PNNL-16356

Pacific Northwest
National Laboratory
Operated by Battelle for the
U.S. Department of Energy

Game Theoretic Evaluation of Threat Detection Problems–The central role of the Adversary

Patrick Heasler, Tom Wood, Barbara Reichmuth

Jan 2007

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

DISCLAIMER

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service,
U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161
ph: (800) 553-6847
fax: (703) 605-6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/ordering.htm

This document was printed on recycled paper.
(9/2003)

PNNL-16356

Game Theoretic Evaluation of Threat Detection Problems–The central role of the Adversary

Patrick Heasler, Tom Wood, Barbara Reichmuth

Jan 2007

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Executive Summary

A wide variety of security problems hinge on the detection of threats and discrimination of threats from innocuous objects. The theory that frames these problems is common among medical diagnostics, radar and sonar imaging, and detection of radiological, chemical, and biological agents. In many of these problems, the nature of the threat is subject to control by a malicious adversary, and the *choice of a reference (or "design basis") threat* is a very difficult, and often intractable, aspect of the problem. It is this class of problems that this report considers.

This report formulates a threat detection problem from a decision theory (i.e. game theoretic) perspective and calculates the optimal strategies for both players. For this problem, containers pass a checkpoint which is monitored by a set of detectors. The adversary desires to introduce a container carrying a threat into the stream of "clean" containers and get it through the checkpoint without detection. The objective of the detector operator is to accomplish the opposite, to find and detain the threat containers. The specific "threat" detection problem we are most interested in evaluating is that of nuclear explosives. However, the framework developed also applies to other threat detection problems, so we present the problem in a more general form.

The decision theoretic formulation most clearly describes the inter-relationships between three components of the problem; (1) the detector capabilities, (2) the player strategies, and (2) the player payoffs. Decision theory provides the best description of detector capability when the checkpoint must account for an intelligent adversary.

The main conclusion is that a proper evaluation of the detectors cannot be done without a game-theoretic formulation. In particular, one has to evaluate the detectors when operated against a "competent" adversary, and a strong case can be made that such an adversary will employ the game-theoretic solution to this problem; this solution is called the "least favorable distribution" (LFD) and should be an integral part in the evaluation of any detection problem involving an intelligent adversary.

Acknowledgments

Financial support was provided by the U.S. Department of Energy under Contract DE-AC05-76RL01830.

Abbreviations and Acronyms

ABBREV	DEFINITION
$\delta(S)$	Algorithm for for transforming detector signal, S, into a Pass/Detain decision on the container.
$\delta_{best}(S)$	The best algorithm from a minimax perspective; the game-theoretic solution for the operator.
$L(i,j)$	Loss to detector operator when he chooses i and adversary j.
LFD	Least Favorable Distribution
P_{LFD}	The least favorable distribution; the game-theoretic solution for the adversary. $P_{LFD}(j)$ is the frequency with which the adversary chooses configuration j.
PND	Probability of Non-Detection $PND = 1 - POD$
POD	Probability of Detection
$f(S,j)$	The distribution of the detector signal, when the adversary presents a container in configuration j. $j = 0$ represents a "clean" container.
S	Detector signal, a vector of all data that the detector produces from the examination of a container.

Contents

Figures

Tables

1.0 The Threat Detection Game

The problem we are concerned with arises from any inspection or detection scenario in which a set of objects (containers) is screened for threats. This includes border-crossing checkpoints, inspection of cargo containers at ports, or inspection of airline passengers or baggage.

It is assumed that a checkpoint has been set up to inspect all containers passing through for threat. The checkpoint is equipped with detectors that remotely examine each passing container and produce a signal that is related to the presence/absence of a specific type of threat in the container (for example, explosives, radioactive material). The checkpoint operator uses this signal to decide whether to let the container pass through or pull it over into secondary inspection. If the container is pulled over for secondary inspection, it is assumed that the status of the container will be determined without error (i.e. the threat, if present, will be found). In contrast, the detectors do not perfectly classify the passing containers; a container containing threat might be missed, resulting in a *non-detection*, while a clean container might be pulled into to secondary inspection, resulting in a *false-call*.

