Nuclear terrorism and virtual risk: Implications for prediction and the utility of models

Robert J. Downes*
Department of War Studies, King's College London

Christopher Hobbs**
Department of War Studies, King's College London

Abstract

Assessing the risk of nuclear terrorism is a challenging task due to the diversity of actors involved, variety of pathways to success, range of defensive measures employed, and the lack of detailed historical record upon which to base analysis. Numerical models developed to date vary wildly in both approach and ultimate assessment: estimates of the likelihood a nuclear terrorist attack differ by up to nine orders of magnitude. This article critiques existing efforts from the standpoint of probability theory, and proposes an alternative perspective on the utility of risk assessment in this area. Nuclear terrorism is argued to be a 'virtual risk' for which it is not possible to meaningfully ascribe a quantitative measure, making numerical estimates of the likelihood of nuclear terrorism misleading. Instead, we argue that focus should be placed on utilising models to identify areas of disagreement as targets for further research, with greater emphasis on understanding terrorist decision-making and adaption in response to nuclear security measures.

Keywords

Nuclear Terrorism; Risk Assessment; Nuclear Weapons; Nuclear Security

Introduction

The fourth and final Nuclear Security Summit, held in the United States in March 2016, capped six years of sustained international effort to reduce the threat of nuclear terrorism. Over this period billions of dollars have been spent on a broad range of nuclear security initiatives. States have minimised their use of sensitive nuclear materials, in particular highly enriched uranium (HEU); increased the physical, information, and human security of nuclear and radiological facilities; developed and implemented new national nuclear security legislation; employed systems to detect nuclear material outside of regulatory control; and made preparations to mitigate the ultimate effects should an incident occur.[1]

Given this level of investment, significant attention has also been directed at measuring the effectiveness of these efforts and whether they have indeed served to reduce the risk of nuclear

---

 *Author's email: robert.downes@kcl.ac.uk
** Author's email: christopher.hobbs@kcl.ac.uk

[1] Sharon Squassoni, 'Outcomes from the 2014 Nuclear Security Summit', Centre for Strategic and International Studies Critical Questions (25 March 2014), available at: http://csis.org/publication/outcomes-2014-nuclear-security-summit} accessed 1 July 2016.

terrorism. Here there are several approaches that can be taken. Narrowly focused assessments might, for example, measure security culture improvements within a specific organisation as a result of targeted education and training programmes. Other analyses are broader and seek to assess the impact of multiple measures, for example the NTI Nuclear Material Security Index, which provides a national level public assessment of both the threat and nuclear security conditions within different states.[2] As the scope of assessment broadens, the complexity of the task increases due, in part, to the heterogeneous and covert nature of possible nuclear terrorist groups, their unique context, and the diversity of the security systems with which they interact.

Despite this complexity a number of mathematical models have been used to produce quantitative estimates of the likelihood of a nuclear terrorist attack. These have an intrinsic appeal to both the public and policymakers, serving to simplify complex problems and providing a seemingly scientific and straightforward way of assessing the effectiveness of programmes or policies.[3] However, serious problems can occur when there is no consensus on the models used or the estimates they produce. This issue is particularly acute in the case of nuclear terrorism, with expert estimates of the annual probability of an attack ranging from one in three billion to more than one in two.[4] Perhaps surprisingly, despite this clear disparity, the use of high-level numerical estimates of nuclear terrorism is widespread. They appear in public discourse, are referenced throughout the academic literature, and can be found in government testimony, thus influencing nuclear security decision-making at the highest levels.[5]

This article discusses the challenge of assessing the risk of nuclear terrorism, highlighting the serious pitfalls that are, we feel, an inevitability when carrying out numerical analysis. Beginning by introducing the process of risk assessment, the article outlines two broad interpretations of probability – the Frequentist and subjective Bayesian approaches – which can be used to produce numerical estimates of the likelihood of nuclear terrorism. Three different categories of risk are then considered, with nuclear terrorism demonstrated to be a 'virtual risk' that defies straightforward quantification. Existing approaches to modelling the likelihood of nuclear terrorism are then critiqued, with a particular focus on event tree analysis, a common and straightforward modelling approach. Finally, alternative perspectives on the utility of risk modelling and the useful insights they can bring when considering nuclear terrorism are presented.

Approaches to risk assessment

Although widely applied, risk is an elusive concept with divergent interpretations across different fields making a precise definition difficult. In general terms, risks are associated with certain events or activities and incorporate three distinct notions: *hazard*, *likelihood*, and *consequence*.[6] A hazard

---

[2] NTI Nuclear Security Index, 'Nuclear Threat Initiative', available at: {http://ntiindex.org/} accessed 10 February 2016.

[3] Theodore M. Porter, *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life* (Princeton: Princeton University Press, 1995), ch. 4.

[4] John Mueller, 'The Atomic Terrorist: Assessing the Likelihood', conference paper, Program on International Security (University of Chicago, 2008), p. 14; Graham T. Allison, *Nuclear Terrorism: The Ultimate Preventable Catastrophe* (London: Macmillan, 2004), p. 15.

[5] Nate Silver, 'Crunching the risk numbers', *Wall Street Journal* (8 January 2010); Graham T. Allison, 'Nuclear attack a worst-case reality?', *The Washington Times* (23 April 2008); Michael Crowley, 'Yes, Obama really is worried about a Manhattan nuke', *Time* (26 March 2014).

[6] Stanley Kaplan and B. John Garrick, 'On the quantitative definition of risk', *Risk Analysis*, 1 (1981), pp. 11–27.

exists as a source of danger; it is what can go wrong. Every hazard has a likelihood of occurring and, if indeed it does occur, a range of consequences will follow.

In essence, a risk analyst collects together a set of hazards, assigns likelihoods to their occurrence and then, assuming the events in question do occur, considers consequences attendant upon the range of possible occurrences. The 'risk' associated with a given system is a function of likelihood and consequence for this collection of potential hazards.[7] When dealing with complex phenomena such as terrorist activity we use the alternative terminology *scenario* as opposed to hazard. While a terrorist group is clearly a hazard, this language emphasises the specific context under consideration, that is, the group together with a range of possible actions that could be undertaken.

In the domain of nuclear terrorism a frequently discussed scenario is the terrorist acquisition of fissile material and subsequent fabrication of an Improvised Nuclear Device (IND).[8] For the purposes of risk analysis the details of such a scenario should be reasonably specific. For example, a group could purchase fissile material on the black market using contacts in Eastern Europe before shipping it to a safe haven in the Middle East through known drug trafficking routes. At a well-equipped facility a scientific team then engineers a crude nuclear device using specialist equipment and other materials that have been legitimately purchased on the open market. Finally, the group transports the IND to the target location, the capital city of a nearby country, and detonates the device.

