http://pubsonline.informs.org/journal/deca/

DECISION ANALYSIS
Vol. 15, No. 3, September 2018, 154–173
ISSN 1545-8490 (print), ISSN 1545-8504 (online)

Insights for Critical Alarm-Based Warning Systems from a Risk Analysis of Commercial Aviation Passenger Screening

Robin L. Dillon,[a] **William J. Burns,**[b,c] **Richard S. John**[d]

[a] McDonough School of Business, Georgetown University, Washington, DC 20057; [b] Decision Research, Eugene, Oregon 97407; [c] California State University, San Marcos, San Marcos, California 92096; [d] University of Southern California, Los Angeles, California 90089

Contact: rld9@georgetown.edu (RLD); william_burns@sbcglobal.net (WJB); richardj@usc.edu (RSJ)

Received: September 26, 2017
Revised: December 2, 2017
Accepted: January 2, 2018
Published Online in Articles in Advance: May 25, 2018

https://doi.org/10.1287/deca.2018.0369

Abstract. The Transportation Security Administration (TSA) recently commissioned a risk analysis of the passenger threat vector in which an adversary gains access through the screening checkpoint. The goal of the project was to provide insights to the TSA to improve both safety and resource allocation as they continue to develop new security procedures in a constantly evolving threat environment. The result was a probabilistic risk model to support the TSA as they plan future safety and resource allocations procedures. Because aviation passenger screening involves highly sensitive information, we discuss the insights gained from the study that are applicable for other highly critical security systems that rely on alarm-based warning technologies to detect anomalies.

Funding: This material is based upon work supported by the U.S. Department of Homeland Security to the USC Center for Risk and Economic Analysis of Terrorism Events [Contract HSHQDC-15-J-00375] and grant [Award Number 2010-ST-061-RE0001-07]. Additionally, part of Robin Dillon's time was supported by the National Science Foundation [Grant CMMI 1757350].

Keywords: risk analysis • commercial aviation security

1. Introduction

Commercial aviation is not the only industry challenged to prevent and detect security threats. In July 2012, three intruders, including an 82-year-old nun, infiltrated the high security Y-12 nuclear facility in Oak Ridge, Tennessee. This failure of security highlights several challenges associated with both modeling and managing critical security situations with alarm-based warning systems. First, regarding the management of the security system, the attackers did set off an alarm, but security officers ignored it because hundreds of false alarms occur at Y-12 every month. Officers stationed inside the uranium-storage facility heard the intruders hammering on the wall, but assumed the sounds were from maintenance workers (Schlosser 2015). The problem in this situation was not with the alarms, which worked as designed. The problem was with the resolution of the alarms in an environment where threats are rare events, and false alarms are common.

Many critical systems face a similar challenge when they rely on alarms for the initial alert, and on trained security personnel for alarm resolution. Additionally, there are interesting challenges in modeling a system like the high security Y-12 nuclear facility prior to an attack for purposes of improving security. One challenge is the various uncertainties related to the detection of system events and the consequences of failed detection. Through testing and "red-teaming"[1] exercises, probabilities associated with how often alarm systems react can usually be estimated. Probabilities associated with human elements are harder to gather or estimate (i.e., the probability that a guard will disregard the alarm and assume the sound was from a maintenance worker). Also, the consequences are difficult to estimate since it had not happened before (i.e., what are the consequences of three intruders infiltrating the high security facility when this has never happened before). In this case, these intruders had no intention to steal anything and instead only wanted to splash blood on the Y-12 facility to mark it as "evil" (Schlosser 2015).

Modeling a broad range of intruders and their intentions in a risk analysis effort makes the modeling task more burdensome as more adversary characteristics

are considered (i.e., highly trained nation-state adversaries versus lone-wolf adversaries, the size of the adversary force, the weapons used, etc.). Very specific scenarios can be crafted, which helps reduce the range of uncertainty in the model, but greater scenario specificity generally increases the complexity of the models. An alternative to modeling many realistic scenarios is not specifying particular scenarios (i.e., size of adversary force) but leaving that as an additional uncertainty that subject matter experts (SMEs) must consider in their probabilistic estimates. While fewer combinations of scenarios are less burdensome, very large uncertainties in outcome distributions can result, providing less precision and greater variability for the decision maker to consider. The point of discussing the Y-12 security breach here is to not only emphasize the broad applicability of probabilistic risk analysis (PRA) for analyzing critical alarm systems, but also to highlight the challenges in applying PRA to critical alarm systems. The remainder of the paper focuses on a case study of the passenger threat vector for commercial aviation, concentrating on the screening checkpoint that faces similar challenges as the security at Y-12.

In this paper, we discuss the development of three components of a probabilistic risk analysis for the passenger screening threat vector: (1) the probability of detection of an adversary during passenger screening for a specific attack device and set of countermeasures; (2) the probability of catastrophic loss of an aircraft, conditional on an adversary navigating the screening with an undetected attack device; and (3) the consequences of catastrophic loss of an aircraft for different aircraft types. We will discuss challenges and insights from a risk analysis of this type, as well as broader implications for modeling and managing critical alarm-based warning systems.

2. Discussion of Challenges

The passenger screening checkpoint at U.S. airports is just one example of a critical alarm-based warning system. Many other critical systems rely on alarms and the resolution of alarms to detect problems. We explore some of the challenges with (1) alarm-based warning systems in general, (2) estimating and representing uncertainty in risk based models of these systems, and (3) unique issues related to commercial aviation screening.

2.1. Challenges with Alarm-Based Warning Systems

When alarm-based warning systems are needed to detect an anomaly (such as X-ray screeners searching bags for banned items or a gas detection system on an oil drilling rig), one perceived option to improve the system is to increase the quantity of detection devices that could produce an alarm. Jackson et al. (2012) refer to this approach as "ladling on," when more and more alarm systems are added to respond to every perceived threat. Increasing the number of detection devices that can alarm is not without costs, including false alarms and in the case of passenger screening, the time spent by passengers undergoing security procedures, and the hassle and privacy sacrifices of travelers undergoing security screening. Given the resources and attention devoted to aviation security in an era in which resource constraints are likely to become ever more important in policy decisions, it is important that we approach aviation security in a rational and defensible way (Jackson et al. 2012).

As the example from the Y-12 plant discussed in the introduction demonstrates, the usefulness of an alarm system is limited if the alarm resolution step is not performed satisfactorily. A problem identified in the Y-12 example and that is well documented in many domains is "alarm fatigue." Alarm fatigue is a significantly reported concern in many fields but especially in the medical community. For example, given the sheer number of potential equipment alarms within the clinical environment, there is a high risk of nurses becoming desensitized to the sound of patient alarms. Additionally, alarms are sometimes disabled, silenced, or ignored. In January 2010, a heart patient died at Massachusetts General Hospital. An investigation into the incident found that 10 nurses on duty that morning could not recall hearing the beeps at the central nurses' station or seeing the scrolling tickertape messages on three hallway signs that would have warned them as the patient's heart rate fell and finally stopped over a 20-minute span. Additionally, investigators said the volume for a separate audible crisis alarm on the patient's bedside monitor was turned off the night before by an unknown person. Investigators did not feel that the one disabled alarm was the single cause of death. Rather, investigators pointed more critically to the desensitization to the alarms that actually sounded as a significant factor in the patient's

death (Kowalczyk 2010). Even before this tragic death, in 2009, the Emergency Care Research Institute (ECRI), a nonprofit research and consulting organization that specializes in medical devices, listed alarms on patient monitoring devices as number two on its top ten list of health technology hazards (Bell 2010, Graham and Cvach 2010).

Other well-known examples exist of alarms being turned off or disabled, including an alarm system designed to alert the crew and prevent combustible gases from reaching potential sources of ignition on the Deepwater Horizon rig. The gas alarm system had been operating in "inhibited" mode for a year to prevent false alarms from disturbing the crew needlessly in the middle of the night (Hilzenrath 2010).

2.2. Challenges in Modeling Uncertainty in Probabilistic Risk Analysis

From a modeling perspective, many low probability, high consequence systems have similar challenges with including uncertainty in a useful way in risk analysis models. One challenge is that data may be available for some pieces, such as alarm rates on TSA detection equipment, or pump failure rates for the nuclear power industry, but data on the economic consequences of another terrorist attack in commercial aviation or the consequences of a loss of coolant event in nuclear power are highly uncertain, actual data are rare, and test data are not available. It has been known for some time that different types of uncertainty characterize different events. One early distinction was classically referred to as risk versus uncertainty (Knight 1921). More recently, the terms epistemic uncertainty and aleatory uncertainty have been used (Paté-Cornell and Dillon 2006). Aleatory uncertainty is the irreducible randomness or variability that remains when the system is well known. The alarm rates on a piece of detection equipment would be an example of aleatory uncertainty. Epistemic uncertainty stems from limitations of fundamental knowledge about a phenomenon of interest. The indirect cost consequences to the U.S. economy following another terrorist attack against a U.S. commercial aircraft is an example of epistemic uncertainty. Realizing that probabilistic risk analyses of critical alarm warning systems will need to include some of both types of uncertainty is important to recognize and capture because it allows the decision maker, who will later use the results, to better understand the source of residual fundamental uncertainties (Paté-Cornell and Dillon 2006). Epistemic uncertainty will commonly be revealed by divergences in the opinions of experts. As will be discussed further in this paper, there are multiple, highly divergent, but still reasonable views on possible indirect economic consequences from another terrorist attack against commercial aviation. Decision makers need to be aware of the potential effect of experts' disagreements on the variability in the results of the risk analysis (Paté-Cornell and Dillon 2006).