1.1 Loss Matrix

Each time a container passes through the checkpoint is a "play" in the threat detection game, and the outcomes of a play can be described by a matrix with the form presented in Table 1.1. The matrix lists the two moves the detector operator can take (represented by the rows of the matrix) versus the moves the adversary can take (represented by the columns of the matrix). The outcome of a play of the game is represented by a cell in the matrix. This cell is filled with the loss associated with that outcome, hence the matrix is called the "loss matrix" for the game. By constructing a loss matrix, we assume the outcomes of the game can be summarized economically; in other words, the outcome can be represented as a numerical loss to one player and corresponding payoff to the other. The specific entries in the matrix represent losses to the checkpoint operator and gains to the adversary. Losses will be measured in dollar terms, but other metrics are possible such as expected lives lost. Both the detector operator and adversary want to choose strategies that minimize their expected losses when playing the game.

Table 1.1. Loss Matrix (i.e. Without Detector Operation)

	Adversary Moves				
Operator	No Threat	Threat			
Moves	Config=0	Config=1	Config=2	...	Config=N
Pass (i=0)	L(0,0)	L(0,1)	L(0,2)	...	L(0,N)
Detain (i=1)	L(1,0)	L(1,1)	L(1,2)	...	L(1,N)

Note that the game as described above is a zero-sum game; in other words, the operator's losses equal the opponent's gains. Obviously, the most realistic and general formulation of the threat detection problem would allow for different losses for each player.

However, with careful definition of the losses, we believe that many threat detection problems can be modeled by a zero-sum game. This statement applies particularly to cases when the adversary is a terrorist who is smuggling the threat to do harm to the detector operator. In this case, the adversary equates any loss to the other player as his gain, and the zero-sum assumption is justified.

1.2 The Adversary's Moves

For each play of the game, the adversary can either choose to;

Not Enter Checkpoint Queue: Do nothing–let a clean container pass through. A clean container is a container not controlled by the adversary and one that does not contain a threat. Consequently, the clean container's configuration and cargo are determined "by chance" and not by the adversary.

Enter Checkpoint Queue: The adversary has control over a threat container, and can as part of his move, choose the container and configure the cargo so as to be difficult for the checkpoint to detect.

In the actual detection problem, there are an infinite number of ways that the threat container might be configured. To simplify the formulation in this report, we will assume that there are a finite number, N, of threat configurations used by the adversary, as illustrated in the loss matrix presented in Table 1.1. So under this framework, the adversary will have the choices; *send Configuration j through, where $j = 0, 1, ...N$, and Config 0 stands for the "do nothing" alternative.*

The term, "threat configuration," might describe actual differences in the threats, but also in how the threats are hidden in the container. For example, one threat configuration may refer to an unshielded nuclear weapon in a container, while another configuration might represent the same nuclear weapon wrapped in lead shielding. The threat configurations must be defined in enough detail so that (1) the economic costs associated with a miss or detection of the threat can be calculated, and (2) the probability of detection of the threat can also be calculated.

In the examples we use in this report, the threat configurations are a discrete representation of a variable that an adversary would have control over. The two variables examined are amount of shielding for a nuclear weapon, and the size of a dirty bomb.

1.3 Operator's Moves

If the operator did not have use of the detectors, the loss matrix presented in Table 1.1 would describe the decision he would have to make. However, the detector produces a signal, let us call it S, and he can use this signal to decide upon his two alternatives. In this framework, his real choice (or move) is transformed into choosing a *decision function*, call it $\delta(S)$, which makes the choice for him. In other words, $\delta(S)$ is a mathematical function or computer program that transforms the detector signal into a pass/detain decision.

The presence of detectors transforms the Loss-matrix into the Risk-matrix presented in Table 1.2.