Given this scenario, a risk analysis would consider the likelihood and consequences of terrorist success. Likelihood of success will depend upon myriad factors including, but not limited to, the financial arrangements of the group, the availability of fissile material of sufficient quality and quantity on the black market, the intelligence and nuclear security arrangements of a number of countries, and the technical ability of the assembled scientists, engineers, and technicians to successfully build a viable nuclear device. The consequences of a successful attack are similarly diverse and include the potential collapse of local governments, strain on regional alliances, and economic turmoil, alongside the significant physical and psychological suffering caused by the blast, fire, and fallout generated by the detonation. This short example highlights some of the complexities found in analysing the risk of nuclear terrorism and represents just one of the many possible scenarios that could reasonably be considered.[9]

When risk analysis is carried out numerically, relationships between a scenario, its likelihood, and its consequence must be specified. Commonly, analysts use the risk = likelihood x consequence equation, or similar.[10] Many quantitative studies concerning the risk of terrorism take this approach.[11] Quantified risk assessment relies on the determination of numerical values for likelihood – the probability of an event occurring, and its consequence – the relative severity of the event. As outlined above such scenarios do not lend themselves to easy quantification. This article explores the

[7] Many potential hazards may fall outside the ambit of our collective knowledge resulting in so-called *Black Swan* events, as discussed by Nassim Taleb in his book of the same name. See Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (London: Peguin, 2008).
[8] Peter D. Zimmerman and Jeffrey G. Lewis, 'The bomb in the backyard', *Foreign Policy* (16 October 2009), available at: {http://foreignpolicy.com/2009/10/16/the-bomb-in-the-backyard/} accessed 1 July 2016.
[9] Charles D. Ferguson and William C. Potter, *The Four Faces of Nuclear Terrorism* (London: Routledge, 2005), p. 5.
[10] Mathematically, this formulation presents risk as the 'expected consequence' of the hazard.
[11] Louis Anthony (Tony) Cox, Jr, 'Some limitations of "Risk = Threat x Vulnerability x Consequence" for risk analysis of terrorist attacks', *Risk Analysis*, 28 (2008), pp. 1749–61.
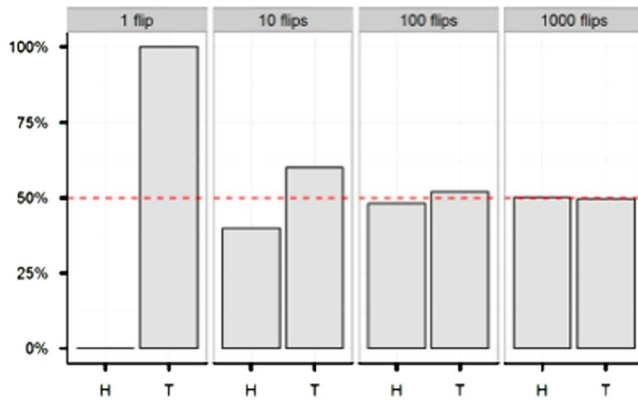
Figure 1. Frequency of heads and tails observed in the flipping of an unbiased coin for an increasing number of flips. The dashed red line indicates 50 per cent.

challenges associated with assessing the likelihood of nuclear terrorist events in quantitative terms. Assessments of the consequences of an act of nuclear terrorism are not addressed here, although have been discussed elsewhere.[12]

Frequentist and subjective Bayesian probability

When determining probabilities for events as part of a quantitative risk assessment, analysts generally adopt either the Frequentist or subjective Bayesian interpretation of probability. The former takes a 'frequency view of probability' in which a large number of independent repetitions of an identical statistical experiment are conducted to form a sample dataset. Probabilities for each possible outcome are calculated by dividing the frequency of each outcome by the total number of events contained in the dataset.[13] The Frequentist approach to probability estimation can be illustrated using a simple example, the flipping of an unbiased coin, a process that has two possible outcomes, heads or tails. Suppose we wish to determine the probability that a coin flip will result in a head. Frequentist probability tells us that we must conduct a large number of coin flips, measure the result each time, and calculate the share of heads and tails for different numbers of flips: 1 flip, 10 flips, 100 flips, etc. The results of such an experiment are given in Figure 1. As the number of flips increases, the proportion of heads approaches 50 per cent. Using this increasingly refined experimental data, the probability of a head resulting from the flip of an unbiased coin is determined to be 50 per cent.

As illustrated by our simple example, a crucial issue in utilising Frequentist probability is sample size. From Figure 1 it is clear that results based on a single flip would be misleading. In general, the larger

---

[12] While consequence analysis typically focuses on a single impact measure, such as economic damage expressed in dollars lost, there is increasing acknowledgement that impacts are multidimensional and should therefore be addressed in such terms. See, for instance, Bruno S. Frey, Simon Luechinger, and Alois Stutzer, 'Calculating tragedy: Assessing the costs of terrorism', *Journal of Economic Surveys*, 21 (2007), pp. 1–24.

[13] B. S. Everitt and A. Skrondal, *The Cambridge Dictionary of Statistics* (4th edn, Cambridge University Press, 2010), p. 174; Alan Hájek, '"Mises redux" – redux: Fifteen arguments against finite Frequentism', *Erkenntnis*, 45 (1996), pp. 209–27. Note that our designation of Frequentist probability could, more properly, be referred to as finite Frequentist probability. This approach demonstrates a wide range of mathematical and philosophical problems as an interpretation of probability despite its numerous enticements.

the sample size, the more confident one can be in using Frequentist probability (subject to certain caveats).[14]

In risk analysis there is often a great deal of useful data available to aid in the evaluation of the likelihood of a scenario, which allows Frequentist probability to be utilised. For example, analysis of past road traffic accident data shows that young male drivers are disproportionately likely to be involved in accidents in the United Kingdom.[15] This analysis relies on a comprehensive dataset to which Frequentist tools can be applied; the probability of an event occurring can be calculated directly from historical information. Various methods can be employed to demonstrate that, while Frequentist probability may change slightly year-to-year, the overall trend remains fundamentally the same. Consequently, young male drivers, who are a greater accident risk than other demographics, pay higher insurance premiums.

Frequentist probability is essentially the mathematical description of a sample dataset to draw inferences concerning the phenomenon under study. Following this approach any analyst provided with the same dataset and utilising the same analytical tools would produce, in theory, identical results – in this sense, Frequentist probability may be thought of as objective.[16] The Frequentist approach contrasts with that adopted by subjective Bayesian probability that *explicitly* incorporates analyst subjectivity.

In general, there is no standardised process for generating subjective Bayesian probability judgements: they are based on the experience of the assessor and represent a *degree of belief* in the likelihood of occurrence of the event at hand.[17] In cases where information is limited and the Frequentist approach to probability cannot be usefully applied, as the dataset is too small to be representative of the phenomenon under study, for example, subjective Bayesian approaches incorporating expert opinion can be used to produce probabilistic judgements.[18] Robert Winkler and Allan Murphy emphasise that the 'subjective [Bayesian] framework does not admit the existence of a "correct" (in the sense of universal) probability'.[19] Instead, subjective probabilistic judgements need only be self-consistent, that is, obey the laws of probability. (For instance, the probability across all possible outcomes of a particular event must sum to 100 per cent.) As this form of probabilistic judgement relies on an investigator's knowledge and prior beliefs, an 'individual making [a subjective Bayesian] assessment will be coherent but [we] cannot force consensus between two different analysts'.[20]

---

[14] Ibid.

[15] Alan F. Williams and Veronika I. Shabanova, 'Responsibility of drivers, by age and gender, for motor-vehicle crash deaths', *Journal of Safety Research*, 34 (2003), p. 527.

[16] Porter, *Trust in Numbers*, ch. 4.

[17] In the context of Bayesian Probability, the same observation can be applied to the generation of the prior probability distribution, although there are guiding principles that can be applied in this case such as the Jeffreys Prior, see D. V. Lindley, 'The use of prior probability distributions in statistical inference and decisions', *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* (Berkeley: University of California Press, 1961), pp. 453–68.

[18] Tim Bedford and Roger Cooke, *Probabilistic Risk Analysis: Foundations and Methods* (Cambridge: Cambridge University Press, 2001), ch. 10.

[19] Robert L. Winkler and Allan H. Murphy, '"Good" probability assessors', *Journal of Applied Meteorology*, 7 (1968), p. 751.

[20] George Apostolakis, 'The concept of probability in safety assessment of technological systems', *Science*, 250 (1990), p. 1359.

