•Quantifying the Frequency of Attack - A Fool's Errand?

N. Siu

The Problem

In a classical risk-informed decision making process addressing multiple threats to a variety of facilities, the frequencies of different attack scenarios are necessary technical inputs to the process. Without explicit provision of these frequencies, the decision makers must make their own informal and generally implicit assessments as to the absolute and relative likelihoods of the threats. However, the fundamental ability of a technical assessment team (even one with strong involvement from the intelligence community) to develop meaningful estimates (including uncertainty) of the attack scenario likelihoods is likely to be a source of significant controversy. This white paper summarizes what appear to me to be the key issues, discusses these issues, and provides recommendations for proceeding.

The Technical Issues and Discussion

I have heard a number of objections to the development of attack scenario frequencies. These objections, which are not necessarily consistent with each other, are both philosophical and practical. These objections and their apparent implications can be summarized as follows.[1]

(1) *Terrorist attacks are not "random events." Rather, they are the product of deliberate planning and action*. Thus, probabilistic models for attacks are not appropriate and the concept of an "attack scenario frequency," which is just a parameter in a particular probabilistic model (the Poisson model), is meaningless.

(2) *Terrorist attacks, as with all human actions, are not predictable*. Thus, the results of any predictive model, including probabilistic ones, are meaningless.

(3) *The available data are too limited to allow timely, meaningful estimates of attack probability model parameters*.

I will address each of these objections in turn. It should be noted that the subject of assessing terrorist attack likelihoods has been strenuously debated over many years. Many of the objectors are very knowledgeable or even expert in the subjects of risk assessment, risk management, and decision making under uncertainty, and their objections need to be taken seriously. The purpose of this paper is to do just that. I hope the discussion will support the development of a path forward for fully risk-informed decision making for NRC-regulated facilities.

[1]The logical implications of the objections are based on my understanding of the objections; any misinterpretations/mis-statements of these implications are my own.

Objection (1): Terrorist attacks are not random events...

From the way this objection is often stated, it appears that some objectors are making a number of implicit premises about random events and their role in probabilistic assessments. At the risk of oversimplification, I summarize them as follows:

(i) The concept of a random event is absolute, i.e., independent of the frame of reference of the observer. Consequently, events can be classified as being random or not in a manner independent of the purpose of the probabilistic assessment.

(ii) Because probabilistic assessments deal with random events,[2] they are not appropriate for events associated with well-defined causes.

Regarding the first premise, it can be argued that the concept of a "random event" is, for the purpose of probabilistic assessments, a relative one. (See Apostolakis [4]; Siu [5] provides a discussion in the context of a particular accident threat.) In particular, it is defined with respect to a specified "model of the world" (to use the terminology of Apostolakis [4]); variations in outcome that are not addressed by the causal mechanisms explicitly addressed in the model are treated as being random (or "aleatory"), regardless of their actual causes. What is addressed in the model, of course, depends a great deal upon the purpose of the assessment.

As an example, in typical probabilistic risk assessments (PRAs) for nuclear power plants (NPPs), the failure of a pump to start on demand is treated as being a random event. It is technically feasible to develop a PRA model that treats pump failure as the deterministic outcome of physical processes affecting the pump. In such an approach, randomness would be introduced when modeling these physical processes. To date, most NPP PRAs do not take this approach because such detail has been judged unnecessary for the decisions requiring PRA support.

I have, perhaps, stated the second premise and its logical implication too superficially. However, just to be clear on the subject, of course many probabilistic assessments (and all PRAs of engineered systems) do treat events that have directly attributable, if often multiple, causes. Probabilistic assessments are performed for scenarios involving such events because there are uncertainties associated with key aspects of the scenarios (e.g., the type and strength of factors affecting the occurrence of constituent events), and there is a need to quantify these uncertainties.

---

[2]According to Webster's [1], the adjective "random" means lacking or seeming to lack a regular plan, purpose, or pattern. Somewhat different flavors are provided by textbooks on probability theory. For example, according to Parzen [2], "A *random* (*or chance*) *phenomenon* is an empirical phenomenon characterized by the property that its observation under a given set of circumstances does not always lead to the same observed outcome (so there is no deterministic regularity) but rather to different outcomes in such a way that there is *statistical regularity*." De Finetti [3] takes a broader (and subjective) point of view on the meaning of random: "...it is simply that of 'not known' (for You), and consequently 'uncertain' (for You), but *well-determined* in itself."

In general, following de Finetti and other subjectivists, probabilistic assessments deal with situations where the truth of important propositions[3] is unknown. These propositions can involve events which may be completely non-deterministic from one perspective (e.g., the perspective of the lay public) but much more deterministic from a different perspective (e.g., the perspective of the attacker). As long as the propositions are philosophically meaningful, the notion of a probabilistic assessment is philosophically meaningful.