Table 1.2. Risk Matrix for the Threat Detection Problem (Detector Results introduced)

Operator Detect Algorithm	Opponent Move: No Contr. Config.0	Threat Config.1	Threat Config.2	...	Threat Config.N
δ_1	$R(\delta_1,0)$	$R(\delta_1,1)$	$R(\delta_1,2)$	...	$R(\delta_1,N)$
δ_2	$R(\delta_2,0)$	$R(\delta_2,1)$	$R(\delta_2,2)$	...	$R(\delta_2,N)$
.	.	.	.	...	.
.	.	.	.	...	.
δ_m	$R(\delta_m,0)$	$R(\delta_m,1)$	$R(\delta_m,2)$	...	$R(\delta_m,N)$

The rows in this matrix represent the different decision algorithms the operator might construct to process the detector signal. In reality there are an infinite number of such algorithms, even though we have only placed m algorithms in the table. The standard decision theory formulation that we will exploit requires that the opponent have a finite number of choices, but makes no such assumption concerning the choices available to the detector operator; the objective is to find the best decision function from all those possible.

The values in the cells of the risk matrix are now expected losses (also called risk) which average loss over the signal distribution. More specifically, the notation $R(\delta, j)$ represents the expected loss when employing decision algorithm δ and is defined by;

$$R(\delta, j) = E(L(\delta(S)|Config = j) \tag{1.1}$$

$$= \int L(\delta(S), Config = j) f(S|Config = j)\, dS \tag{1.2}$$

From this equation, we see that the detector signal distributions $f(S|Config = j)$ must be known in order to calculate the desired risk values.

1.4 Description of Detector Operation

As one can see from the discussion in the previous section, the detector operating characteristics enter into the loss calculations through the signal distributions represented by the terms $f(S|Config = j)$. If the detector is effective at finding the threat, there will be a difference between the "no threat" distribution $f(S|config = 0)$ and the threat distributions $f(S|config = i)$ $(i > 0)$.

In order to construct a solution for the threat detection problem described in this report, one therefore requires the signal distributions, $f(S|config = j)$ for the detector. These distributions may be constructed by using the detector physics or empirically through the statistical analysis of actual detector data. In most cases, a combination of both techniques is used to produce these distributions.

It should be noted that the detector signal, as represented by the variable S, is meant to represent more than a single value. A typical detector might produce hundreds of measurements–for

example a gamma-detector might produce an entire energy spectrum, which would be represented as a vector of values. The raw detector measurements are typically corrected for biases through various calibration procedures and then summarized into detection statistics (or features). The variable S is meant to represent these summarized statistics.

The detector signal, S, might also represent the results of two different detectors that are being employed concurrently at the checkpoint. For example, the containers might be inspected by both a gamma ray detector and radiography, and in this case $S = (S_1, S_2)$ with S_1 representing gamma results and S_2 radiography.

1.5 Constraint on FCP

An important constraint for the threat detection problem is associated with the secondary inspection facility. At most checkpoints, the secondary is not designed to allow 100% inspection of the traffic stream. Typically the secondary inspection facility is limited in that it can only process a small percentage of the incoming traffic, (call this fraction, *the secondary inspection capability*). Therefore, it is not possible to select a detection algorithm that exhibits a false call probability (FCP) greater than the secondary inspection capability.

To incorporate this constraint into the threat detection problem in the easiest manner possible, one formulates two problems, solves both, and then selects the solution from the two. The first problem formulation is the problem as described above without the constraint. If one works out the solution to this problem and the associated detection algorithm happens to have a FCP less than the secondary inspection capability, then it is the solution for the constrained problem.

If the unconstrained solution violates the constraint then we require the decision algorithm to have a FCP equal to the secondary inspection capability. This constraint produces a simplification in the Risk matrix; the risks associated with the adversary's "no threat" move become constant and one only needs to calculate the risks associated with threat moves.

2.0 Specific Examples of This Problem

In this section, we present two important examples of the threat detection problem and discuss how each might be formulated in the framework presented above. The chief differences between these examples are the costs that are placed in the loss matrix and the moves available to the adversary.