Mathematically, Bayesian probability is a procedure for revising and updating probability estimates in light of new evidence. This requires the assessor to first specify their beliefs about an event in quantitative terms as a prior probability distribution. Bayes' Theorem is then applied to derive the posterior probability distribution taking into account new evidence or data conditioned on this prior. This mathematical theory 'has two main elements: the use of the laws of probability as coherence constraints on rational degrees of belief … and the introduction of a rule of probabilistic inference' based on the revision of these rational degrees of belief using new data.[21] In this article reference to the Bayesian approach to probability indicates the former aspect of the theory – coherence constraints on rational degrees of belief. For clarity, from this point we designate this as the subjective Bayesian interpretation of probability. This is the approach adopted by the quantitative studies of nuclear terrorism that are considered later in the article.

It must be noted that interpretations of probability are not neatly divided into objective and subjective camps. It is uncontroversial to note that subjectivity wends its way into many scientific activities perceived as being objective: readers with an interest in this are directed to Matthews' discussion on the subject which, in part, deals with subjectivity in Frequentist methods of statistical inference.[22] As one simple example, Aven et al. have described that to 'define the frequentist probability … we have to construct a population of similar situations. This can however be done in many different ways. There is no objective approach for making this mental construction.'[23] Furthermore, our distinction between Frequentist and subjective Bayesian interpretations has been carried out with a view to expounding the challenges of probabilistic estimation relating to nuclear terrorism in particular. We are careful to work only with subjective Bayesian interpretations as defined above and, again, emphasise that we do not refer to the general mathematical Bayesian interpretation of probability.

These caveats aside, the crucial difference between subjective Bayesian and Frequentist probability is neatly expressed by Michael Goldstein: 'We have moved away from a traditional view of [Frequentist] analysis, which attempts to express what we may learn about some aspect of reality by analysing an individual dataset. Instead, the [subjective] Bayesian analysis expresses our current state of belief based on combining information from the data in question with whatever other knowledge we consider relevant.'[24] While Frequentist probability is appropriate for a wide range of investigations, there are many problems that are not easily dealt with by analysing past frequency of occurrence or through an extensive sampling process (as in the case of flipping a coin). Daniel Kahneman and Amos Tversky offer three pertinent examples: 'What are the chances that *this* 12-year-old boy will grow up to be a scientist? What is the probability that *this* candidate will be elected to office? What is the likelihood that *this* company will go out of business?'[25] Such questions are effectively unique: there are no realistic datasets that will shed light on such questions and, therefore, Frequentist tools are inapplicable. These are cases where a subjective Bayesian approach to probability can still enable quantitative probabilistic judgements to be made, albeit subject to certain limitations.

---

[21] William Talbott, 'Bayesian Epistemology', *The Stanford Encyclopedia of Philosophy*, Winter (2016).

[22] Robert A. J. Matthews, 'Fact Versus Fiction: The Use and Abuse of Subjectivity in Scientific Research', European Science and Environment Forum Working Paper, No. 2 (1998).

[23] Terje Aven, Ortwin Renn, and Eugene A. Rosa, 'On the ontological status of the concept of risk', *Safety Science*, 49 (2011), p. 1077.

[24] Michael Goldstein, 'Subjective Bayesian analysis: Principles and practice', *Bayesian Analysis*, 1 (2006), p. 408.

[25] Daniel Kahneman and Amos Tversky, 'Subjective probability: a judgement of representativeness', in C.-A. Staël von Holstein (ed.), *The Concept of Probability in Psychological Experiments* (Dordrecht: D. Reidel Publishing Company, 1974), p. 44, emphasis added.

When seeking to analyse the risk associated with a rare or unique event for which data is either extremely limited or does not exist, the subjective Bayesian approach is often adopted. Analysts provide their own subjective Bayesian judgement concerning the probability of certain events. This analysis incorporates a wide range of data pertaining to similar but ultimately different situations. Probabilities for these similar events are estimated using a Frequentist approach and are used to create what Stanley Kaplan and B. John Garrick call a *bureau of standards* – an analyst's calibration scale against which a particular scenario can be compared.[26] Analysts use the bureau of standards as a guide when expressing a subjective Bayesian probability of occurrence for a given scenario. However, there is no formal process for drawing upon the bureau of standards and, hence, the extent to which different bureau components affect probability judgements will vary between analysts. Following Kaplan and Garrick, a Bayesian analysis then consists of two steps:[27]

1. **Constructing a bureau of standards.** A good analyst will have high *substantive goodness*, that is, demonstrable and relevant experience in the domain in which assessments are being made;[28]

2. **Drawing on the bureau of standards to produce a probability estimate.** A good analyst will have high *normative goodness*, that is, expertise in probabilistic assessment with which to encode subjective assessments in a coherent quantitative manner that obeys the laws of probability.[29]

While normative goodness can be enhanced through expert calibration training,[30] we may be faced with a scenario for which there exists no relevant frequency data to guide probabilistic extrapolation. The bureau of standards may be effectively empty or the data contained within of only tangential relevance to the scenario under consideration. An analyst will therefore have low substantive goodness when dealing with such a scenario. This crucial gulf between relevant frequency data within our bureau of standards and the scenario in which we are expressing a degree of belief is referred to in this article as the *knowledge gap*.[31] Note that, while normative and substantive goodness are concepts taken from an established literature, the term knowledge gap has been coined here for explanatory purposes.

For clarity, we make a clear distinction between the notions of substantive goodness and knowledge gap. The former relates to the properties of the assessor while the latter relates to the scenario in question. A scenario with a large knowledge gap is ill suited to probabilistic assessment. Regardless of the qualities of the analyst the nature of the scenario militates against the formation of a relevant bureau of standards. It is therefore impossible for any analyst to have high substantive goodness in such a context. For a scenario with a large knowledge gap, all analysts operate with low substantive goodness. No action on the part of the analyst can improve this situation.

While there are intrinsic challenges in applying subjective Bayesian methods, they have been successfully utilised in many areas, for example, in designing risk prediction models for treating common diseases.[32]

---

[26] Kaplan and Garrick, 'On the quantitative definition of risk', p. 18.

[27] Ibid.

[28] Winkler and Murphy, '"Good" probability assessors', p. 752.

[29] Ibid.

[30] Bedford and Cooke, *Probabilistic Risk Analysis*, pp. 191–217.

[31] An alternative name for the knowledge gap could be *epistemic goodness*, the extent to which relevant knowledge about the hazard or scenario under consideration is knowable.

[32] Andrew H. Briggs, Ron Goeree, Gord Blackhouse, and Bernie J. O'Brien, 'Probabilistic analysis of cost-effectiveness models: Choosing between treatment strategies for gastroesophageal reflux disease', *Medical Decision Making*, 22 (2002), p. 291; Naresh A. Dewan, Christopher J. Shehan, Steven D. Reeb, Lisa S. Gobar,

Table 1. A comparison of Frequentist and Bayesian approaches to probability.

Approach	Frequentist	Bayesian
Data input	Large high quality dataset	Limited or no dataset
Opinion input	No	Yes
Output	Quantitative description of dataset	Quantitative description of degrees of belief
Judgements	Objective (analysts produce same result)	Subjective (varies between analysts)

A simple description of the major differences between Frequentist and Bayesian approaches to probability estimation is summarised above in Table 1.

