Unexampled Events, Resilience, and PRA

Steve Epstein
ABS Consulting, Tokyo, Japan
sepstein@absconsulting.com

Abstract

"Why isn't it loaded? Are you afraid of shooting yourself?"
"Of course not. These weapons don't go off accidentally. You have to do five things in a row before they'll fire, and an accident can seldom count higher than three ... which is a mystery of probability that my intuition tells me is rooted at the very base of physics. No, it's never loaded because I am a pacifist."
-- Field Marshall Strassnitzky of the First Hussars of the Belvedere during WW I [1]

1 INTRODUCTION

I hope that my contribution to this symposium will be read as an essay (literally "an attempt"), a literary genre defined as the presentation of general material from a personal and opinionated point of view; my contribution should not be read as an academic or technical piece of writing, well balanced and well referenced. Instead please consider this as an attempt to clarify, mostly for my self, some general ideas and observations about unexampled events, resilience, and PRA.

I would like to discuss three topics:

1. Some ideas about rare, unfortunate juxtapositions of events, those called unexampled events, leading to accidents in well-tested, well-analyzed man/machine systems;
2. Some ideas about resilience and the relationship between unexampled events and resilience;
3. Safety and PRA: focusing not on the numbers, but on the act of doing the PRA itself, several times, to increase the ability of an individual or organization to respond to unexampled events resiliently.

2 UNEXAMPLED EVENTS

From the PRA point of view, there are three senses of unexampled events. The first sense is that of an extraordinary, never before thought of, challenge to the normal, daily flow of a system or an organization, such as three hijacked airplanes concurrently flown into well-known American buildings. The second sense is a juxtaposition of seemingly disparate events, practices, and influences, usually over time, which as if from nowhere suddenly create a startling state, as the operators at Three Mile Island discovered. The third sense of an unexampled event is one whose probability is so small that it warrants little attention, even if the consequences are severe, because though one can postulate such an event, there is no example, or evidence, of such an event ever occurring.

From my point of view, I believe that inattention to all three types of unexampled events that can lead to severe consequences poses a grave danger to the safety of vigilantly maintained, well-tested systems and organizations. Of special concern to me is the question if PRA can help an organization respond to significant unwanted, unexampled events.

I would like to discuss my notions of unexampled events, not solely as an intellectual exercise (which does have its own beauty and purpose), but also as the first steps towards (1) looking at safety in a different way, and (2) defining what it means to be resilient to unexpected impacts on safety.

In the spirit of Kaplan [2], an event can be defined by three attributes: a scenario, a likelihood function, and a consequence. Mathematically, risk analysts express this as $\boldsymbol{e = <s, l, c>}$. The letter ***s*** is a description of the event, the scenario; ***l*** is the likelihood of the event actually occurring, perhaps some measure like the odds a bookmaker gives; and ***c*** are the consequences of the event, sometimes a measurable entity like money or deaths/year, but often a list or description.

As an example, let ***e*** be an event that entails the release of toxic chemicals into the environment. So $\boldsymbol{e = <s, l, c>}$, where ***s*** is ***"the canister of toxic fluid was dropped from 2 meters"***, ***l*** is the judgment ***"not very likely"***, and ***c*** is the list ***"drum breaks, release of chemical on floor, cleanup necessary, no deaths"***.

It is easy then to imagine a set of events big $\boldsymbol{E = \{e_i\}}$, where big ***E*** might be defined as ***"the set of all events where canisters of toxic fluid are dropped"***. Some of the little $\boldsymbol{e_i}$ in big ***E*** are of special interest to the PRA analyst: the events where canisters leak, where workers are injured, where toxic fumes get into the ventilation system. To practice our art, or deception, we try to enumerate all of the events we can think of which lead to the consequences of interest, somehow give odds for the occurrence of each event, and present the results in such a way that decisions can be made so as to prevent, or lessen the impact of, the unwanted consequences.

Continuing the example of toxic chemicals, suppose by some black art we can assemble all of the events which lead to deaths by inadvertent chemical release, measure the likelihood of each event as a probability between 0 and 1, and then finally assign a number of deaths that could result from each event. We could then plot the results, such as the curve idealized in Figure 1 (in most situations the more severe consequences actually do have a lower probability of occurrence than less severe, but the curve is usually more jagged).

