear together in software packages in various ways, though only very recently. The rules set forth in the semi-quantitative LOPA (Center for Chemical Process Safety, 2001) help reduce bias and structure critical barriers while the qualitative risk bowtie helps preserve the 'big picture' in an accessible visual format.

The risk bowtie is a qualitative assessment of barrier sufficiency and is typically developed through a workshop following a qualitative process hazard assessment performed by subject matter experts. It is an intuitive model to communicate the causes and consequences of unwanted events and the barriers to deployed to prevent and mitigate them, showing event sequence from left to right. The top event (sometimes call the 'materially unwanted event') is shown at the center of the bowtie and is discussed as the loss of control of an identified hazard. Threats must penetrate barriers before causing the top event (i.e., the barriers must fail).

Despite its visual utility, the risk bowtie is not a test of barrier sufficiency and does not prioritize the relative importance of barriers.

LOPA uses simplifying rules to evaluate independent layers of protection against a hazard scenario. It is a semi-quantitative approach that applies to a single cause-consequence pair when the failure mode is well understood. LOPA is not a substitute for detailed quantitative risk analysis when the failure mechanisms and barriers are overly complex or poorly understood, or where precise risk-return tradeoff comparisons are needed to support decisions.

The initiating event (typically an annual frequency) is multiplied by conditional modifiers (probability of ignition, fatality etc.) to arrive at the frequency of an unmitigated consequence. The frequency of the mitigated consequence is calculated by multiplying the probability of failure-on-demand (PFD) of identified independent protection layers (IPLs), analogous to the barriers in the bowtie. Sources of failure data include generic databases, actual site measurements, expert consensus, and manufacturer predictions.

IPL failures are typically described in order of magnitude units, which recognizes the usual paucity of comparable data and makes the analysis inherently conservative. The LOPA literature lays out basic criteria for IPLs. The IPL must be effective in preventing the consequence when it functions as designed, independent of the initiating event and the components of any other IPL already claimed for the same scenario, and auditable, i.e., the assumed effectiveness in terms of consequence prevention and PFD must be capable of validation in some manner (by documentation, review, testing, etc.). These criteria can be applied to existing risk bowtie barriers.

LOPA is usually summarized in a worksheet calculation without detailed symbolic visual representation. This might be a useful format for a technical audience but increases the perception of being a theoretical exercise. In addition, a lack of transparency with respect to subjectivity applied can hinder decision making up and across the organization**.** Adding LOPA elements to the bowtie and ensuring the bowtie barriers meet IPL requirements provides for more intuitive comprehension and makes the logic more accessible by decomposing the mathematics back onto a time-based event flow diagram.

INCORPORATING EXPERT JUDGMENT

A successful decision-making process or framework relies on quality inputs, careful attention to the limitations of the assumptions, thoughtful framing of the problem and its elements, and incorporates the perspectives—consistent or not—of observers and stakeholders.

In our case, the problem is initially framed starting with the existing bowtie and assessing that the proposed barriers meet LOPA IPL independence, effectiveness, and auditability requirements. The bowtie is then seeded with default initial data (e.g. a basic process control system has a default probability of failure on demand of 10^{-1}, or 10%) which depends on a number of assumptions. This generates a default case, or a rebuttable presumption. The default case is an important element of the framework as it avoids starting in a vacuum, provides a basis from which to de-rate or up-rate the values based on workshop participation expertise or more specific data, yet avoids an overly rosy “we can’t think of anything that would go wrong here” bias among the workshop participants.

A workshop is held to review the default case with site and ideally other experts. The aim is a consistent and systematic approach which explicitly harnesses and documents subjectivity but does not provoke over-analysis. Specific goals include:

- Validation or improvement of the bowtie
- Adjustment of IPL strengths in half order-of-magnitude units (i.e. factor of $3 \approx 10^{1/2}$) in relation to reference class incidence data
- Identification of IPL improvements, both low-hanging fruit and more complex process or other improvements that may require the investment of resources.

- Agreement on next steps, which may be execution of mitigation actions to improve the IPLs or further evidence gathering or analysis where required to have an actionable conclusion. These steps should be auditable or at least amenable to clear monitoring.

