2010

Is Screening Cargo Containers for Smuggled Nuclear Threats Worthwhile?

Jason R. W. Merrick
Virginia Commonwealth University, jrmerric@vcu.edu

Laura A. McLay
Virginia Commonwealth University

© 2010 INFORMS. This is the author's version of a work that was accepted for publication as Merrick, J. R. W. and McLay, L A. (2010) Is Screening Cargo Containers for Smuggled Nuclear Threats Worthwhile? Decision Analysis 7(2): 155-171. http://dx.doi.org/10.1287/deca.1100.0171

Is Screening Cargo Containers for Smuggled Nuclear Threats Worthwhile?

Jason R. W. Merrick

Virginia Commonwealth University

Richmond, VA, 23284, jrmerric@vcu.edu

Laura A. McLay

Virginia Commonwealth University

Richmond, VA, 23284, lamclay@vcu.edu

Abstract

In recent years, Customs and Border Protection (CBP) has installed radiation sensors to screen cargo containers entering the United States. They are concerned that terrorists could use containers to smuggle radiological material into the country and carry out attacks with dirty bombs or a nuclear device. Recent studies have questioned the value of improving this screening system with new sensor technology. The cost of delays caused by frequent false alarms outweighs any reduction in the probability of an attack in an expected cost analysis. We extend existing methodology in three ways to demonstrate how additional factors affect the value of screening investments. We examine the effect that screening has in discouraging terrorists. We model multiple levels of screening. Finally, we consider additional objectives beyond cost. We find that the conclusion about screening depends on key inputs to the probability model (reflecting uncertainties) and to the value function (reflecting the stakeholders' fundamental objectives).

Keywords: Applications: Terrorism; Probability: Applications; Multiple Objective Decision Analysis.

1. Introduction

Terrorist groups are trying to obtain radiological material to attack targets inside the United States with radioactive dispersal devices, or dirty bombs (Gardner 2003). One method for smuggling radiological material into the United States is a container. Containers and container ships carry about 90% of non-bulk goods worldwide (Ebeling 2009). 20.4 million containers enter the United States each year (Bonner 2005). US Customs and Border Protection (CBP) screens these containers looking for smuggled radiological materials. A radiation portal monitor (RPM) scans the container for radioactive sources. If the RPM detects radiation then CBP performs further screening or physically inspects the container. Additional inspection uses gamma ray and x-ray scanners. CBP personnel use radioactive isotope identification devices for a physical inspection. We study these critical screening and inspection decisions using decision analysis and extend previous work to determine which inputs to the decision structure, probability model, and value model are critical in determining whether and how to screen containers.

Decision analysis has been used widely in counter-terrorism decisions, including ranking critical infrastructures to protect (Haimes et al. 2002, Apostolakis and Lemon 2005), learning from multiple sources of intelligence information (Paté-Cornell 2002), protecting commercial airlines against missiles (von Winterfeldt and Sullivan 2006), determining which bridges to protect (Leung et al. 2004), and determining strategies for using Potassium Iodide treatments after exposure to radioactive iodine from a nuclear incident (Feng and Keller 2006). Game theory has also been widely applied, including general strategy and defense modeling (Kunreuther and Heal 2002, Zhuang and Bier

2007, Bier et al. 2007b, Zhuang et al. 2007) and specific applications, such as protecting power transmission systems (Bier et al. 2007a) and commercial airlines (Heal and Kunreuther 2005). Decision analysis has proved useful in modeling the sequential nature of these decisions and the uncertainties inherent in them. Game theory has proved useful in modeling the action and counter-action of the terrorist and counter-terrorist, including the ability to catch terrorists with secret defenses and deter terrorists with revealed defenses. Paté-Cornell and Guikema (2002) and Bier (2007) combine decision analysis and game theory to model both uncertainty and the action and counter-action of the opposing sides. Rios Insua et al. (2009) give an overview of this work and the issue of modeling terrorist decisions directly with game theory versus modeling the uncertainty about their actions with a probability distribution.

Several studies have looked at improving the screening process for containers with operations research methods. Lewis et al. (2003) look at the inspection problem faced by major "hub" container ports, like Singapore. These ports must screen containers as they pass through on the way to their destination port. Lewis et al. develop algorithms for minimizing delays while inspecting a given percentage of containers. Ramirez-Marquez (2007) consider a combination of different types of sensors, developing inspection strategies that minimize inspection cost while maintaining stated detection rates. Madigan et al. (2007) develop algorithms for selecting the optimal sequence of sensors and the optimal signal threshold for each sensor beyond which inspectors should take further action. Bakir (2008) considers a more specific problem, how best to protect the US-Mexican border. The options include screening on the Mexican side of the border and implementing other general security measures. Bakir's analysis includes the choice

on the US side of the border between new Advanced Spectroscopic Portals (ASP) and existing RPMs.

Bakir (2008) provides a wealth of probability and cost estimates. For instance, we would have to screen 10 billion containers to detect four true threats. However, we would also detect 249,999,996 naturally occurring radioactive material (NORM) sources, which are regular goods that give off radiation. Examples include ceramic tile and irradiated iron. Inspecting containers because of NORM sources is costly and delays the cargo. Bakir (2008) concludes that CBP should not implement new ASP technology. We examine whether even the original screening investments are worthwhile. We consider several extensions to Bakir's analysis and examine whether these extensions could potentially change the conclusions about screening.

Firstly, we consider the effect that screening has in deterring nuclear smuggling in containers. Radiological material is hard to get. Container transportation is only one avenue for smuggling it into the country; other avenues include freight vessels, air cargo, and private boats and aircraft. Terrorist groups will choose the method and route that give the best chance of getting the radiological material through. If CBP screens containers and does not screen the other avenues, terrorists will be less likely to use a container. However, the reverse of this argument is also true. If CBP does not screen containers, terrorists will be more likely to choose them. Thus, screening provides a deterrence effect. We examine the effect of deterrence on the screening decision and determine how large the effect needs to be before screening becomes cost-effective.

Secondly, we consider multiple levels of screening. Bakir (2008) assumes that if CBP screens a container and an alarm sounds, then they will perform an inspection. Does

CBP learn enough from an alarm to justify inspection? Examining the expected costs suggests the answer is no. The probability of a false alarm is much higher than the probability of a threat. One option for obtaining more information is to perform multiple levels of screening. We analyze this option with existing screening technology. To model multiple levels of screenings accurately, we consider the occurrence of both threats and NORM sources and their relative probabilities of causing alarms. We extend the analysis to determine how many times CBP should screen a given container.

Lastly, we include additional objectives in the screening decision. Bakir (2008) performs a thorough expected-cost analysis. Keeney (2007) suggests other objectives for counter terrorism decisions. We include objectives relevant to the screening decision in a value model. We find that the number of levels of screening depends heavily on the parameters of the value model that reflect stakeholder preferences.

We note that the techniques used herein are well developed. However, these simple extensions shed new light on this critical decision. Section 2 outlines the part of the analysis in Bakir (2008) we focus on, namely the decision whether to screen containers entering the United States. In Section 3, we extend the analysis to examine how large the deterrence effect must be before screenings become worthwhile. We consider the decision to perform multiple levels of screening in Section 4. Section 5 adds additional objectives to the analysis through a multi-attribute value function. We draw conclusions in Section 6 and provide recommendations for further research on this problem.

2. The Basic Screening Decision

Figure 1 shows a decision tree for deciding whether CBP should screen containers. One option is to screen the container with current radiation portal monitors and then decide whether to inspect it. The other option is to decide whether to inspect the container without screening. The decision is applicable to all containers entering the US. We use parameter values from Bakir (2008) in our decision. In the following, we use the term threat for radiological material smuggled in a container.

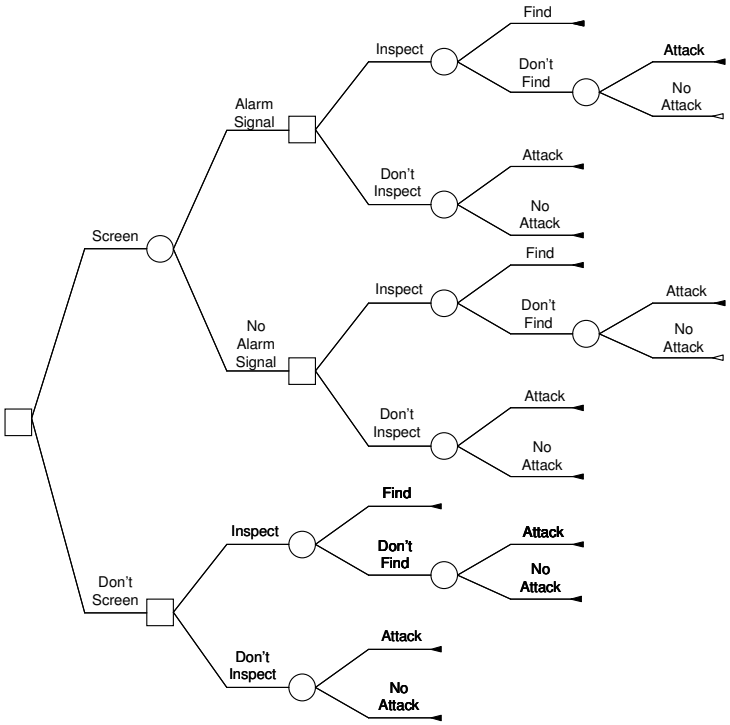

Figure 1. A decision tree for the container-screening problem with one screening.