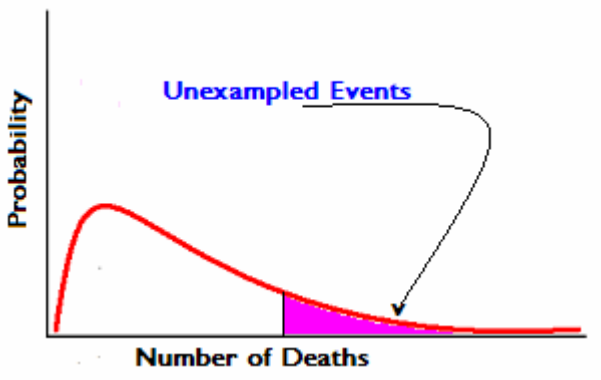

Figure 1

Each point on the curve represents an event in terms of number of deaths and probability (there may be many events with the same probability, of course, and in that case methods exist for combining their probabilities in a suitable manner). The extreme right part of the graph, the tail, is the home of unexampled events, what are called the outliers, the low probability/high consequence events. There are also unexampled of events in the tail on the left part of the graph (not a big tail here), and are usually of less interest because of the less severe consequences.

Of course, we have made one questionable assumption: we assumed that we have assembled all of the possible events that can lead to death by accidental release of toxic chemicals. Obviously, there will always be events not imagined and juxtapositions of circumstances not considered. But let us assume, for the moment, that in vigilantly maintained, well-tested systems, projects and organizations, these unconsidered unexampled events are of low probability. Later, I will add force to this assumption.

A rigorous definition would here give a method to locate the point on the graph where an event becomes unexampled. But this is not easily done, nor perhaps can be done. An unexampled event is a normative notion, depending on cultural influences, personal history, and the events under scrutiny: is 1 out-of 1000000 a limit for an unexampled event? However, if we look at how PRA characterizes risk, we can find an interesting connection with risk and the point where an event becomes an unexampled event.

In general, a PRA gives the odds for unwanted events occurring: the risk. In general, organizations that use the results of PRA (regulatory agencies, governments, insurance corporations) decide where they will place their marker on the graph and say, “Here is where the number of deaths is tolerable since the probability of the event occurring is 1 out-of 1000000; I’ll bet the event won’t happen.”

By the act of picking a point, one coordinate on the probability axis, one coordinate on the consequence axis, a decision maker makes an operational, normative definition of accepted, unwanted events; she decides what risks are acceptable in this situation. One can see that the shaded area of unexampled events in Figure 1 and the shaded area of unwanted, but accepted, events in Figure 2 bear a strong resemblance.

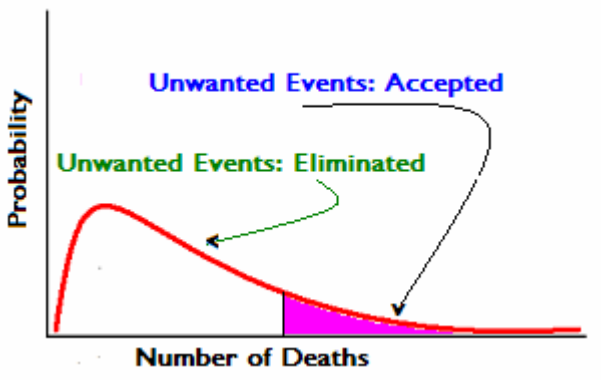

Figure 2

Clearly, PRA bets that decreasing focus on low probability/high consequence events, will not impact total safety of the situation.

Imagine a complex, dangerous situation, such as the excavation and disposal of 500,000 chemical weapons. Assume that in the design and operations of this system there is a very high degree of reliability of equipment, that workers and managers are vigilant in their testing, observations, procedures, training, and operations so as to eliminate the unwanted events in the white area of the curve in Figure 2. Given that an accident does occur in this situation, will the superior vigilance and performances postulated lower the probability that it the accident is severe? Surprisingly, at least to me 15 years ago, the answer is "No."

