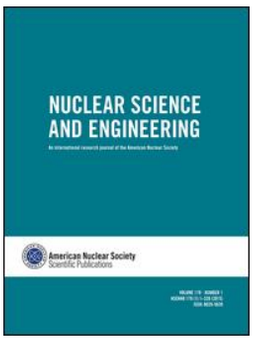

Nuclear Science and Engineering

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/unse20

Possible Does Not Mean Useful: The Role of Probability of Attack in Security Risk Management

Gregory D. Wyss & Adam D. Williams

To cite this article: Gregory D. Wyss & Adam D. Williams (2022): Possible Does Not Mean Useful: The Role of Probability of Attack in Security Risk Management, Nuclear Science and Engineering, DOI: 10.1080/00295639.2022.2129224

To link to this article: https://doi.org/10.1080/00295639.2022.2129224

NUCLEAR SCIENCE AND ENGINEERING

DOI: https://doi.org/10.1080/00295639.2022.2129224

Possible Does Not Mean Useful: The Role of Probability of Attack in Security Risk Management

Gregory D. Wyss* and Adam D. Williams

Sandia National Laboratories, Albuquerque, New Mexico 87185-0650

Received May 11, 2022
Accepted for Publication September 12, 2022

Abstract — *Providing adequate security to civilian nuclear materials and facilities exemplifies the long-standing, dynamic challenge of addressing the potential for facility damage under operational uncertainty. Estimating attack likelihood with enough precision to be useful and actionable for security risk management is philosophically, scientifically, and practically challenging. In response, this paper discusses the conceptual and analytical shortcomings of various approaches to calculating the likelihood of attack as a foundational element of security risk management. From these shortcomings emerge a set of characteristics that can guide the creation of alternative concepts that provide more robust and actionable security risk management approaches better aligned with the Ievolutionary growth in civilian nuclear facilities. Such broader conceptions would support movement from traditional interpretations of probability of attack toward more nuanced and complex depictions to enhance security risk management.*

Keywords — *Security, risk, uncertainty, risk management.*

Note — *Some figures may be in color only in the electronic version*

I. INTRODUCTION

Successfully operating nuclear facilities and managing nuclear materials is a difficult endeavor. Yet, the potential benefits—clean energy, baseload power, technical sophistication—incentivize the pursuit of nuclear energy–related projects in the face of multifaceted challenges. Decision makers and stakeholders who manage nuclear facilities are continuously required to balance many competing sources of potential damage and loss, such as safety, equipment reliability, product quality, and security.

Providing adequate security to civilian nuclear materials and facilities exemplifies the long-standing, dynamic challenge of addressing the potential for facility damage under operational uncertainty. Consider advances in adversary capabilities as an example of more traditionally familiar external drivers of security, while uncertainty or changes in regulatory standards or budget availability are examples of less familiar internal drivers of security uncertainty. The civilian nuclear sector's continued path toward advancements (e.g., small modular reactors) and evolution (e.g., advanced reactors) also introduce new drivers of security, including increased digitization and susceptibility to increased personnel radicalization.

Quantifying, analyzing, and designing security to reduce (or eliminate) the potential for damage or loss is a complex endeavor and requires addressing undesired events that have not occurred and outcomes of

*E-mail: gdwyss@sandia.gov

undetermined magnitude. Coincidentally, for security, the range of possible undesired events and associated outcomes are related to previously mentioned sources of uncertainty. Yet, because resources are limited, decision makers are often forced to prioritize which potential undesired events they should spend resources to prevent, necessarily leaving others unaddressed.

In response, the concept of risk is a popular approach to compare and prioritize among potential undesired events and to develop security solutions to address and mitigate the various drivers of security uncertainty. However, several important nuances are present in the concept of risk when it is applied to security, especially relating to estimating scenario likelihood in general, and the likelihood of attack in particular. This is a subject of intense debate in both the risk and nuclear security communities. In response, this paper examines the nuances of security risk estimation in order to identify how they may complicate proposed U.S. Nuclear Regulatory Commission (NRC) risk-informed security methods and requirements

II. CONTEXTUALIZING THE CONCEPT OF RISK IN NUCLEAR SECURITY

Given its ubiquity, there is a surprisingly lack of consensus on the appropriate application of this concept of risk. For example, the Society for Risk Analysis—the "multidisciplinary, interdisciplinary, scholarly, international society … in risk analysis"—does not have a single definition for risk. Rather, this body offers a list of seven descriptions in its official glossary,[1] all of which are qualitative in nature, including that risk is

1. the possibility of an unfortunate occurrence
2. the potential for realization of unwanted, negative consequences of an event
3. exposure to a proposition (e.g., the occurrence of a loss) of which one is uncertain
4. the consequences of the activity and associated uncertainties
5. uncertainty about and severity of the consequences of an activity with respect to something that human's value
6. the occurrences of some specified consequences of the activity and associated uncertainties
7. the deviation from a reference value and associated uncertainties.

Each of these definitions revolve around the concept of an unknown future state subject to a range of undesired outcomes, which is a fitting description for the nuclear security domain. Yet, the absence of a consensus framework for quantifying these unknown future states subject to a range of undesired outcomes has resulted in (at best) confused and (at worst) inaccurate risk mitigations, another fitting description of the nuclear security domain.

Historically, quantifying "risk metrics" is often traced back to Blaise Pascal's 1662 *Logica sive Ars Cogitandi*, in which he offers a characterization of how to judge future events, namely,

> in order to decide what we ought to do to obtain some good or avoid some harm, it is necessary to consider not only the good or harm in itself, but also the probability that it will or will not occur, and to view geometrically the proportion all these things have when taken together.[2]

Mathematically, "to view geometrically the proportion" commonly implies multiplication, so Pascal's ideas led to the popular, and long-standing, quantitative notion that the risk of an event is the product of the undesired event's probability and outcomes (more commonly called consequences), often expressed as

$$R_e = P_e \cdot C_e \,. \tag{1}$$

Pascal argues that one should have greater fear of events with a greater R_e, so a decision maker would naturally want to prioritize prevention for undesired events with a higher R_e. In terms of nuclear security, this matches the intuition to focus limited resources on mitigating the highest security risks. Though mathematically tractable and conceptually appealing, evolving thought in risk science has found that the equivalences produced by this simple but powerful formula do not always correspond to human understanding or observations, especially when comparing rare but extreme events with common but relatively benign events.[3]

Consider Fig. 1, for example. In this notional example, four different undesired events are quantified in boxes [A] through [D]. The horizontal category, labeled "frequency of occurrence," relates to the probability element discussed in the preceding paragraph. For a nuclear security example, consider the different detection abilities from a highly reliable versus a low-quality sensor. The vertical category, labeled "severity of consequence," relates to the consequence element discussed in the preceding paragraph. For a nuclear security example, consider the difference between an adversary accessing a gate on a perimeter fence versus the door to the reactor control room. While the results of boxes [A] and [D]