When a container arrives at a port, CBP has the decision whether to screen it. If CBP screens the container, then either the alarm sounds, or it does not. Then CBP has the decision to inspect the container or to allow it through. Note that we are not assuming that CBP will automatically inspect the container if an alarm sounds. Instead, we consider

the conditions under which the information gained from an alarm is sufficient to justify inspection. CBP must make this same inspection decision if they choose not to perform screening, but without the additional information provided by the RPM. If CBP inspects the container then there is a chance that a threat will be found. If CBP does not inspect the container, or if they find nothing during an inspection, they allow it to pass through, and we observe whether there is an attack. We calculate the probability of an attack from the probability that there is a threat inside the container and the probability that CBP does not find the threat during an inspection. It also incorporates the probability that law enforcement agencies do not stop the attack once inside the country.

Several critical parameters and events drive the decision. We denote the event that a threat is inside any given container by T and the event that an alarm sounds by S. We use the over-score notation for the complement of each event. F denotes the event that CBP finds a threat during the inspection of a container and A denotes the event that an attack occurs. In terms of probabilities, we must define:

- the probability of a threat being inside a given container (p_T)
- the probability of an alarm or signal if there is a threat in the container (true alarm, $p_{S|T}$)
- the probability of an alarm or signal if there is not a threat in the container (false alarm, $p_{S|\bar{T}}$)
- the probability of finding the threat upon inspection ($p_{F|T}$)
- the probability of an attack if there is a threat in a container that is passed through ($p_{A|T}$)

In terms of costs, we must define:

- the average cost of screening each container (c_S), the average cost of inspecting a container (c_I)
- the cost of containing and disposing of radiological materials discovered in a container (c_F)
- the cost of an attack (c_A)

Table 1 shows three values of each parameter, a base value, as well as a minimum and maximum value. We base the values in Table 1 on those given in Bakir (2008).

Table 1. The base, minimum, and maximum values of the parameters for the container decision problem.

Parameter	Minimum	Base Value	Maximum
c_S	\$0	\$6	\$40
c_I	\$0	\$600	\$1,000
c_F	\$0	\$100,000	\$500,000
c_A	\$10 Billion	\$40 Billion	\$100 Billion
p_T	$5\text{x}10^{-11}$	$5\text{x}10^{-10}$	$5\text{x}10^{-9}$
$p_{S\|T}$	0.6	0.8	1
$p_{S\|\bar{T}}$	0.01	0.025	0.1
$p_{F\|T}$	0.1	0.9	1
$p_{A\|T}$	0.1	0.5	1

The General Accounting Office (GAO) performed an analysis of the cost and efficacy of container screening. Their report (GAO 2006) states that the cost of installing and maintaining radiation detection equipment over the next 10 years is $1.2 billion. Bakir (2008) uses this value as their cost basis. During the 10-year life cycle, 204 million containers will enter the country, if current cargo volumes remain constant. This gives an average cost of approximately $6 per container. Our use of an average cost does not allow for discounts if CBP implements screening more widely. However, we use a minimum value of $0 to test if this cost is a driving factor in the screening decision. Recently, the GAO has updated the total cost to $3.1 billion over the next ten years (GAO 2008). We use a maximum value of $40 per container to account for further increases and to allow for the use of new technology. The cost of inspection includes both direct costs and the cost of cargo delays. Bakir (2008) estimates this cost to be $600 per container. The maximum value for this is $1,000 per container, and we use $0 as the minimum value again.

The containment and removal of radiological material found during an inspection will require specialized personnel and equipment. Previous studies have not considered such costs. We estimate this cost at $100,000, but with a range from $0 to $500,000. This cost actually implies a penalty for finding radiological material during the inspection. However, it is a real cost and is much lower than the cost of an attack.

We assume that the attack uses a radioactive dispersal device (RDD). A more common name for an RDD is a dirty bomb. An RDD uses a standard explosive device to scatter radioactive material. It does less damage than a nuclear bomb, but it is less difficult to build. De Rugy (2007), Bronskill and Bailey (2007), and Rosoff and von

Winterfeldt (2007) each assessed the total cost of an RDD attack. Bakir aggregated their estimates to obtain a base value of $40 billion and a range from $10 billion to $100 billion.

Bakir (2008) estimates a probability of 0.1 that there will be an attempt in the next 10 years to smuggle radiological material inside a container entering the US. 204 million containers will enter the US in that period. The probability that there is a threat inside any given one is $5x10^{-10}$. We vary this probability between $5x10^{-11}$ to $5x10^{-9}$. Bakir estimates the probability of an attack at 0.855 given that a threat is successfully smuggled into the country. Discussion with intelligence officials suggests that this could be high. We set the probability to 0.5. The probability is the same if we do not inspect, or if we inspect and do not find the threat ($p_{A|T} = p_{A|T,\bar{F}}$). There can be no attack if inspectors find and contain the threat, so $p_{A|T,F}$ is zero. We vary the attack probability from 0.1 to 1. Bakir estimates the probability of finding a threat hidden inside a container during the inspection to be 0.9. We vary this probability from 0.1 to 1.

Bakir (2008) gives a true alarm rate of 80% and a false alarm rate of 2.5% based on current technology. The minimum values considered are 60% true alarms and 1% false alarms. The maximum values considered are 100% true alarms and 10% false alarms. Given these input values, we must calculate the probabilities required for the decision tree. Firstly, we must calculate the probability of an alarm signal (p_S). We use the law of total probability, conditioning on whether there is a threat in the container, specifically $p_S = p_{S|T}\, p_T + p_{S|\bar{T}}\, p_{\bar{T}}$. The base value of p_S is $3.13\ x10^{-4}$.

The remaining probabilities also depend on the probability of a threat ($5\text{x}10^{-10}$ at the base value). We must find the probability of a threat given the outcome of the screening, specifically given an alarm signal

$$p_{T|S} = \frac{p_{S|T}\, p_T}{p_{S|T}\, p_T + p_{S|\bar{T}}\, p_{\bar{T}}}$$

and given no alarm signal

$$p_{T|\bar{S}} = \frac{p_{\bar{S}|T}\, p_T}{p_{\bar{S}|T}\, p_T + p_{\bar{S}|\bar{T}}\, p_{\bar{T}}}.$$

At the base parameter values, $p_{T|S}$ is $1.6\text{x}10^{-8}$ (alarm signal) and $p_{T|\bar{S}}$ is $1\text{x}10^{-10}$ (no alarm signal).

We must calculate the probability of an attack (p_A), the probability of an attack given an alarm signal ($p_{A|S}$), and the probability of an attack given no alarm signal ($p_{A|\bar{S}}$). These probabilities can be calculated using $p_A = p_{A|T}\, p_T$ ($2.5\text{x}10^{-10}$ at the base values), $p_{A|S} = p_{A|T}\, p_{T|S}$ ($8.05\text{x}10^{-9}$ at the base values), and $p_{A|\bar{S}} = p_{A|T}\, p_{T|\bar{S}}$ ($5.00\text{x}10^{-11}$ at the base values).

Lastly, we must calculate the probability of finding a threat upon inspection (p_F), the probability of finding a threat upon inspection given an alarm signal on screening ($p_{F|S}$), and the probability of finding a threat upon inspection given no alarm signal on screening ($p_{F|\bar{S}}$). These probabilities can be calculated using $p_F = p_{F|T}\, p_T$ ($4.5\text{x}10^{-10}$ at the base values), $p_{F|S} = p_{F|T}\, p_{T|S}$ ($1.45\text{x}10^{-8}$ at the base values), and $p_{F|\bar{S}} = p_{F|T}\, p_{T|\bar{S}}$ ($9.00\text{x}10^{-11}$ at the base values).

Figure 2 shows the solution of the decision tree in Figure 1 using the base values for the probabilities and costs. The optimal decision is not to screen. Figure 2 shows that the probability of a threat increases when an alarm sounds, but not enough to justify screening. Bakir points out that the main reason for this is the rate of false alarms. The probability of an alarm is 0.025 at the base values. This includes both true alarms and false alarms, but how many of each? The probability of a threat is only 5×10^{-10} and 90% of these would lead to an alarm. This means that four in 10 billion containers would contain a threat that causes an alarm, while 249,999,996 of them would contain a NORM source that causes an alarm. Thus, most alarms are false. Figure 2 shows that even when the alarm sounds, CBP should not inspect the container, as the alarm is likely to be false. If it is not worth inspecting containers when the alarm sounds then it is not worth screening. We could also say that the expected value of information from screening is less than the cost of screening.