Subjective Bayesian and Frequentist approaches in security studies

Both subjective Bayesian and Frequentist approaches have been used to make predictions within the field of security studies. For example, Aaron Clauset et al. apply a Frequentist approach to explore the severity of terrorist acts.[33] The authors take as their dataset the National Memorial Institute for the Prevention of Terrorism database from 1968 to 2006 containing approximately thirty thousand acts of terrorism. Eleven thousand of these resulted in at least one person being injured or killed. By calculating frequency-severity distributions for this dataset, the authors identify a 'scale invariant' inverse power law relationship between the frequency of terrorist events and their severity in terms of death and injury. The strength of this analysis and its applicability to future acts of terrorism rests on two key assumptions. Firstly, that each event in the dataset is independent of any other, as required by the Frequentist approach. Analogous to the earlier coin-flipping example, where it was assumed that each flip is independent of every other flip. Secondly, that the dataset is sufficiently representative of terrorism as a phenomenon. This enables the authors to state, based on their analysis, that there is a 'global pattern in the frequency statistics of terrorist attacks' characteristic of terrorism itself due to the comprehensive nature of their underlying dataset.[34]

Given the complex real-world problems encountered in security studies, a lack of directly relevant data can often pose a serious barrier to the use of Frequentist analysis. In these situations subjective Bayesian approaches can be used to analyse effectively unique scenarios. For example, they have been used to estimate the likelihood of terrorist attacks under a variety of circumstances 'when existing information is vague or uncertain yet may lead experts to deduce probabilities that can later be updated as new information becomes available'.[35] Mixed methods may be employed as the mathematical theory of Bayesian probability 'allows for incorporation of subjective probability judgments into assessments that may include frequentist calculations'.[36] However, although the subjective Bayesian approach allows for the incorporation of subjective opinion when real-world data is lacking, care must be taken when quantitative probabilistic judgements are made. Some analysts have urged caution when discussing the role subjective Bayesian probabilities have played in

Walter J. Scott, and Kay Ryschon, 'Likelihood of malignancy in a solitary pulmonary nodule: Comparison of Bayesian analysis and results of FDG-PET scan', *Chest*, 112 (1997), pp. 416–22.

[33] Aaron Clauset, Maxwell Young, and Kristian S. Gleditsch, 'On the frequency of severe terrorist events', *Journal of Conflict Resolution*, 51 (2007), pp. 58–87.

[34] Ibid., p. 64.

[35] Henry H. Willis, Tom LaTourrette, Terrence K. Kelly, Scot Hickey, and Samuel Neill, *Terrorism Risk Modeling for Intelligence Analysis and Infrastructure Protection* (RAND, 2007), p. 5.

[36] Ibid.

predicting terrorist threats, believing that these 'should not be discussed in terms of probability because historical data does not exist with which to perform actuarial calculations of event frequencies'.[37]

Nuclear terrorism as virtual risk

To date the malicious use of nuclear and radiological materials by non-state actors has been limited and there have been no large-scale incidences of nuclear terrorism. As a result it is not possible to adopt a Frequentist approach to risk estimation. When examining the likelihood of a major terrorist attack such as the detonation of an IND by a terrorist group, it is therefore necessary to consider the application of subjective Bayesian methods. However, as emphasised above, the validity of this approach for the scenario under consideration hinges upon the size of the knowledge gap. It must be possible for an analyst to construct a bureau of standards with high substantive goodness for the resulting probability judgements to be anything other than highly speculative. In exploring the extent of the knowledge gap for nuclear terrorism it is illustrative to draw on John Adams's three kinds of risk typology wherein risks are categorised as *directly perceptible*, *perceived through science*, and *virtual*.[38]

Directly perceptible risks are typically managed 'instinctively and intuitively', alleviating the need for a formal risk assessment.[39] As a simple example, when a person experiences hazardous sensory input, the withdrawal reflex attempts to arrest bodily damage by retreating from the causal sensory stimuli, without the need for thought or higher brain function. A second class of risks are only perceptible thanks to the development and application of scientific theory and tools, requiring an interpretive third party (the scientist) to mediate the relationship between individual and hazard – these risks are perceived through science. Water-borne diseases such as cholera are canonical examples, being evidentially accessible only to medico-scientific personnel with domain-specific expertise.[40] Applying sophisticated statistical techniques to large volumes of pertinent data, highly-trained individuals can identify at-risk sub-communities, behaviours, and paths to disease propagation. The collection, collation, and analysis of data clearly lies within the purview of technical specialists, although the implications are often far-reaching with consequences for entire populations. Accordingly, the evidence-based (often in the Frequentist sense) development of policies to mitigate against disease outbreak and manage risks is often placed in the hands of government agencies and the broader healthcare community, rather than devolved to the individual.[41]

[37] Ibid.

[38] John Adams, 'Risk and morality: Three framing devices', in Richard Victor Ericson and Aaron Doyle (eds), *Risk and Morality* (University of Toronto Press, 2003), pp. 87–104.

[39] Ibid., p. 87.

[40] Mary K. Kindhauser (ed.), *Communicable Diseases 2002: Global Defence Against the Infectious Disease Threat* (World Health Organisation, 2002), available at: {http://apps.who.int/iris/bitstream/10665/42572/1/9241590297.pdf} accessed 8 July 2016.

[41] In the United Kingdom, the response to mid-nineteenth century cholera epidemics followed precisely this pattern. Popularly attributed to the revolutionary statistical epidemiology of Dr John Snow, who painstakingly identified a contaminated communal water pump handle as the source of a central London outbreak, the realisation that cholera spread via the faecal-oral route forced the authorities to act. Legislative changes followed and regulations governing public health were disseminated by the General Board of Health that required local authorities to 'provide dispensaries operating around the clock with sufficient medical aid to treat cholera patients' among other stipulations, see John Snow, *On the Mode of Communication of Cholera* (John Churchill, 1855); Donald Cameron and Ian G. Jones, 'John Snow, the Broad Street pump and modern

A similar situation obtains in the security domain when forecasting risks arising from conventional terrorism. There exist substantial and complex datasets upon which quantitative analysis can be performed to make predictions and to inform risk-minimising actions at the tactical, operational, and strategic levels, particularly in the distribution of resources for defensive systems. For example, the installation of metal detectors at airports has been shown to reduce instances of aircraft hijacking, although the overall effect of this policy is effectively neutral once substituted terrorist activities are taken into account.[42] While it may appear obvious that metal detection technology will decrease hijacking, a rigorous cost-benefit analysis is needed to determine whether such a preventive policy delivers sufficient risk-minimising value compared with alternatives. The authority and skills required to conduct such a study are beyond the purview of officials at individual airports, and the financial, legislative, political, and social implications of such a policy generally require decision-making at both the national and international level.[43]

Virtual risks are a final category in this typology and are characterised by a paucity of data – it is not possible to claim to have sufficient information to meaningfully ascribe probabilities and, as such, no analyst can attain high substantive goodness. The knowledge gap between calibration data and the event in which an analyst is expressing a degree of belief is simply too great. In comparison with risks perceived through science, the critical difference is the size of the knowledge gap, which is considerably larger for virtual risks. A characteristic example of a virtual risk is the invasion of Earth by hostile aliens. Famously, American astronomer Frank Drake proposed his eponymous equation in the early 1960s, which, through a probabilistic argument, has 'helped guide speculation about the likelihood of intelligent extra-terrestrial life' contacting humanity.[44] However, extant frequency data is non-existent. No analyst could make a meaningful judgement about this risk because, from humanity's current vantage, the relevant data is (currently) unknowable. Information is increasingly available concerning the existence of planets outside the solar system that could, in theory, support life and some scientists have attempted to quantify the likelihood of existence of intelligent extra-terrestrials. However, even with access to a comprehensive catalogue of life-supporting planets throughout the galaxy, there exists no precise way of estimating whether they are populated by intelligent beings, let alone whether such beings could cross the vastness of space should they actually exist.[45] If such an estimate could be made, we would still lack a crucial piece of the puzzle: are the intentions of the alien species likely to be hostile? Even with a means of estimating the distribution of intelligent life across the galaxy, any statement concerning the alien's attitude towards humanity can be nothing but highly speculative without additional sources of data. The large numbers of unknowns in this case militate against the formation of a meaningful quantitative likelihood estimate and, hence, there is 'no realistic way to evaluate one prediction against another'.[46]

Virtual risks are extremely difficult to analyse and any attempt at quantification represents ungrounded numerical speculation. As Adams makes clear, for virtual risks the 'veneer of scientific

epidemiology', *International Journal of Epidemiology*, 12 (1983), pp. 393–6; S. L. Kotar and J. E. Gessler, *Cholera: A Worldwide History* (Jefferson, NC: McFarland and Company, 2014), pp. 151–60.