A second consideration when modeling uncertainty in a PRA is the level of specificity to include for events. The more specific an event regarding adversary type, number of adversaries, type of weapon (e.g., metallic versus nonmetallic bomb assembled or not assembled), airport target, time of day, etc., the more precise SMEs can be when providing probability estimates. The trade-off is that greater event specificity requires more scenarios that must be considered, which adds to the assessment burden if historical data are not available and the only data available is from SMEs. The challenge is in choosing the minimal level of specificity required to provide the insights needed, and there is little research on best practices for this. For example, in considering detection rates for different types of screening equipment at the passenger checkpoint, the type of weapon and the method of concealing it were critical, so we considered metallic explosive devices versus nonmetallic explosive devices versus metallic guns either on the adversary's body or in a carry-on. This resulted in six weapon configurations to consider. We left more specific details to the experts to include in their uncertainty estimates, such as level of bomb assembly at screening, accuracy and calibration of detection equipment, target airport, time of day, traveler load, etc. Because significant details were not specified, the results were expected to include large uncertainty bounds, capturing the unspecified details of the attack that are a source of uncertainty faced by system-wide strategic decision makers. Ideally, the probability assessments will be updated over time as more knowledge is gained by TSA experts. As part of any updates to the analysis, the level of specificity of events could be revisited.

2.3. Challenges with Commercial Aviation Security

Many researchers, including Elwell (2008), have warned of the risk of recurrent attacks in commercial aviation. Elwell (2008, p. 11) states, "Terrorists will visit aviation again and again . . . . [They] are learning from their mistakes and from the mistakes of others . . . their attack methodologies [are] chang[ing] quickly, adapting to new security measures." To mitigate this threat, over time, more and more measures have been adopted to enhance commercial aviation security. As we discover new threats, travel and screening rules are modified, but changes have come with an increasing burden on passengers in terms of time, convenience, and invasiveness of screening.

The use of risk analysis has become a central element in decision making for the TSA and the preferred approach for setting priorities. Risk analysis explicitly considers the probability that a damaging event will occur and the consequences of it occurring. We explore both of these components in a risk analysis of the passenger threat vector with a focus on checkpoint screening. It is important to realize that there are other threat vectors to commercial aviation (e.g., checked baggage, stand-off missile, insider threat, etc.) and that an adaptive adversary can shift from one threat vector to another. In addition to the threat, decision making informed by risk analysis should include consideration of the costs of mitigation alternatives. In the case of commercial aviation, costs can vary significantly among individual travelers. One person may be significantly offended or inconvenienced by a use of a particular countermeasure or search technique, while that same search may not bother another person at all (Jackson et al. 2012). Additionally, the amount of time a passenger has allocated to traverse a checkpoint screening will influence the perceived inconvenience of delays and wait times.

3. Modeling and Analysis

Over several months, we had the opportunity to interact with TSA management and SMEs to develop a probabilistic risk model of an adversary moving past the passenger screening checkpoint and into the "sterile" area. The effort involved estimating (1) the probability of detection of an adversary during passenger screening for a specific attack device and set of countermeasures, (2) the probability of catastrophic loss of an aircraft conditional on an adversary navigating the screening with an undetected attack device, and (3) the consequences of catastrophic loss of an aircraft for different aircraft types.

3.1. Probability of Detection of an Adversary During Passenger Screening

The starting point for an assessment of risk is an understanding of the pathways for attacking the aviation system to cause damage and disruption. Early in our study for the TSA, we chose six means of attacking an aircraft by one of the passengers traversing the checkpoint screening. These attack devices were (1) a metallic improvised explosive device (IED) hidden on the adversary's body, (2) a nonmetallic IED hidden on the adversary's body, (3) a metal gun hidden on the adversary's body, (4) a metallic IED hidden in carry-on luggage, (5) a nonmetallic IED hidden in carry-on luggage, and (6) a metal gun hidden in carry-on luggage. Metallic versus nonmetallic weapons carried on the body or in luggage are important distinctions in assessing the ability of alarm equipment to identify the weapon. It was assumed an adversary would choose only one device. The question of whether the device was assembled or not at the screening checkpoint was unspecified and left to the SMEs to incorporate as uncertainties in their estimates. Representative available countermeasures were selected, although not all countermeasures are available at all airports (e.g., canines, Sachs 2016), to offer a rich array of combinations. The countermeasures considered were (1) canines, (2) behavioral detection officers (BDOs), (3) explosive trace detection (ETD) devices, (4) advanced imaging technology (AIT) devices, (5) walk-thru-metal-detectors (WTMD), (6) pat-downs, and (7) an x-ray for carry-on baggage. We assumed that all passengers walked through an AIT or a WTMD and put their carry-on luggage through an x-ray. Thirty-two combinations of countermeasures were constructed ranging from what might be offered during expedited screening (x-ray, WTMD) to complex combinations (x-ray, canines, BDO, ETD, WTMD, pat-down regardless of an alarm). It was assumed that a pat-down would be used to resolve any alarms for a person and that a bag search would be used to resolve alarms for carry-on luggage. A pat-down without a prior alarm could be planned as a

screening measure. It was also assumed that a passenger receives no more than one pat-down, regardless of the number of alarms triggered.

Checkpoint screening detection probabilities were calculated considering the capability of each of these thirty-two countermeasure combinations to discover an adversary bringing one of the three attack devices through the checkpoint, either on their person or in carry-on items. For example, a walk-through metal detector is better at detecting metallic weapons than nonmetallic explosives.

Computing detection probabilities required that experts provide three types of assessments. These assessments were based on standard expert elicitation techniques, developed for risk analyses in the nuclear power industry and published in NUREG-1150 (U.S. Nuclear Regulatory Commission 1990). The landmark NUREG-1150 study used expert judgment to assess failure frequencies, failure modes, recovery actions, accident progression, and source term behavior, among other phenomena, to characterize the risk associated with severe accidents in operating reactors. As detailed in the remaining section, every attempt was made to acknowledge the limitations of relying on expert assessments (Morgan 2014), since this was the only source of information available given the constraints of this study.

SMEs were identified by TSA management for the project. Experts were identified and recruited who could provide information regarding the alarm and resolution rates of the various countermeasures and interdiction rates. Across the entire process, seven different SMEs provided input, but not all SMEs provided input for all estimates. SMEs only provided data for questions they were qualified to answer. Prior to the elicitation meetings, the process was described in a group setting to multiple SMEs. The final elicitations were conducted with each SME individually. Each provided his or her best assessment of the 10% probability, the 50% probability, and the 90% probability of specific alarm events. It is important to understand that the uncertainties were expected to include sources of variation related to unspecified details of the attack scenario, such as location and size of the airport, time of day, passenger loads, staffing and training, etc., as well as unresolvable uncertainty regarding performance of equipment and personnel. As explained to the SMEs, the 50% estimate represented an estimate of the likelihood of the event such that the SME was equally likely to believe the probability was greater than or less than the stated estimate. The extreme 10% (90%) estimate represented an estimate of the likelihood of the event such that the SME believed the probability was nine times more likely to be greater than (less than) the stated estimate. The probability of alarm was conceptualized as a Bernoulli parameter $p(\text{alarm} \mid \text{weapon})$. Each SME was asked for median (50%) and extreme (10% and 90%) estimates on this Bernoulli parameter. In essence, SMEs provided three points on a secondary probability distribution for the probability of an alarm for each alarm system, conditional on the presence of the specified weapon.

The TSA SMEs first provided probabilities that a countermeasure would alarm in the presence of each of the six attack devices. Second, given that the countermeasure alarmed, they provided probabilities that the attack device would be detected during the alarm resolution procedure. This probability was also conceptualized as a Bernoulli parameter $p(\text{detection} \mid \text{alarm})$. Third, they described relevant standard operating procedures (SOPs) for checkpoint screening, allowing the analysis to account for dependencies among countermeasures used in combination. Table 1 describes the alarm events asked of SME's, and Table 2 describes a simplified version of the questions asked regarding the resolution of alarms.[2]

In the next step, we calculated parameters for beta probability distributions for each SME's assessments. Standard beta distributions were used since the distribution covers the entire range of probability from 0 to 1 which can compensate for experts' overconfidence (Morgan 2014) and are fully determined by two points on the cumulative distribution function. For all the parameters, we calculated three possible beta distributions using all three pairs of estimated points, and for each parameter we used the distribution that included the median and the most extreme estimate (i.e., the longest tail) in an effort to best characterize the skewed tail of the distribution. While there are indeed multiple possible approaches for aggregating the SMEs' cumulative distribution estimates, we focused on three particular options: (1) average the 10%, 50%, and 90% estimates across experts at all three

Table 1. Countermeasure Alarm Probability by Attack Type

	Metallic IED on body	Nonmetallic IED on body	Weapon (metal gun) on body	Metallic IED in carry-on	Nonmetallic IED in carry-on	Weapon (metal gun) in carry-on
AIT alarms						
WTMD alarms						
X-ray alarms						
BDO alarms						
ETD of hands alarms						
Canine alarms						

points for each distribution; (2) calculate the beta distributions for each expert and then average the distributions using simulation; or (3) construct decision trees for each expert independently and average the final overall detection probabilities. We followed the second option where the selected beta distributions for each expert were used in a Monte Carlo simulation to generate estimates that aggregated the data and uncertainties provided by the individual experts, providing not only a point estimate (mean) of the probability of the target event but a distribution reflecting the SMEs' uncertainty about that probability value. It is worth noting that the third option was precluded since we did not have complete assessments for an entire event tree from each expert, as individual experts answered only questions that were within their area of expertise.

Using the distributions for each of the previously described events aggregated for all SMEs, detection rates were then determined for sets of countermeasures. Two important assumptions were made at this stage in the analysis:

- Each countermeasure rate was independent of the others in the set (i.e., having an ETD alarm did not change the alarm rate of an AIT).
- At most one pat-down will be conducted, regardless of the number of independent alarms.

With those two assumptions, the resulting detection rates were simulated for each set of feasible countermeasures considered.