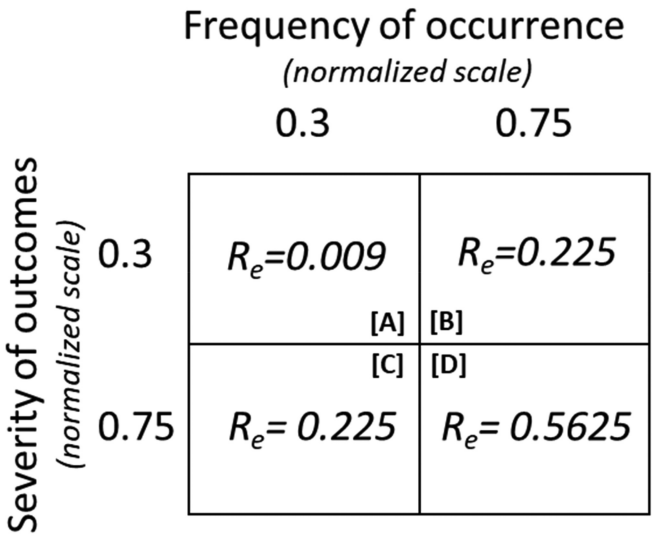

Fig. 1. Notional example highlighting false equivalences in popular approach to quantifying risk.

match intuition (namely, that resources should be prioritized to prevent [D] over [A]), boxes [B] and [C] present a conundrum. More specifically, if R_e of undesired event [B] is the equal to R_e of undesired event [C], then they are commonly assumed to be equally important. Yet, the mitigations for [B] that focus on addressing the frequency of the undesired event are not likely applicable to mitigations for [C] that focus on addressing the severity of the undesired event.

When considered in determining "total risk" across scenarios, such false equivalences (e.g., R_e for [B] and [C] in Fig. 1 being equivalent) can generate nonsensical results. Consider, for example, deriving total risk as the summation of risks R_e for statistically independent loss events, which is a statistically valid computation for aggregating risk across all known potential undesired outcomes. Computing total risk in this manner, however, exacerbates false equivalences and produces equivalent total risk for both highly reliable high-consequence outcomes (box [C] in Fig. 1) and much less reliable low-consequence outcomes (box [B] in Fig. 1). Such results are not aligned with humans' risk perceptions.[3]

To counter these concerns, the seminal argument by Kaplan and Garrick[4] holds that an event's risk should be viewed as a table of triplets consisting of the following for all relevant scenarios—description, likelihood, and its expected consequences—but that likelihood and consequences should not be multiplied. Instead, risk is represented as a complementary cumulative distribution function (also called a risk curve, an exceedance frequency curve, or a Farmer curve, see Fig. 2), which makes explicit both the likelihood and consequence elements of risk, thus eliminating the troubling multiplication-induced risk equivalences. But even here, to compare risk among various potential undesired outcomes, a statistically independent estimate of an event's likelihood is critical to understanding its risk. For random events, such as equipment failures or those induced by nature, achieving statistical independence is possible and validated in practice. For security-related undesired events, where a malevolent human is making deliberate decisions about whether to initiate an attack and choosing only among attack pathways they believe to be advantageous, statistical independence cannot be achieved because these decisions cannot conceivably be considered random.

For additional clarity, consider the spirited debate in risk analysis regarding the effectiveness of one popular security-based interpretation of the geometrical proportionality between frequency and severity, the threat, vulnerability, consequence (TVC) framework. Commonly used in terrorism risk analysis, this framework replaces the probability P from Eq. (1) with the terms threat T and vulnerability V). On one side of the debate, risk scholars contend that the simplistic TVC construct for quantifying risk is appropriate for addressing adaptive adversaries and can be rigorously conducted with high accuracy.[5] On the other side of the debate, risk scholars question if "TVC models are useful in general, or usually, for correctly assessing attack risks or setting priorities" and describe them as "simplistic, unvalidated, and low-performing" risk models.[6] Further, dissenting risk scholars argue the TVC-type risk models

1. do not address size differences in information partitions between the adversary planning their actions and analysis attempting to ascertain their actions

2. cannot account for the self-defeating element of conditioning terrorism risk analysis on judgments of that adversary's actions

3. do not ask the right questions or elicit relevant information for predicting risks

4. have not demonstrated better results than random models

5. have ill-defined key terms

6. neglect important conditionality between the mathematical terms

7. omit crucial information needed to predict and manage risks.

At the heart of this debate is the drive to develop and deploy risk analysis techniques to "provide useful

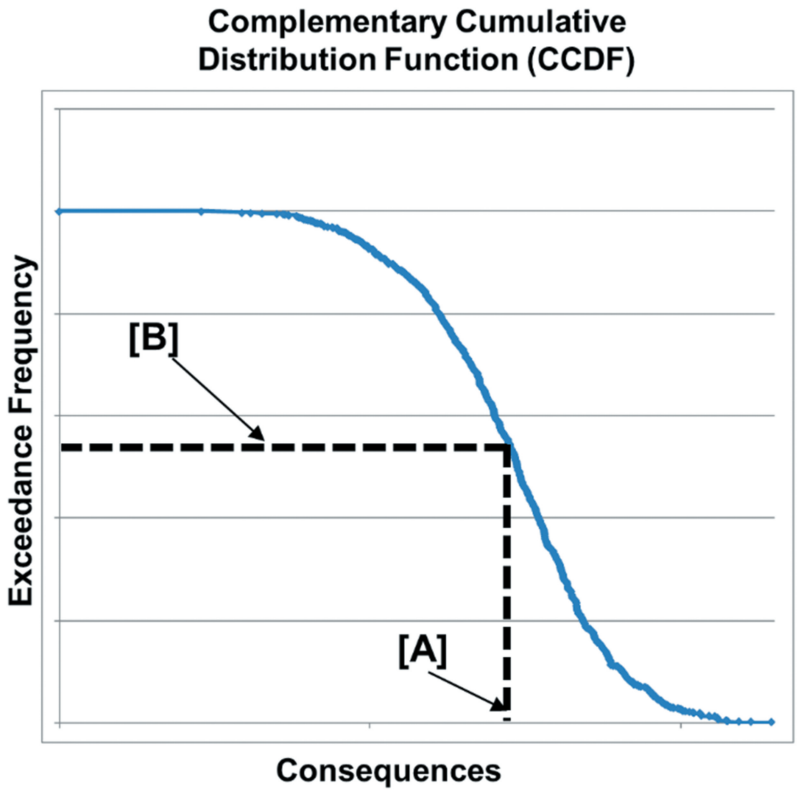

Fig. 2. Notional complementary cumulative distribution function (also known as a risk curve, an exceedance frequency curve, or a Farmer curve), wherein [A] if wondering about this level of consequence, then [B] represents the likelihood of occurrence for a scenario with consequences that are at least that bad.

insights guiding effective allocation of limited defensive resources" (Ref. 7, p. 194), which is a particularly germane challenge for nuclear security. Yet, consider how this debate in the larger risk analysis community characterizes similar discussions regarding security for nuclear materials and facilities. In November 2020, the International Nuclear Materials Management professional society convened leading risk academics and nuclear security practitioners for a workshop on Quantifying the Likelihood of Attack, ultimately taking the position that

> [q]uantifying the likelihood of an adversarial attack is *not substantially different from estimations/quantifications that have been performed for other rare events in nuclear safety analyses*. Expert judgments are currently used for physical security purposes and, through appropriate guidance, could be extended to include likelihood of attacks using quantitative or semi-quantitative methods.[8] (emphasis added)

In terms of the larger risk analysis community debate, this workshop suggests that many of the attendees sit on the pro-TVC framework side. Yet, if the logical and analytical issues sitting on the other side of the debate are valid, there is a need to better explore their implications for adequate security for civilian nuclear facilities.