In 1991 I was attempting a PRA of the software for the main engines of the NASA space shuttle. I was working with two more experienced colleagues, Marty Schuman and Herb Hecht. While planning tasks for the next day, Herb mentioned to me to pay close attention when doing a code walk-down on a software module which was seldom invoked during operations. The reason, he said, was " … infrequently executed code has a higher failure rate than frequently executed code."

I had suddenly been awakened from my dogmatic slumbers.

Herein I summarize Herb Hecht's ideas (first contained privately in "Rare Conditions – An Important Cause of Failures" , then in [3]) from my copy of the first paper, now covered with the coffee stains and tobacco ashes of time. Like many artful and calm insights, Herb's thesis is immediately grasped, perceived almost as a tautology; I had all of the same data available, but I had not seen the connection between them. Perhaps he was not the first to think of such an idea, however I believe that he was the first to see its implications for well-tested systems:

1. In well-tested systems, rarely executed code has a higher failure rate than frequently executed code;
2. consequences of rare event failures in well-tested systems are more severe than those of other failures;

3. given that there is a failure in a well-tested system, significantly more of the failures are caused by rare events;
4. inability to handle multiple rare conditions is a prominent cause of failure in well-tested systems.

In short, we have tested out all of the light stuff and what we are left with are rare accidents with severe consequences in any well-tested software system.

How does this apply to other well-tested, vigilantly maintained systems, with well trained staff and enlightened management, good operating procedures in place; do Herb Hecht's observations about software systems apply to a process plant or nuclear facility? I believe that they do.

Look at figure 2 again. Our hypothetical chemical weapon disposal facility has calculated the risk of the unwanted events, and assigned a point to represent the risks they are willing to accept, the magenta area. The white area represents the unwanted events that the facility wants to entirely eliminate. By exceptional planning, maintenance, reliability of equipment, human factors, training, and organizational development skills, the facility is successful. The known and the easy problems are vanquished. What is left are the events in the magenta area, the accepted risks, the unexampled, rare events. So if there is a failure, chances are the failure is an unexampled event.

Moreover, Herb Hecht's study makes the following observation: all of the software which failed from three rare events, also failed, perhaps less severely, from two rare events, and three-quarters of the software which failed from two rare events, also failed, perhaps less severely, from one rare event.

What this means at my postulated facility is that if unwanted events and their consequences are actively guarded against, and equipment is vigilantly maintained, barriers in place, and staff prepared to prevent these events, and if indeed symptoms of unwanted events begin to occur, then there is a good chance that if we are on a failure path, it is the start of a severe accident scenario, out there in the tail of Figure 1. Perhaps more failures will occur to compound the situation and form a scenario which may have never been thought of, or previously dismissed as being improbable, and there are no procedures, nor experience nor training to aid in recovery. Chances are that this is not a simple or known situation; the first rare event failure has a good probability of being a harbinger of a severe accident scenario.

3 RESILIENCE

I would like to step away from unexampled events for a moment and look at resilience, with an eye as to how it applies to the occurrence of an unexampled event with severe consequences.

Resilience can be defined as a technical term: the capability of a strained body to recover its size and shape after deformation caused especially by compressive stress. It can also be defined for general usage: an ability to recover from or adjust easily to misfortune or change. Both definitions imply a reaction, not an action, on the part of a material, individual, or, perhaps, an organization to an impact or stress. The technical definition also has an operational aspect: resilience of a material, its coefficient of restitution and spring like effect, can only be determined by experiment. Resilience is something that cannot be measured until after the fact of impact. Perhaps one could try to prepare for acting resiliently in a given situation (a proactive measure of presilience?). Predicting resilience maybe easier, but entails knowing the essential properties of materials, and *mutatis mutandis*, individuals or organizations, which make them resilient, and they must confirmable in principle.

In the hypothetical chemical weapons disposal facility, I have postulated extraordinary activities to eliminate unwanted events; good work rules and work habits have been codified, proper procedures have been installed, proper surveillance and technical oversights groups are in place. In short, the facility has institutionalized what were before successful reactions, resilient responses, to known accidents, as standard operating procedures.