Using these two assumptions and other relevant SOP data provided by the PRAs, event trees were constructed for each combination of the six attack devices and the 32 countermeasures. Constructing these event trees was not a simple task, but the insight gained was very helpful. In particular, we learned the following:

1. The value of adding a countermeasure depends on the set of countermeasures already in use, and the various attack mode threats.

2. Countermeasure detection rates are limited by the effectiveness of the resolution.

We discuss both of these insights in more detail next.

3.1.1. Insight: The Value of Adding a Countermeasure Strongly Depends on What Other Countermeasures Are Being Used and What the Threats Are. From the SMEs' assessment of the probability of alarms and the probability of detection resolution, we constructed estimates of the probability of detection for different types of weapons and mode of transport for all 32 different portfolio combinations of countermeasures. Figure 1 provides an example of what results could look like for three attack methods.

Table 2. Simplified Alarm Resolution by Attack Type

	Metallic IED on body	Nonmetallic IED on body	Weapon (metal gun) on body	Metallic IED in carry-on	Nonmetallic IED in carry-on	Weapon (metal gun) in carry-on
Detection from pat-down						
Detection from ETD of TSO officer's gloves following a pat-down						
Detection from bag check						

Figure 1. (Color online) Example Portfolio Detection Rates for Different Sets of Countermeasures

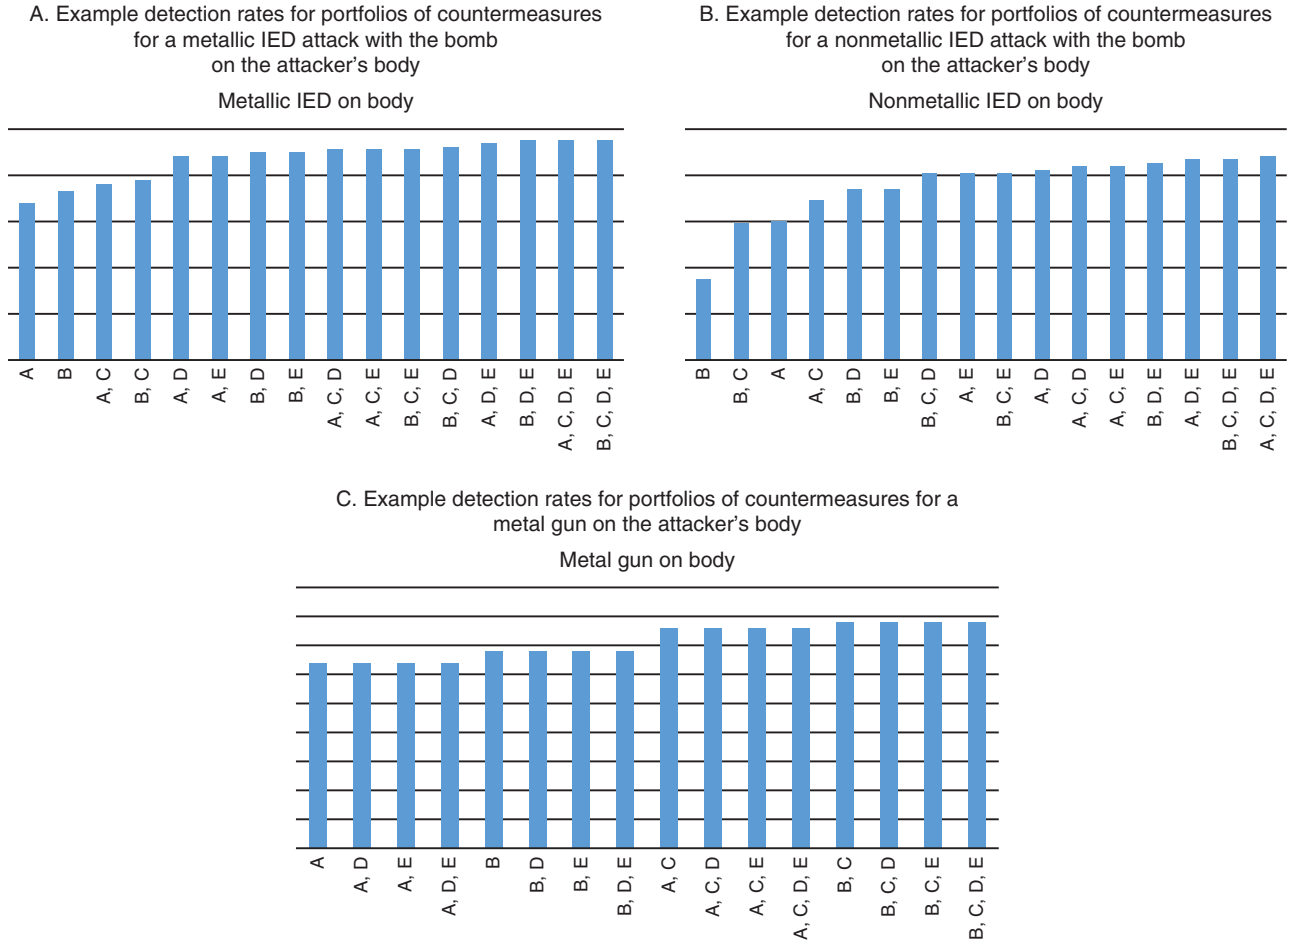

From this analysis, we are able to draw a number of conclusions regarding the effects of adding additional countermeasures to currently existing configurations of countermeasures:

1. The improvement in detection from adding a particular countermeasure is nonadditive. Detection improvement for individual countermeasures is contingent on the entire configuration of countermeasures in an interactive manner. As can be seen in various places in Figure 1, depending on the countermeasures already included and the attack type, sometimes additional countermeasures do not result in improvements in detection.

2. The specific nature of countermeasure detection improvement, including interactive effects among countermeasures, is highly dependent on the particular attack threat, including both weapon type and mode of transport through screening. Some countermeasures are not designed to detect some items (e.g., metallic versus nonmetallic weapons), and some countermeasures can have redundant detection capabilities for particular attack threats. In Figure 1, if countermeasure B is better than countermeasure A at detecting metal items, then for some attack modes countermeasure B is better than A but for others the reverse is true.

3. Despite the nonadditive, interactive effects described previously, there is an overall effect of marginally decreasing reduction in detection probabilities as more countermeasures are added to the screening configuration. This non-additive, diminishing return of adding countermeasures was observed across all attack types and modes of transport considered. "Ladling on" (Jackson et al. 2012) is not always

the best strategy. This appears as plateaus are reached in each chart in Figure 1.

4. The differential value added in many instances was small in comparison to the large uncertainties in overall probability estimates. In the first panel of Figure 1, the difference between the portfolios {A, E} and {B, E} was 2%. While that suggests that {B, E} could be preferred to {A, E} for this attack type, this 2% difference was small in comparison to the overall uncertainty described in the model. This large overall uncertainty will be discussed in more detail in Section 3.2.1.

In summary, blind expenditure of resources to provide particular countermeasures across all airports or screening areas within an airport is unlikely to provide optimal detection probability. Resource allocations should occur at a local level, with full consideration of the countermeasures currently in place, and an assessment of the relative threats along the passenger threat vector. It is important to understand which types of threats current sets of countermeasures can detect and choose countermeasures to fill any gaps strategically. Additionally, if multiple sets of countermeasures (with uncertainty) are roughly equal in effectiveness, decision makers should consider the possibility of thwarting some terrorist learning by changing portfolios over time.

3.1.2. Insight: Countermeasure Detection Rates Are Limited by the Effectiveness of the Resolution. Consider an event tree for a simplified example where an adversary with a nonmetallic IED on his body traverses a walk-through metal detector (WTMD). Figure 2 denotes the probability of the WTMD alarming given that a bomb is present as the probability p and the probability that the alarm is successfully resolved to detect the weapon as the probability q.

Table 3. Probability of Detecting Weapon Given One Is Present

WTMD Probability of alarm (%)	Alarm resolution—Pat-down probability of detection (%)							
	30	40	50	60	70	80	90	100
30	9	12	15	18	21	24	27	30
40	12	16	20	24	28	32	36	40
50	15	20	25	30	35	40	45	50
60	18	24	30	36	42	48	54	60
70	21	28	35	42	49	56	63	70
80	24	32	40	48	56	64	72	80
90	27	36	45	54	63	72	81	90
100	30	40	50	60	70	80	90	100

In the risk analysis for TSA, p and q were provided by SMEs. In this example, we will consider a range of possible values between 30% and 100% for both probabilities as shown in Table 3.

From the data provided in Table 3, we can observe that as the probability of the WTMD's alarm rate increases, the probability of detecting the weapon increases, but it can only increase to the level of the probability of alarm resolution, i.e., the effectiveness of the pat-down in detecting the weapon. To further explore this insight, consider adding a second screening countermeasure, a near-perfect explosive trace detection (ETD) device with an alarm effectiveness of 99%. This scenario is shown in Figure 3.