Because providing security requires decision makers to understand and manage a dynamically uncertain risk landscape, estimating attack likelihood with enough precision to be useful and actionable for security risk management is philosophically, scientifically, and practically challenging. The remainder of this paper discusses the conceptual and analytical shortcomings of various approaches to calculating the likelihood of attack as a foundational element of security risk management. Augmented with representative examples, this paper then offers a set of characteristics that can guide the creation of alternative concepts that provide more robust and actionable security risk management approaches better aligned with the (r)evolutionary growth in civilian nuclear facilities.

III. EVOLUTION OF SECURITY RISK IN THE CIVILIAN NUCLEAR DOMAIN

The historical development of nuclear security, as outlined by Sandoval et al.[9] traces the evolution from simplistic "guards, guns, and gates"–based solutions through

revolutionary (at the time) systems engineering–based approaches to the cutting-edge modeling and simulation capabilities of today. This history highlights three important trends describing security risk management. One, core security principles and concepts have remained relevant over last 70+ years of civilian nuclear operations. Two, the analytical approaches have evolved and grown commensurate with underlying logical advancement and trends, increasing the complexity of implementing security solutions. Three, the most common conception of security risk in the civilian nuclear sector, however, has not evolved in the same manner.

III.A. Historical Conceptual Origins of Security Risk Concepts

Some of the longest-standing tenets of security relate to designing defensible facilities around the concepts of (1) understanding the tools, tactics, and capabilities that adversaries might be able to use, and (2) incorporating features to either render these capabilities less effective or to inflict an expectation of unacceptable casualties for the attackers. These tenets are consistent with commonly accepted drivers of uncertainty in security risk. By logical extension, as adversary capabilities have evolved, security designs have been abandoned, modified, or rebuilt to address new adversary capabilities. In response, adversaries have continually searched for new, novel capabilities or defensive vulnerabilities to increase the success of their next malicious act. Upon observing new adversary capabilities or having a defensive vulnerability exposed, security designs would again be abandoned, modified, or rebuilt. This basic (and reactive) cyclic paradigm describes the majority of security risk management over time, including common practices in the civilian nuclear sector.

A second approach to security risk management—a feature-based security design paradigm—also has ancient roots. In this paradigm, observed new adversary capabilities or exposed defensive vulnerabilities are the basis for a list of security features (or best practices) required for adequate defense. Such requirements may (or may not) be adequate to mitigate the flexibility and rapidity with which adversaries have been able to adjust tactics and to creatively identify vulnerabilities. As security has become more formalized in the civilian nuclear sector, features on such lists have become design and construction requirements,[a] which ultimately suggest that the presence of these features/requirements makes the facility "secure enough."

Though often visualized as two ends of the same security risk continuum, the reality is that most applied security risk management programs consist of a blend of these approaches. The logic of each of these security risk management paradigms has implications for addressing uncertainty in an effort to prevent undesired events. For example, consider how each paradigm addresses uncertainty in adversary actions. For the adversary capability–based paradigm, uncertainty is addressed, ensuring timely updating in security design. For the features list–based paradigm, uncertainty is addressed in the presence of the identified features for mitigation or protection. Where the former approach tends to require more resources and is favored by higher-security facilities, the latter approach is often associated with facilities of lesser security needs. Though these two approaches bound a range of clear concepts for security risk, regulators and practitioners have struggled to find consensus within this range of concepts for civilian nuclear security risk management.

III.B. Analyzing Security Risk: ERDA-7 and Its Progeny

Before the mid-1970s, civilian nuclear power safety was ensured using design-basis criteria, which included both required equipment lists and deterministic accident definitions against which facility designs must be demonstrated safe.[10] Based on this experience and familiarity, nuclear security was defined in a similar fashion by a combination of security design criteria and a design-basis threat (DBT), a list of adversary capabilities against which the facility is to be defended. As the need for more formal and official security designs at civilian nuclear facilities increased, so did the need to develop more robust security risk understanding and management. Cutting-edge efforts to improve safety at civilian nuclear facilities in the mid-1970s, led by Rasmussen's groundbreaking 1974 study (WASH-1400) that introduced probabilistic risk assessment[11] (PRA), were similarly borrowed by the nuclear security community. The foundational principles of PRA follow the philosophy espoused by Pascal, where risk (for a scenario) is computed as the product of the frequency of the initiating event F_{IE}, the conditional probability that consequences occur given that the initiating event occurs $P_{C|IE}$, and the magnitude of the scenario's consequence should it occur C_S. Mathematically, this describes the expected value of a scenario's risk.

Shortly after, a modified version of the WASH-1400 risk model was first proposed for nuclear domestic

[a] A variant of this method is the standard design paradigm, where an exemplar exists for something that has proven very secure, and exact copies are built elsewhere, even where differences in terrain, weather, or adversary capabilities render the standard design vulnerable in unexpected ways.

safeguards, another term describing security in the United States,[12] known as the Energy Research and Development Administration (ERDA)-7 proposal. ERDA-7 took Rasmussen's equation and redefined its components for security applications. F_{IE} became the likelihood of attack (probability or frequency as appropriate). $P_{C|IE}$ became the conditional probability that the attack is successful for the adversary. C_S remains defined as the consequence should the successful attack occur

Over the years, the ERDA-7 proposal has been subject to reintroduction and modification.[13,14] In some cases, the attack likelihood term is treated qualitatively, quantitatively, or removed altogether. Removing it altogether leads to the computation of conditional risk values where the risk that would exist given that an attack of a specific type was to occur.[15,16] For example, in 1986 the U.S. Department of Energy (DOE) published a *Master Safeguards and Security Agreements (MSSA) Guide* that introduced a version of security risk that "… deviates from classical risk evaluation by basing risk calculations on the assumption that the adversary event occurs."[17] Yet, subsequent critical reviews stated that an approach like ERDA-7, based on traditional risk assessment, should not be used for security risk.[18,19] For this reason, security risk management methods used by the DOE and NRC rely on performance-based standards for the effectiveness of security systems,[20,21] as well as addressing consequences where possible, instead of only in ERDA-7-like risk quantification.

Despite never being formally adopted, ERDA-7 has a strong legacy in security risk discussions. Perhaps the most commonly used recent example of ERDA-7-like thinking is the oft-quoted conceptual "equation" describing nuclear security risk in terms of the TVC framework, where threat and vulnerability are considered an unconditional and a conditional probability, respectively. This description parallels a common understanding for safety risk in the nuclear power sector. It provides a clear framing for addressing the challenge of protecting critical assets against malicious acts. However, it also promotes common misperceptions about security risk that can confuse security risk management efforts at civilian nuclear facilities.