2.1 Smuggling a Nuclear Weapon

For this example, we assume that a terrorist desires to smuggle a nuclear weapon past the checkpoint to detonate in a large city. One important characteristic of a nuclear weapon (at least one that a terrorist could manufacture) is its relatively constant size, which translates into a relatively constant loss when it is successfully detonated. Thus, for this example, we assume that if he gets through this checkpoint, he will succeed and cause a trillion dollars worth of damage ($\$10^{12}$).

Table 2.1. Loss Matrix for a Nuclear Weapon

	Adversary Moves				
Operator Moves	No Threat Config=0	Threat Config=1	Threat Config=2	...	Threat Config=N
Pass (i=0)	\$0	$L_{miss} = \$10^{12}$	$\$10^{12}$	...	$\$10^{12}$
Detain (i=1)	$L_{fcp} = \$10^{3}$	\$0	\$0	...	\$0

This assumption produces a loss matrix with the form illustrated in Table 2.1–the loss associated with each bomb configuration is the same. The different bomb configurations listed in the loss matrix are meant to describe different methods the adversary might use to "hide" the bomb from the detector, which we will assume is a gamma-ray detector. It is assumed that the adversary can employ different amounts of shielding (up to a certain maximum, dictated by weight constraints). Thus, the configurations in the loss matrix range from "No shielding" (Configuration N), to "Maximum Shielding" (Configuration 1).

The adversary's choice of container configuration therefore only affects the detector signal distribution, $f(S|config = j)$. Figure 2.1 provides an example of how the adversary's choice might affect these distributions for a gross-count gamma detector. For multi-spectral detectors, the effect to the distribution are much more complicated.

When the operator inspects a clean container, a cost of \$1000 is incurred when a container is detained and sent through secondary. This represents the cost due to lost time and manual inspection.

2.1.1 Constraint on the Threat Frequency

Nuclear weapons are difficult devices to manufacture, especially for a clandestine organization such as a terrorist network. We would expect such an organization's deployment strategy to be

Figure 2.1. Detector Signal Distribution for various Adversary Choices

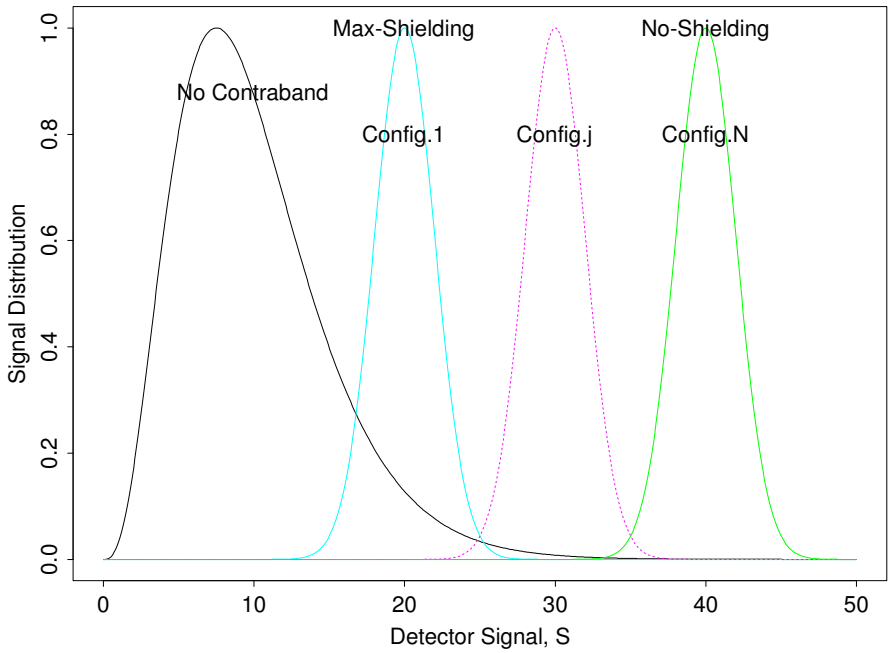

constrained by this fact. For the detection game under consideration, this would translate into a constraint on the frequency for which the adversary can choose anything but $Config = 0$.