As can be seen in Table 4, the probability of the WTMD alarming is irrelevant because the ETD alarms

Figure 2. (Color online) Alarm and Detection Rates for a WTMD

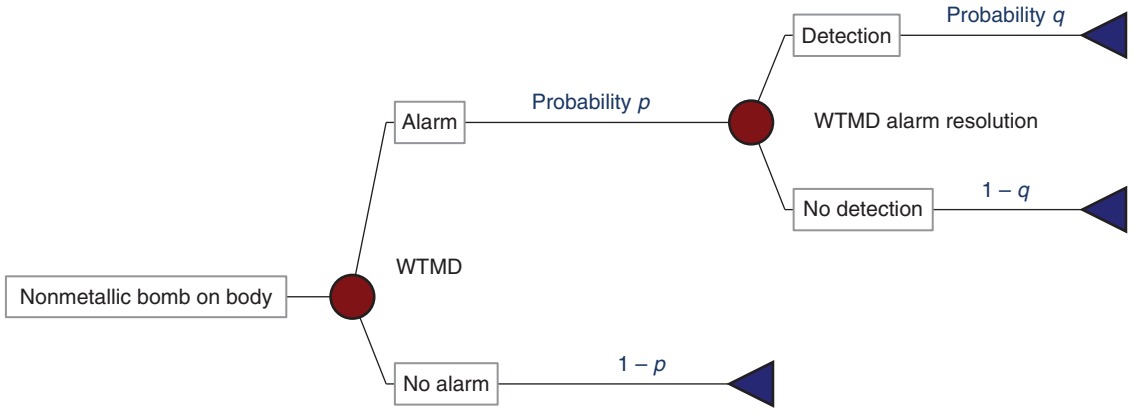

Figure 3. (Color online) Alarm and Detection Rates for a System That Includes a 99% Effective ETD

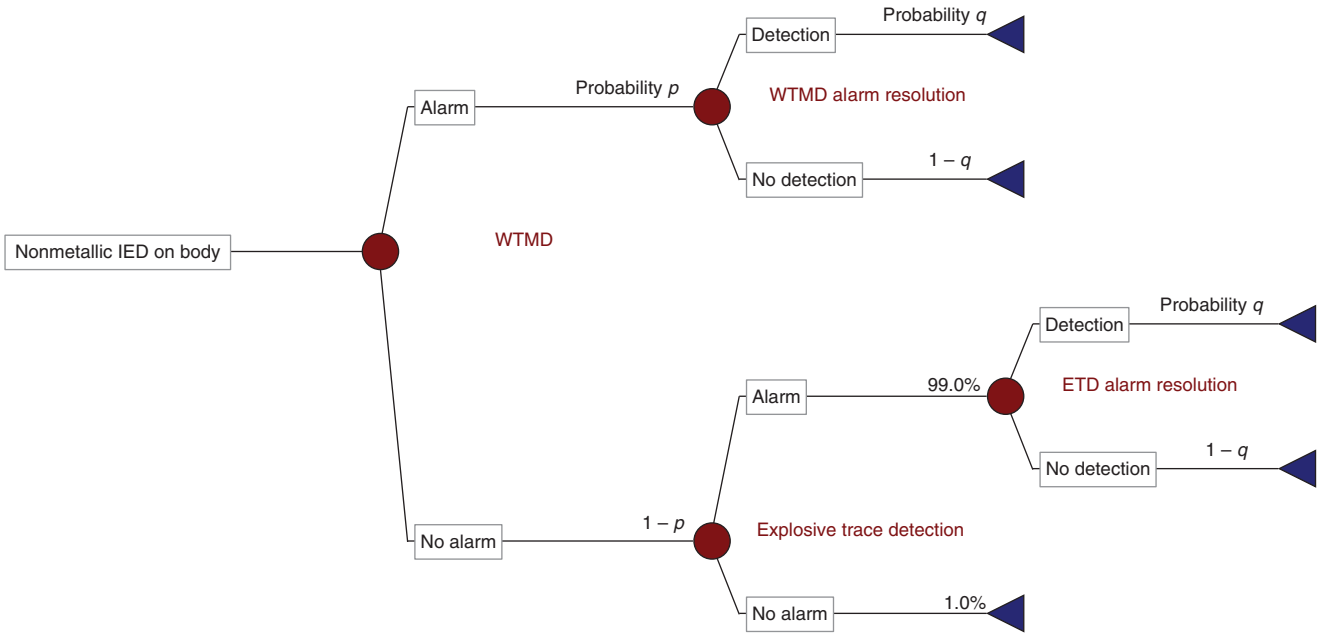

with great accuracy. Regardless of the accuracy of the ETD alarm, the maximum detection rate is constrained by the probability of the alarm resolution, i.e., the maximum detection rate for weapons can be no better than the pat-down rate, even if the detection technology warns us with nearly perfect accuracy.

An important conclusion from this part of the analysis is that adding more detection technologies will not be helpful unless the screening staff is effective in their resolution pat-downs by identifying the weapons that the alarms are detecting.

Table 4. Probability of Detecting Weapon Given One Is Present Given a WTMD and a 99% Effective ETD

Probability of WTMD alarm (%)	Alarm resolution—Pat-down probability of detection (%)							
	30	40	50	60	70	80	90	100
30	30	40	50	60	70	80	89	99
40	30	40	50	60	70	80	89	99
50	30	40	50	60	70	80	90	100
60	30	40	50	60	70	80	90	100
70	30	40	50	60	70	80	90	100
80	30	40	50	60	70	80	90	100
90	30	40	50	60	70	80	90	100
100	30	40	50	60	70	80	90	100

We next focused on an adversary progressing undetected through a screening checkpoint into the secure area of the airport.

3.2. Probability of Catastrophic Loss of an Aircraft Given an Adversary Passes the Screening Checkpoint Undetected

Once the screening checkpoint is passed, the individuals most likely to interact with an adversary, i.e., passengers, airline employees, and crew, are not typically trained to detect weapons or to intervene with adversaries. However, individuals may "see something, say something" (Department of Homeland Security 2016), as demonstrated on previous flights in which individuals have proven effective in interdicting attackers. For example, passengers tackled and restrained Umar Farouk Abdulmutallab on Northwest Flight 253 (bound for Detroit from Amsterdam on Christmas Day, 2009) and put out the fire before the plastic explosives hidden in his underwear could detonate.

To prevent an attack once the adversary has cleared the screening checkpoint, one of the following scenarios is required to occur:

1. A trained law-enforcement officer such as a federal air marshal, a traveling off-duty law-enforcement

Table 5. Probabilities of Interdicting Adversary After Security Checkpoint Screening and Corresponding Probability of Successful Attack

Law enforcement interdicting in gate area

Probability of interdiction in sterile area

Passenger/Crew Interdicting in gate area

	10%	20%	30%	40%	50%	60%	70%	80%	90%
90%	91%	92%	93%	94%	95%	96%	97%	98%	99%
80%	82%	84%	86%	88%	90%	92%	94%	96%	98%
70%	73%	76%	79%	82%	85%	88%	91%	94%	97%
60%	64%	68%	72%	76%	80%	84%	88%	92%	96%
50%	55%	60%	65%	70%	75%	80%	85%	90%	95%
40%	46%	52%	58%	64%	70%	76%	82%	88%	94%
30%	37%	44%	51%	58%	65%	72%	79%	86%	93%
20%	28%	36%	44%	52%	60%	68%	76%	84%	92%
10%	19%	28%	37%	46%	55%	64%	73%	82%	91%

Law enforcement interdicting in cabin given no interdiction in gate area

Probability of interdiction in aircraft cabin

Passenger/Crew Interdicting in cabin given no interdiction in gate area

	10%	20%	30%	40%	50%	60%	70%	80%	90%
90%	91%	92%	93%	94%	95%	96%	97%	98%	99%
80%	82%	84%	86%	88%	90%	92%	94%	96%	98%
70%	73%	76%	79%	82%	85%	88%	91%	94%	97%
60%	64%	68%	72%	76%	80%	84%	88%	92%	96%
50%	55%	60%	65%	70%	75%	80%	85%	90%	95%
40%	46%	52%	58%	64%	70%	76%	82%	88%	94%
30%	37%	44%	51%	58%	65%	72%	79%	86%	93%
20%	28%	36%	44%	52%	60%	68%	76%	84%	92%
10%	19%	28%	37%	46%	55%	64%	73%	82%	91%

Probability interdicting in cabin given no interdiction in gate area

Probability of successful attack

Probability Interdicting in gate area

	10%	20%	30%	40%	50%	60%	70%	80%	90%
90%	9%	8%	7%	6%	5%	4%	3%	2%	1%
80%	18%	16%	14%	12%	10%	8%	6%	4%	2%
70%	27%	24%	21%	18%	15%	12%	9%	6%	3%
60%	36%	32%	28%	24%	20%	16%	12%	8%	4%
50%	45%	40%	35%	30%	25%	20%	15%	10%	5%
40%	54%	48%	42%	36%	30%	24%	18%	12%	6%
30%	63%	56%	49%	42%	35%	28%	21%	14%	7%
20%	72%	64%	56%	48%	40%	32%	24%	16%	8%
10%	81%	72%	63%	54%	45%	36%	27%	18%	9%

officer, or on-duty airport security is in the vicinity of the adversary and detects signs of an attack from the adversary, and interdicts the adversary in the gate area.

2. Untrained passengers and airline employees and crew in the adversary's vicinity in the gate area detect signs of an attack from the adversary and interdict the adversary in the gate area.

3. Additional TSA security screening in the gate area (e.g., canine units, etc.) detects signs of an attack from the adversary and interdicts the adversary in the gate area. (Note that additional security could be available due to additional information about a risky traveler or insider threat, or could be available due to random patrol patterns.) *We do not explore this option in this analysis.*

4. The targeted flight has been assigned a federal air marshal who detects signs of an attack on the aircraft in time to interdict.

5. The targeted flight has Federal Flight Duty Officers (FFDOs), i.e., armed pilots, and the FFDOs are in a position to interdict when the adversary attempts to hijack the aircraft by overtaking the cockpit.