III.C. Current State of Security Risk Practice

Current security risk management practices are a mixture of all the methods described previously, generally captured in best practices and required design feature lists for both physical security and cybersecurity. Often, higher-security facilities focus more on adversary capabilities while facilities with lesser security needs focus more on prescriptive lists of required security features. As such, high-consequence facilities tend to bolster these approaches with more robust analytical methods. These can include approaches that account for current or future adversary capabilities, like the threat-based planning and capability-based planning[22] methods. Security designers and analysts are encouraged to think about threat, vulnerability, and consequence, sometimes using a qualitive view of ERDA-7 to frame discussions. There is still a tendency to quantify ERDA-7 approaches to security. Consider what the U.S. Department of Homeland Security (DHS) uses for terrorism risk assessment and other security risk management activities. Specifically, the U.S. Coast Guard developed the Maritime Security Risk Analysis Model[23] (MSRAM) for port security to identify and prioritize critical infrastructure, key resources, and high-consequence scenarios using the ERDA-7 framing. Such examples are commonly used to defend the inclusion of quantified ERDA-7 approaches for nuclear security risk management.

In the commercial nuclear sector, the Design Evaluation Process Outline[15] (DEPO) methodology created by Sandia National Laboratories is widely considered the state of the art for nuclear security. The DEPO methodology is popular for its relative ease of use and probabilistic modeling paradigm, which concludes when equilibrium is achieved between the assumed level of security risk and security system effectiveness.[15] DEPO also borrows both an ERDA-7 approach and a timeline-based model for describing risk for nuclear safety. The accident timeline is replaced by comparing timelines for adversaries to achieve a malicious act versus the response force protective actions. Consistent with the ERDA-7 legacy, DEPO measures security as probabilistic influences on this timeline (e.g., the probability that a particular sensor alarms when an adversary enters a prohibited area or the probability that the response force is able to intercept the adversary) to describe threat and vulnerability.

While the current state of practice for security risk management for civilian nuclear facilities is dominated by the combination of prescriptive security features and adversary capability estimates, the risk-related elements of these approaches are built on a conceptual foundation framing security risk in terms of ERDA-7.

IV. WHAT'S MISSING IN SECURITY RISK … AND WHAT'S NEEDED TO MOVE FORWARD?

As previously mentioned, a range of influences are poised to challenge the efficacy of the traditional

conceptions of security risk management in civilian nuclear facilities. The most common security risk concepts for commercial nuclear applications—the features list, adversary-basis, and ERDA-7 framing—each suffer from philosophical, scientific, and practical gaps and shortcomings. Yet, these gaps and shortcomings also identify crucial next steps in the evolution of nuclear security risk management.

Features list–based methods for security risk management are attractive for their philosophical simplicity and practical clarity. In these methods, the definitive answer to the question "how secure is secure enough?" is embedded in the list of features itself. Yet, these advantages carry several key gaps and shortcomings. Philosophically, features lists are inherently challenged by the inability of those responsible for developing the lists to perfectly, adequately, and comprehensively identify what may be needed to counter any specific adversary act. In addition, any accepted features list causes a conceptual imbalance by effectively mapping a static security posture against a dynamic threat. The problem is compounded when features lists are generic (i.e., they apply to more than one facility), because a generic list cannot address the unique security challenges posed by characteristics that differ between facilities.

These philosophical shortcomings result in the scientific challenge of incompleteness. Features list–based approaches cannot account for possible inadequacy to defend against the actual threat or for adversary adaptation to defeat the features. Further complications can arise from inadequate implementation if the listed features do not operate as expected. At best, this would fail to provide effective security. At worst, this could manifest in unexpected and perhaps even unidentified susceptibilities in the security posture of civilian nuclear facilities. If risk management means understanding the possibility of experiencing an undesired state, then features list–based methods provide one mechanism for defining those undesired states. Yet, the associated uncertainties and incompleteness challenge the accuracy and appropriateness of these descriptions of security risk, thereby limiting their effectiveness in supporting security risk management.

The DBT-based methods are similarly attractive and popular because they build upon the logic of features lists with a focus on countering specified adversary capabilities. In these methods, the answer to the question "how secure is secure enough?" is related to how accurately adversary capabilities are defined and how they are expected to be defeated. Because these approaches do not require any specific security features and conceptually map to expected adversary capabilities, they provide more flexibility for addressing security risk management. Like features lists, DBT-based approaches also struggle with philosophical shortcomings. For example, adversaries typically can adapt their capabilities [in both tactical (near-term) and strategic (long-term) timeframes] more quickly than a DBT can be (temporally or permanently) updated, representing a slightly more advanced version of mapping a static solution to a dynamic problem. An additional challenge for DBT methods stems from selecting reasonable and credible attack scenarios. Conceptually, the logic of DBT methods discounts a simple attack scenario only slightly outside the DBT (e.g., a scenario using just one beyond-DBT capability), yet includes other, potentially highly complex attack scenarios that are fully within the DBT, even if the former is more credible. The logic underpinning these methods also tends to bias designers toward scenarios at the maximum capability of the DBT, while lesser attacks may not be fully explored or expected. If conservatism shifts security focus to the maximum DBT capability, then lesser adversarial acts may increase their chances of success, particularly using stealth or deceit strategies.

Furthermore, where DBT-based methods provide a more flexible mechanism for framing security risk, these philosophical gaps yield additional sources of potential analytic incompleteness. As written, DBT methods require that security risk management "defeat all credible attacks by DBT-included adversaries." While a seemingly simple and straightforward task, defining credible is a difficult and subjective task that incorporates adversary characteristics into scenario complexity (and uncertainty) to define (un)desired boundaries for security risk.

The family of ERDA-7–related methods heavily borrows from advances in nuclear safety in an attempt to more directly manage security risk through quantification. In these approaches, a version of the simplified risk equation is used to compute risk and compare against numerical thresholds to answer the question: "How secure is secure enough?" One of the most liberal simplifications of ERDA-7–related methods is treating the risk variables in the TVC framework as statically independent. Here, consider that a major goal of many security posture upgrades at civilian nuclear facilities is to increase the difficulty of adversary tasks, thereby reducing the adversary's likelihood of success and thus driving the adversary to decide not to attack. This concept of "adversary decides not to attack" is deterrence and is not well captured when TVC framework elements are considered independent. Similarly, because security posture upgrades affect both T and V, there is both conditionality

(as opposed to independence) and nonlinearity within the ERDA-7 security risk equation. For example, if a facility becomes more secure (decrease the *V*), observations and experience suggest that the probability of attack will decrease (decrease in *T*) in response. Conversely, if a facility becomes less secure (increase the *V*), observations and experience suggest that the probability of attack will increase (increase in *T*") in response. This simplified approach to security risk assessment cannot capture similar critical elements observed in adversary behavior, including adaptation, threat shifting, or deterrence.[24,25]

This misrepresentation of the canonical risk equation also manifests in scientific shortcomings that result in analytic gaps. First, consider the common failure to adequately define the terms in this simplified risk equation.[19] One key definitional concern is the point that conditionality of the computed variables is frequently neglected, or at least left unstated, leading to statistically meaningless results. For example, if the characteristics of the adversary for which likelihood is asserted in *T* would overwhelm vulnerability *V*, then the value of *V* should approach unity, whereas a threat that would have great difficulty overcoming this vulnerability may require a value of *V* that approaches zero.[19] Second, ERDA-7–based security risk representations seem based on the premise that depictions of probability of attack must behave similarly to the initiating event frequency in safety risk equations, which is a mathematical instantiation based on an incomplete philosophical argument.