This is no small or trivial accomplishment. For anyone who has experienced a facility like this, a large nuclear power plant or the bigger than life oil platforms in the North Sea, the attention and concern given to safety is impressive, which Figure 3 represents as the white part of the graph, where the need for resilient reactions has been transformed into the need to strictly follow standard operating procedures. The magenta area in the tail of the graph represents unexampled, unwanted events, where resilience to stress is unknown and untested.

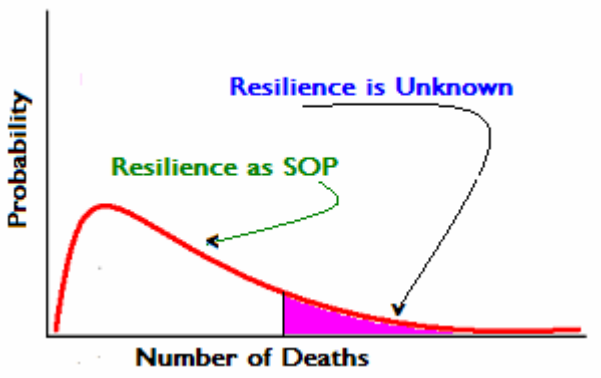

Figure 3

At this point, I wish that I were qualified to analyze the essential properties of resilience. This is probably the domain of psychology, a discipline in which I am formally unschooled. But I have had serious experience in situations where resilient reactions were needed during the 10 years I worked and lived in Israel. Without too much explanation, let me list some of the attributes which I believe are necessary, but not sufficient, properties of resilience:

1. **Experience** – nothing is second to experience with a system and adversity;
2. **Intuition** – intuition can give the best evidence for something being the case;
3. **Improvisation** – knowing when to play out of book;
4. **Expecting the unexpected** – not to be complacent
5. **Examine preconceptions** – assumptions are blinders;
6. **Thinking outside of the box** – look at it once, then look at it again;
7. **Taking advantage of luck** – when it happens, assent immediately.

This list is not exhaustive, but a witnessed list of underlying traits for successful resilient response to novel, critical situations, from removing a tractor from a sea of mud to acts of heroism. In all cases, my sense was that individuals or groups that showed the traits enumerated above, were much different than individuals or groups whose work entailed following strict protocols, procedures, and rules. In critical situations there were sometimes clashes of these two different cultures. Resilience in an individual meets with no internal resistance, however in a group, those that follow a rule, and those who improvise a tune, can find themselves at odds.

It is tempting to make comparisons between resilience and adaptation in the Darwinian sense. Natural selection, the *modus ponens* of evolutionary theory, makes the entailment of survival from adaptation: "natural selection is the claim that organisms enjoying differential reproductive success will be, on the average, those variants who are fortuitously better adapted to changing local environments, and that those variants will then pass their favored traits to offspring by inheritance" [4]. For me, the word "fortuitously" is key. An organism does not decide that the trait of wider stripes will be a better adaptation to a change in the environment, but by chance those organisms with wider stripes proliferate, as do the stripes.

Is resilience in response to an unexampled event a type of adaptation? I will carry the metaphor a bit longer to gain some insight. Unwanted, unexampled events are experienced as changes to the environment, albeit bad ones. The exact skills that may be necessary so as to rebound from the situation cannot be known ahead of time. No procedure is made for the unexampled by definition. However, if individual characteristics fortuitously exist that can aid in resilient response to critical and sudden unexampled events, I believe that severe consequences may be dampened and perhaps stopped. Can we plan what traits are needed ahead of time? Perhaps the list I presented above is a start to understand what underlies resilient response.

Does resilience apply to groups as well as individuals? Darwin believed quite strongly that natural selection applied only to individual organisms, not to groups, species, or clades. In 1982, Stephen J. Gould and Niles Eldredge proposed the theory of punctuated equilibrium to explain the long periods of no change in the fossil record of a population, then suddenly a flurry of speciation:

> "A new species can arise when a small segment of the ancestral population is isolated at the periphery of the ancestral range. Large, stable central populations