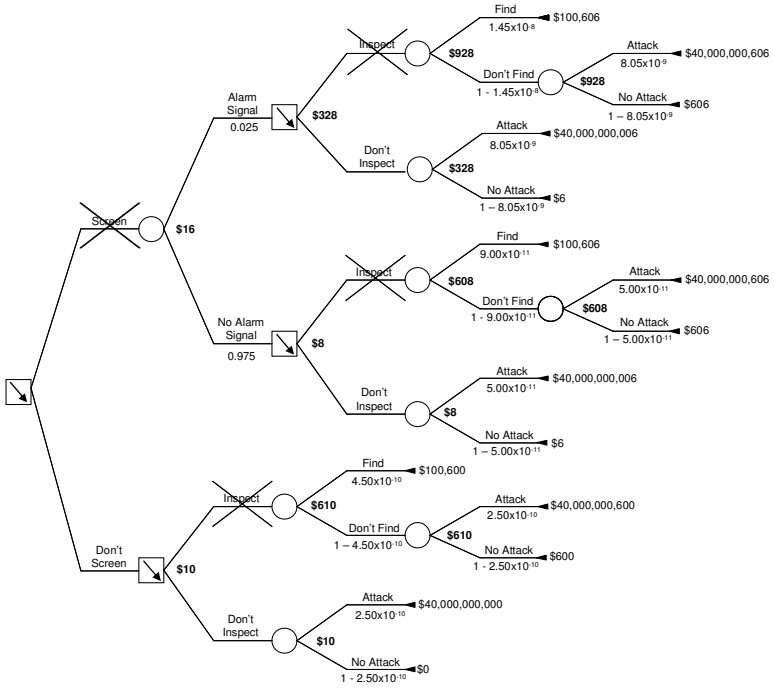

Figure 2. The solution for the container-screening problem with one screening.

3. The Effect of Deterrence

Thus far, we have implicitly assumed that terrorist groups will attempt to smuggle radiological material in containers, whether CBP screens them or not. This is because we have assumed that the probability of a threat is the same in either case. However, it is reasonable to assume that screening will influence terrorist groups' decisions. We hope that it will deter them from using containers for smuggling radiological material. CBP is the decision maker in this situation and to CBP the actions of terrorist groups are uncertainties. Thus, we model the deterrence effect of increased screening as a change in the probability of a threat.

Bakir (2008) estimates a specific change in the probability. The change represents the decision maker's beliefs about how screening will affect the terrorist group's decisions. Estimating the exact value of this change is a difficult judgment task for intelligence experts. We could consider it unknowable, rather than just unknown. Instead, we introduce a parameter to represent the change in threat probability. Bakir subsequently performed a sensitivity analysis on the change in probability. We suggest that one should start with a sensitivity analysis to find the parameter value at which the optimal decision changes. Then the decision maker only needs to judge which side of this value the parameter lies on, rather than attempting to estimate the parameter directly.

We let d be a deterrence multiplier that captures the ratio of the threat probabilities when CBP screen containers and when it does not. Given that we are modeling a decision without intelligence information specific to a container, our decision reflects the choice to screen all containers. Thus, the terrorists know about the choice to screen before initiating their attack. If we denote the decision to perform the screening

process by P, then $p_{T|\bar{P}} = dp_{T|P}$. Bakir (2008) used the values of the threat probability in Section 2 for the choice to keep existing screening technology. Thus, we use these values for $p_{T|P}$ when CBP screens containers. In Figure 1, the upper branch of the screening decision node represents the choice to screen and then to decide whether to inspect. The lower branch represents the choice not to screen, but still to decide whether to inspect. d represents the change in threat probability for the upper branch in comparison to the lower branch. Thus, we know that $d \geq 1$ as screening with possible inspection will have more of a deterrence effect than just possible inspection. Using this assumption, we obtain similar expressions for the probability of an attack and the probability of finding a threat upon inspection, namely $p_{A|\bar{P}} = dp_{A|P}$ and $p_{F|\bar{P}} = dp_{F|P}$.

Figure 3 shows the expected cost of screening and the expected cost of not screening. The expected cost of screening is higher than the expected cost of not screening when $d = 1$. This is the situation represented in Figure 2. The expected cost of not screening increases linearly with d. The expected costs of screening and not screening are equal when $d = 1.6$ and the expected cost of not screening exceeds that of screening for $d > 1.6$. Thus, our analysis shows that if the probability that a terrorist group will attempt such nuclear smuggling increases by 60% or more if we do not screen, then the lower cost option is to screen. Note that this calculation uses the base values of the parameters.

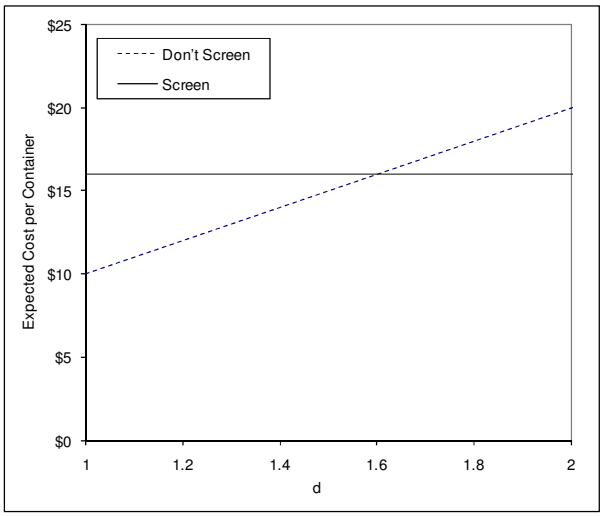

Figure 3. The expected cost of screening and not screening for varying levels of deterrence multiplier d

In summary, adding deterrence to our analysis can change the final recommendation. Thus, actually finding smuggled radiological materials is not the only reason to screen containers. This approach to modeling deterrence is simple for the basic screening decision, but it will be useful as we extend the model. The reader will note that the screening decision in Figure 1 has only two alternatives: screen or not. In practice, random inspections are often used when not screening. Randomly inspecting a percentage of containers would provide a deterrence effect and reduce the cost compared to the choice of inspecting all containers. However, this would mean that we would need to model a deterrence effect that depends on the proportion of containers screened, an even more complex judgment task. Thus, decision analysis is limited in this respect compared to game theory. Future work will address this limitation.

4. The Effect of Multiple Screenings

Screening may be economically viable if we take the effect of deterrence into account. Another way to improve the value of screening is to improve the accuracy of the

screening process, or more specifically to reduce the false alarm rate. CBP can reduce the false alarm rate with improved technology, and such efforts are ongoing. They can also reduce it with current technology by performing multiple levels of screening. However, the results of two screenings of the same container are dependent. We develop a model for the distribution of the results of multiple screenings of a container.

4.1 A Closer Look at False Alarms

NORM sources cause most false alarms in nuclear screening. Suppose we screen a container and get an alarm, but there is not a threat in the container. If we screen it again, there is a high probability that we will get a second alarm. Truly random false alarms from extraneous background sources of radiation or equipment errors are rare. False alarms from NORM sources inside the container (such as ceramic tiles, irradiated iron, and medical equipment) are frequent. Thus, we must extend the probability model from Section 2 to model multiple levels of screening.

Let N denote the event that there is a NORM source in a given container. p_N is then the probability of this event and $p_{S|N,\bar{T}}$ denotes the probability of an alarm given only a NORM source inside the container. We assume that screening equipment detects NORM sources as often as it does threats, so our baseline value for $p_{S|N,\bar{T}}$ is 0.8. From Section 2, 2.5% of containers give a false alarm and 80% of NORM sources cause an alarm. This implies that 3.125% of containers would have to contain a NORM source or $p_N = 0.03125$ ($0.03125 \times 0.8 = 0.025$). We also extend our notation from $p_{S|T}$ to $p_{S|T,\bar{N}}$ to represent the probability of an alarm given that there is a threat in the container, but no NORM source. If there is both a threat and a NORM source in the container then each of

these sources can send off gamma-rays and so each has a chance of independently causing an alarm. We calculate $p_{S|T,N} = p_{S|T,\bar{N}} + p_{S|\bar{T},N} - p_{S|T,\bar{N}} p_{S|\bar{T},N}$ assuming the independence of alarms caused by NORM sources and alarms caused by threats. We assume that without a threat or a NORM source, the probability of an alarm ($p_{S|\bar{T},\bar{N}}$) is zero. This assumption ignores false alarms from background radiation or equipment errors, but does not affect our calculations.

We must first find the probability of alarms caused by the combinations of threats and/or NORM sources. We can use simple conditional probability calculations to find the probabilities of an alarm given the contents of the container, namely

$$
\begin{aligned}
p_{S,T,\bar{N}} &= p_{S|T,\bar{N}} p_{T,\bar{N}|P} = p_{S|T,\bar{N}} p_{T|P}(1-p_N) = 3.88\times10^{-10} \\
p_{S,\bar{T},N} &= p_{S|\bar{T},N} p_{\bar{T},N|P} = p_{S|\bar{T},N}(1-p_{T|P})p_N = 0.025 \\
p_{S,T,N} &= p_{S|T,N} p_{T,N|P} = (1-(1-p_{S|T,\bar{N}})(1-p_{S|\bar{T},N}))p_{T|P}p_N = 1.5\times10^{-11}.
\end{aligned}
$$

We can then calculate the probability of a threat and/or a NORM source given an alarm, specifically

$$
\begin{aligned}
p_{T,\bar{N}|S} &= \frac{p_{S,T,\bar{N}}}{p_{S,T,\bar{N}} + p_{S,\bar{T},N} + p_{S,T,N}} = 1.55\times10^{-8} \\
p_{\bar{T},N|S} &= \frac{p_{S,\bar{T},N}}{p_{S,T,\bar{N}} + p_{S,\bar{T},N} + p_{S,T,N}} = 0.999999984 \\
p_{T,N|S} &= \frac{p_{S,T,N}}{p_{S,T,\bar{N}} + p_{S,\bar{T},N} + p_{S,T,N}} = 6\times10^{-10}.
\end{aligned}
$$

The overall probability of a threat given a single alarm is then

$$
p_{T|S} = p_{T,\bar{N}|S} + p_{T,N|S} = 1.61\times10^{-8},
$$

the same value we obtained in Section 2. However, it is noticeable that the probability of a NORM source is almost one. We did update our probability of a threat from the prior

value of 5×10^{-10} to the posterior value of 1.61×10^{-8}. Thus, we do increase our belief that there is a threat in the container an alarm sounds, but we believe more strongly that a NORM source caused the alarm.