Even if the WASH-1400 and ERDA-7 conditionality among the risk equation's terms is retained, critical elements of adversary behavior (adaptation, threat shifting, or deterrence) cannot be captured. This is because probability of attack is assumed to be the independent variable, whereas adversary decision-making behavior shows that this variable is strongly dependent on both the attack's likelihood of success and expected outcomes. For example, the hidden dependencies that must be addressed while eliciting the probability of attack include estimating the likelihood that some known or unknown individual or group will exist during some specified future time period, as well as estimating that that particular adversary group will

1. decide that an attack can achieve an outcome or consequence that they desire

2. understand and validate an exploitable vulnerability or pathway to plan a viable attack

3. obtain the weapons, tools, skills, and information required to accomplish the attack

4. decide that the attack's likelihood of failure, potential losses, and risks are acceptable

5. decide that the costs and sacrifices required to accomplish the attack are acceptable

6. decide that this is the best opportunity to accomplish a desired objective by comparison to all other known opportunities at this or any other facility or location.

Each and every one of these factors must be accounted for if the desired probability of attack result is a joint probability capturing all of these conditions, many of which are interdependent. The arguments against heavy reliance on using probability of attack are longstanding, with Rasmussen, the primary author of WASH-1400 and father of probabilistic risk assessment, articulating that

> [f]or a number of reasons, however, I do not believe that safeguards [e.g., security] risks can be quantified using these [PRA] procedures … the basic assumptions in the [reactor safety study] RSS methodology is that failures are basically random in nature … allows one to estimate a system failure by an appropriate combination of the failure rates of its parts. In the case of deliberate human action, as in imagined diversion scenarios, such an assumption is surely not valid.[20]

Last, semi-quantitative interpretations of ERDA-7 where likelihood (e.g., probability of attack) and consequence are viewed as subjective measures converted to corresponding ordinal values for risk calculation have been shown to often lead to wildly inaccurate risk results.[26] Further, the negative impacts on security risk from inaccurate results are compounded by the often-large uncertainty associated with any probability of attack number, which relates to four different types of unknowns. First, the uncertainty about known or unknown adversaries must consider what world events or personal events might cause new adversaries to manifest or existing adversaries to cease. Second, conditionalities 1, 4, 5, and 6, listed previously, all depend on each adversary group's value set, so that making changes in these values over time further contributes to uncertainty in the probability of attack. Third, conditionalities 2, 3, and 6 depend on each adversary's opportunities to obtain the information and resources necessary to identify and exploit a facility's vulnerabilities. Again, these opportunities may change over time in uncertain, and sometimes radical and sudden, ways as facility information leaks (e.g., WikiLeaks) and adversarial tool availability changes

(e.g., wide availability of unmanned aerial vehicles). For these reasons, ERDA-7–based interpretations of the probability of attack as independent likely yields extraordinarily uncertain results, particularly if assessed over an extended future time period using Bayesian methods.

If these uncertainties are neglected, the security risk results are likely inaccurate. Yet, if the uncertainties are fully represented, the security risk results are likely not to be statistically significant, and therefore not actionable. In short, the apparent analytic benefit of quantifying probability of attack to create a numerical criterion for "how secure is secure enough" is largely negated by security risk distributions whose uncertainty can span several orders of magnitude.

Even assuming a reasonable approach for addressing the philosophical and scientific shortcomings were to be developed, there still exists a range of practical challenges to ERDA-7–type approaches. Consider how local or geopolitical events can significantly increase or decrease adversary motivation to attack in very short order, which presupposes a need to reevaluate the probability of attack with similar frequency. In addition, security risk centered on the probability of attack assumes high-level, timely, and accurate knowledge of a (at least mostly) complete set of adversaries. For domains with higher frequencies of adversary actions, these assumptions are more reasonable. The MSRAM is one example of successfully using an ERDA-7–type approach,[27] in large part driven by the ease and frequency of attacks on port facilities. In contrast, the ERDA-7 approaches the DHS uses in its terrorism risk assessments are plagued with such broad uncertainties that it is almost impossible to draw from them statistically significant risk management insights.[28] Moreover, reliance on probability of attack as a numerical security risk criterion can also cause confusion between regulators and operators, as demonstrated by the popular terms risk based and risk informed.[29] Though the former is making decisions in relation to established PRA thresholds that define acceptable levels of risk and the latter is making decisions in which insights from PRA are integrated into a broader process, ERDA-7 approaches often conflate them as equivalent.

One variant of ERDA-7–type approaches that has been explored to better account for the challenges inherent to security risk management revolves around the concept of conditional risk. In other words, these ERDA-7 variants attempt to overcome the many problems establishing probability of attack by computing risk assuming that an attack occurs. While this approach does mitigate some of the previously discussed issues with probability of attack, it simply ignores the large span of uncertainty, and yet, still fails to capture the interdependence between threat, vulnerability, and consequence. Treating likelihood in this way fundamentally and completely misses the basic fact that the conditionality in security risk is different from, and essentially the opposite of, independence in the safety risk equation. For example, conditional risk assessment is often colloquially and wrongly described as "setting probability of attack to 1.0." While conditional risk actually describes the risk if the specified attack occurs, the tradition of setting probability of attack to unity is a statement of belief that the specified attack will occur. In a practical sense, a conditional risk approach describes security risk as equivalent between a facility with a high vulnerability and low consequence and a facility with a low vulnerability but high consequence, a security version of the false equivalencies in Fig. 1. In addition, despite the perception it can make security more uniform, these potential misinterpretations of conditional risk actually increase the difficulty in comparing security risk across organizations, facilities, and locations.

While each of these traditional approaches to security risk has been successful (at least to a degree), the evolution necessary to effectively mitigate twenty-first century malicious threats to nuclear facilities can start by exploring the philosophical, scientific, and practical needs for security risk management.

V. ADVANCING SECURITY RISK

Successful security risk management requires methods that both produce accurate and actionable results and have sound philosophical and scientific foundations. As the preceding section made clear, there is currently no method that satisfactorily meets these criteria in general, let alone for civilian nuclear facilities where attacks are very rare events, consequences could be very large, and both likelihood and consequences are extraordinarily uncertain. Thus, overcoming the gaps and shortcomings discussed in the previous section introduces an opportunity for framing next-generation security risk management.

Despite exploring the challenges associated with the probability of attack for nearly 50 years, the risk community is still debating the same philosophical and scientific questions. There is a need to better address the interdependencies within security risk in order to effectively manage security responses to intentionality, human choices, deterrence, and threat shifting. Furthermore, security risk management methods must deal with not

just the adversary, but the potential for multiple adversaries to target a facility. Each such adversary will have differing attack objectives, some of which the defenders may not know exist, let alone understand their capabilities, motives, and intent. More specifically, what is needed are new approaches to quantifying risk that trade simplistic calculations for probability of attack for descriptions that capture how the likelihood that an adversary will select a specific attack to initiate. This depends on factors, such as whether

1. the proposed attack would lead to an outcome (consequence) that some adversary desires
2. this potential adversary believes that they can successfully execute the attack or can adapt sufficiently to achieve success in an attack
3. this potential adversary considers their expected cost to be acceptable
4. this potential adversary does not have other attack options that are significantly superior to this proposed attack.