> exert a strong homogenizing influence. New and favorable mutations are diluted by the sheer bulk of the population through which they must spread. They may build slowly in frequency, but changing environments usually cancel their selective value long before they reach fixation. Thus, phyletic transformation in large populations should be very rare—as the fossil record proclaims. But small, peripherally isolated groups are cut off from their parental stock. They live as tiny populations in geographic corners of the ancestral range. Selective pressures are usually intense because peripheries mark the edge of ecological tolerance for ancestral forms. Favorable variations spread quickly. Small peripheral isolates are a laboratory of evolutionary change." [4]

To continue the Darwinian metaphor: we should expect the operations of a vigilantly maintained, well-tested, well surveilled system to proceed flat and normally, with no signs of change or need of resilience, most of the time, then suddenly punctuated by critical challenges. Under great environmental stress, such as an unexpected, unforeseen accident, we can expect only a small, isolated group to respond resiliently and to "speciate" from the larger group.

And as those challenges are met, or not met, then the standard operating procedures of the system are changed to incorporate the resilient reactions that mitigated the situation; and another period of stasis will be entered.

3 PRA

My focus in this essay has been on well-tested, well-analyzed, vigilantly maintained systems, unexampled events, and resilience. I have tried to show that (1) unexampled events have an increased probability of severe consequences in these systems, and (2) resilience to respond to unexampled events is a trait that may be antithetical to the mindset that must run these systems without incidents. I would now like to focus on the implications to PRA.

PRA is the discipline of trying to quantify, under uncertainty, the risk or safety of an enterprise. To briefly state my views:

Quantification, or measuring, the risk/safety of a situation is not the goal of a PRA. Nor is it necessary to "quantify" with numbers (one could use colors). The act of trying to measure the risk involved is the source of knowledge. The acts of trying to assign values, combining them, questioning their verisimilitude, building the model are the great treasure of PRA: the key to the treasure is the treasure itself.

Uncertainty is not some noisy variation around a mean value that represents the true situation. Variation itself is nature's only irreducible essence. Variation is the hard reality, not a set of imperfect measures for a central tendency. Means and medians are the abstractions.

Too often risk is defined as ***risk = likelihood * consequence*** and ***safety = 1-risk***. I disagree with this. Risk is likelihood **and** consequence, not a simple multiplication with safety as the additive inverse of risk. Risk and safety are normative notions, changing with situations and expectations, and must be assessed accordingly.

Modern PRA began after Three-Mile Island with the publication of WASH-1400, "(The) Reactor Safety Study", a report produced in 1975 for the USNRC by a committee of specialists under Professor Norman Rasmussen. It considered the course of events that might arise during a serious accident at a (then) large modern light water reactor, and estimated the radiological consequences of these events, and the probability of their occurrence, using a fault tree/event tree approach.

The proposed event trees and fault trees were very small with respect to the number of system and events they modeled. The mathematics was approximate and the data little more than reliability studies, the initiating events well known possible disturbances to normal plant operations. However, these easy methods gave operators and managers the first feel for the safety of the plant as a measurement, certainly one step in knowledge.

Times have changed, but the methods have not. Nuclear plant PRA models are orders of magnitude larger than envisioned by Rasmussen. The models are so large that they are neither reviewable nor surveyable. The results calculated are severe approximations with no knowledge of the error factors involved. Reliability techniques are used for human actions in questionable ways. The numerical results of the PRAs are highly suspect, and yet the desiderata of the activity.

The focus of these PRAs is almost entirely on known system disturbances as initiating events, and static, sequential views of accident emergence and progression. As a result, procedures, training, regulations, and methods of operation were put in place to guard and watch out for the known disturbances. Risk models were used not for their insights, but for the quantitative results offered, thus never exploring novel failure modes of the facilities, totally missing the ability to postulate unexampled events and strange system and extra-system influences/interactions/background.

The result is that the attention of the risk analysts is not on unexampled events. Given that symptoms of system failure occur, attention will not be on the tail of the distributions where unexampled events reside. There will be little experience in the organization for imagining scenarios that change critical assumptions, have slightly different symptoms, or include multiple failures. Moreover, the standard operational culture is focused on the procedures and rules for dealing with known disturbances and standard ways of solving problems. And rightly so, since without this focus on the checklists, procedures, and protocol controllable situations can easily escalate out of control, and the daily safety of the facility impacted.