The use of half-order-of-magnitude allows for a greater degree of granularity to up-rate or de-rate initial estimates than the conservative full order of magnitude (factor of 10) common in LOPA methodology. It also reflects emerging best practice in enterprise/strategic risk management (Pergler, 2020), where a factor of ~3x in the likelihood (or impact) of major risks is worth reflecting in their relative prioritization, yet more precise estimation may be unfeasible or completely unreliable for risks following an approximately power-law distribution.

The workshop operates on a "grounded consensus" model. Participants bring their expertise and evidence to support changes to the assumptions. Where there is initial disagreement, drivers of difference are explored using a modified version of the Delphi method. In this approach, consensus is sought through discussion, rather than merely voting or averaging. Where stable dissensus persists (e.g. two or more of ~10 continue to hold a divergent view) the more conservative estimate is utilized, however actions to support evidence-based reassessment of it may be agreed on as part of the next steps.

The Delphi method was initially developed at the RAND Corporation in 1964 and involved several rounds of voting on predictive macroeconomic and other estimates in the context of strategic scenario development. Participating experts provided and voted on not only specific estimates but their rationale, which was anonymized and shared as input into the next round. The initial objective was to avoid the biasing of group consensus by the most vocal or prestigious of the experts. More recent applications especially in the risk management context have focused more on free-flowing exploratory discussion after each round to identify salience biases and hidden pools of information, and the use of technology to speed up cycle time to minutes rather than days. As a result, expert groups tend to converge to a stable consensus reflecting all information available to them in 2-3 Delphi rounds in many cases, or a "stable dissensus" reflecting continuing differences of opinion as to the relevance or completeness of information (Freeman & Pergler, 2008)

It is helpful to impose certain rule-based constraints on the level of allowed flexibility afforded to the workshop to deviate from the default case. For example, one could set a lower limit to the failure of a human reaction to an abnormal situation under stress at 10^{-1}. These constraints are broadly addressed in the literature on LOPA generic data (Center for Chemical Process Safety, 2015). Such constraints help address potential group biases arising from overconfidence or sunflower management, as well as demonstrably indicate the desire for conservatism and provide pragmatic circuit breakers for discussions that can be perceived as individually evaluative: "we only allow a limit of 10^{-1} for such a process unless additional IPLs are implemented in conjunction" cuts off a perceived implied conflict of "why don't you just trust me when I say my people do their job?"

The direct output of the workshop updates the default case as follows:

1. A "consensus base case" reflecting the current state with grounded, agreed-on, and documented changes versus the default case
2. Typically, one or more future state options are added to arrive at "consensus base case with proposed mitigations".
3. A placeholder is added for a "one-year verification follow-up" to track progress.

The results can be compared to corporate risk tolerance if applicable, or more readily to similar risk evaluations with the same assumptions at other assets. They provide a consistent and grounded input to resource allocation decisions within the asset or across assets to provide maximum overall risk reduction.

Importantly, the process of creating the above output typically already builds immediate consensus around a certain number of highly effective risk mitigation steps that are low-hanging fruit to put in place. These form part of the next steps and may include ensuring that certain processes are applied as described, documented, and monitored, or that simple process changes are just put in place. This contributes to the actual and perceived utility of the risk discussion. However, it is important to differentiate what mitigation is actually demonstrably already in place (and should be in the consensus base case) versus what is aspirational, if easy to implement, and should be part of one or more of the future case options and subject to monitoring and follow-up.

EXAMPLE: WORKER EXPOSURE TO ARSINE GAS IN A CLEANED STORAGE TANK

This worked illustrative example builds upon a fragment of the generalized risk bowtie for worker exposure to arsine proposed by Mitsui, Koehler and Krysa (2019).

Arsine gas can be generated when acid comes into contact with arsenic metal. Arsenic metal can be present from upstream reduction processes or generated from arsenic solutions by a metal more reactive than a specific potential. Confined spaces with residual arsenic bearing solution are areas workers could drop metals tools or walk inside with active metal dust on their boots and generate very toxic arsine gas and be exposed.