We use Bayesian updating to learn from multiple levels of screening (Bernado and Smith 1994). Suppose we screen a container n times. Evidently, alarms from multiple screenings of a given container are dependent; each alarm depends on the contents of the container. If there is a threat and/or a NORM source, then each screening has a high probability of sounding an alarm. If there is no threat or NORM source, then the probability of an alarm is (effectively) zero for each screening. We assume that the occurrences of alarms from n screenings are exchangeable events when conditioned on the contents of the container (Bernado and Smith 1994, chapter 4). Let K be the number of alarms that occur during n screenings. The conditional exchangeability assumption implies that the distribution of K given the contents of the container is binomial (Bernado and Smith 1994, page 223), specifically

$$\begin{aligned} K \mid T, \bar{N} &\sim Binomial(n, p_{S\mid T,\bar{N}}) \\ K \mid \bar{T}, N &\sim Binomial(n, p_{S\mid \bar{T},N}) \\ K \mid T, N &\sim Binomial(n, p_{S\mid T,N}), \end{aligned}$$

where $p_{S\mid T,\bar{N}}$, $p_{S\mid \bar{T},N}$, and $p_{S\mid T,N}$ are specified above. We calculate the joint probabilities of alarms and container contents by

$$\begin{aligned} p_{T,\bar{N},K} &= \binom{n}{K} p_{S\mid T,\bar{N}}^{K} (1-p_{S\mid T,\bar{N}})^{n-K}\, p_{T\mid P}(1-p_{N\mid T}), \\ p_{\bar{T},N,K} &= \binom{n}{K} p_{S\mid \bar{T},N}^{K} (1-p_{S\mid \bar{T},N})^{n-K} (1-p_{T\mid P})\, p_{N\mid \bar{T}}, \\ p_{T,N,K} &= \binom{n}{K} p_{S\mid T,N}^{K} (1-p_{S\mid T,N})^{n-K}\, p_{T\mid P}\, p_{N\mid T}. \end{aligned}$$

Again, we calculate the posterior probabilities of the contents of the container given K alarms as

$$p_{T,\bar{N}|K} = \frac{p_{K,T,\bar{N}}}{p_{K,T,\bar{N}} + p_{K,\bar{T},N} + p_{K,T,N}},$$

$$p_{\bar{T},N|K} = \frac{p_{K,\bar{T},N}}{p_{K,T,\bar{N}} + p_{K,\bar{T},N} + p_{K,T,N}},$$

$$p_{T,N|K} = \frac{p_{K,T,N}}{p_{K,T,\bar{N}} + p_{K,\bar{T},N} + p_{K,T,N}}.$$

Figure 4 shows the possible posterior probabilities for five levels of screening. It is clear from examining Figure 4 that we are confident that there is no threat if there is no alarm. If we observe one or more alarms, we become certain that there is a NORM source in the container. The prior probability of a threat is much lower than the prior probability of a NORM source. The posterior probability of a threat does increase after observing one or more alarms. However, it does not increase sufficiently to become confident that a threat caused the alarms instead of a NORM source.

The likelihood functions for K alarms given the contents of the container are driven by the probability of an alarm given the radioactive source, namely $p_{S|T,\bar{N}}$ and $p_{S|\bar{T},N}$. Current radiation sensors are as likely to give an alarm for a NORM source as they are for a threat. The base values of these probabilities are equal. Thus, inspectors cannot differentiate between an alarm from a threat and an alarm from a NORM source. We begin with a higher prior probability of a NORM source. We end with a higher posterior probability of a NORM source when we observe alarms. So what is the effect of screening equipment that can differentiate between threats and NORM sources?

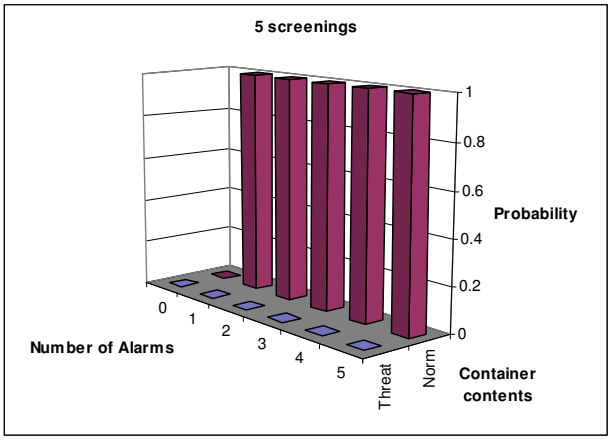

Figure 4. The posterior probability of a NORM source and a threat given the number of alarms for different levels of screening

Figure 5 repeats the calculations in Figure 4 with three different values of $p_{S|\bar{T},N}$, namely 0.1, 0.05, and 0.01. These values illustrate the effect of differentiating between threats and NORM sources. We are confident that there is a threat only when $p_{S|\bar{T},N}$ = 0.01 and there are four or five alarms. An increase in the accuracy of the screening equipment for detecting threats ($p_{S|T,\bar{N}}$) gives the same conclusion.

These calculations allow us to draw two conclusions. The first is that we must differentiate between threats and NORM sources before we can derive any certainty that there is a threat in the container. The second is that we need multiple levels of screening to become certain that there is a threat. This is because our prior probability of a threat is much lower than our prior probability of a NORM source. However, the cost of an attack is much greater than the cost of an inspection that only reveals a NORM source. Thus, we must determine how many levels of screening are cost effective and again consider the deterrence effect.

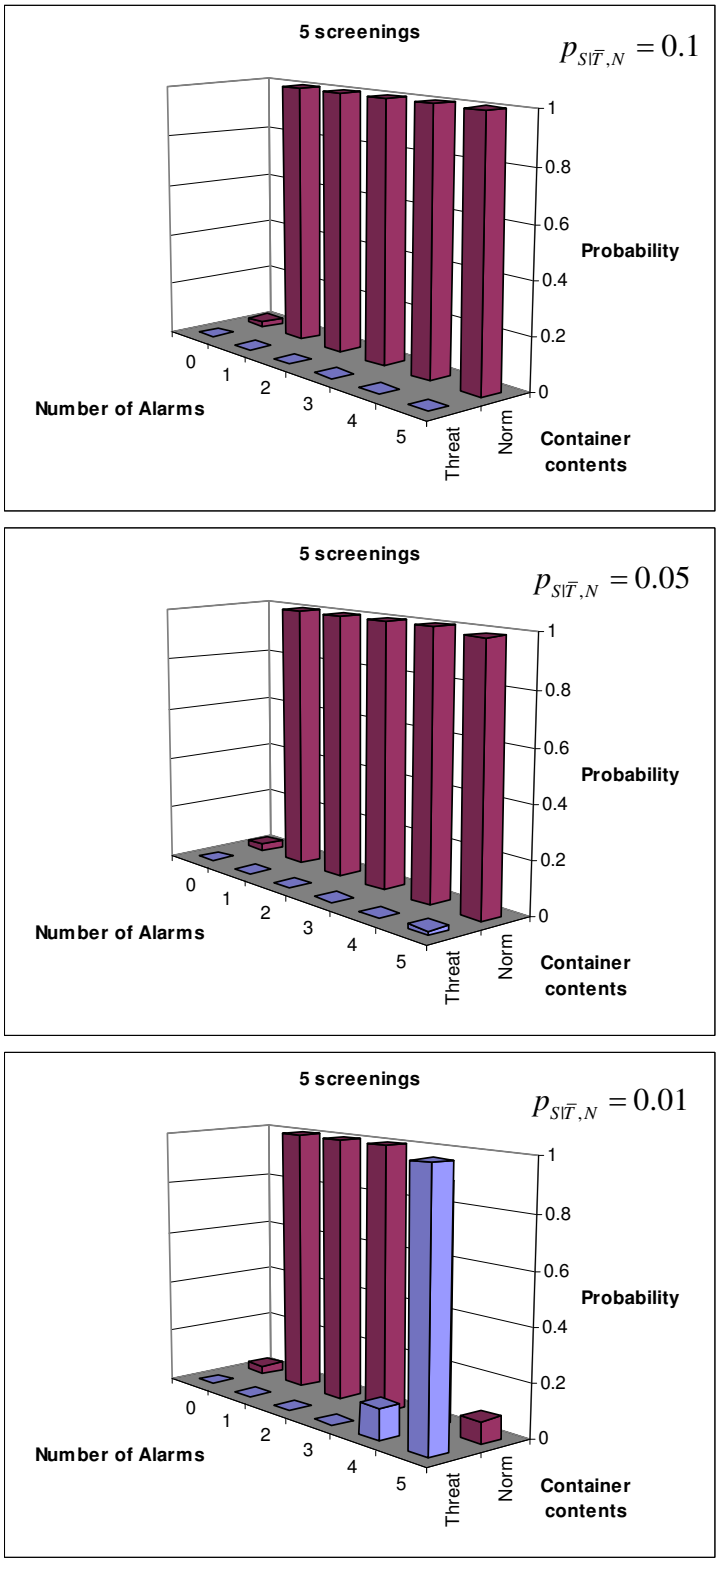

Figure 5. The posterior probabilities of threats and NORM source for varying probabilities of alarms given NORM sources given five levels of screening.