A second culture is also needed. To restate a central theme in this essay, in well-tested, etc., systems, given that there is an accident, chances are the level of consequence is

high and that the causes had not been modeled in the PRA. The second culture, to be prepared for the unexampled event, must play with the model, question assumptions, run scenarios, and understand the uncertainty. When initial indications or symptoms that a system may be going astray, the second culture moves away from the probable and into the possible.

This can be visualized by using the typical Color Risk Matrix used by many, including me, to present risk analysis results. Here is an example of two 6x5 risk matrices:

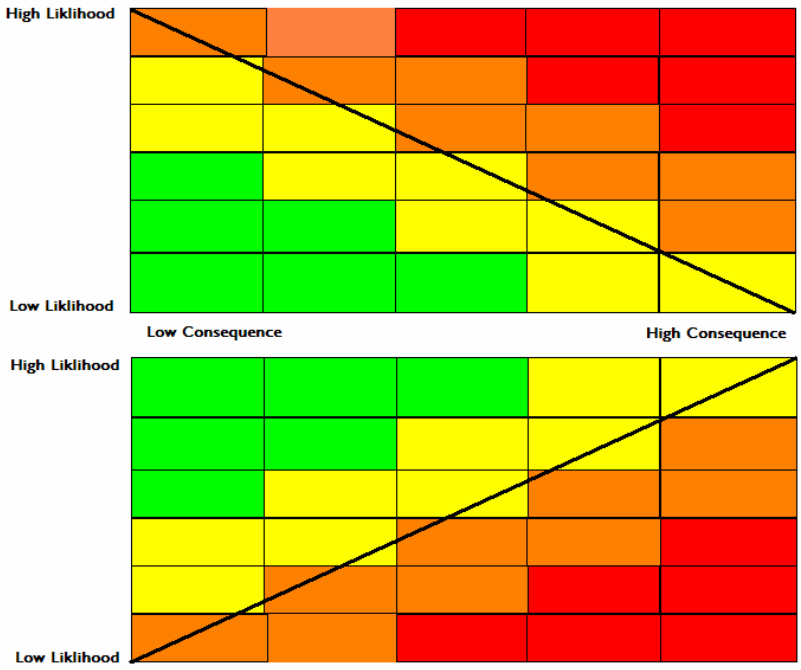

Figure 4

In this type of matrix, colors represent risk, with the order usually being like a traffic light: red, orange, yellow, and green (from high risk to low). The two dimensions represent consequence and likelihood as marked in Figure 4.

The upper matrix is the typical risk matrix for the standard operating culture, focusing on the area above the diagonal. The lower matrix is the typical risk matrix for the second culture, focusing on the area below the diagonal. Note how the two matrices are rotated.

Can these two cultures coexist? Can one of the cultures "proactively presilient"? I do not know the answers at all. But I do know, that without them both, we can be assured of accidents with higher levels of consequence than not.

Safety is connected not only to risk, but also to expectation. It is a normative notion. In operations like a nuclear power plant or a chemical weapons disposal facility, which are of the well-tested etc. category, I expect the rare events to be guarded against, also. I

weight consequence more heavily than likelihood to calculate safety in the well-tested etc.

> ... It's in words that the magic is--Abracadabra, Open Sesame, and the rest--but the magic words in one story aren't magical in the next. The real magic is to understand which words work, and when, and for what; the trick is to learn the trick.
> ... And those words are made from the letters of our alphabet: a couple-dozen squiggles we can draw with the pen. This is the key! And the treasure, too, if we can only get our hands on it! It's as if--as if the key to the treasure is the treasure!
>
> John Barth in Chimera

REFERENCES

[1] Helprin, Mark (1991). A Soldier of the Great War, Harcourt and Brace, pg. 546

[2] Garrick, B.J. & Kaplan, S. (1981). On the Quantitative Definition of Risk. *Risk Analysis Vol. 1, No. 1*.

[3] Crane, P. & Hecht, H. (1994). Rare Conditions and their Effect on Software Failures. *Proceedings Annual Reliability and Maintainability Symposium*.

[4] Gould, S.J (2004). The Structure of Evolutionary Theory, Harvard Press