[42] Walter Enders and Todd Sandler, 'The effectiveness of antiterrorism policies: a vector-autoregression-intervention analysis', *The American Political Science Review*, 87 (1993), pp. 829–44.

[43] Preventive Security Measures, Convention on International Aviation, Annex 17, 'Security: Safeguarding International Civil Aviation Against Acts of Unlawful Interference' (International Civil Aviation Organisation, 2014).

[44] Mark J. Burchell, 'W(h)ither the Drake equation?', *International Journal of Astrobiology*, 5 (2006), pp. 243–50.

[45] SETI Institute, 'The Drake Equation', available at: {http://www.seti.org/node/434} accessed 2 February 2016.

[46] Burchell, 'W(h)ither the Drake equation?', p. 249.

authority imparted by quantified probability often can withstand little scratching'.[47] However, this does not necessarily stop analysts from conducting numerical analysis, drawing on what little prior knowledge and experience may be available. Numerical assessment of virtual risks is likely to be characterised by strong disagreement among analysts resulting from the considerable extrapolation from incomplete or partially relevant data. According to Adams it is common for 'reputable scientists [and analysts to] contend with one another'.[48] Estimates produced diverge to the point that a coherent picture regarding purported risks cannot be divined and 'convictions, prejudices, and superstitions' form the basis for analysis.[49] If quantitative tools are applied, this disagreement will manifest in significant disparities between and understandings of numerical estimates.

Under Adams's typology nuclear terrorism fits squarely within the virtual risk category. In terms of prior experience there are no large-scale incidences upon which to draw. Although relevant information does exist, this is often limited, incomplete, open to interpretation, or only applicable to a small part of the overall puzzle. For example, when considering whether a terrorist group could acquire nuclear material, an essential part of many nuclear terrorism scenarios, it is instructive to consider past thefts. Here, though, there are just a handful of malicious cases in the public domain upon which to draw in analysing precisely how nuclear materials can be stolen, an activity about which it is difficult to generalise due to the diversity of states, facilities, measures, materials, and actors involved. This lack of sufficient detail limits the calibration of probability estimates regarding success or failure of potential adversaries.[50]

Nevertheless, attempts to deduce probabilities either directly from or referencing the historical record are not uncommon. A widely-read study by Matthew Bunn illustrates these issues clearly while offering a nuanced understanding of the challenges facing analysts engaging in numerical assessment. Bunn assesses the probability of a terrorist group choosing to pursue black market acquisition of nuclear material as 'fairly large', equating to an approximate 30 per cent annual likelihood, and assigns the probability of a group being successful in an attempt at 20 per cent.[51] These estimates are based upon the observation that '[b]oth Aum Shinrikyo and al Qaeda have pursued this method of acquisition.'[52] In this case the bureau of standards is extremely limited, compromising just two groups, neither of which were successful in procuring nuclear material – a limitation acknowledged by the author.

Assessing terrorist intentions to carry out acts of nuclear terrorism is similarly challenging. While some analysts point to terrorist statements regarding Weapons of Mass Destruction, others pass this off as empty rhetoric, arguing that the same groups will be deterred from acquiring and using such weapons for fear of alienating their political base, diminishing their meagre financial resources, and

---

[47] Adams, 'Risk and morality', p. 92.

[48] Ibid.

[49] Ibid.

[50] There are likely to have been many more unrecorded or unreported cases, with political or other reasons acting as a constraint upon national authorities from sharing such sensitive information. The International Atomic Energy Agency's Incident and Trafficking Database (ITDB) is a case in point: event reporting is voluntary and therefore the completeness of the database is open to question. While statistical tools can be applied to the ITDB, this must be done in the knowledge that the results are inherently limited for this reason.

[51] Matthew Bunn, 'A mathematical model of the risk of nuclear terrorism', *The Annals of the American Academy of Political and Social Science*, 607 (2006), pp. 103–20.

[52] Ibid., p. 113.

risking their long-term survival through reprisal action by the states they target.[53] These statements cannot be interpreted in an unambiguous manner and offer scant basis for numerical reasoning.

Efforts to quantify the risk of nuclear terrorism

Despite intrinsic challenges in assessing the risk of nuclear terrorism, a substantial literature has developed over the past two decades driven by the belief that 'the danger of high-end terrorism is growing'.[54] Analysts such as Brian Jenkins observe a shift in terrorist strategy: formerly terrorists wanted 'a lot of people watching, *not* a lot of people dead … [whereas increasingly] terrorists want a lot of people watching *and* a lot of people dead'.[55] The geographical spread of nuclear materials and expertise has also increased which, in the eyes of many, has served to lower the technical barriers groups would need to overcome to perform acts of nuclear terror.[56]

Contemporary studies are largely focused on providing a qualitative analysis of the relative risk between different options open to terrorist groups.[57] Some of these are generic, while others focus on the intentions and efforts made by specific terrorist groups.[58] Also commonly explored are psychological factors and structural aspects of groups that might pursue nuclear terrorism, the routes they might take, and how they could be thwarted.[59] Within this literature there is a clear division between analysts that believe an act of nuclear terrorism to be an imminent threat and those who see such an attack as a highly unlikely event. Heated exchanges between scholars from both ends of this spectrum have arguably served to polarise this field, contributing to an impasse within the analytical community.[60] Numerical estimates have been mobilised in support of different viewpoints, with Graham Allison and John Mueller assessing the annual likelihood of nuclear terrorism as 'more likely than not' (that is, upwards of one in two) and more than one in three billion, respectively.[61]

In modelling terrorism and, in particular, nuclear terrorism, event tree models have been widely utilised. This conceptually simple method has been applied to assess the reliability of nuclear reactors and is commonly used in the context of Probabilistic Safety Analysis or Probabilistic Risk Analysis for large engineering projects.[62] The development of probabilistic tools for risk

---

[53] Ferguson and Potter, *The Four Faces of Nuclear Terrorism*, pp. 33, 194–5.

[54] Ibid., p. 4.

[55] Brian Michael Jenkins, 'The new age of terrorism', in Brian Michael Jenkins (ed.), *McGraw-Hill Homeland Security Handbook* (RAND, 2006), ch. 9.

[56] Benjamin Cole, *The Changing Face of Terrorism* (London: I. B. Tauris, 2010).

[57] Ferguson and Potter, *The Four Faces of Nuclear Terrorism*, ch. 1.

[58] Rolf Mowatt-Larssen, 'The Armageddon test: Preventing nuclear terrorism', *Bulletin of the Atomic Scientists*, 65 (2009), pp. 60–70.

[59] Brecht Volders and Tom Sauer (eds), *Nuclear Terrorism: Countering the Threat* (London: Routledge, 2016); Michael Levi, *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2009).

[60] After receiving criticism for 'alarmist' views on the issue, Peter Zimmerman famously included a section headed 'John Mueller: Pollyanna?' in his paper, 'Do we really need to worry? Some reflections on the threat of nuclear terrorism', *Defence Against Terrorism Review*, 2:2 (2009), pp. 1–14. In *Achieving Nuclear Ambitions: Scientists, Politicians, and Proliferation* (Cambridge: Cambridge University Press, 2012), Jacques Hymans notes that Allison 'cites – without irony – an analysis conducted by science fiction writer Tom Clancy' as part of an argument showing terrorists or even individuals could self-produce fissile materials for inclusion in an IND.