The requirement to consider multiple diverse adversaries requires repackaging security risk to frame it in terms of asset and facility exposure. This perspective also acts as a mechanism for addressing basic adversary decision processes that were simplified away when the probability of attack was assumed independent. In many ways, the probability of attack is likely the most dependent variable in the ERDA-7-like risk formulation. This cannot be overemphasized; probability of attack for one scenario, or one facility, is strongly dependent on the characteristics of all other attack scenarios that are available to the adversary at this and all other facilities. In response, next-generation security risk management approaches must more clearly address this fundamental difference between security risk from safety risk.

By extension, it is important to remember that pursuing a Bayesian approach to quantify these elements related to strategic security risk (e.g., probability of attack) using an ERDA-7-like method brings extraordinary uncertainty. Propagating such large levels of uncertainty risk results makes producing statistically significant risk management insights nearly impossible. For example, inadequately managing the uncertainty of such basic adversary behaviors, like adaptation and threat shifting, into security risk assessment may provide decision makers with an inaccurate picture of adversary behavior by failing to model the very phenomena that form the basis for security risk.[28]

It is also possible that lurking just beneath the surface of the desire to quantify security risk is the philosophical notion that it is not possible to effectively manage what cannot be quantified. Yet, the definitions of risk from the Society for Risk Analysis[1] are all qualitative, suggesting an appropriateness in exploring alternate forms of risk management to overcome these philosophical barriers to security risk quantification. Part of the solution to this dilemma may lie in a revised concept of measurement. Hubbard[26] claims that measurement should not be considered as assigning a number to something, but rather as reducing its uncertainty in a meaningful or actionable way. Under this philosophy of measurement, the components of security risk, or their surrogates, can be qualitatively measured in a meaningful way for use in security risk management.

If the related timing dynamics contribute to philosophical confusion and significantly large uncertainties, then security risk management should seek to reconcile differences between near-term (tactical) versus long-term (strategic) time horizons. For example, risk managers may have credible adversary information to clarify a near-term (e.g., matter of days, weeks, maybe months) likelihood of attack and deploy tactical security solutions. In contrast, strategic decisions to address security risk (like facility construction details that may be in place for decades) are made against the reality that dramatic changes to adversary motivation and intent may occur over the years these construction details remain in place. Now consider scenarios in which adversaries execute and abort decision-action loops many times while probing the security system over the long term (strategic timeframe), possibly learning and adapting after each iteration. Such examples reinforce the need for philosophical approaches capable of addressing the security risk implications of both time domains, even if via different risk management methods.

One recent attempt leveraging such a new paradigm of security risk, Sandia National Laboratories's Risk Informed Management of Enterprise Security (RIMES) methodology, demonstrated the robustness of qualitative methods for security risk management over the strategic time horizon at facilities with very low probability of attack and very large uncertainties.[30] The RIMES method uses a relative rating scale comparing attack scenarios on the basis of their difficulty and consequences, providing the analyst with a risk landscape of potential attacks to be managed. It captures elements of important adversary behaviors, like human choices, adversary adaptation, and threat shifting, precisely because it does not attempt

to quantify probability of attack to produce numerical estimates for security risk. Trading the traditional notion of probability of attack for a function of difficulty, consequence, and scenarios suggests the benefit of shifting the basis of security risk management from "risk is a quantified and uncertain singular value" to "risk is a rigorous description of an uncertain undesired future state."

Building on this broader philosophical foundation, the scientific approach to security risk should evolve past the traditional, if mostly false, dichotomy between safety and security. For example, one of the most widely marketed characteristics of advanced and small modular reactors are varying layers of passive safety, which are mechanisms to maintain cooling and/or containment that require little or no external power sources or human involvement. While this logic supports a reduced risk of an accidental radiological release, passive safety is also claimed to reduce security risk. Yet, while the passive safety features may reduce the chance for a release at the reactor itself, the passive safety features now represent an expanded set of assets exposed to potential adversary actions. Here, the passive safety mechanisms may not reduce the security risk, but instead simply change how security risk manifests. Indeed, if they are easy for an adversary to defeat, they could potentially even increase security risk.

Moving beyond quantifying probability of attack for security risk, this perspective of security risk management provides an opportunity to use a broader set of scientific approaches to address security risk over both strategic and tactical timelines. Consider, for example, different conclusions that can be drawn from frequentist and Bayesian interpretations of probability of attack. For frequentist inference, probability of attack relates to deriving the frequency of adversary attack from a known data set, and larger data sets increase mitigation of related uncertainties. For Bayesian inference, probability of attack relates to belief in the specific attack occurring and evidence supporting prior beliefs of the specific attack. While both approaches to deriving statistically backed insights for complicated questions have long traditions, neither, at least in their most common applications, adequately capture security risk. Despite current limitations in both frequentist and Bayesian statistical approaches, cutting-edge advances in approach and applications are worth exploring to enhance security risk management.

Some of the underlying logical framework from these scientific approaches can be leveraged to better address the complexity of identifying, measuring, and describing security risk. While both frequentist and Bayesian approaches provide mechanisms for describing security risk uncertainty, evolutions of the latter have been shown to address interdependencies similar to what is observed in nuclear security.[31] Enhancing the ability to incorporate interdependencies will help reduce uncertainties related to adversary actions. Consider, for example, how better scientific approaches inclusive of interdependencies can better address

> the totality of the security and exposure characteristics when evaluating whether specific proposed changes to [asset or facility] exposure frequency will significantly affect an asset's security risk.[24]

This is but one example of how new (or nontraditionally interpreted) scientific approaches can mitigate the challenges surrounding perceptions of probability of attack as an independent variable.

An expanded range of scientific approaches invites opportunities for exploring new analytical tools and approaches for security. For example, RIMES (previously described) and the expanding use of Bayesian belief networks (e.g., Ref. 32) illustrate current efforts to describe risk more accurately in practice by using novel applications of traditional risk assessment approaches. In RIMES, the focus on quantified risk is replaced with a focus on comparative risk management as a function of attack difficulty, consequence, and scenario. Similarly, Bayesian belief networks extend the traditional use of prior information and new evidence into a graphical network that mathematically describes many of the interdependencies that had traditionally been simplified away.

On the other hand, an improved understanding of security risk, and by extension, better implementation of security risk management, can also be generated from lessons learned from other domains. Here, the complex systems and hazards analysis disciplines stand out as rife with useful insights. As a representative example, consider the application of multiple objective decision analysis (MODA) and systems-theoretic process analysis (STPA) to nuclear security. Born out of systems engineering, MODA approaches are useful for organizing multiple (often conflicting) objectives, capturing explicit value trade-offs, integrating facts with value preferences, and aggregating these dependent factors in an accountable manner. In Ref. 33, the many interdependent influences on security risk are modeled as attributes impacting potential access by a nonstate actor (e.g., scale of nuclear infrastructure in a country) and attributes impacting the effectiveness of security implementation and culture (e.g., subject matter expert assessment of physical protection

system capabilities). By explicitly addressing interdependence, MODA is well suited to enhance nuclear security risk management.