6. Passengers and/or crew onboard the targeted aircraft could detect signs of an attack and interdict.

We explored five of these six scenarios (excluding scenario 3 as previously noted) with SMEs who provided their best assessment of the likelihood of interdiction events in the gate area and on the aircraft. Unlike at the checkpoint screening where there are clear alarm and resolution procedures, there are not similar procedures in the gate area. Instead of continuing to consider separate alarm and resolution rates, we focused on an interdiction rate that would combine alarm and resolution rates.[3] SMEs were identified who could provide information regarding interdiction rates. In this stage of attack for gate and cabin interdiction, four SMEs provided estimates. The final elicitations were conducted with two SMEs individually and with the remaining pair together. Using the same procedure as in the previous assessment, each was asked for the 10% probability, the 50% probability, and the 90% probability of specific interdiction events in the gate area and the aircraft cabin. The events considered different individuals in the gate area who may or may not identify various weapon types (metallic IED on body, nonmetallic IED on body, metal gun on body, metallic IED in carry-on, nonmetallic IED in carry-on, and metal gun in carry-on). Individuals who were considered candidates to detect the weapon included FAMs officers, law enforcement officers (on-duty at the airport or traveling), passengers, and crew members. SMEs first considered how often each of these groups would be in the sterile airport area, including at the gate, and then considered how likely it was for each to interdict a weapon conditional on the type of weapon and location of weapon (attacker body or carry-on luggage). While passengers and crew are at the gate 100% of the time and may be primed to "see something, say something," nevertheless they are not trained in methods of behavioral detection or weapons interdiction. Relying on untrained individuals to detect something will result in numerous false alarms that security resources need to be allocated to resolve.

The next step of the analysis required calculating parameters for beta probability distributions for each SME's assessments of the effectiveness of each of the three groups (FAMS, LEOs, passengers/crew) interdicting an attacker who has successfully passed the TSA screening checkpoint. The same calculation procedure described earlier was used for estimating beta distributions using the median and the long-tail estimate of the distribution. As before, the calculated distributions were utilized in a Monte Carlo simulation to aggregate experts and to generate overall estimates and distributions for the probability of interdiction. This task was then repeated for the cabin of the aircraft assuming the weapon was not detected in the sterile area of the airport. SMEs were asked to provide their best assessment of the 10% probability, the 50% probability, and the 90% probability of interdiction of an attack in the cabin of an aircraft given that the adversary has successfully navigated the checkpoint screening and the sterile gate area. In addition to passengers, crew, federal air marshals, and law enforcement officers, some flights have armed pilots (FFDOs).

Table 5 combines these four scenarios for detection and demonstrates the implications of various values for detection if a weapon gets through the screening checkpoint. This table could be used to examine the successful interdiction rates using the means of the beta probability distributions assessed from SMEs.

Two examples are highlighted in Table 5. In the Black scenario, if we assume a 20% chance of law enforcement interdicting an adversary in the gate area, and a 50% chance that an airline employee or passenger interdicts an adversary given that the adversary is not interdicted by law enforcement, there is a 60% chance that the adversary is stopped in the gate area. Assuming similar chances once in the cabin of the aircraft, there could be an additional 60% chance that the adversary is stopped once onboard. Overall, for this scenario with greater than 50% chance of intervention in the gate area and then the cabin area, there is then a 16% chance (or 4 in 25) of a successful attack if the adversary successfully traverses the checkpoint screening. In the Gray scenario, we assume the same probabilities as in the Black scenario, with the exception of more law enforcement in the gate area such that there is a 60% chance that law enforcement will interdict an adversary in the gate area. In this Gray scenario, the overall chance of a successful attack from an adversary who successfully traversed the checkpoint screening, is cut

in half to 8% (or 2 in 25). Other interdiction scenarios can be explored using Table 5 and the means of the probability distributions provided by SME's regarding interdiction rates. Steps taken to increase trained individuals in the sterile area will also help with adversaries who have insider access, a scenario that was not considered in this study.

From examining interdiction scenarios post-checkpoint, two insights were gained:

1. Since there is great variation in the effectiveness of individuals with no training to detect or interdict terrorists who have successfully navigated the screening checkpoint, the range of uncertainties is large for each interdiction probability, and these large uncertainties can propagate to very large uncertainties for the overall probability of success (from the adversaries' perspective) of the attack.

2. A risk-based approach is a useful way to strategically allocate resources.

Next we discuss both of these insights in more detail.

3.2.1. Insight: The Range of Uncertainties Regarding the Probability of Outcomes When People's Behaviors Are Involved Can Be Large. The Black and Gray scenarios in Table 5 described only the mean value of a probability distribution. From our experience with the TSA study, when the probability distributions are considered, the plausible ranges of possible outcomes were broad. There were significant uncertainties at each step of the security process, and these uncertainties continued to propagate through the analysis. We provide Figure 4 as a demonstration of a practical way to convey results with significant uncertainty to a decision maker. To understand Figure 4, the median value is shown by the line in the box. The 25th percentile and 75th percentile are shown by the box, and the 5th and 95th percentiles are shown by the lines. As imagined in the Black scenario, the range of possible values between the 5th and 95th percentiles are between 5% and 60%. The equivalent range for the Gray scenario is 2% to 40%. Typically, decision makers would feel very differently about the probability of a successful attack being 5% versus 60% and 2% versus 40%. (Figure 4 is a demonstration of a graphical tool to convey important information about the range of possible likely outcomes.) While this example used illustrative numbers, this example demonstrates the types of results available to the decision maker from this type of modeling and the broad range of uncertainties that may need to be considered when we do not know how untrained individuals will respond when faced with a terrorist adversary that has successfully traversed the screening checkpoint.

3.2.2. Insight: A Risk-Based Approach Is a Useful Way to Strategically Allocate Resources. If different risks could be assigned to different flights, such a score could be used to assign limited resources to flights

Figure 4. Distributions of Attack Success Probability for Two Post-Screening Detection Scenarios

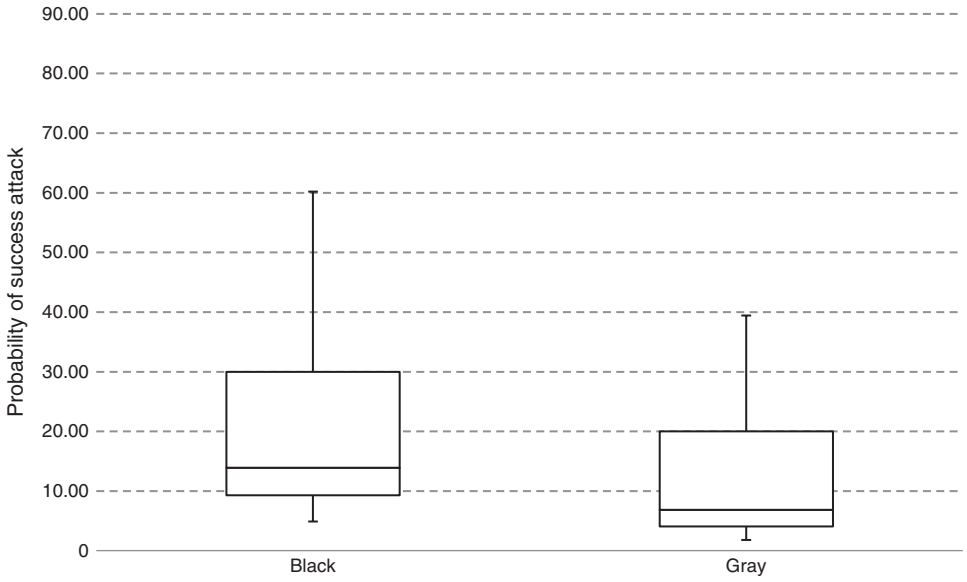

post-checkpoint screening. This may be considered risk-informed decision making since many other factors are involved, but in our study we refer to this as risk-based decision making and so continue to use this term. We did not consider the strategic allocation of additional security resources in the sterile area in our study, but this should be considered in future studies.

Without some method of identifying and prioritizing flights post-checkpoint screening, determining clear rationale for assignments of extra security resources is challenging. Factors that differ among flights include origin city (constant within an airport being considered), destination city, aircraft type, airline, time of day, departure gate, screening areas (for airports with multiple areas), traveling armed passengers (i.e., federal and local law enforcement), federal flight duty officers, federal air marshals, and preidentified risky passengers. We discuss each of these factors briefly:

—The origin city and destination city may have symbolic value (e.g., New York City, Washington, DC), but with an adaptive adversary, flights can be chosen for more random reasons. Abdulmutallab is believed to have targeted Flight 253 to Detroit because the ticket was cheaper than tickets to either Houston or Chicago (Caulfield 2011). Thus, within an airport, all flights have the same origin city, and prioritizing on a perceived higher value destination city has questionable value with an adaptive adversary. Risks associating with passengers from connecting flights (i.e., different initial city of origin) are discussed with another factor, prescreened passengers, in the following section.

—The origin or destination city may also influence the type of aircraft available because only some U.S. airports have wide-body aircraft and others are predominantly regional jets. The aircraft type will be considered further in the discussion of estimating consequences since larger aircraft with more people would result in more fatalities in a successful attack.

—The departure gate and the airline should have little influence on risk because gates are interchangeable and airlines have similar procedures.

—The time of day and different screening areas could affect the probabilities of a successful attack if screening procedures are different in different areas or if more or less crowded security can affect screening, but standard operating procedures are in place to reduce the chances of inconsistent screening.

—Traveling armed passengers (i.e., federal and local law enforcement), federal flight duty officers, and federal air marshals are all knowable to the TSA in advance and are unlikely to be knowable to the adversaries unless discovered at flight time. But as with anyone's travel plans, these individuals' travel plans could change.

—Identifying in advance risky passengers through the information collected when flight reservations are made is clearly the most effective factor in determine a risk for a particular flight. This would include detailed information about passengers whose travel originates outside the United States and who have connected to a domestic flight.

3.3. Consequences of the Catastrophic Loss of an Aircraft

The costs of security are justified when they offset the magnitude of potential losses. In the case of warning systems for commercial aviation, the costs of security accumulate very quickly because of the number of travelers in the system every day. For example, approximately 780,000 passenger flights take place monthly across the United States.[4] When we consider the number of travelers who use commercial aviation and assume even modest wait times, with reasonable values for business and leisure travelers' time, costs of security can easily reach into the billions of dollars per year. Traveling passengers and airline organizations do not have an inexhaustible tolerance for inconvenience and other security costs. For these reasons, we developed a framework to estimate the impact of a catastrophic loss of the aircraft, since it is the avoidance of this cost that justifies the inconvenience and security costs.