For simplified illustration this example focuses on the preventive side of the materially unwanted event and associated hazard. Mitigating IPLs such as PPE, evacuation, and emergency services are not discussed. We define the initiating event to be a worker entering the storage tank and the enabling condition as the probability of there being dissolved and reducible arsenic in the storage tank solution. We further define the materially unwanted event to be the worker's exposure to hazardous levels of arsine gas in the tank. In the absence of any controls the event sequence is as follows:

- Enabling condition: Dissolved arsenic in storage tank
- Initiating event: One or more workers enter storage tank
- Materially unwanted event: Worker contacts dissolved arsenic solution with active metal, generating arsine gas, and is exposed

The enabling condition is modeled to depend on several factors:

- The probability that excessive arsenic enters the process feedstock and goes undetected
- The probability that the process is designed or controlled such that undetected arsenic enters the chemical process and can be dissolved into the tank liquid
- The probability that the process medium can dissolve sufficient arsenic
- The probability that any downstream independent verification of the initial arsenic detection in has failed

Each of these can be assigned a probability outcome based on known or estimated excessive arsenic levels in feedstock, detection intervals for sampling and assaying, and chemical related factors. In order for the worker exposure to arsenic to take place, the worker must enter the tank in the presence of the enabling condition and the following potential IPLs must fail:

- IPL1 - Failure of procedural control: The worker brings active metal into the tank despite being trained on the procedure prohibiting this.
- IPL2 - Failure of ventilation control: The tank mechanical ventilation fails and/or active testing for arsine gas during service fails, and/or the tank entry is from the top. Top entry tanks pose a greater risk than side entry tanks since arsine is heavier than air and would accumulate. This is an illustration of a proposed composite IPL with points based on a truth table that could also be built up via fault or event tree.
- IPL3 – Failure of procedural control: The worker enters an unclean tank which could contained dissolved arsenic solution. Special considerations apply if two procedural controls are to be credited as IPLs under the same scenario.

- Conditional modifier: Probability that the worker contacts the arsenic solution with active metals via tool drop or boot dust etc., resulting in generation of and exposure to arsine gas. We refer to this as a ‘luck factor’ and in practice it would scale with the number of workers in the tank.

  The ‘event flow’ showing the enabling conditions and IPLs is shown in Figure 1.