4.2 How Many Levels of Screening Should We Perform?

Figure 6 shows an influence diagram for deciding the number of levels of screening to perform. It represents a one-time decision, not a sequential one. Although deciding whether to screen again based on the results of earlier screenings may be desirable, it is not practical. CBP passes the containers through the RPMs on a rail car or a truck bed. Constructing a pathway to divert a container if an RPM does not give an alarm would use valuable space at the container port. We assume that CBP passes the containers through each RPM one after another.

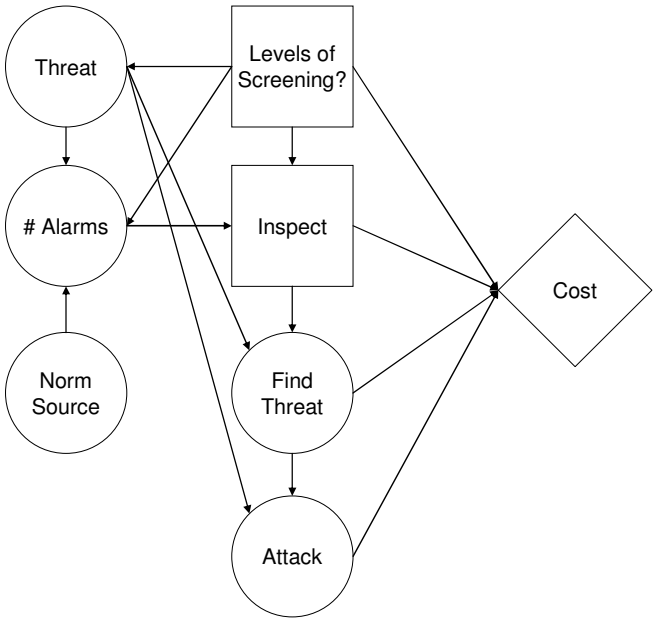

Figure 6. An influence diagram of the decision of how many levels of screening to perform.

We also assume that CBP does not use the same RPM each time. This would take too much time and would produce dependencies in the results beyond those modeled in Section 4.3. We must have multiple RPMs to perform multiple levels of screening. This

allows for the possibility of different types of screening equipment that can search for different types of threats and filter out NORM sources. We do not examine such possibilities here though.

Figure 7 shows the expected cost of no screening and one through five levels of screening. The deterrence multiplier (d) varies from 1 to 5. We assume that the level of deterrence is the same for any number of levels of screening. Increasing the number of screening machines would add equipment at the location of the screening. However, it is not obvious that this would increase the caution of terrorists attempting to smuggle radiological material. We must also admit that this is a simplifying assumption. As we have already stated, it is difficult, if not even unwise, to estimate the exact value of d . Introducing additional deterrence parameters for two through five levels of screening would exacerbate the estimation problem.

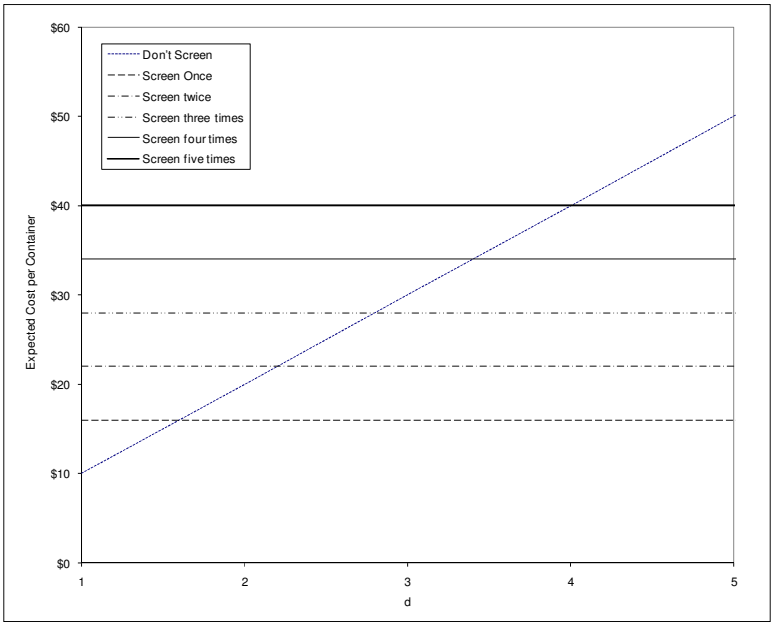

Figure 7. The expected cost of not screening and of performing from one to five levels of screening for varying levels of deterrent multiplier

The expected cost for one level of screening and no screening in Figure 7 are identical to those in Figure 3. These lines still intersect at d =1.6. However, the expected cost of two through five levels of screening increases with the additional cost of screening equipment. Screening is not economically viable if $d < 1.6$. If $d \geq 1.6$, then CBP should perform only one level of screening when we set the parameters to their base values. The purpose of this screening is purely deterrence, as the container is not inspected even with an alarm signal. In fact, the container is not inspected even with five alarm signals as the posterior probabilities in Figure 4 show that the cause is most likely a NORM source. In the next section, we will examine whether this changes with $p_{S|\bar{T},N}$ at 0.01 .

In summary, at our base parameter values, we do not screen if we believe that the deterrence effect is small ($d < 1.6$) and we screen only once if we believe the deterrence effect is large enough ($d \geq 1.6$).

4.3 Sensitivity to Parameter Values

The baseline analysis of multiple levels of screening leaves several critical questions. What if we can avoid alarms for NORM sources? Does this affect the economic calculation? Would additional levels of screening be advisable if the equipment was better at detecting a threat? What other parameters affect the decision?

Table 2. The decision to inspect and threat probability based on the number of levels of screening performed and the number of alarms observed.

Number of Alarms	Number of Levels of Screening					
	0	1	2	3	4	5
0	Don't Inspect $5.00\text{x}10^{-10}$	Don't Inspect $1.00\text{x}10^{-10}$	Don't Inspect $2.00\text{x}10^{-12}$	Don't Inspect $4.00\text{x}10^{-13}$	Don't Inspect $8.00\text{x}10^{-13}$	Don't Inspect $1.60\text{x}10^{-13}$
1		Don't Inspect $1.28\text{x}10^{-6}$	Don't Inspect $2.59\text{x}10^{-7}$	Don't Inspect $5.22\text{x}10^{-8}$	Don't Inspect $1.05\text{x}10^{-8}$	Don't Inspect $2.13\text{x}10^{-9}$
2			Don't Inspect $1.02\text{x}10^{-4}$	Don't Inspect $2.07\text{x}10^{-5}$	Don't Inspect $4.18\text{x}10^{-6}$	Don't Inspect $8.44\text{x}10^{-7}$
3				Inspect 0.0081	Inspect 0.0017	Inspect $3.34\text{x}10^{-4}$
4					Inspect 0.3960	Inspect 0.1169
5						Inspect 0.9813

As we have seen in Section 4.1, we obtain higher posterior probabilities of a threat from multiple alarms if NORM sources do not trigger alarms as often. To test this effect we re-created the analysis in Figure 7 setting $p_{S|\bar{T},N} = 0.01$. The change did not affect the expected costs for not screening and one level of screening. The expected costs for two through five levels of screening decrease when $p_{S|\bar{T},N} = 0.01$, but not enough to justify these levels of screening. Thus, even with the lower value of $p_{S|\bar{T},N}$, the decision remains not to screen if $d < 1.6$ and to screen once if $d \geq 1.6$. However, Table 2 shows the result of the inspection decision for zero to five levels of screening with the corresponding posterior threat probability with $p_{S|\bar{T},N}$ at 0.01. We can see that CBP should not inspect the container if they screen less than three times, no matter how many

alarms they get. If they screen three times, they should only inspect if they get alarms each time. If they screen four times, they should only inspect if they get three or four alarms. If they screen five times then they should only inspect if they get three, four, or five alarms. Thus, the inspection decisions are affected by $p_{S|\bar{T},N}$ for three or more levels of screening, but three or more alarms occur so infrequently that this does not affect the overall screening decision.

Holding $p_{S|\bar{T},N}$ at 0.01, we also increased the probability of detecting a threat in a container $p_{S|T,\bar{N}}$ to its maximum value. Again, we do not change our decision from that obtained when these parameters are set to their base values. Thus, improving the sensitivity of screening devices does not change our conclusions about screening.

We now consider the sensitivity of this analysis to other parameters. Setting the prior threat probability to its maximum level changes the ranking of alternatives, as does setting the cost of screening to its minimum level. Setting the probability of an attack if the threat gets through to its maximum level and setting the cost of an attack to its maximum level both change the expected costs, but do not change the ranking of alternatives. The cost of inspection, the cost of finding a threat on inspection and the probability of finding a threat upon inspection do not significantly affect the expected cost of each alternative.

Figure 8 shows the same calculations as Figure 7, but with the threat probability ($p_{T|P}$) set at its maximum value, ten times the base value. Figure 8 shows that screening either four or five times is better than not screening with a higher threat probability. The size of the deterrent effect does not matter. The best alternative is five levels of screening.

In fact, this is true for values of $p_{T|P}$ below the maximum value. When the threat probability reaches 6.36 times the baseline value, five levels of screening becomes the best alternative.