Likewise, STPA's creation to enhance hazard analysis for complex systems provides a structured way to address external influences (e.g., increasing adversary capabilities), internal influences (e.g., existing gaps between operational and protective force personnel), and interactive influences (e.g., increasingly restrictive security budget and resource constraints) on security risk.[34] STPA's broader view of causality and complexity, as demonstrated in the growing literature on the topic, suggests that applying STPA to nuclear security could help shift risk management away from preventing failures and toward enforcing security constraints. STPA's systems and control theoretic foundation position is a technique capable of better mitigating current challenges in all elements of security risk management.

As new approaches for analyzing security risk are improved and provide better capabilities to identify, characterize, and mitigate uncertainties of security risk, the ability to effectively communicate security risk management across stakeholders will also improve. Today, visualizing security risk in terms of independent probabilities of attack is clear and straightforward, a clarity for regulation and decision making that supports a desire for a clear and straightforward quantitative risk standard. Tomorrow, visualizing security risk will be based on new approaches, models, and frameworks that provide both equal levels of clarity and a more comprehensive description against which to make regulatory and budgetary decisions about where to spend the next security risk management dollar. So, just as nuclear safety has transitioned from a purely risk-based to a risk-informed paradigm, so too can security risk management as it evolves beyond singular probabilities of attack.

VI. CONCLUSIONS, INSIGHTS, AND IMPLICATIONS

Quantitative PRA has been used to help ensure the safety of nuclear facilities for almost half a century. Its benefits have included the following:

1. the ability to discover unexpected safety challenges
2. a mechanism by which to grade and prioritize nuclear safety research
3. methods by which the significance of plant events as potential precursors to safety issues are determined
4. the development of risk-informed regulation
5. a mechanism whereby facility design and retrofit proposals can be objectively discussed and adjudicated between facility and regulator.

The desire to replicate these benefits in the security arena is large, and the ERDA-7 proposal provided a mathematical roadmap for doing so. However, using PRA to manage security risk in support of responsible operation of nuclear facilities presents substantial philosophical, scientific, and practical challenges, many of which can be traced back to the faulty, yet popular, assertion that the probability of attack can be treated as a probabilistically independent variable, parallel to the initiating event frequency in a safety PRA.

Philosophically, a key problem is that the likelihood that some presumably rational adversary will decide to initiate any attack, let alone a specific attack, is not a random event. Rather, it is a human decision based on motivation, intent, cost-benefit calculations, and (likely) limited information or knowledge, each of which leads to uncertainty for the adversary. It is possible to estimate an unconditional likelihood of attack in a Bayesian sense because the probability estimate embodies our understanding and beliefs about uncertain future events. However, for potential major attacks, the analyst attempting to estimate this likelihood must grapple with many information and knowledge limitations themselves, including those limitations stemming from the difficulty of accounting for the list of hidden dependencies described in Sec. IV. In addition to all the previously noted decision factors, Bayesian approaches must also address uncertainties related to all relevant conditionalities and hidden dependencies (also, as described in Sec. IV). For these reasons, the philosophical assumptions supporting ERDA-7–based approaches should be reexamined and scrutinized to ensure the high fidelity and acceptably low uncertainty of any attempts to quantify the probability of attack.

Scientifically, the lack of statistical independence for the probability of attack means that the math embodied in ERDA-7, and in the more contemporary TVC framework, is invalid. All of the terms in these equations are interdependent: C and V depend on the characteristics of T; the likelihood embodied in T depends on its required characteristics in addition to the uncertainties inherent in previously described philosophical issues. Both formulations, as generally implemented, ignore important conditionalities. For this reason, the fundamental aspects of human behavior, such as deterrence and threat shifting, cannot be modeled using these methods and equations.

The importance of these behaviors cannot be overstated; it is generally the goal of the defenders of high-consequence facilities to deter adversaries from attacking and to do so by reducing vulnerabilities and/or consequences for attacks. A tool that cannot address the fundamental aspects of the underlying system's behavior should be viewed with suspicion as a risk management tool.

There are practical consequences resulting from the philosophical and scientific issues with these methods. If the true uncertainties in attack likelihood are propagated through the ERDA-7 and TVC framework models, the risk uncertainties often span several orders of magnitude and drown out statistically significant risk insights. If, however, these uncertainties are not propagated, the risk results may provide decision makers with the wrong inferences upon which to execute security risk management. Furthermore, as conditions change, either at a facility or among the adversary community's decision criteria, the computed probabilities become invalid and must be recalculated, rendering the risk assessment process time consuming and in need of constant revision. This is undesirable as a basis for making long-term security risk management decisions, such as new construction that is intended to be in service for decades.

ERDA-7 has been the subject of research for nearly a half century, and yet the questions and challenges are the same today as when it was first proposed. The likelihood of attack issues make these quantitative methods questionable for long-term strategic security risk management, but that does not mean that likelihood of attack is useless. Indeed, it is critical for short-term tactical security risk management, particularly for smaller target footprints and/or over a short, finite time span. For example, if intelligence information indicates that there is an elevated likelihood of attack at one facility, or a class of facilities, prudent risk managers make a tactical decision to rapidly increase security to deter such an attack over the duration of the elevated threat. As a risk indicator for smaller target footprints and shorter time spans, qualitative estimates of attack likelihood are generally adequate for tactical security risk management. So, if it is not beneficial for security risk management over the span of years/decades and unnecessary for tactical security risk management over the span of days/weeks, perhaps the quantification of attack likelihood may be unwarranted.

In support of exploring new directions for the future of security risk management, this paper offers a few suggestions. Cutting-edge advances in applying Bayesian approaches to incorporate interdependencies in risk discussions[32] and conditionalities in risk monitoring for complex systems[35] show promise for improving security risk management. Similarly, insights from implementing novel [e.g., RIMES (Ref. 36)] and non-traditional [e.g., MODA (Ref. 33) and STPA (Ref. 37)] analysis techniques seem germane to overcoming current challenges to nuclear security management. Regardless of the path forward, security risk management must address the basic elements of adversary behavior in order to be useful. Here, if security risk focuses on preventing "undesired states, behaviors, or outcomes," then this paper suggests a broadening of the term security risk beyond simple probabilistic single-value descriptions. A broader concept would support movement from traditional interpretations of probability of attack toward more nuanced and complex depictions of security risk. In their current form, ERDA-7 and similar methods require the estimation of a quantitative attack likelihood. While such estimation is possible in a Bayesian sense, it is not useful for long-term security risk management.

Acknowledgment

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the DOE's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in this paper do not necessarily represent the views of the DOE or the U.S. government.

Disclosure Statement

No potential conflict of interest was reported by the authors.

ORCID

Adam D. Williams http://orcid.org/0000-0001-8159-9737

References

1. "Society for Risk Analysis Glossary," Society for Risk Analysis (2018).
2. G. PITA, "Setting the Framework for Risk Assessment—Blaise Pascal," *Risk Science and Engineering* (April 7, 2017); https://riskmodeling.org/rse-blog/2017/4/7/setting-

a-framework-for-risk-assessment-blaise-pascal (current as of May, 11, 2022).