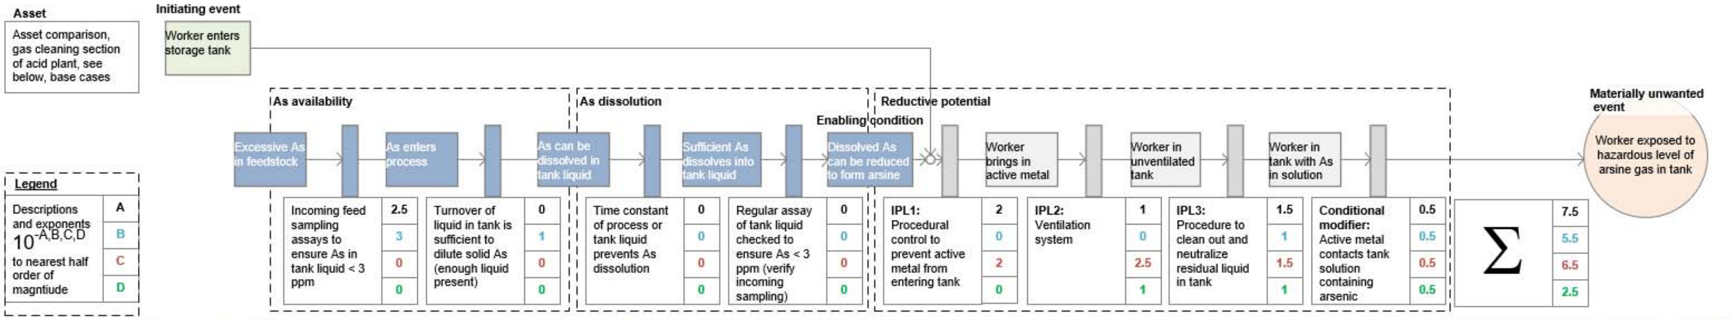

	Initiating event	Enabling condition probability build-up				IPL probability of failure on demand (PFD)				Materially unwanted event	
Description	**Initiating event frequency (IEF)** (tank-entries/year)	Human or mechanical/ automated sampling	A function of how much As present, how much deports to product, rate at which it is blowdown from process	Kinetics, solid phase transition to dissolved (time), not all solid arsenic dissolves	Human or mechanical/ automated sampling	Human procedure; active metal can include tools, ladder, boots with metal dust	Tanks can be side entry or top entry, atmosphere tested before entry, mechanical ventilation	Human procedure	Worker drops a tool, places a ladder, walks on surface with arsenic bearing solution etc.	**Total indicative likelihood** (events/tank-entry)	**Frequency of mitigated consequence** (events/year)
A. Default starting data	1 tank x 1 entry/year	**$10^{-2.5}$** Assume excess As always present and feed sampling each day, 1 failure/year = 1/365 days ~ 0.3% of year Arsenic enters process ~ $10^{-2.5}$ Sampling failure rate from generic data: Human error during task performed greater than once per week: 1 failure/ year	**10^{0}** Process specific Placeholder assumption 100% probability = 10^{0}	**10^{0}** Process specific Placeholder assumption 100% probability = 10^{0}	**10^{0}** Site specific Consider sampling frequency relative to incoming feed batching, other process factors, and earlier sampling etc. Placeholder assumption 100% probability = 10^{0}	**10^{-2}** Generic data: Human error during task performed less than once per month: 0.01 failures/year 1 tank-entry/year therefore failure rate = 0.01/1 = 10^{-2}	**10^{-1}** Mechanical vent = 1 *PFD truth table* *Sum:* *Top entry = 0* *Bottom entry = 1* *Not testing = 0* *Test before enter = 1* *No ventilation = 0* *Mechanical vent = 1* *Allow max total 2.5*	**$10^{-1.5}$** Generic data: Human error during task performed less than once per month: 0.01 failures/year 1 tank-entry/year therefore failure rate = 0.01/1 = 10^{-2} De-rate 0.5 OOM for unpleasant task. = $10^{-1.5}$	**$10^{-0.5}$** Assume: 30% chance, probability ~ $10^{-0.5}$ Depends on tools 'Luck factor'	**$10^{-7.5}$**	1 tank-entry/year x $10^{-7.5}$ events/tank-entry **= $10^{-7.5}$ events/ year**
B. Smelter 1, gas cleaning section of acid plant	3 tanks x 1 entry/year	**10^{-3}** Excessive As enters 30% of time ($10^{-0.5}$), daily test	**10^{-1}** True, but no active control, 10% failure	**10^{0}** As will dissolve in weak acid, 100% failure	**10^{0}** No	**10^{0}** No procedure	**10^{0}** All top entry, no ventilation, no checking	**10^{-1}** True, but does not recognize arsine risk	**$10^{-0.5}$** 30% luck factor	**$10^{-5.5}$**	3 tanks x 1 entry/year $10^{-5.5}$ events/tank-entry **~ 10^{-5}**
C. Smelter 2, gas cleaning section of acid plant	3 tanks x 1 entry/year	**10^{0}** Excessive As enters 100% of time, no controls	**10^{0}** No controls	**10^{0}** As will dissolve in weak acid, 100% failure	**10^{0}** Yes, but too infrequent	**10^{-2}** Site data shows 1% procedure failure rate	**$10^{-2.5}$** All side entry, mech. ventilation, checking	**$10^{-1.5}$** Site data shows 1% failure rate + de-rate	**$10^{-0.5}$** 30% luck factor	**$10^{-6.5}$**	3 tanks x 1 entry/year x $10^{-6.5}$ events/tank-entry **~ 10^{-6}**
D. Smelter 3, gas cleaning section of acid plant	10 tanks x 1 entry/year	**10^{0}** Excessive As enters 100% of time, no controls	**10^{0}** No controls	**10^{0}** As will dissolve in weak acid, 100% failure	**10^{0}** Yes, but is above the limit most of the time	**10^{0}** No procedure	**10^{-1}** All side entry, no mech. ventilation, no checking	**10^{-1}** Site data shows 3% failure rate + de-rate	**$10^{-0.5}$** 30% luck factor	**$10^{-2.5}$**	10 tanks x 1 entry/year x $10^{-2.5}$ events/tank-entry **~ $10^{-1.5}$**

Figure 1. Worker exposure to arsine gas in a cleaned storage tank – Three asset base case comparison

MULTI-ASSET COMPARISON

Figure 1 illustrates the proposed visual framework with potential enabling conditions and IPLs shown on the top row with interim events in between. The initiating event frequency (tank-entries/year) is shown on the left side, with the probability of failure for defined step shown to the right. The final estimated consequence frequency (worker exposures/year) is shown at the far right. Going down the rows four cases are shown: Row A - Default starting case, Row B – Smelter 1 consensus base case, Row C – Smelter 2 consensus base case, and Row D – Smelter 3 consensus base case. In certain cases, some of the listed potential enabling conditions and IPLs may be unknown to the workshop participants and present an immediate opportunity for risk mitigation at little to no cost.