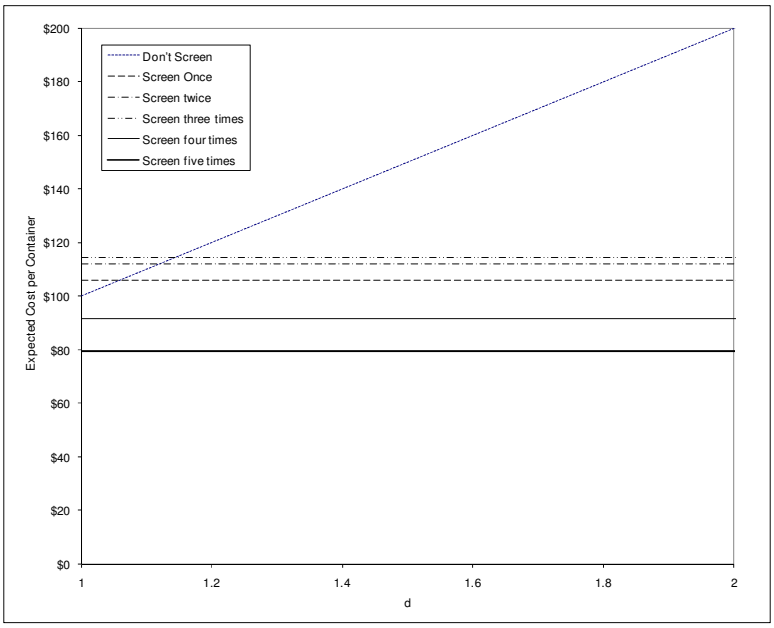

Figure 8. The expected cost for zero to five levels of screening for varying levels of deterrent effect with the maximum threat probability

The other parameter that changes the ranking of alternatives is the cost of screening. Figure 9 shows the same calculations as Figure 7, but with free screening. The optimal decision in Figure 9 is five levels of screening. Again, five levels of screening are best for costs above the minimum (free). Below 69¢ per container, it is better to perform five levels of screening even if there is no deterrence effect.

The cost of an attack (c_A) changes the expected cost values, but does not change the ranking of alternatives. The same is true for the probability of an attack ($p_{A|T}$). However, these parameters do change the effect of deterrence on the decision. The baseline values of these parameters are c_A = \$40 billion and $p_{A|T} = 0.5$. At these values

one level of screening is preferred if $d \geq 1.6$. If we increase c_A to \$100 billion, then we only require that $d \geq 1.24$. If we decrease c_A to \$10 billion, then not screening is preferred unless $d \geq 3.4$. For the probability of an attack, the maximum and minimum values are $p_{A|T} = 1$ and $p_{A|T} = 0.1$ and these correspond to deterrence effects in the range of $d \geq 1.3$ and $d \geq 4$ respectively for one level of screening to be preferred to no screening.

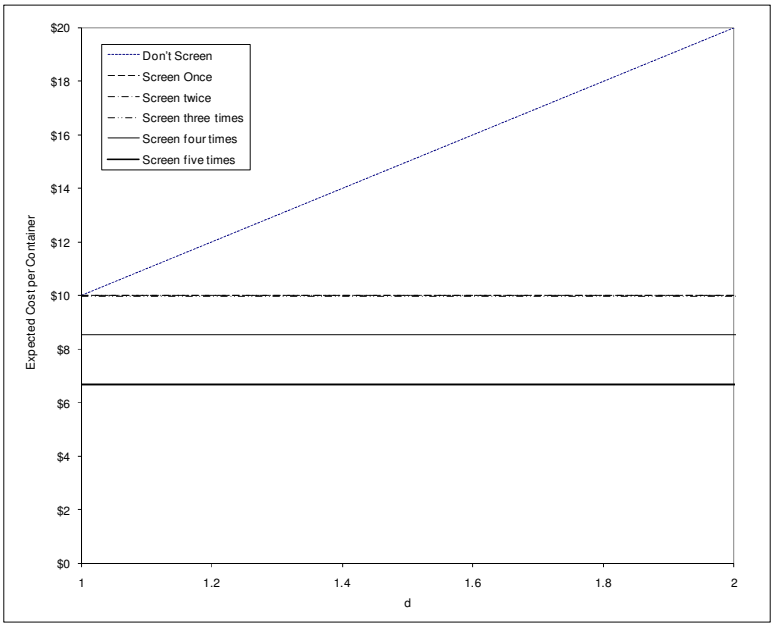

Figure 9. The expected cost for zero to five levels of screening for varying levels of deterrent effect with free screening

In summary, the threat probability and the attack cost are the most influential parameters in this decision. Set at their maximum levels, screening five times is the option with the lowest expected cost. However, we will see in the next section that we should consider attributes other than cost in making this decision.

5. Considering Additional Objectives

Keeney (2007) proposes that decision makers should consider values and objectives in the analysis of counterterrorism decisions. Keeney developed fourteen attributes for evaluating the outcomes of such decisions. Six attributes are costs, namely

- direct and indirect costs to individuals
- direct and indirect costs to businesses
- direct and indirect costs to the government.

Three other attributes are counts of the effect of the attack on individuals, namely

- the numbers of deaths caused by the attack
- the number of victims disabled by the attack
- the number of jobs lost because of the attack

These measure the effect of the attack on the victims, both directly and indirectly. The five other attributes are less obvious, including

- the effect on recruitment of future terrorists
- the number of dollars flowing to support terrorism
- the level of political support for US actions to counter terrorism
- the number of US citizens experiencing limits to their freedom
- the number of citizens experiencing fear and despair

Thus far, we have measured the effects of an attack by their cost. Bakir (2008) gives a range for the cost of an attack. This range is an aggregation of estimates from previous studies. Each study considers attributes suggested by Keeney. For instance, the total costs in Rosoff and von Winterfeldt (2007) include the number of fatalities, the number of cancers caused by radiation exposure, the cost of evacuation, the cost of

decontamination, and the cost of lost business and lowered property values. It is possible to model each of these attributes individually and perform a value tradeoff. Rosoff and von Winterfeldt (2007) implicitly makes a value trade-off in the assumption that each human life is equivalent to a cost of $5 million. However, as Bakir (2008) aggregates multiple studies to get the total cost, we do not know the corresponding values of each attribute. Instead, we include all but the last five of Keeney's attributes above in the one total cost attribute.

Do we need to include any of the remaining five attributes? Some terrorism-related decisions affect support for US policies and limit the freedom of citizens. However, screening containers is not a politically sensitive issue and its effect is commercial, not individual. Successful attacks cause fear and despair, but we can represent this through an increase in the cost of an attack. We have already performed such a sensitivity analysis.

This leaves us to consider effects on the recruitment of future terrorists and monetary support for terrorism. One may think that for our decision, only a successful attack will affect these attributes. However, if CBP finds radiological material in a container during an inspection, then this could negatively affect terrorist support in terms of both money and recruitment. Thus, we should also consider changes in support for terrorists caused by the outcomes of the screening decision.

5.1 A Simple Two Attribute Value Function for the Screening Decision

Our analysis will use two attributes, the total cost of an attack (x_1) and the effect on support for terrorist groups (x_2).We use a constructed scale (Clemen and Reilly 2001, ch.

4) for terrorist support. If an attack is successful, we set $x_2 = +1$. This represents an increase in terrorist support. If CBP finds a threat upon inspection, thereby stopping the attack, then we set $x_2 = -1$. This represents a decrease in terrorist support. If terrorist support is neither increased nor decreased from the current level of support, then we set $x_2 = 0$.

We assume a linear-additive value function (Keeney and Raiffa 1976, ch. 3; Kirkwood 1997, ch. 4), where

$$v(x_1, x_2) = w_1 v_1(x_1) + w_2 v_2(x_2)$$

and $w_1 + w_2 = 1$. A value function represents the decision maker's preferences over these two attributes. We could also use a utility function and represent the decision maker's aversion to risk. The simple additive form above is unlikely to be sufficient to model risk aversion over these two attributes. Thus, we use a value function to keep the form simple for this demonstration of the approach. The value function for x_2 is simple and represents our objective to minimize terrorist support. We set $v_2(+1) = 0$, as this is the worst outcome. We set $v_2(-1) = 1$, as this is the best outcome. Lastly, we set $v_2(0) = a$. We use the parameter a to show whether preferences are stronger for avoiding an increase in terrorist support or for achieving a decrease in terrorist support. If a is above 0.5 then a increase in terrorist support from a successful attack is more important than a decrease in terrorist support from finding smuggled radiological material. If a is below 0.5 then a decrease in terrorist support from finding smuggled radiological material is more

important than the increase in terrorist support from a successful attack[1]. At the extremes, if a is 1 then an increase in terrorist support is important, but a decrease is not. If a is 0 then a decrease in terrorist support is important, but an increase is not.

For x_1, we use a proportional scoring value function (Clemen and Reilly 2001, ch. 15), specifically

$$v_1(x_1) = \frac{x_1^{\max} - x_1}{x_1^{\max} - x_1^{\min}},$$

where $x_1^{\max}$ =\$100,000,000,630 (the maximum cost of a successful attack even after inspection and five levels of screening) and $x_1^{\min}$ =0 (the cost if we do not screen or inspect and either there is no threat or the attack is not successful). Thus $v_1(x_1^{\max}) = 0$ and $v_1(x_1^{\min}) = 1$ as we wish to minimize the total cost.