This is not to dismiss Objection (1) out of hand. Aside from the data concerns raised under Objection (3), there are challenging questions regarding the fundamental models used in many probabilistic assessments. For example:

(iii) Even accepting that terrorist attacks are uncertain events for which probabilistic assessments are philosophically meaningful, it is unlikely that such events will obey the assumptions underlying the Poisson model. Therefore, is the concept of an "attack scenario frequency" appropriate?

The Poisson model, which provides a means for calculating the probability of observing a certain number of events over a specified period of time,[4] is based on a set of five assumptions (e.g., see Parzen [2]). For modeling purposes, the two key ones are:

- the events are independent (i.e., the observation of an event in one time interval does not affect the probability of occurrence of an event in a subsequent time interval), and

- the event generating process is stationary (i.e., the probability of an event occurring in a time interval of fixed duration Δt does not change over time).

It can be seen that, for terrorist attacks, these assumptions are probably incorrect. The important questions that need to be addressed are then:

(iv) When (or under what conditions) is the Poisson model good enough?

(v) If the Poisson model is not good enough, what model(s) should be used?

---

[3]A proposition is a logical statement which can either be true or false. For example: "NPP X will be attacked tomorrow in manner Y." The statement "NPP X has significant vulnerabilities to terrorist threat Y" is not a proposition until the terms "significant" and "vulnerabilities" are clearly defined.

[4]According to the Poisson model

$$P\{r\ events\ in\ time\ T|\lambda\} = \frac{(\lambda T)^r}{r!}e^{-\lambda T}$$

where λ, the event occurrence frequency, is the single model parameter. It can be shown that the Poisson model (which deals with counts of events) is mathematically equivalent to the exponential model for event occurrence times:

$$P\{time\ of\ next\ occurrence\ \leq t\} = 1 - e^{-\lambda T}$$

These questions can be partially addressed by technical analysis, e.g., through discussion of the mathematical properties of the Poisson model and alternatives (e.g., non-homogeneous Poisson models) and examination of available intelligence data. However, the answers are also dependent on the notion as to what is "good enough," are therefore dependent on decision problem at hand and the characteristics of the decision maker, and are therefore not likely to be generic.

Without trying to develop a full answer in this paper, it is useful to make two observations.

First, the Poisson model does not assume or predict that events are evenly spaced in time. On the contrary, it predicts the occurrence of clusters of events, separated by relatively long gaps. (This follows from the fact that the event "inter-arrival times," i.e., occurrence times, are exponentially distributed - shorter occurrence times are more likely than longer ones.)

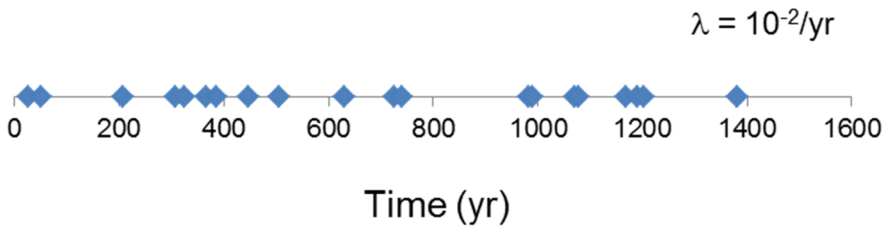

Figure 1. Schematic of the characteristic time signature of a Poisson process

Thus, for example, the simple observation of a flurry of events in a relatively short time interval does not, by itself, imply that a Poisson model for event occurrence is incorrect. One has to probe into the causes of the events to see if there are significant dependencies between events. And even if there are dependencies, the Poisson model may still provide numerical results that are deemed to be adequate for decision making purposes. For example, the model is the basis of the concept of return periods, which is used in the management of risks associated with such natural hazards as earthquakes and floods, despite recognized sources of dependencies between events.

Second, as more information becomes available, improved probabilistic models that do not derive from the Poisson model can certainly be developed. Knowledge of a crack in a pump rotor can be used to develop a more mechanistic, predictive model of pump performance on the next demand. Knowledge of ground saturation and reservoir levels following a major rainstorm can be used to develop a more mechanistic, predictive model of flooding following an approaching storm. However, in long-term planning situations, such detailed information and, therefore, a basis for causal modeling, is lacking. It seems that, in such situations, there is a place for such non-mechanistic models as the Poisson.

Objection (2): Terrorist attacks are not predictable...

The superficial response to this objection is that, while acknowledging that deterministic predictions of attacks are not generally possible (certainly not with information we could expect to have access to), probabilistic predictions, i.e., identified scenarios with associated likelihoods, are certainly possible and meaningful. Moreover, decisions based upon such predictions (perhaps more often implicit than explicit) are being made all the time. This response is a direct consequence of the subjectivist definition of probability [3,4], in which probability is a measure of an individual's (e.g., a decision maker's) degree of belief in the truth of a proposition, based upon the information the individual has available at the time.[5]

A more difficult, follow-on question is:

(vi) To what extent should such important decisions as the ones we (as an agency or as a nation) are faced with rely upon personal probabilities?