In Row A the default starting case is populated with generic data and assumptions. For example, the probability that workers fail to adhere to a procedure forbidding metal tools is assigned a failure rate of 1% from generic LOPA data (this default case could be updated with less generic data by a company that has, or over time develops, its own internal or industry dataset).

In Row B – Smelter 1 consensus base case, the workshop determines lot assaying of arsenic feed is performed daily and that site data shows that on average 30% of feed deliveries contain excessive arsenic. They have basic control on the turnover of liquid in the tank. From site data there are 3 tanks and 1 entry per year per tank. There is no procedure to forbid workers from bringing in active metals and no ventilation system. A basic procedure exists to clean out the tank upon worker entry and a 30% chance of dropping a tool or having active contact the tank liquid is estimated across all sites. The resulting indicative frequency of worker exposures per year is 10^{-5}, or one event per 100,000 years, or 0.001% chance of occurrence per year. There no convincing way to interpret if this frequency is reasonable on an absolute basis due to lack of recorded incidents at the site or sites with similar controls. However, an insight from Row B is that daily sampling provides the highest amount of protection and at relatively low cost compared to potential mechanical interventions. The calculated risk could be further be reduced by implementing a procedure to forbid active metals from entering the tank, which would involve no capital costs and minimal changes to operating practices.

In Row C for the Smelter 2 consensus base case, there are no controls on arsenic since the site expects excess arsenic in the feed at all times, however, there are strict controls on active metals in the tank and strict adherence to tank cleaning procedures, with a separate dedicated tank cleaning team highly aware of arsine risks. The procedural success is documented by site logs over the last three years. In addition, there is a regularly maintained and tested tank ventilation system. From site data there are 3 tanks and 1 entry per year per tank. The resulting indicative frequency of worker exposures per year is 10^{-6}, or one event per roughly 1,000,000 years. As with Smelter 1, there no convincing way to interpret if this frequency is reasonable on an absolute basis due to lack of recorded incidents at the site. However, it could be argued the site is one order of magnitude safer than Smelter 1 if there is confidence in the precision of the assumptions applied.

In Row D for Smelter 3, there are no controls on arsenic and the amount is known to be over a safe limit for arsine generation. There 10 tanks and 1 entry per year, which is a ½ order of magnitude higher than the other sites in terms of initiating event frequency. The tanks are side entry, which provide some protection against heavier-than-air arsine gas accumulating in the tank, and the procedural controls on tank cleaning are deemed satisfactory on site, but random inspection for general cleaning procedures over the last five years suggests poor cleaning in 3% of cases. The resulting indicative frequency of worker exposures per year is $10^{-1.5}$, or one event per roughly 30 years, which equates to a 33% chance of occurrence over the next ten years if the absolute value is taken at face value.

Even if the absolute frequency for Smelter 3 is deemed too aggressive, there is a clear set of basic deficiencies relative to Smelter 1 and Smelter 2 which appear qualitatively reasonable. This suggests that Smelter 3 has the highest risk of exposing workers to arsine relative to the other sites and any planning efforts should accelerate improvements at Smelter 3 to at least bring the relative risk in line with the others. In parallel, the additional controls with relatively low implementation cost should be applied to all sites once workshop participants agree the methods would be effective and practical. Finally, where there are critical unanswered questions a plan should be enacted to collect the data that would allow for an absolute estimate of consequence likelihood so that it can be compared to an absolute corporate risk tolerance.

SINGLE-ASSET REVIEW

Figure 2 demonstrates the framework focused on a single asset, in this case Smelter 3, following the site comparison. Row A shows the default case as before, while Row B shows the consensus base case discussed above. Row C shows the ‘consensus base case with proposed mitigations’, while Row D leaves a blank for annual verification of the implementation of the proposed mitigations.

In Row C the workshop group estimates that a procedural control to prevent active metal from entering the tank could be applied with at least the same success rate as the tank cleaning procedure. In addition, the Smelter 3 site personnel agree to ventilate their tank and implement regular air quality testing during tank servicing. Ventilation will require some capital spending, which will need to be prioritized during the next budget cycle, and could include forced air respirators. The group will reconvene in one year to re-assess the risk following mitigations and populate Row D.