The analysis now depends on the value of w_1 (as $w_2 = 1 - w_1$) and a. In the next section, we will perform a sensitivity analysis to determine how w_1 and a affect the optimal decision.

5.2 Results of the Two Attribute Analysis

Figure 10 shows a strategy chart indicating the alternative with highest expected value at each combination of w_1 and a between 0 and 1. Let us first consider the extreme combinations of these two parameters. We will then discuss what happens at the non-

[1] Imagine a decision maker faced with a choice of A. going from an increase in terrorist support from a successful attack ($x_2 = +1$) to the current levels ($x_2 = 0$) or B. going from the current level ($x_2 = 0$) to a decrease in terrorist support ($x_2 = -1$) by finding the threat by inspection. If the DM prefers A $\succ$ B then the preference difference (from an increase to the current level) $a > 0.5$. If the DM prefers B $\succ$ A then the preference difference (from the current level to a decrease) $1 - a > 0.5$, so $a < 0.5$.

extreme combinations. In this analysis, we set $p_{S|\bar{T},N} = 0.01$ as discussed in Sections 4.1 and 4.3.

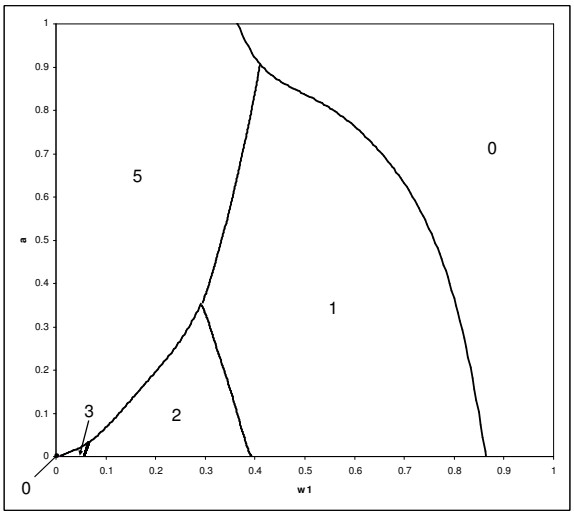

Figure 10. A strategy chart showing the optimal number of levels of screening for values of w_1 and a

When $w_1 = 1$, the analysis is the same as the purely cost based analysis. When $w_1 = 0$, we are only basing the decision on the importance of increases and decreases in terrorist support. When $a = 0$ (and $w_1 = 0$) then only decreases in terrorist support from finding smuggled material are important. When $a = 1$ (and $w_1 = 0$) then only increases in terrorist support from a successful attack are important. If $0 < w_1 < 1$ and $0 < a < 1$ then each of these outcomes affects the analysis and their interplay causes interesting dynamics in the optimal decision.

There are two dynamics at play in Figure 10. The first occurs as the value of increases in terrorist support from successful attacks becomes high enough (to the top-left of Figure 10). We reach a point where the best decision changes from no screening to

five levels of screening. When a is high enough and w_1 is low enough, the decision maker's preference against increases from terrorist support from successful attacks is high. Thus, stopping the attacks is most important and the best option is five levels of screening.

Below this threshold of a and w_1, the decision is based on how important it is to decrease terrorist support by finding smuggled threats. When $w_1 = 1$, we are considering only total cost in the analysis. Then we obtain the same result as our purely cost-based analysis and choose not to screen (we are not considering the effect of deterrence yet). As we scan from the right of Figure 10 to the bottom-left, the optimal number of levels of screening increases. For a and w_1 near zero, we choose higher numbers of levels of screening.

One final observation is an effect at the origin. When $a = 0$ and $w_1 = 0$, the only outcome that matters is decreasing terrorist support by finding smuggled threats. In this special case, the best decision is to inspect every container, as cost is not important ($w_1 = 0$). However, if we increase either a or w_1 by as little as 0.0001, we choose five levels of screening. Thus, this is a special case at the origin only.

The analysis represented in Figure 10 assumes no deterrence effect (or $d = 1$). Recall for the purely economic decision, if $d \geq 1.6$ then the best decision would be to perform one level of screening. Next, we re-created the strategy plot from Figure 10, but setting $d = 2$. Figure 11 puts the strategy plots for $d = 1$ and $d = 2$ side by side for comparison. The values of w_1 and a for which the best decision is not to screen when $d = 1$, switch to screen once when $d = 2$ due to the effect of deterrence. The best

decisions for other values of w_1 and a are mostly unaffected by the switch from $d = 1$ to $d = 2$, except near the origin.

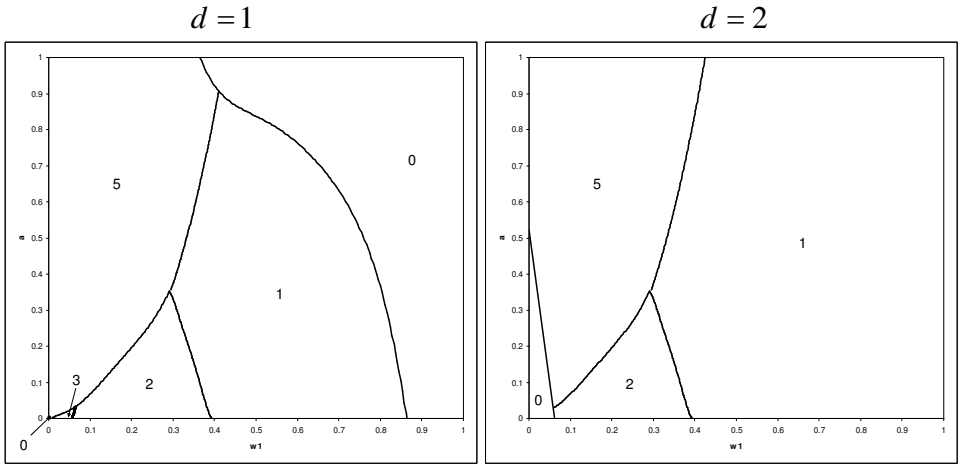

Figure 11. A strategy chart showing the optimal number of levels of screening for values of w_1 and $v_2(0) = a$ when $d = 1$ vs. $d = 2$

Recall that with $d = 1$ we had a special case at the origin (when $a = 0$ and $w_1 = 0$) where we would choose no screening. This area increases when $d = 2$. Close to the origin, the value of decreasing terrorist support by finding nuclear threats is high. With an increased deterrence effect, the threat probability increases when not screening. This makes the chance of finding radiological material upon inspection higher. Thus, when $d = 2$, the region where the optimal decision is to inspect every container covers a larger area around the origin, rather than just at the origin with no deterrent effect ($d = 1$). We note, however, that while this decision near the origin is of mathematical interest, setting $a = 0$ and $w_1 = 0$ is not realistic. Thus, this effect is of little practical interest.

For realistic values of w_1 and a, Figure 11 indicates that screening (at least one level) should be performed even if the deterrence effect is small ($d = 2$). This suggests that we should not just look at the cost of the screening process and the cost of attacks. We should also consider the second-order effects of outcomes on future terrorist plans.

In Section 4.3, we examined the influence of key parameters on the purely cost-based decision. Do these parameters still influence our decision when we consider two attributes?

Figure 12 shows a comparison of the strategy charts for two different numerical values of $p_{S|T,\bar{N}}$, the probability of an alarm when there is a threat. On the left, we use the base case value ($p_{S|T,\bar{N}} = 0.8$). On the right, we use perfect detection ($p_{S|T,\bar{N}} = 1$). The threshold for five levels of screening moves to the right when $p_{S|T,\bar{N}} = 1$. If each of the five levels of screening is more accurate, then this threshold is reached for more combinations of w_1 and a. Below this threshold, fewer levels of screening are required at a given combination of w_1 and a when $p_{S|T,\bar{N}} = 1$; each level of screening is better at finding threats.

The remaining parameters do not have as much effect on the strategy plot. The threshold beyond which five levels of screening are optimal moves to the right if we increase the threat probability ($p_{T|P}$). The same is true for increases to the cost of an attack (c_A). These changes make attacks either more likely or more costly, making five levels of screening more attractive. Changes in the probability of an alarm from a NORM source ($p_{S|\bar{T},N}$) and changes in the cost of inspection (c_I) do not affect the strategy chart.

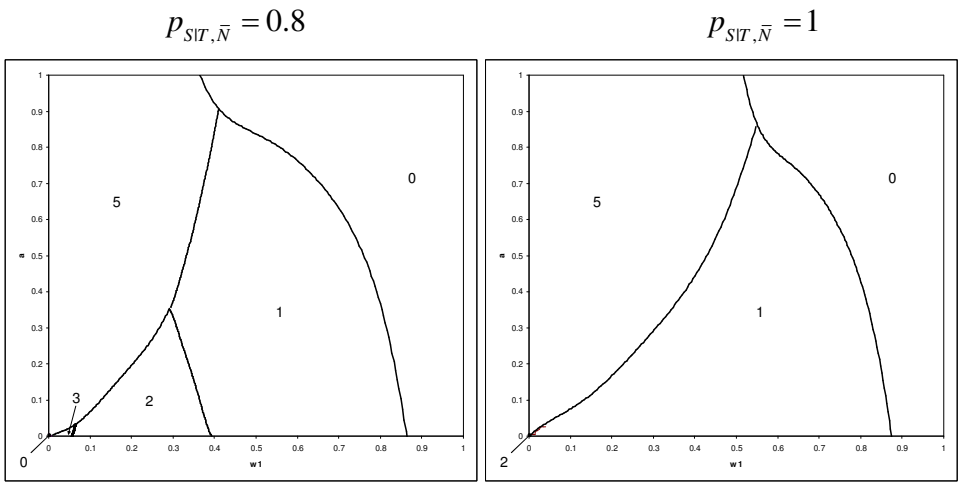

Figure 12. A strategy chart showing the optimal number of levels of screening for values of w_1 and $v_2(0) = a$ when $p_{S|T,\bar{N}} = 0.8$ vs. $p_{S|T,\bar{N}} = 1$

6. Conclusions and Further Research

Bakir (2008) recommends that Customs and Border Patrol does not implement new advanced spectroscopic portal technology, unless the threat probability or the cost of an attack is significantly higher than current best estimates. We examined the investments for the original radiation portal monitor technology and found the same to be true. However, we performed several extensions of the analysis that can change this conclusion.