This is, of course, a policy question that needs to be answered by policy makers. However, it is useful to acknowledge that there are a variety of approaches to decision making under uncertainty that make more or less explicit use of personal probabilities. Classical multi-attribute utility theory, which relies upon an explicit statement of decision maker probabilities and preferences (as measured by "utilities"), is only one approach. Two alternatives are the min-max approach, and the game theoretic approach.

In the min-max approach, the decision maker selects the decision option that has the smallest possible negative consequence. For example, in the highly simplified case of a choice between the two options shown in Figure 2, the decision maker will choose Option 2, even if p_1 is close to unity.

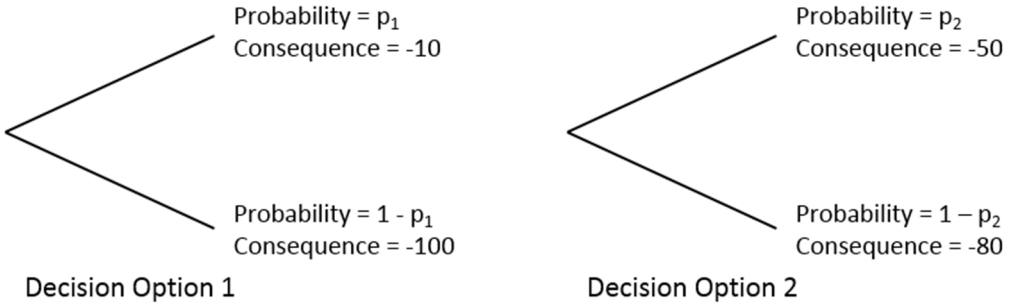

Figure 2. Example Decision Options

[5]In the subjectivist school, probabilities must conform to certain rules (the calculus of probabilities) to be "coherent." These rules ensure that, in the face of strong data, coherent probabilities will be consistent with the data. However, in situations where the data are weak or their relevance is uncertain, there is no external standard for the "correctness" of a probability.

It can be seen that personal probabilities (of the decision maker, or of the analysts providing information to the decision maker) will enter when defining the scenarios for which consequences will be assessed (e.g., it will likely be decided that some scenarios are "credible" and others aren't), but do not explicitly enter the formal decision making process.

From a safety philosophy perspective, a problem with the min-max approach is that it runs counter to the advances we've made in safety assessment over the last 25 years; it requires designating scenarios as being "credible" or "incredible," and then treats all credible scenarios essentially as being equally likely. From an implementation perspective however, the decision maker is not required to explicitly assess probabilities nor to explicitly express personal or corporate values in weighting consequences (e.g., to numerically specify how risk-aversion concerns should be dealt with, if at all).

In the game-theoretic approach, the decision problem is put in the context of a conflict between adversaries. It is explicitly recognized that the attacker can change objectives and plans based upon the actions of the defender, that the defender can, in turn, respond to the attacker, and so forth. The situation gets more complicated when one or both sides has imperfect information about the other. The question for both sides is how to optimize the use of resources.

In one consulting company's approach aimed at insurance applications [6], the relative probabilities of various attack modes at different target locations are assessed and used, in combination with the probability of successful defense and the consequences of a successful attack, to estimate the relative risk for each target modeled. The attack mode probabilities are based upon considerations of a number of factors believed to be important to the attackers:

- weapon prioritization (including considerations of reliability, effectiveness, and detectability);
- target attractiveness (including considerations of economic and human loss consequences, symbolic value, and name recognition); and
- cost effectiveness (including considerations of target hardening and uncertainties in the degree of hardening)

The attack mode probabilities also take into consideration recognized tactics adopted by the attackers (e.g., randomized target selection).

With estimates of relative risk, defensive strategies can be developed to optimize the use of defensive resources. Although not explicitly discussed in Ref. 6, presumably the approach can be applied in long-term planning (e.g., in searching for solutions that cover a broad range of planning scenarios), or in short-term response (e.g., in redistributing resources to react to specific, new intelligence).

The game-theoretic approach described above is quite similar to the classical decision approach supported by current PRAs.[6,7] Although it acknowledges a malicious, intelligent adversary, it still relies upon assessments of probability and preference. Where it differs from current PRAs is that it does not require an assessment of the absolute probabilities of attack. This is because the approach is aimed at an optimal distribution of resources among facilities subject to a single class of threat - a terrorist attack. It is not aimed at distributing resources among facilities in a manner that is optimal from both safety and security considerations, nor is it aimed at answering the question as to whether anything needs to be done at all.