It is worthwhile to note the application of the half-order of magnitude principle throughout, and the resultant implied precision of the estimates. It is not important to estimate whether poor cleaning occurs in 2, 3, or 4% of the cases, or whether the chance of dropping a tool is really 30% versus 25% or even 55%. As a result, there is also no expectation that the overall estimate is in any way more precise than half or even one order of magnitude. However, the output is still a grounded estimate useful for clarifying the relative level of risk across assets, and for documenting the importance of certain fairly easy to implement additional barriers (and the irrelevance of potential others).

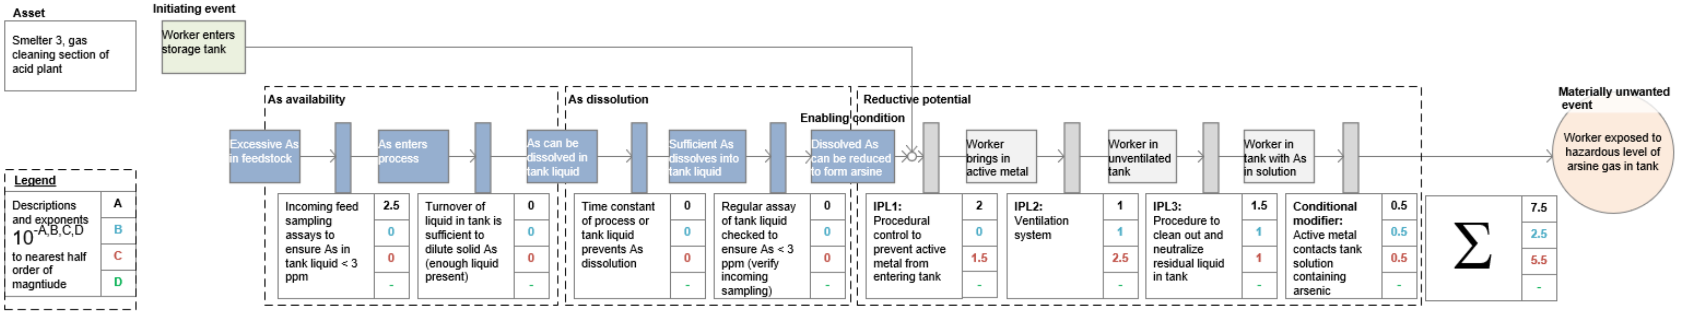

	Initiating event	Enabling condition probability build-up				IPL probability of failure on demand (PFD)				Materially unwanted event	
Description	**Initiating event frequency (IEF)** (tank-entries/year)	Human or mechanical/ automated sampling	A function of how much As present, how much deports to product, rate at which it is blowdown from process	Kinetics, solid phase transition to dissolved (time), not all solid arsenic dissolves	Human or mechanical/ automated sampling	Human procedure; active metal can include tools, ladder, boots with metal dust	Tanks can be side entry or top entry, atmosphere tested before entry, mechanical ventilation	Human procedure	Worker drops a tool, places a ladder, walks on surface with arsenic bearing solution etc.	**Total indicative likelihood** (events/tank-entry)	**Frequency of mitigated consequence** (events/year)
A. Default starting data	1 tank x 1 entry/year	**$10^{-2.5}$** Assume excess As always present and feed sampling each day, 1 failure/year = 1/365 days ~ 0.3% of year Arsenic enters process ~ $10^{-2.5}$ Sampling failure rate from generic data: Human error during task performed greater than once per week: 1 failure/ year	**10^{0}** Process specific Placeholder assumption 100% probability = 10^{0}	**10^{0}** Process specific Placeholder assumption 100% probability = 10^{0}	**10^{0}** Site specific Consider sampling frequency relative to incoming feed batching, other process factors, and earlier sampling etc. Placeholder assumption 100% probability = 10^{0}	**10^{-2}** Generic data: Human error during task performed less than once per month: 0.01 failures/year 1 tank-entry/year therefore failure rate = 0.01/1 = 10^{-2}	**10^{-1}** Mechanical vent = 1 *<u>PFD truth table</u>* *Sum:* *Top entry = 0* *Bottom entry = 1* *Not testing = 0* *Test before enter = 1* *No ventilation = 0* *Mechanical vent = 1* *Allow max total 2.5*	**$10^{-1.5}$** Generic data: Human error during task performed less than once per month: 0.01 failures/year 1 tank-entry/year therefore failure rate = 0.01/1 = 10^{-2} De-rate 0.5 OOM for unpleasant task. = $10^{-1.5}$	**$10^{-0.5}$** Assume: 30% chance, probability ~ $10^{-0.5}$ Depends on tools 'Luck factor'	**$10^{-7.5}$**	1 tank-entry/year x $10^{-7.5}$ events/tank-entry **= $10^{-7.5}$ events/ year**
B. Consensus base case	10 tanks x 1 entry/year	**10^{0}** Excessive As enters 100% of time, no controls	**10^{0}** No controls	**10^{0}** As will dissolve in weak acid, 100% failure	**10^{0}** Yes, but is above the limit most of the time	**10^{0}** No procedure	**10^{-1}** All side entry, no mech. ventilation, no checking	**10^{-1}** Site data shows 3% failure rate + de-rate	**$10^{-0.5}$** 30% luck factor	**$10^{-2.5}$**	10 tanks x 1 entry/year x $10^{-2.5}$ events/tank-entry **~ $10^{-1.5}$**
C. Consensus base case with proposed mitigations	10 tanks x 1 entry/year	**10^{0}** Excessive As enters 100% of time, no controls	**10^{0}** No controls	**10^{0}** As will dissolve in weak acid, 100% failure	**10^{0}** Yes, but is above the limit most of the time	**$10^{-1.5}$** Implement procedure Site data 3% failure	**$10^{-2.5}$** All side entry, add mech. vent, checking	**10^{-1}** Site data shows 3% failure rate + de-rate	**$10^{-0.5}$** 30% luck factor	**$10^{-5.5}$**	10 tanks x 1 entry/year x $10^{-5.5}$ events/tank-entry **~ $10^{-4.5}$**
D. One-year verification follow-up	TBD	TBD	TBD	TBD	TBD	TBD	TBD	TBD	TBD	TBD	TBD