[61] Allison, *Nuclear Terrorism*, p. 15; Mueller, 'The Atomic Terrorist', p. 14.

[62] William Keller and Mohammad Modarres, 'A historical overview of probabilistic risk assessment development and its use in the nuclear power industry: a tribute to the late Professor Normal Carl Rasmussen', *Reliability Engineering and System Safety*, 89 (2005), pp. 271–85.

analysis and their application to nuclear engineering has a long and complex history, which has reflexively shaped nuclear safety, a process culminating in the famed 1975 WASH-1400 study.[63] When applied to terrorism, a group's behaviour is broken down into a sequence of actions or decisions, each of which has an associated probability. An initiating event, often the decision of a terrorist group to undertake an attack, is followed by all possible realisations of this decision point and all subsequent contributory decisions laid out in a sequence or tree.[64] Each tree ends at a terminal event; in the case of nuclear terrorism this could be the successful detonation of an IND. As 'the probability of each event is displayed conditional on the occurrence of events that precede it in the tree, the joint probability of the intersection of events that constitute a sequence (or "scenario") is found by multiplication' along the sequence in question.[65] This is a trivial calculation that further supports the widespread application of event tree models by limiting their technical content. Event trees have been the subjects of intense academic study and a growing literature concerning their construction and operation is now available and accessible to the lay reader.[66]

In the context of nuclear terrorism, an attack scenario is broken down into constituent sequential steps. These are then analysed from the perspective of the adversary, with the probability of a terrorist group completing each step considered in turn.[67] A subset of these steps might include an insider first defeating facility security systems to acquire nuclear material, bypassing detection systems to illicitly remove the material from the facility, and overcoming technical challenges to fabricate a viable device, before transporting the device to a target, and successfully detonating. The probability of this scenario occurring is determined by multiplying the probabilities of success of each sequential step along a given tree, from the initiator to the terminus.

In the aforementioned study carried out by Bunn, an event tree model is used to explore the risk of nuclear terrorism, 'mak[ing] explicit the assumptions about the key factors affecting the risk and provid[ing] a tool for assessing the effectiveness of alternative policies'.[68] It estimates the global likelihood of an act of nuclear terrorism as '29 percent probability … in the next decade' and identifies the security of fissile material as a key chokepoint upon which policy efforts should be focused.[69]

However, despite their conceptual and operational simplicity, there are serious problems with the way in which event tree models have been applied to terrorism and, in particular, nuclear terrorism. These stem largely from the way they serve to simplify terrorist group's decision-making processes, failing to treat terrorists as intelligent and adaptive adversaries, with probabilities regarding group behaviour required as model inputs as opposed to outputs. These weaknesses are discussed in more detail below.

---

[63] *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants* (US Nuclear Regulatory Commission, 1975).

[64] Barry C. Ezell et al., 'Probabilistic risk analysis and terrorism risk', *Risk Analysis*, 30 (2010), pp. 575–89.

[65] Elisabeth Paté-Cornell, 'Fault tree vs. event trees in reliability analysis', *Risk Analysis*, 4 (1984), pp. 177–86.

[66] For a particularly accessible introduction to the theory, see ibid.

[67] Bunn, 'A mathematical model of the risk of nuclear terrorism', pp. 104–6; Mueller, 'The Atomic Terrorist', p. 14. Levi stands in contrast to this approach, for example, by analysing the set of defensive measures as a whole, that is, as a layered defensive system.

[68] Bunn, 'A mathematical model of the risk of nuclear terrorism', p. 103.

[69] Ibid.

Terrorists are intelligent and adaptive

Event trees were designed to study large engineering projects incorporating many interdependent physical components. While this type of system may be complicated in structure, the relevant characteristics of individual components (for example, failure rates) can be readily described by static probability distributions and easily encoded into the relevant event tree branches. This is not the case when describing terrorist decision-making, which is a dynamic process wherein groups adapt their tactics and strategy in response to various external and internal stimuli, such as the introduction of new security measures by state adversaries or changes in the makeup of group leadership, respectively. Studies have shown that terrorist groups 'learn from experience, adapting their strategies and practices in response to information and feedback'.[70] In the case of hijacking, terrorist groups responded to metal detection systems at airports by substituting their activities for others, such as increased hostage taking and assassination attempts, which ultimately proved as costly in terms of civilian lives lost.[71] This is problematic not least because, from a decision-makers perspective, understanding both intended and unintended impacts of different choices is essential in the evaluation of any potential policy change.

A characteristic example in the domain of nuclear security is the installation of border monitoring systems designed to detect nuclear and radiological materials outside of regulatory control. The existence of such a system will impact upon the nascent trafficker's decision to illicitly transport materials across detector-enabled borders, particularly if attention is drawn to the existence and capabilities of deployed systems.[72] Effective modelling tools should recognise the responsive nature of terrorist or criminal groups to this information and their ability to adapt dynamically to a changing environment.[73] Modelling efforts to assess the probability of nuclear terrorism, and those utilising event trees in particular, have not incorporated this essential dynamic component accounting for the interplay between terrorist groups and their adversaries. Bunn clearly acknowledges this difficulty in his study, stating that 'intelligent and adaptive adversaries may react to security upgrades not by giving up but by increasing their capabilities'. However, this action-reaction dynamic is not captured by extant models as this functionality is largely absent from the basic event tree modelling toolkit.[74]

Probabilities regarding terrorist behaviour should be model outputs, not inputs

Numerical estimates of the likelihood of nuclear terrorism typically ascribe probabilities to the decision to engage in an act of nuclear terrorism and the likely success of such an endeavour. Considered through the event tree lens, this situation applies to each step of the nuclear terror process. When determining these probabilities, subject matter experts are currently required to predict how adversaries

---

[70] Michael Kenney, 'From Pablo to Osama: Counter-terrorism lessons from the war on drugs', *Survival*, 45 (2003), pp. 187–206.

[71] Enders and Sandler, 'The effectiveness of antiterrorism policies'.

[72] For instance, via a joint US-Russia blue ribbon ceremony inaugurating the first US Second Line of Defense border monitoring system in place at Moscow's Sheremetyevo airport, see Lara Cantuti and Lee Thomas, 'Second Line of Defense Program', paper presented at The Institute of Nuclear Materials Management (US Department of Energy, 1999), p. 3.

[73] David P. Morton, Fend Pan, and Kevin J. Saeger, 'Models for nuclear smuggling interdiction', *IIE Transactions*, 39 (2007), pp. 3–14.

[74] One notable exception is the study by Morton, Pan, and Saeger where resource allocation is dynamically modelled in the face of a (potentially) unknown number of adversaries.

will behave *a priori*, providing such figures as model inputs. As emphasised by the US National Research Council in their report on the misapplication of event tree analysis in the bioterrorism domain, for this approach to be valid, 'the subject-matter experts must grasp nuances of alternatives and outcomes and render opinions founded on an analysis of the entire decision process' of a terrorist group.[75] For an analyst making a probabilistic assessment, an event tree thus presents a problem: in order to assess the behaviour of the group at each step it is necessary to take account of the entire decision-making process of the group, including scenario-specific tactical choices and broader strategic intentions. However, the insights into the entire decision-making process of nuclear terrorist groups are precisely the desired outputs of the modelling efforts themselves.