It is worth noting that Ref. 6's approach has apparently been applied to a wide variety of diverse facilities.

[6]This isn't surprising, as the apparent principal behind the technical approach described in Ref. 6 has a strong background in PRA.

[7]It is not clear from the description provided in Ref. 6 if the game-theoretic framework has been fully implemented. For example, it is not clear if attacker and defender moves and countermoves are explicitly modeled as a part of the search for a (time-dependent) optimum. Such an approach might be necessary if the sequencing of moves is important.

Objection (3): Data are too limited to allow timely, meaningful probabilities...

There appear to be two concerns underlying this objection:

(a) the frequent lack of specific, accurate information regarding adversary objectives, resources, plans, and knowledge/skills/abilities; and

(b) the compartmentalization of what information that is available resulting from the application of the "need to know" principle.

Regarding the first concern, I have no doubt that global attack probabilities are being generated (in some fashion) and are being used. After all, there surely is some basis for the threat condition levels announced by Homeland Security. Furthermore, activities are underway to develop more localized attack probabilities (e.g., in terms of relative attractiveness of different facilities) for use in actual decision making. (Ref. 6 provides only one example; similar, perhaps less sophisticated, approaches are being developed for dams [7], and perhaps other facilities. Note that I do not have independent confirmation of the extent to which the approaches described in Ref. 6 and Ref. 7 are actually used.)

There is always, of course, the possibility that the assessments will be wrong (e.g., due to an incorrect understanding of attacker intentions), and that their results could prompt wrong decisions. And, as mentioned in the previous section, there are alternative decision making processes that rely less heavily on assessments of attack probability. However, it is hard to envision a practical decision making approach that would completely ignore considerations of attack likelihood,[8] and such considerations should clearly be influenced by what information is available to the intelligence community.

Regarding the second concern, there may very well be insufficient information within NRC to assess the likelihood of attacks on each of the facilities/processes it regulates. (Although expert panels can be convened, there is always the question of the panel makeup, and there is the additional question as to how the decision maker(s) will use the results of the panel.) This does not obviate the need to assess the likelihood. Rather, it identifies a need for support from appropriate federal agencies, and the need to develop a process to gain and use that support. Such a process might, for example, involve the development of attack probabilities by other federal agencies (with appropriate input from NRC), and the provision of these probabilities to NRC to support NRC decision making.

[8]I am assuming that, in a comprehensive homeland security strategy, some differentiation of potential targets is necessary.

Concluding Remarks and Recommendations

I believe that the assessment of attack probabilities is both a meaningful and useful exercise, and should be pursued. In examining the objections to such an assessment, I have found no compelling philosophical or practical arguments to support a long-term course that would avoid assessing probabilities.

In the near term, the structured development of conditional probabilities for use in relative assessments appears to be both technically feasible and accepted in a number of applications. Such an approach can cover a broad range of facility types, and may very well be completely sufficient for many of NRC's decision making needs. The key potential barrier to success is organizational: processes should be established to allow the development of probabilities that are based upon the best information available to the U.S. intelligence community. It is conceivable that these probabilities could be developed by other federal agencies and provided to NRC on an as-needed basis. Note that the probabilities, and the decision making framework employing these probabilities, may need to be dynamic (to reflect, in a timely fashion, changes in intelligence assessments).

In the longer term, the development of absolute probabilities (which will allow balancing of accident and security concerns, as well as decision making as to whether facilities are safe enough) appears to be technically feasible. Some technical work should be done to address potential barriers. In particular, the adequacy of the Poisson model for attack occurrence should be assessed. If the model is found to be inadequate for decision making purposes, then an alternate model should be developed. Processes for development of the actual attack probabilities should be developed in a manner similar to that used for the development of relative attack probabilities.

References

1. *Webster's Third New International Dictionary of the English Language, Unabridged*, P.B. Gove, ed., G. & C. Merriam Company, Springfield, MA, 1969.

2. E. Parzen, *Modern Probability Theory and Its Applications*, Wiley, New York, 1960.

3. B. de Finetti, *Theory of Probability: A Critical Introductory Treatment*, Wiley, New York, 1974.

4. G. Apostolakis, "The concept of probability in safety assessments of technological systems," *Science*, 250, 1359-1364(1990).

5. N. Siu, "Uncertainty analysis and pressurized thermal shock: an opinion," white paper, U.S. Nuclear Regulatory Commission, September 3, 1999, ADAMS Accession No: ML992710064.

6. "Understanding and managing terrorism risk," Risk Management Solutions, Inc., www.rms.com, 2002.

7. M. Chavira, "United States Bureau of Reclamation Security Risk Analysis (RSRA)," *Proceedings of PSAM 6, International Conference on Probabilistic Safety Assessment and Management*, San Juan, Puerto Rico, June 23-28, 2002.