3. P. SLOVIC, "Perception of Risk," *Science*, **236**, *4799*, 280 (1987); https://doi.org/10.1126/science.3563507.
4. S. KAPLAN and B. J. GARRICK, "On the Quantitative Definition of Risk," *Risk Anal.*, **1**, 11 (1981); https://doi.org/10.1111/j.1539-6924.1981.tb01350.x.
5. B. EZELL and A. COLLINS, "Letter to the Editor," *Risk Anal.*, **31**, *2*, 192 (2011); https://doi.org/10.1111/j.1539-6924.2010.01562.x.
6. G. G. BROWN and L. A. COX JR., "How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts," *Risk Anal.*, **31**, *2*, 196 (2011); https://doi.org/10.1111/j.1539-6924.2010.01492.x.
7. G. G. BROWN and L. A. COX JR., "Making Terrorism Risk Analysis Less Harmful and More Useful: Another Try," *Risk Anal.*, **31**, *2*, 193 (2011); https://doi.org/10.1111/j.1539-6924.2010.01563.x.
8. J. RIVERS, A. LINTEREUR, and M. DURBIN, "INMM Workshop on the Quantification of the Likelihood of Attack," International Atomic Energy Agency (2020).
9. J. S. SANDOVAL, A. D. WILLIAMS, and R. ROSANO, "The History of Guns, Gates, and Guards in Nuclear Security," *Proc. for the Institute for Nuclear Materials Management 61st Annual Meeting* (2020).
10. G. D. WYSS, "The Accident that Could Never Happen: Deluded by a Design Basis," in *Learning from a Disaster: Improving Nuclear Safety and Security After Fukushima*, E. D. BLANDFORD and S. D. SAGAN, Eds., Stanford University Press, Redwood City, California (2016).
11. "WASH-1400: The Reactor Safety Study: The Introduction of Risk Assessment to the Regulation of Nuclear Reactors," U.S. Nuclear Regulatory Commission (1975).
12. C. A. BENNETT, W. M. MURPHEY, and T. S. SHERR, "Societal Risk Approach to Safeguards Design and Evaluation," Energy Research and Development Administration (1975).
13. C. J. UDELL et al., "Risk Evaluation System for Facility Safeguards and Security Planning," *Proc. of the 30th Annual Meeting of the Institute of Nuclear Materials Management*, Deerfield, Illinois (1989).
14. C. J. UDELL, J. A. TILDEN, and R. T. TOYOOKA, "Short-Form Risk Evaluation Method," presented at the 34th Annual Mtg. of the Institute of Nuclear Materials Management, Scottsdale, Arizona (1993).
15. M. L. GARCIA, *The Design and Evaluation of Physical Protection Systems*, 2nd ed., Butterworth-Heineman, Burlington, Massachusetts (2008).
16. B. E. BIRINGER, R. V. MATALUCCI, and S. L. O'CONNOR, *Security Risk Assessment and Management: A Professional Practice Guide for Protecting Buildings and Infrastructures*, John Wiley & Sons, Inc., Hoboken, New Jersey (2007).
17. J. SANDOVAL, "The Use of the PRA Risk Equation in DOE Security: A Chronological History," SAND2014-1003C, Sandia National Laboratories (2014).
18. J. M. RICHARDSON, "Comprehensive Safeguards Evaluation Methods and Societal Risk Analysis," Sandia National Laboratories (1982).
19. L. A. COX JR., "Some Limitations of "Risk = Threat X Vulnerability X Consequence" for Risk Analysis of Terrorist Attacks," *Risk Analysis*, 28, *6*, 1749 (2008).
20. N. C. RASMUSSEN, "Probabilistic Risk Analysis: Its Possible Use in Safeguards Problems," *Proc. of the 17th Annual Mtg. of the Institute of Nuclear Materials Management*, Deerfield, Illinois (1976).
21. B. H. GARDNER, W. K. PAULUS, and M. K. SNELL, "Determining System Effectiveness Against Outsiders Using Assess," *Proc. of the 32nd Annual Mtg. of the Institute of Nuclear Materials Management*, Deerfield, Illinois (1991).
22. T. BALASEVICIUS, "Is It Time To Bring Back Threat-Based Planning?," The Mackenzie Institute, April 7, 2016; https://mackenzieinstitute.com/2016/04/is-it-time-to-bring-back-threat-based-planning/ (current as of May 11, 2022).
23. B. DOWNS, "Maritime Security Risk Analysis Model: USCG Presentation to Area Maritime Security Committee," Presentation to the Critical Infrastructure Protection Workshop, Center for Homeland Defense and Security, June 2008.
24. G. D. WYSS, "Asset Exposure, Attack Opportunity, and Security Risk," Sandia National Laboratories (2021).
25. G. WYSS et al., "Risk-Based Cost-Benefit Analysis for Security Assessment Problems," presented at the IEEE 44th Annual Int. Carnahan Conf. on Security Technology, San Jose, California (2010).
26. D. W. HUBBARD, *The Failure of Risk Management: Why It's Broken and How to Fix It*, Wiley (2009).
27. "Coast Guard: Security Risk Model Meets DHS Criteria, but More Training Could Enhance Its Use for Managing Programs and Operations," U.S. Government Accountability Office (2011).
28. "Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change," National Research Council (2008).
29. E. ZIO and N. PEDRONI, "Risk-Informed Decision-Making Processes: An Overview," Foundation for an Industrial Safety Culture, Toulouse, France (2012).
30. G. D. WYSS et al., "A Method for Risk-Informed Management of Enterprise Security (RIMES)," SAND2013-9218P, Sandia National Laboratories (2013).

31. N. U. I. HOSSAIN et al., "A Framework for Modeling and Assessing System Resilience Using A Bayesian Network: A Case Study of an Interdependent Electrical Infrastructure System," *Int. J. Crit. Infrastruct. Prot.*, **25**, 62 (2019); https://doi.org/10.1016/j.ijcip.2019.02.002.

32. P. GEORGE and V. RENJITH, "Evolution of Safety and Security Risk Assessment Methodologies Towards the Use of Bayesian Networks in Process Industries," *Process Saf. Environ. Prot.*, **149**, 758 (2021); https://doi.org/10.1016/j.psep.2021.03.031.

33. S. CASKEY, B. EZELL, and R. DILLON-MERRILL, "Global Chemical, Biological, and Nuclear Threat Potential Prioritization Model," *J Bioterrorism and Biodefense*, **4**, 125, (2013); https://doi.org/10.4172/2157-2526.1000125.

34. A. D. WILLIAMS, "System Security: Rethinking Security for Facilities with Special Nuclear Materials," *Trans. Am. Nucl. Soc.*, **109**, *1*, 1946 (2013).

35. R. MORADI et al., "Integration of Deep Learning and Bayesian Networks for Condition and Operation Risk Monitoring of Complex Engineering Systems," *Reliab. Eng. Syst. Saf.*, **222**, 108433 (2022); https://doi.org/10.1016/j.ress.2022.108433.

36. F. A. DURAN et al., "Risk-Informed Methodology for Enterprise Security: Methodology and Applications for Nuclear Facilities," *Proc. 52nd Annual Mtg. of Institute of Nuclear Materials Managemnt*, Palm Desert, California (2013).

37. G. POPE, "Risk Management Using Systemic Theoretic Process Analysis (STPA)," LLNL-CONF-776677, Lawrence Livermore National Laboratory (2019).