For example, to assess the probability of a terrorist group successfully procuring fissile material through a specified channel it is necessary to understand the group's overall strategy regarding their planned act of nuclear terror. Given the significant resources that must be accorded to a nuclear terrorist endeavour, any group engaging in such activity may well formulate a clear (albeit malleable) plan of action, taking into account their own capabilities, the barriers they will face, their strategic aims, and their own perceptions regarding success or failure. This analysis, fully synthesised, will form a basis for the plan of action a group will likely adopt that will guide subsequent behaviour. This leads to the aforementioned logical inconsistency in the application of event tree models: in order to determine the probability of, for example, a terrorist group attempting to acquire fissile material through the black market, it is necessary to understand the group's broader strategy pertaining to acts of nuclear terror, which itself is exactly what the model attempts to determine. As the US National Research Council study makes clear, '[f]or decision problems as complex as those motivating [this study], the assessment of the probabilities that adversaries will choose courses of action should be the *outputs* of analysis, not required *input parameters*.'[76]

Event trees overly simplify nuclear terrorism

Event tree models are overly prescriptive in structural terms to the point at which their use militates against a realistic representation of terrorist behaviour. Between the initiator-terminus extremes, events are laid out in a linear chain with a prescribed and non-negotiable ordering. This rigidity goes against the evidence available in both the nuclear terror domain and in broader discussions of highly-engineered systems designed and produced by illicit groups. That terrorists will undertake a clear, linear chain of actions to achieve their goals is an *a priori* assumption that is not compatible with available albeit limited evidence.[77] This does not mean that terrorist groups operate without a clear overarching strategy. Rather, it suggests that strategic aims can be fulfilled in a non-linear way, and that the best-laid plans often go awry.

In the application of event trees, the characteristics of terrorist groups that would pursue nuclear terrorism are considered to be generic. This is an unrealistic simplification highlighted by Michael Levi who makes clear that '[r]ather than assuming a single model of skill and capability building, an intelligent defensive strategy will prepare to take advantage of a wide range of

---

[75] US National Research Council Committee on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis, 'Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change' (National Academies Press, 2008), p. 27.

[76] Ibid.

[77] Randy Borum, 'Understanding terrorist psychology', in Andrew Silke (ed.), *The Psychology of Counter-Terrorism* (London: Routledge, 2010).

terrorist approaches.'[78] This statement is reinforced by a cursory examination of two high-profile groups that have in the past considered the possibility of nuclear terrorism, Aum Shinrikyo (now defunct) and al-Qaeda. The former, a doomsday cult, was based largely in a First World country with attendant security and intelligence services;[79] the latter operates transnationally, in some cases in fragmented states without effective security forces, at other times receiving direct state support.[80] These two groups and others willing to engage in nuclear terrorism differ significantly in terms of their aims, motivations, structures, financial arrangements, and openness to external influence, and so warrant a distinct assessment taking these differences into account.

This situation obtains in numerous quantitative studies of nuclear terrorism. For instance, in one survey-based study, 75 per cent of respondents reported the black market route as the most likely pathway for terrorist acquisition of nuclear material.[81] However, this judgement of terrorist group behaviour is conditional upon a wide range of group- and scenario-specific assumptions and factors, none of which are made explicit in the analysis, which thus renders the figure meaningless as a descriptor of terrorist behaviour.

Risk models and nuclear terrorism: a new perspective

The preceding sections outline the difficulties associated with numerical risk assessments of nuclear terrorism. Efforts to date are centred on overly simplistic models that fail to take account of terrorist intelligence and adaptation in response to measures designed to mitigate the risks of terrorism, and offer wildly divergent predictions. In operation, model input requirements demand analysts to make judgements of terrorist behaviour that, themselves, should be the output of well-defined modelling processes. However, by adopting an alternative perspective on the utility of risk modelling it is possible to obtain useful insights relating to nuclear terrorism.

Conventional wisdom suggests that modelling has two purposes, to explain or predict some portion of the real world.[82] Probabilistic risk models are predictive in that they attempt to anticipate future events. Those negative events considered most likely to occur become the focus for measures designed either to diminish the likelihood of occurrence or to mitigate negative consequences. This has been the case to date when modelling nuclear terrorism, with emphasis placed on determining probabilities for the purpose of policy formation. The crucial obstacle here is that nuclear terrorism is a virtual risk and, therefore, attempts to apply subjective Bayesian probabilistic estimates are subject to a large knowledge gap. However, models can perform a far wider range of roles beyond these two basic functions.[83] Among this panoply are two connected functions, to structure thinking and to provide a locus for discussion among relevant stakeholders.[84]

---

[78] Levi, *On Nuclear Terrorism*, p. 49.

[79] Ian Reader, *Contemporary Religious Violence in Japan: The Case of Aum Shinrikyo* (Honolulu: University of Hawaii Press, 2000).

[80] Colin Flint and Steven M. Radil, 'Terrorism and counter-terrorism: Situating al-Qaeda and the global war on terror within geopolitical trends and structure', *Eurasia Geography and Economics*, 20 (2009), pp. 150–71.

[81] Richard G. Lugar, 'The Lugar Survey on Proliferation Threats and Responses', Office of United States Senator Richard G. Lugar (2005), p. 16.

[82] Galit Shmueli, 'To explain or predict', *Statistical Science*, 25 (2010), pp. 289–310.

[83] Mary S. Morgan and Margaret Morrison (eds), *Models as Mediators: Perspectives on Natural and Social Sciences* (Cambridge: Cambridge University Press, 1999).

[84] Peter McBurney, 'What are models for?', in Massimo Cossentino, Michael Kaisers, Karl Tuyls, and Gerhard Weiss (eds), *Multi-Agent Systems* (Springer Lecture Notes in Computer Science, 2012), pp. 175–88.

Models to structure thinking

Models offer a formal environment to structure thinking. In the context of nuclear terrorism, they allow analysts to make a range of simplifying assumptions so that important characteristics of terrorist behaviour become amenable to investigation. By codifying these steps mathematically a formal model can 'make explicit the assumptions about the key factors affecting the risk', offering the analyst a framework in which to operate.[85] The model is 'itself the starting point for future discussion and hence shapes those discussions', acting as a medium through which analysis can occur.[86] As Elisabeth Paté-Cornell observes in the context of nuclear terrorism: 'Because we illustrate our model using fictitious numbers, the importance of this work is not so much in the specific ranking of countermeasures that it suggests as in the framework for reasoning that it provides.'[87] As the 'numbers' inputted into any model of nuclear terrorism will be 'fictitious' in the sense that nuclear terrorism is a virtual risk, this is a positive way of utilising models developed to date.

For example, Bunn's model has been highly influential in this respect. By adopting a supply-chain perspective of nuclear terrorism, with the attendant breakdown of terrorist decision-making into a linear chain, the model focuses the analyst's attention on the different options open to terrorist groups. These are explicitly linked to policy options that may decrease the likelihood of acquisition of nuclear weapons or fissile material. The power of this approach is that 'by breaking a large, and at first glance, intractable question into a series of smaller individual questions, an estimate can be obtained for the overall question.'[88] This is a significant contribution to the debate, which both emphasises the importance of policy-relevant contributions and legitimises investigation into this contentious subject. This is, of course, subject to the proviso that model outputs are not treated as inviolable and that analysts utilising such methods acknowledge this essential truism.