Figure 2. Worker exposure to arsine gas in a cleaned storage tank – Single asset base case with proposed mitigations and follow up

CONCLUSIONS

The proposed risk reduction decision framework is an adapted layer of protection analysis (LOPA) and risk bowtie applied using a modified Delphi method and half order of magnitude heuristics in a workshop setting. The framework can be used at individual sites on individual risks where good data and a well understood failure mode supports an absolute estimate of consequence frequency. It allows one to identify material weak spots in barriers in place and build consensus around it. The framework aims for transparency to show where subjectivity is being applied and its impact. Value can be clarified and documented through event frequency reduction of specific mitigation actions. The framework is most useful when comparing similar risk scenarios. Reliance on the same assumptions allows for productive relative analysis and proposal of mitigation actions to reduce absolute risk, without necessarily being able to estimate the absolute reduction within a target level of precision. This approach allows for cost-benefit trade-offs to avoid overinvestment and identification of problem risk 'hotspots' among several assets. The framework is expected to be applicable in the metallurgical setting, as well as more broadly for catastrophic operational risks in any industry sector where process safety is paramount and risk aggregation and prioritization across sites and assets is relevant for mitigation decisions.

REFERENCES

Center for Chemical Process Safety (2001). *Layer of Protection Analysis: Simplified Process Risk Assessment*. Wiley.

Center for Chemical Process Safety (2015). *Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis*. Wiley.

Freeman, F., and Pergler, M., (2008). *Probabilistic Modeling as an Exploratory Decision-Making Tool*. McKinsey Working Papers on Risk #6.

Kahneman, D. (2002). *Maps of Bounded Rationality: A Perspective on Intuitive Judgment and Choice*. Nobel Prize Lecture. Retrieved from https://www.nobelprize.org/uploads/2018/06/kahnemann-lecture.pdf

Mitsui, A., Koehler, T., and Krysa, B., (2019). *The Utility of the ISO 45001:2018 Framework to Generate Risk Profiles and Mitigation Strategies for the Worker/Asrine* [sic] *Gas Exposure Risk*. Proceedings of the 58th Conference of Metallurgists Hosting the International Copper Conference.

Pergler, M (2020). *Be Your Own Risk Manager*. Unpublished, in progress.

UK Health and Safety Executive (HSE) (2009). *A review of Layers of Protection Analysis (LOPA) analyses of overfill of fuel storage tanks, report RR716*. Retrieved from https://www.hse.gov.uk/research/rrpdf/rr716.pdf 2019