When considering the loss of an aircraft from a terrorist attack, economic impacts are typically attributed to those costs that follow directly from the event versus those costs that follow indirectly as a result of public and institutional response to the event. Direct impacts include loss of life and injuries, property damage, cleanup, rescue, and short-term business interruption. These losses have relatively straightforward actuarial interpretations (e.g., value of a statistical life, depreciation of an aircraft, temporary loss of revenue while affected businesses are closed). Indirect impacts involve broader and more long-term economic costs

to the U.S. economy. These impacts ripple and possibly amplify through the economy as the public, media, business and government react to the attack. Because these indirect losses depend heavily on the context of the attack, they are much more difficult to estimate. Several studies suggest that indirect impacts can be one or two orders of magnitude larger than direct impacts (Burns and Slovic 2007; Giesecke et al. 2012, 2015; Rose et al. 2016) and hence become an important consideration when planning security measures. Direct impacts are well understood, but indirect impacts are not.

Depending on the degree of success of the attack, the loss of life can be substantial, and so it is natural to inquire how economic consequences might depend on loss of life. In the case of direct impacts, most costs are proportional to the loss of life. However, because indirect impacts are largely context dependent, loss of life may be only one of many factors to consider. In the case of commercial aviation, the potential for loss of life depends on the size of aircraft. For purposes of illustration, we discuss three aircraft types: regional, narrow body, and wide body aircraft. Regional jets are aircraft that transport 20–100 passengers, narrow body aircraft transport 101–200 passengers, and wide body transport 201–300 passengers, based on a survey of commercial aircraft currently in use. Focusing on different aircraft types is important for at least two reasons: (1) Regional jets are found in larger numbers at our nation's smaller airports, while wide body aircraft are found only at our country's busiest airports and are used primarily for international travel; hence calculating economic losses for different aircraft allows a comparison of possible impacts across our nation's airports. (2) Aircraft passenger capacity itself might be an important factor. For example, 100 casualties on a regional jet means there were no survivors, whereas 100 casualties on a wide body aircraft means that over half of the passengers survived the attack. Since the narrative following these two scenarios is likely to be different, the indirect economic consequences are likely to be different as well. Hence, we explore possible indirect impacts due to loss of life across a range of 1 to 300 passengers.

We draw upon a study (Rose et al. 2016) sponsored by the TSA to provide the indirect estimates we need. This study involved a longitudinal experiment that provided survey inputs to a national computable generalized equilibrium (CGE) economic model that estimated the indirect effects of an attack. Researchers created a scenario of an attack on a commercial passenger aircraft leaving Los Angeles International Airport (LAX). Shortly after takeoff, a bomb detonates killing two passengers and injuring ten others, but the plane manages to land safely. A national survey panel reviewed a video newscast of this hypothetical attack and were asked to imagine they had just heard this report on the news. They were then questioned about current plans to fly and whether, after hearing this news, they would either continue to fly as planned, switch to another mode of transportation, or cancel their travel plans entirely. Based on these inquiries, it was estimated that about 7% of planned airline trips would be canceled in the first year following such a terrorist attack. A CGE model using data from the survey estimated the economic losses to the U.S. economy to be about \$8 billion (2012 dollars) during this first year primarily caused by reported fear. The indirect impacts of this attack were determined to be about nine times larger than the direct impacts.

The study as described involved only two deaths. To explore the relationship between loss of life and economic impact, we represent economic loss as a power function:

$$y = kx^{\alpha}. \tag{1}$$

Using the standard power function equation and the parameters from the aforementioned study, \$8 billion for an attack with two deaths ($k = \$ \times$ billion), and $x = z/2$, where z is the number of deaths for a particular attack, this yields the following power function:

$$y = 8(z/2)^{\alpha}. \tag{2}$$

When $\alpha = 0$, economic losses are \$8 billion for all number of deaths and when $\alpha = 1$, economic losses are linear in loss of life.

In Table 6, we consider a range of possible lives lost (1–300 passengers) and vary α for values in the range of 0.0–0.5. Notice that the estimates for the indirect impacts for the first year after an attack range from \$5.7 billion (1 death, $\alpha = 0.5$) to \$98 billion (300 deaths, $\alpha = 0.5$). To put these estimates in perspective, consider that the annual budget for the TSA in 2012 was \$8.1 billion, and the estimated economic loss for the attack involving two deaths was \$8 billion. Even

Table 6. The Indirect Economic Impacts ($ Billion) of an Attack During First Year Varying Loss of Life and α

Loss of life	Values of α for power function										
	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
1	8.0	7.7	7.5	7.2	7.0	6.7	6.5	6.3	6.1	5.9	5.7
50	8.0	9.4	11.0	13.0	15.2	17.9	21.0	24.7	29.0	34.1	40.0
100	8.0	9.7	11.8	14.4	17.5	21.3	25.9	31.5	38.3	46.5	56.6
150	8.0	9.9	12.3	15.3	19.0	23.5	29.2	36.3	45.0	55.8	69.3
200	8.0	10.1	12.7	16.0	20.1	25.3	31.8	40.1	50.5	63.5	80.0
250	8.0	10.2	13.0	16.5	21.0	26.7	34.1	43.4	55.2	70.3	89.4
300	8.0	10.3	13.2	17.0	21.8	28.0	36.0	46.2	59.4	76.3	98.0

an attack with only two deaths and the plane landing safely is equivalent to approximately 99% of the TSA's annual operating budget. Conversely, according to Rose (2009), the economic losses attributed to the attacks on September 11 were approximately $109 billion over a multiyear period. If we assume the indirect impacts are about nine times larger than the direct impacts (Rose et al. 2016), then the indirect effects of the attacks on September 11 were about $98 billion over multiple years. Therefore, it appears reasonable to vary α over the range from 0.0–0.5.

Similar studies suggest indirect impacts depend of the context of the attack. Rose et al. (2016) also estimated the indirect impact of a hypothetical attack on the passenger security checkpoint area at LAX in which 20 people are killed and 98 people are injured to be $12.2 billion during the first year. The indirect impacts were found to be 73 times that of the direct impacts. In a different domain, Giesecke et al. (2012) used survey input and CGE modeling to estimate the indirect impacts of a dirty bomb attack on the financial district of Los Angeles. In this attack, 180 people were killed, and the researchers estimated the indirect impacts would be $2.6 billion the first year and $15.8 billion over a 10-year period. In this study, the indirect impacts were 15 times that of the direct impacts. Clearly, indirect impacts depend on context, of which loss of life is but one factor.

From Table 6, we can make a few observations and generalizations. First, if $\alpha = 0$, then the economic impacts of a single attack in the first year are the same regardless of loss of life. One interpretation is that perceptions of risk are based more on the fear that terrorists just successfully detonated a bomb on a commercial aircraft than on the concern of how many people were just killed. If this were empirically true, adversaries would have an incentive to attack the most vulnerable airports. These might be smaller airports with limited screening resources, or airports that based on adversary surveillance, appear to deviate in a predictable manner from standard security operating procedures.

Second, it would be very unlikely that the indirect economic impacts during the first year following a single attack would exceed 50% of the economic impacts of the attacks on September 11, calculated over multiple years. Hence, from Table 6, reasonable values of α for this single attack should be between 0.0 and 0.35.

Third, within this bounded range for α, economic impacts range from $6.3 billion to $46 billion. This suggests that at the low end a successful attack would exceed 70% of the TSA annual budget, and at the high end would be over five times that of the TSA's annual budget. Hence, any successful attack, regardless of loss of life, would represent a significant economic event.

Fourth, we see that for any value of α within the range of 0.0 to 0.35, more than 50% of the maximum economic impacts occur for loss of life from 1–50 passengers, and in some cases this contribution is over 90%. In many instances, the economic impacts for the loss of a regional jet versus wide body aircraft are not much different. This suggests that the loss of a smaller aircraft taking off from one our nation's smaller airports represents an economic threat comparable to that for the nation's largest airports. This threat is accentuated by the fact that an adversary can readily transfer to a larger aircraft without going through additional security at most of our country's airport hubs.

Fifth, the values in Table 6 do not reflect possible inflection points in loss of life. These would be values for loss of life that represent a disproportional change in economic impact, perhaps due to their psychological significance. For example, a value of 0 would indicate everyone survived, whereas values of 100, 200, and 300 deaths for regional, narrow body and wide body aircraft, respectively, indicate that no one survived the attack.

From modeling the consequences of an economic loss from a successful attack as a power function of number of deaths, two insights were gained:

1. Indirect economic impacts greatly exceed direct impacts because fear drives trip cancellations.

2. Since it is hard to predict how individuals will respond to a future terrorist attack, the range of uncertainties is large for the consequences with many contributing factors that could be further researched.

Next we discuss both of these insights in more detail.

3.3.1. Insight: Indirect Impacts Greatly Exceed Direct Impacts Because Fear Drives People's Responses to Terrorism. It is possible to explore the connection between loss of life and indirect economic impacts due to fear, for example, in the following way. To estimate consequences in Table 6, the power function in Equation (2) was used. However, when estimating the economic impacts of an attack, canceled trips due to security concerns were the primary factor driving the economic impacts. Economic losses were found to be approximately linearly related to the percentage of canceled trips. Recall that in the Rose et al. (2016) attack scenario, there were two deaths resulting in a 6.9% trip cancellation and \$8 billion economic loss. We can specify this economic impact as a function of change in cancelled trips. For example, indirect economic impacts can be expressed as

$$y = 8(p\%/6.9\%), \tag{3}$$

where 6.9% is the percentage of canceled trips resulting from an attack with two deaths and \$8 billion of economic impacts. Likewise, $p\%$ is the predicted numbered of canceled trips for some number of deaths z from the power function. Setting these expressions equal,

$$y = 8(z/2)^{\alpha} = 8(p\%/6.9\%) \tag{4}$$

and solving for α:

$$\alpha = \ln(p\%/6.9\%)/\ln(z/2). \tag{5}$$

By examining how the percentage of canceled trips $p\%$ caused by fear vary as the number of deaths z, reasonable values of α can be obtained.