Firstly, if we believe that an attempt to smuggle radiological material or nuclear devices into the United States is less likely if we screen containers (a deterrence effect), then screening can become the recommended decision. This alone could justify screening. Bakir gives point estimates of the deterrence effect. We parameterize it and find the values at which it changes the decision. This allows the decision maker to perform a simpler judgment task. They need only to decide if they believe the deterrence

effect is smaller or larger than this change-point. They do not need to estimate its actual value.

Secondly, we consider the main reason that screening has such a high cost, false alarms. CBP can reduce the false alarm rate by improving screening technology. We show that they can also reduce false alarms by performing multiple levels of screening. We find that multiple levels of screening do improve the ability to find threats in a container. However, the recommended decision remains screening once if the level of the deterrence effect is high enough and not screening if the deterrence effect is low in the expected cost analysis.

Our final extension considers objectives beyond those included in Bakir's analysis. We develop a two-attribute value function that includes the total cost and changes in support for terrorist groups. We see that the recommendation can vary from zero to five levels of screening depending on our preferences over these two attributes. We have not elicited the parameters of the value model from decision makers and could not publish them if we had. Instead, we show the recommended decision for different values of the parameters and allow the decision maker to decide which to use. However, this analysis shows that total cost is not sufficient for full evaluation of the decision.

This paper shows that the decision to screen containers entering the United States is complex. The effect of deterrence alone may well justify screening. Moreover, if we consider attributes other than the cost, if we consider multiple levels of screening, and if we consider improved technology, then an increase in the level of screening can be justified. We have also shown that the decision depends heavily on key parameters. There is considerable uncertainty about these parameters, so the next step is to model this

uncertainty explicitly and to model the decision maker's attitude to risk and uncertainty through a utility function rather than a value function (as done in Section 5).

We should note a limit to our analysis, however. We model the effect of deterrence as a decrease in the probability that terrorists will smuggle radiological material into the United States inside a container. This treats the actions of terrorists as uncertain quantities (Bedford and Cooke 2001, ch. 2) even though they are clever enemies. Their decisions about how to smuggle radiological material into the United States are not as simple as a choice between smuggling inside a container and abandoning their attack. If we make the container choice less appealing, then terrorists will look for other methods. Work to model terrorist decisions, US government decisions, and multiple methods of entering the United States is ongoing. Parnell et al. (2009) propose an influence diagram where the decisions of the US government and a terrorist organization are explicitly included, but with opposing objectives. Rios Insua et al. (2009) also propose a combination of game theory and decision analysis that could solve this problem. We could model the deterrence effect of random inspections with either method. These approaches appear promising for further consideration of this critical decision.

Acknowledgments

Developed partially under grants from the U.S. Department of Homeland Security's Domestic Nuclear Detection Office under Grant Award Number 2008-DN-077 ARI001-02 and the National Science Foundation (CBET-0735735). The work was done at Virginia Commonwealth University. 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.

References

Apostolakis, G. E., D. M. Lemon. 2005. A screening methodology for the identification and ranking of infrastructure vulnerabilities due to terrorism. *Risk Analysis* **25**:2 361-376.

Bakir, N. O. 2008. A decision tree model for evaluating countermeasures to secure cargo at United States southwestern ports of entry. *Decision Analysis* **5**(4) 2304-248.

Bedford, T. M., R. M. Cooke. 2001. *Probabilistic Risk Analysis: Foundations and Method.* Cambridge University Press, Cambridge UK.

Bernado, J., A. Smith. 1994. *Bayesian Theory*. Chichester, UK: Wiley.

Bier, V.M. 2005. Game-theoretic and reliability methods in counter-terrorism and security. A. Wilson, N. Limnios, S. Keller-McNulty, Y. Armijo, eds., *Mathematical and Statistical Methods in Reliability*, *Series on Quality, Reliability and Engineering Statistics*. World Scientific, Singapore, 17-28.

Bier, V. M., E. R. Gratz, N. J. Haphuriwat, W. Magua, K. R. Wierzbicki. 2007a. Methodology for identifying near-optimal interdiction strategies for a power transmission system. *Reliability Engineering & System Safety* **92** 1155-1161.

Bier, V.M., S. Oliveros, L. Samuelson. 2007b. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *J. Public Economic Theory* **9**:4 563-587.

Bonner, R. C. 2005. Statement of Robert C. Bonner—Hearing before the permanent Subcommittee on Investigations, Senate Committee on Homeland Security and Governmental Affairs. http://hsgac.senate.gov/_files/STMTBONNERCBP.pdf.

Bronskill, J., S. Bailey. 2007. Dirty bomb would cause panic, cost billions: Study. TheStar.com. http://www.thestar.com/article/231670.

Clemen, R. T., T. Reilly. 2001. *Making Hard Decisions with Decision Tools Suite*, 1st Edition. Duxbury Press, Pacific Grove, California.

de Rugy, V. 2005. *Is port security making us safer?* American Enterprise Institute for Public Policy Research, Washington, D.C.

Ebeling, C. E. 2009. Evolution of a box. *Invention and Technology* **23**(4) 8-9.

Feng, T., L. R. Keller. 2006. A multiple-objective decision analysis for terrorism protection: potassium iodide distribution in nuclear incidents. *Decision Analysis* **3**:2 76-93.

GAO. 2006b. Combating nuclear smuggling: Challenges facing U.S. efforts to deploy radiation detection equipment in other countries and in the United States. GAO-06-558T, U.S. Government Accountability Office, Washington, D.C.

Gardner, F. 2003. Al Qaeda "was making dirty bomb." *BBC News Online*.

Haimes, Y. Y., S. Kaplan, J. H. Lambert. 2002. Risk filtering, ranking, and management framework using hierarchical holographic modeling. *Risk Analysis* **22**:2 383-397.

Keeney, R. L. 2007. Modeling values for anti-terrorism analysis. *Risk Analysis* **27**(3) 585-596.

Keeney R. and Raiffa H. 1976. *Decisions with Multiple Objectives, Preferences, and Value Tradeoffs*. John Wiley and Sons: New York.

Kirkwood, C. W. 1997. *Strategic Decision Making, Multiobjective Decision Analysis with Spreadsheets.* Duxbury Press: Belmont CA.

Kunreuther, H., G. Heal. 2002. Interdependent security. *J. Risk & Uncertainty* **26**:2-3 231-249.

Leung, M., J. H. Lambert, A. Mosenthal. 2004. A risk-based approach to setting priorities in protecting bridges against terrorist attacks. *Risk Analysis* **24**:4 963-984.

Lewis, B. M., A. L. Erera, C. C. White III. 2003. Optimization approaches for efficient container security operations at transshipment seaports. *Transportation Research Record* **1822** 1-8.

Madigan, D., S. Mittal, F. Roberts. 2007. Sequential decision-making algorithms for port of entry inspection: Overcoming Computational Challenges. *Proceedings of the 2007 IEEE Intelligence and Security Informatics Conference*, New Brunswick, NJ, 1-7.

Parnell, G. S., C. M. Smith, D. I. Moxley. 2009. Intelligent adversary risk analysis: a bioterrorism risk management model. Submitted to *Risk Analysis*.

Paté-Cornell, E. 2002. Fusion of Intelligence Information. *Risk Analysis* **22**:3 445-454.

Paté-Cornell, E., S. Guikema. 2002. Probabilistic modeling of terrorist threats:
A systems analysis approach to setting priorities among countermeasures. *Military Operations Research* **7**:4 5-20.

Ramirez-Marquez, J. E. 2008. Port-of-entry safety via the reliability optimization of container inspection strategy through an evolutionary approach. *Reliability Engineering and System Safety* **93** 1698-1709.

D. Rios Insua, J. Jesus and D. Banks, 2009. Adversarial risk analysis. J*ournal of the American Statistical Association* **104**(486) 841-854.

Rosoff, H., D. von Winterfeldt. 2007. A risk and economic analysis of dirty bomb attacks on the ports of Los Angeles and Long Beach. *Risk Anal*. **27**(3) 533–546

von Winterfeldt, D., T. M. Sullivan. 2006. Should we protect commercial airplanes against surface-to-air missile attacks by terrorists? *Decision Analysis* **3**:2 63-75,

Zhuang, J., V. M Bier. 2007. Balancing terrorism and natural disasters – defensive strategy with endogenous attacker effort. *Operations Research* **55**:5 976-991.

Zhuang, J., V.M. Bier, A. Gupta. 2007. Subsidies in interdependent security with heterogeneous discount rates. *The Engineering Economist* **52**:1 1-19.