Risk models as loci of discussion

Modelling can catalyse stakeholder community engagement with a topic of investigation in a shared manner, with models 'focusing debate and highlighting the basis for disagreements'.[89] In the domain of nuclear terrorism, assessments have typically been carried out by individual analysts, while group efforts to determine likelihoods have focused on simplistic elicitation of overall probabilities.[90]

These two approaches can be brought together through the shared utilisation of probabilistic models. This can 'tame the complexity of the domain … enabl[ing] stakeholders to jointly explore relevant concepts, data, system dynamics, policy options, and the assessment of potential consequences of policy options, in a structured and shared way'.[91] Areas of both agreement and disagreement can be identified through this process which, in turn, sets an agenda for future research in the nuclear terrorism domain to determine why this is the case. Identifying areas of agreement allows for the simultaneous identification of relevant evidence (the bureau of standards) and scrutiny

[85] Bunn, 'A mathematical model of the risk of nuclear terrorism', p. 103.
[86] Burchell, 'W(h)ither the Drake equation?', p. 244.
[87] Elisabeth Paté-Cornell and Seth Guikema, 'Probabilistic modeling of terrorist threats: a systems analysis approach to setting priorities among countermeasures', *Military Operations Research*, 7 (2002), p. 1.
[88] Burchell, 'W(h)ither the Drake equation?', p. 244.
[89] Bunn, 'A mathematical model of the risk of nuclear terrorism', p. 117.
[90] For instance, Lugar, 'The Lugar Survey on Proliferation Threats and Responses'.
[91] McBurney, 'What are models for?', p. 184.

of extant policy advice. By contrast, areas of disagreement can be investigated further so that the underlying reasons for divergent views can be clearly understood. Evidence used in support of divergent views can be critically considered and a comprehensive assessment of possible credible perspectives undertaken, again with a view to establishing whether extant policy advice has been well-rendered.

Bunn has partially argued for an approach of this kind, encouraging readers who 'disagree with some of the numbers … to use the model with numbers of their own, to develop and analyze their own risk assessments'.[92] The crucial subsequent step is to critically analyse the results of this activity across the expert community. According to Kaplan and Garrick, shared background knowledge is key to consistent and mutually intelligible risk assessment;[93] policymakers seeking advice on nuclear terrorism risks would do well to note that members of the nuclear terrorism analytical community *themselves* do not expect anything approaching agreement when discussing numerical risk assessment of nuclear terror events. Understanding the extent of and reasons for this divergence is thus an essential activity in which modelling can play an important role.

How to use risk models

Accepting the notion that models *can* act as a structured locus for discussion, it is natural to ask precisely *how* to use a model in this way. Fortunately, established methods already exist – when assessing the risk of events with low occurrence rates or for which adequate data is unavailable recourse to expert opinion for probabilistic estimation is a well-studied activity.

When considering nuclear terrorism, the most important component of these methods is the explicit use of expert *groups* rather than individuals. This observation stems from the belief that 'expertise is unlikely to reside in a single expert'.[94] Efforts to date by individuals are essentially expert self-elicitation exercises, in which no attempt is made to canvass opinion widely and hence control for inevitable personal bias in probabilistic estimation.

As for the form risk models might take, extending them beyond static event chains could allow for the consideration of more complex feedback and response mechanisms, and the explicit inclusion of dynamic policy choices and adversary behaviour. For example, the formalism offered by State-Transition models or (semi-)Markov decision processes incorporate multiple states with associated transition probabilities, subject to a range of policy choices in the case of the latter – different policies result in different transition probabilities.[95] This may offer a more realistic alternative to static and linear event chains although, to date, the specifics have not been investigated in the risk assessment or terrorism context beyond a small number of speculative studies.[96]

---

[92] Bunn, 'A mathematical model of the risk of nuclear terrorism', p. 108.

[93] Kaplan and Garrick, 'On the quantitative definition of risk'.

[94] Anthony O'Hagan et al., *Uncertain Judgements: Eliciting Experts' Probabilities* (Hoboken, NJ: Wiley, 2006), p. 25.

[95] Ronald A. Howard, *Dynamic Probabilistic Systems, Volume II: Semi-Markov and Decision Processes* (Mineola, NY: Dover, 2007), p. 965; for instance, William M. Miller, 'A state-transition model of epidemic foot-and-mouth disease', *New Techniques in Veterinary Epidemiology and Economics*, 1 (1976).

[96] See, for example, Ransom Weaver et al., 'Modeling and Simulating Terrorist Decision-Making: A 'Performance Moderator Function' Approach to Generating Virtual Opponents', Proceedings of the 10th Conference on Computer Generated Forces and Behavioural Representation (2001).

Another potential avenue of investigation would involve treating an event tree model as a tool for classifying key decision points in terrorist behaviour, and then applying a range of alternative modelling techniques to the study of each. This approach acknowledges that different elements of terrorist nuclear weapon development are best modelled individually. For example, in modelling nuclear material acquisition, significant effort has been expended in modelling non-nuclear black and parallel markets. A number of modelling techniques or (in)formal analogies could be explored to aid in the determination of key parameters for black market acquisition of nuclear material. In this sense, an event tree model classifies events which can then be investigated in a multitude of different ways as befits their unique characteristics, offering a framework under which outputs of these different approaches can be brought together to investigate the larger problem under consideration. Crucially, this could alleviate the criticism of event tree modelling applied to nuclear terrorism arising from the US National Research Council study into bioterrorism risk assessment, that analysts are required to input probabilities that themselves should be the output of modelling efforts into event tree models.[97]

Finally, it is important to stress that the above discussion should be caveated by the intrinsic challenges in formalising a process as complex as nuclear terrorism. Capturing this within any kind of mathematical model may serve to shape the resultant discussion in a way that obscures features of fundamental importance. Even when less emphasis is placed on the numbers themselves, the very act of modelling serves to constrain the way analysts can approach complex issues. As Martha Lampland has argued, the utility of processes whereby false or provisional numbers are produced is highly context dependent.[98] The efficacy of such processes will also evolve over time. Consequently, probabilistic modelling should be viewed as one tool among many available to analysts – and is not something that should be used to the exclusion of other quantitative or qualitative approaches.

Conclusion

Risk estimation in the area of nuclear terrorism is an extremely challenging task. Despite this there are no shortage of efforts to quantify the likelihood of its occurrence. While numerical estimates have an intrinsic appeal and offer a potential means of benchmarking nuclear security efforts, great care must be taken in their production. Here is it important that analysts understand the strengths and weaknesses of the models that they employ and caveat their conclusions accordingly. However, to date the majority of studies have over-reached and under-caveated their conclusions, through employing inappropriate models and using data that is far from representative of nuclear terrorism. This article has attempted to demonstrate, through outlining different approaches to probability estimation, that nuclear terrorism should be considered as a virtual risk that escapes simple quantification. That does not mean, however, that models cannot be used, rather that their ultimate purpose must be reconsidered. Models of nuclear terrorism can be used to structure thinking and to serve as loci of discussion, highlighting areas of disagreement and, hence, serving to direct future research efforts. When employing them to this end, use should be made of expert groups over individuals and models should look to incorporate feedback and response mechanisms. It should also be recognised that the development of unifying models for the nuclear terrorism phenomenon as a whole is unrealistic. Instead individual models should be constructed and applied to different components as appropriate. This nuanced understanding is necessary to ensure the questionable results of modelling efforts in the nuclear terrorism domain are not oversold to the policymaking community.

[97] US National Research Council, 'Risk assessment'.
[98] Martha Lampland, 'False numbers as formalizing practices', *Social Studies of Science*, 40 (2010), pp. 377–404.

Acknowledgements

Biographical information

Rob Downes is a MacArthur Fellow in Nuclear Security working at the Centre for Science and Security Studies, within the Department of War Studies at King's College London. Trained as a mathematician, he received his PhD from University College London in 2014.

Dr Christopher Hobbs is Co-Director of the Centre for Science and Security Studies (CSSS) and a Senior Lecturer in Science and Security within the Department of War Studies at King's College London. A physicist by background, he works on nuclear security, non-proliferation, and intelligence issues.