For example, suppose survey respondents are asked whether they would cancel their travel plans following an attack with 100 deaths. Suppose further we find the sample cancellation percentage to be 14% (about two times larger than for two deaths). Computing α, we have

$$\alpha = \ln(14\%/6.9\%)/\ln(100/2) = 0.18. \tag{6}$$

This corresponds to an economic loss of \$16.2 billion. Assessing trip cancellation in relation to loss of life is an easier judgment than going directly from loss of life to economic loss and will be a significant factor in the indirect impacts that occur following a terrorist event.

3.3.2. Insight: There Is a Large Uncertainty Regarding How the Nation Will Respond to Another Terrorist Attack. There are multiple factors that are not well understood that will affect the consequences. For example, does an attack that kills all 100 passengers on a regional jet evoke the same level of public reaction as an attack that kills 100 passengers on a narrow body aircraft carrying 200 passengers? To what extent does even one survivor change the narrative of the event and most probably the economic impacts? We assumed in Table 6 that 100 passengers killed was the important variable without considering the number of passengers that survive.

The consequences model also does not consider contextual factors other than loss of life. Other factors that could be relevant to a public's response include airport size, flight origination and destination, TSA or airline negligence, or type of passengers (e.g., children, celebrities, political officials) on board.

Additionally, in all of the analyses, we assumed that a successful attack on a commercial aircraft was clearly defined. However, more ambiguous scenarios are possible. Suppose many passengers are killed during an attack, but the plane lands safely. The impact of attack location is also unclear. Is an attack in the sterile area of an airport as consequential as an attack while the plane is in the air? Is an attack in the public areas of an airport as consequential as an attack in either the sterile areas or while the plane is in the air? What role does loss of life play on public reactions and behavior? How will a successful attack that kills no one be construed? Some researchers refer to such an event as a near-miss (Dillon and Tinsley 2008, Dillon et al. 2014). The public and media could portray this near-miss as demonstrating that commercial aviation is resilient to attacks. On the other hand, the public could construe this event as a tragedy that almost happened. Public attribution would affect demands for increased security measures. Additionally, with everyone surviving this attack, there would be some tales of heroism that could significantly change the narrative in support of the resilience to attack perspective.

Public reactions to terrorist attacks in the United States are not well understood, partly due to their relatively low frequency; hence, there is great uncertainty associated with the consequences of either a successful or "near miss" terrorist attack on aviation.

4. Discussion

From our study, we identified six insights of importance to the TSA in managing resources to protect the passenger vector. These insights are relevant to other critical alarm-based warning systems. We review each in terms of their general applicability beyond airport screening systems.

Insight 1: *The value of adding a countermeasure strongly depends on what other countermeasures are being used and what the threats are.*

Consider cyber attacks, where countermeasures include antivirus software, firewalls, intrusion detection, encryption, etc. (Kumar et al. 2008). Organizations should recognize that each component is part of a portfolio interacting with other components, and it is critical to understand the value of portfolios of different types of security measures in the context of different threats. Countermeasures in cybersecurity focus on threat prevention, detection, and recovery. With increasing threats from ransomware, many organizations have found gaps in their countermeasure portfolio regarding recovery. Organizations responsible for critical alarm-based warning systems should consider both the threats to their systems and their current countermeasures before "ladling on."

Insight 2: *Countermeasure detection rates are limited by the effectiveness of the resolution.*

Continuing with examples from the cybersecurity field, reports following the Target data breach in November–December 2013 described how their malware detection tool from FireEye provided warnings about a possible intrusion but were ignored (Albanesius 2014). Organizations should emphasize a safety culture and follow through with appropriate actions that maintain vigilance in response to critical warnings (Tinsley et al. 2011). Regarding the TSA in particular, the public needs to recognize that thorough pat-downs to resolve system alarms are a critical part of the security process (Martin 2017).

Insight 3: *The range of uncertainties regarding the probability of outcomes when people's behaviors are involved can be large.*

While alarm systems can be tested and calibrated, human components of the system are inherently less predictable. Repeating the example from Y-12, it would be difficult to predict that a security guard would misconstrue hammering as maintenance rather than a threat. Decision makers should understand that even the best constructed risk analysis models will have large uncertainties associated with possible outcomes when human actions are a critical step in the resolution process.

Insight 4: *A risk-based approach is a useful way to strategically allocate resources.*

Many agencies within the U.S. Department of Homeland Security (DHS), including the TSA, describe the risk of terrorist attacks as a function of three components: threat, vulnerability, and consequences (Ezell et al. 2010). Because of the difficulty of assessing threat, and the dire consequences that would likely follow a successful attack on an airline, the TSA initially sought to reduce the nation's vulnerability by screening all passengers with the same intensity regardless of their individual risk level. In 2011, the TSA expanded its risk-based security strategy to include an expedited screening program (TSA Precheck). Passengers were prescreened by the precheck system, so that passengers deemed lower risk could be expedited. Precheck allows TSA to allocate resources in a way that increases security and is more cost effective in the long term. It also improved passengers' experience going through the screening checkpoint area. As the TSA looks to the future, an understanding of the effectiveness of different countermeasure combinations could be used in conjunction with real-time individual passenger threat assessment to determine the appropriate combination of countermeasures to achieve a desired level protection across all passengers against a wide array of attack devices.

TSA's risk-based approach to security has been expanded throughout the DHS, and now includes numerous programs designed to screen U.S. travelers re-entering the United States from other countries. These systems all follow the same approach of requiring U.S. travelers to provide background information that allows a vetting process to assign a level of risk. Examples include Global Entry, NEXIS (travel across the U.S.–Canadian border), and SENTRI (travel across the U.S.–Mexican border). Such risk-based approaches

are common at industrial, military, and government sites, where employees and contractors are routinely vetted and issued credentials that allow minimal screening for entry to sensitive areas, allowing security resources to focus on unknown guests and visitors seeking admission.

The application of such risk-based approaches at public venues, such as shopping areas, entertainment halls, and sports arenas and stadiums, is a promising next direction. Hall et al. (2008) provide a risk assessment model for sporting events to identify and evaluate potential threats. Similarly, researchers at the Department of Homeland Security's Center of Excellence for Command, Control, and Interoperability Center for Advanced Data Analysis (CCICADA 2013) review best practices for risk assessment and management of sporting venues. Such risk-based security systems could be modeled on the TSA Secure Flight system that assigns risk level to travelers based on a variety of databases available. Such a system for public events could incorporate information related to the ticket purchase, much like the TSA obtains information about travelers from the airlines. Adoption of such risk-based approaches to security beyond airports and border crossings is likely to optimize the effectiveness of security allocated by allowing minimal screening for those who are vetted, and greater scrutiny of those who are unknown.

Insight 5: *Indirect impacts greatly exceed direct impacts because fear drives people's responses to terrorism.*

Kasperson et al. (1988) were among the first to offer an explanation as to why the indirect economic, social, and political impacts of certain hazards tend to far exceed what would be expected based on an event's direct impacts such as loss of life. The authors' social amplification framework suggests that hazards that inspire dread and whose outcomes are highly uncertain, poorly understood, or difficult to control are likely to cause a great deal of media attention, public concern, and possible institutional reaction. A number of studies (Pidgeon et al. 2003, Slovic 2010) support this explanation. According to Robinson et al. (2010) societal fear and anxiety are among the most important consequences of terrorism. Terrorist acts seek to inspire fear by causing the public to believe that they have little power to protect themselves and by eroding public confidence that societal institutions can effectively protect them. Hence, the context of an event plays a defining role in the magnitude of its indirect impacts. For example, in the immediate aftermath of an attack on a commercial airline, the public would naturally show great concern. However, the long-term impacts will be largely determined by what this event signals about the TSA's ability to handle the risk of future attacks.

Insight 6: *There is a large uncertainty regarding how the nation will respond to another terrorist attack.*

There is considerable uncertainty surrounding public response to future attacks. As previously mentioned, public reaction depends on the context of the attack and what it signals for future risk. Opinions among researchers differ on how the public might adjust to such risk. For example, Stewart and Mueller (2017), in a review of a wide array of terrorist attacks, acknowledge that economic impacts depend largely on indirect losses due to public reaction (e.g., reduction in flying) and to institutional reactions (e.g., increased security measures, military action). However, the authors caution against using public reaction to the attacks on September 11 as a benchmark. They observe that while the attacks on September 11 had a singular impact on air travel and tourism, subsequent attacks on transportation (e.g., London Underground, 2005, 52 deaths; Madrid commuter train, 2004, 191 deaths) have had only moderate to small impacts on travel and tourism despite significant loss of life and injuries. They note the public adapted to these new risks by simply switching tourist destinations, not by canceling or appreciably altering their travel plans.

Other research suggests there can be conditions that induce a much stronger public response. Burns and Slovic (2007) systematically varied contextual factors to determine their influence on perceptions of risk, fear, and willingness to resume normal activities following an attack. They found that factors such as attack mechanism and possible negligence in security significantly altered reported perceptions of risk and behavioral intentions. Their conclusion is that context matters, and they predict that when an event suggests a risk is not under control public reaction will be significant. Consider for example, the recent attacks in France in which terrorists decided to focus on soft targets in Paris 2015 (130 deaths) and in Nice 2016 (86 deaths). Public reaction was pronounced, security measures were significantly increased throughout the country, and terrorist

risk became a focal point in the French presidential election in 2017. It is not unreasonable to think that recent attacks in Europe and the United States influenced the U.S. presidential election and the current administration's call for travel restrictions and a wall along the border between the United States and Mexico.

5. Conclusions

Learning should be possible for the TSA, as well as terrorists, especially given possibilities for development of new technologies and methods and the acquisition of knowledge related to ongoing operations. Updating risk analyses is a good way of consolidating such learning, and this work provided the TSA with a framework to consider security and resource allocations at the time of the study and in the future. As a contribution, this paper describes this method for conducting a PRA of a critical component of commercial aviation security, the passenger screening. Also, from the analysis, important results relevant to strengthening passenger screening are provided that include the following: (1) developing training and incentives for transportation security officers (TSOs) to perform alarm resolution procedures that always conform to standard operating procedures (SOP); (2) always combining countermeasures in portfolios that address the individual limitations of certain countermeasures against specific attack devices; (3) realizing that if multiple portfolios are relatively equivalent in effectiveness (given uncertainties), then changing portfolios over time can help thwart terrorists; and (4) recognizing the overall effect of marginally decreasing reduction in detection probabilities as more countermeasures are added to the screening configuration. Finally, broader insights regarding the value of additional countermeasures, the reality of large uncertainties, the challenges of risk-based approaches, and the role of indirect outcomes are provided to organizations that rely on critical alarm-based warning systems and that may or may not face risks from future terrorist attacks.

Acknowledgments

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security or the National Science Foundation. The authors also acknowledge the cooperation of the TSA, who provided both the experts and support for conducting probabilistic elicitations.

Endnotes

[1] A red team is an independent group that challenges an organization to improve its effectiveness.

[2] For Tables 1 and 2, the actual questions were more specific but providing any more detailed version of the questions would expose TSA SOPs.

[3] More research is needed to understand to what extent alarms from passengers are attended to and how many are false alarms, but we assume that if a weapon is detected, interdiction will occur.

[4] Homeland Security, U.S. Department of Transportation, Transportation System Sector-Specific Plan, 2015.

References

Albanesius C (2014) Target ignored data breach warning signs. *PC Magazine*. Accessed August 1, 2017, https://www.pcmag.com/article2/0,2817,2454977,00.asp.

Bell L (2010) Monitor alarm fatigue. *Amer. J. Critical Care* 19(1):38.

Burns WJ, Slovic P (2007) The diffusion of fear: Modeling community response to a terrorist strike. *JDMS: J. Defense Model. Simulation: Appl., Methodology, Tech.* 4(4):298–317.

Caulfield P (2011) Christmas 2009 "underwear bomber" targeted Detroit because it was the cheapest flight: Report. *New York Daily News* (March 24), http://www.nydailynews.com/news/national/christmas-2009-underwear-bomber-targeted-detroit-cheapest-flight-report-article-1.118654.

CCICADA (Command, Control and Interoperability Center for Advanced Data Analysis) (2013) BEST PRACTICES in Anti-Terrorism Security for Sporting and Entertainment. Report submitted to the U.S. Department of Homeland Security. Accessed August 11, 2017, https://www.safetyact.gov/externalRes/refdoc/CCICADA%20BPATS.pdf.

Department of Homeland Security (2016) Website for DHS "See something, say something" campaign. Accessed May 21, 2018, https://www.dhs.gov/see-something-say-something.

Dillon RL, Tinsley CH (2008) How near-misses influence decision making under risk: A missed opportunity for learning. *Management Sci.* 54(8):1425–1440.

Dillon RL, Tinsley CH, Burns WJ (2014) Evolving risk perceptions about near-miss terrorist events. *Decision Anal.* 11(1):27–42.

Elwell R (2008) The threat to aviation by terrorist acts. *The Log* (British Airline Pilots Association, West Drayton, UK), 11.

Ezell BC, Bennett SP, Von Winterfeldt D, Sololowski T, Collins AJ (2010) Probability risk analysis and terrorism risk. *Risk Anal.* 30(4):575–589.

Giesecke JA, Burns WJ, Rose A, Barrett A, Griffith M (2015) Regional dynamics under adverse physical and behavioral shocks: The economic consequences of a chlorine terrorist attack in the Los Angeles financial district. Nijkamp P, Rose A, Kourtit K, eds. *Regional Science Matters: Studies Dedicated to Walter Isard* (Springer, Berlin), 319–350.

Giesecke JA, Burns WJ, Barrett A, Bayrak E, Rose A, Slovic P, Suher M (2012) Assessment of the regional economic impacts of catastrophic events: CGE analysis of resource loss and behavioral effects of a RDD attack scenario. *Risk Anal.* 32(4):583–600.

Graham K, Cvach M (2010) Monitor alarm fatigue: Standardizing use of physiological monitors and decreasing nuisance alarms. *Amer. J. Critical Care* 19(1):28–35.

Hall S, Marciani L, Cooper WE, Rolen R (2008) Introducing a risk assessment model for sport venues. *Sport J.* Accessed May 21, 2018, http://thesportjournal.org/article/introducing-a-risk-assessment-model-for-sport-venues.

Hilzenrath D (2010) Technician: Deepwater Horizon warning system disabled. *Washington Post* (July 23), http://www.washingtonpost.com/wp-dyn/content/article/2010/07/23/AR2010072305419.html.

Jackson BA, LaTourrette T, Chan EW, Lundberg R, Morral AR, Frelinger DR (2012) *Efficient Aviation Security: Strengthening the Analytic Foundation for Making Air Transportation Security Decisions* (RAND Corporation, Santa Monica, CA).

Kasperson RE, Renn O, Slovic P, Brown HS, Emel J, Goble R (1988) The social amplification of risk: A conceptual framework. *Risk Anal.* 8(2):177–187.

Knight FA (1921) *Risk, Uncertainty, and Profit* (Houghton Mifflin Company, Boston).

Kowalczyk L (2010) Alarm fatigue linked to patient's death. *Boston Globe* (April 3), http://archive.boston.com/news/local/massachusetts/articles/2010/04/03/alarm_fatigue_linked_to_heart_patients_death_at_mass_general/?page=1.

Kumar RL, Park S, Subramaniam C (2008) Understanding the value of countermeasure portfolios in information systems security. *J. Management Inform. Systems* 25(2):241–279.

Martin H (2017) TSA quietly launches new "enhanced" pat-down procedure. *Los Angeles Times* (March 6), http://www.latimes.com/business/la-fi-tsa-pat-down-20170306-story.html.

Morgan MG (2014) Use (and abuse) of expert elicitation in support of decision making for public policy. *Proc. Natl. Acad. Sci.* 111(20):7176–7184.

Paté-Cornell ME, Dillon RL (2006) The respective roles of risk and decision analyses in decision support. *Decision Anal.* 3(4):1–13.

Pidgeon N, Kasperson RE, Slovic P, eds. (2003) *The Social Amplification of Risk* (Cambridge University Press, Cambridge, UK).

Robinson LA, Hammitt JK, Aldy JE, Krupnick A, Baxter J (2010) Valuing the risk of death from terrorist attacks. *J. Homeland Security Emergency Management* 7(1):14.

Rose AZ (2009) A framework for analyzing the total economic impacts of terrorist attacks and natural disasters. *J. Homeland Security Emergency Management* 6(1):1–26.

Rose AZ, Avetisyan M, Rosoff H, Burns WJ, Slovic P, Chan O (2016) The role of behavioral responses in the total economic consequences of terrorist attacks on U.S. air travel targets. *Risk Anal.* 37(7):1403–1418.

Sachs A (2016) Don't mind the wet nose: TSA enlists more dogs to screen passengers. *Washington Post* (January 21), https://www.washingtonpost.com/lifestyle/travel/dont-mind-the-wet-nose-tsa-enlists-more-dogs-to-screen-passengers-for-explosives/2016/01/20/26e11d98-b983-11e5-829c-26ffb874a18d_story.html?utm_term=.338197fb17ea.

Schlosser E (2015) Break-in at Y-12. *New Yorker* (March 9), https://www.newyorker.com/magazine/2015/03/09/break-in-at-y-12.a.

Slovic P, ed. (2010) *The Feeling of Risk: New Perspectives on Risk Perception* (Earthscan, London).

Stewart MG, Mueller J (2017) *Are We Safe Enough? Measuring and Assessing Aviation Security* (Elsevier, Cambridge, MA).

Tinsley CH, Dillon RL, Madsen PM (2011) How to avoid catastrophe. *Harvard Bus. Rev.* (April):1–9. [Reprint IR1104G.]

U.S. Nuclear Regulatory Commission (1990) Severe accident risks: An assessment for five U.S. nuclear power plants. *NUREG*-1150 (U.S. Nuclear Regulatory Commission, Washington, DC).

Robin L. Dillon is a professor in the McDonough School of Business at Georgetown University. She seeks to understand people's decision making following near-miss events. She has received research funding from the National Science Foundation, NASA, the Department of Defense, and the Department of Homeland Security. She has a B.S./M.S. from the University of Virginia and a Ph.D. from Stanford University. For AY 2017–2018, she is on an Interagency Personnel Agreement (IPA) assignment at the National Science Foundation as the program director for the Humans, Disasters, and Built Environment program.

William J. Burns completed his Ph.D. at the University of Oregon in decision science and has held positions at the University of Iowa and the University of California, Davis. He is currently a research scientist at Decision Research (Eugene, Oregon), associated with the Center for Risk and Economic Analysis of Terrorism Events (CREATE at USC) and a part-time faculty member at California State University, San Marcos and the University of Redlands. His work has been funded by the Department of Homeland Security and the National Science Foundation.

Richard S. John is an associate professor in psychology at the University of Southern California (USC) and an associate director at the Center for Risk and Economic Analysis of Terrorism Events (CREATE). His research focuses on normative and descriptive models of human judgment and decision making and methodological issues in the application of decision analysis and probabilistic risk analysis. He has a Ph.D. in quantitative psychology from USC, an M.S. in applied mathematics from USC, and a B.S. in applied mathematics (summa cum laude) from the Georgia Institute of Technology.