For example, the checkpoint might expect to see no more than, say, one nuclear device every 10 years. If the checkpoint was inspecting a million containers a year, this would place a one-in-ten-million constraint on the frequency with which the adversary could send a threat container through, a very severe constraint.

The constraint can be dealt with in the same manner as described for the FCP constraint placed on the detector operator. That is, first ignore the constraint and determine the solution for the unconstrained problem. If this solution obeys the constraint, then it is also the solution to the constrained problem.

If the constraint needs to be included in the problem, we accomplish this by altering the definition of "move" for the game. If the adversary can manufacture one weapon in say, 10 years, we define a move as the placement of the weapon somewhere in a 10-year traffic stream of containers. This changes the adversary's choices to configuration $j = 1, 2, ..N$, with an average risk per move of;

$$R(\delta, j) = \$10^3(1 - T_d)Pr(\delta(S) = 1|Config = 0) + \$10^{12}T_dPr(\delta(S) = 0|Config = j) \qquad (2.1)$$

where T_d is the *threat density*, the rate at which the adversary can produce a nuclear weapon.

2.1.2 Comments on the Loss Matrix

We close our discussion of this example with an important comment about the loss matrix. One could argue that the loss matrix presented in Table 2.1 is unduly pessimistic from the checkpoint operator's perspective. This loss matrix assigns a *zero* dollar value (i.e. a status-quo outcome) when the checkpoint operator happens to actually detect a weapon. When the weapon is detected, it is true that no loss to the city occurs, so from this perspective a zero-loss assignment makes sense.

However, one could argue that other benefits would accrue to the checkpoint operator; he would gather substantial information on the terrorist network that might allow him to put it out of the bomb building business for a substantial time period. At the very least, one could argue that the detector operator has captured a nuclear weapon, and the cost of this weapon should be credited to the operator.

If one accepted that this perspective had merit, then one would enter a negative number in the lower, right-hand cells of the loss matrix to represent this loss. For example one could easily argue that the capture of a terrorist's container represents a gain of ten million dollars (or loss of -10^7).

This alteration in the loss matrix changes the solution considerably. With the original loss matrix, detectors would have to achieve a probability of non-detection below 10^{-9} to have a deterrent effect[a], an unattainably small number for real detectors. If the altered loss matrix were considered applicable, detectors with PND's below 10^{-5} would achieve a deterrent effect, and this performance is possible for a good detector.

In fact, one can work out a simple equation that describes what PND our detectors should achieve, if we desire a deterrent effect. The formula is

$$PND < \frac{L_{fc} - L_{detect}}{L_{miss} - L_{detect}} \tag{2.2}$$

where L_{fc}, L_{detect}, and L_{miss} represent the losses associated with a false call, a detection, and a bomb that was missed.

2.2 Smuggling a Dirty Bomb

In this example, we assume a terrorist desires to smuggle a dirty bomb through the checkpoint and detonate it in a city. In contrast to the nuclear device, the adversary can choose the size of the explosive device–with the size being proportional to the amount of damage the device would cause; the adversary might choose to smuggle a small device through the checkpoint, which is hard to detect, but would produce less damage than a larger device. We will assume that this is the principal choice (or move) that the adversary must make.

In this case, the primary loss matrix would resemble the matrix presented in Table 2.2. To con-

(a) Having a *Deterrent Effect*, means that the adversary's optimal solution includes Config 0 with positive probability

struct this matrix, we have assumed that loss associated with each configuration is proportional to the size, X_i, of the dirty bomb. The "size" of the bomb, as represented by the variable, X, is measured in Kg. The loss matrix assumes that loss is directly proportional to the bomb size. We also assume that the detector operator receives a payoff of $100,000 whenever he detects a bomb.

Table 2.2. Primary Loss Matrix for a Dirty Bomb

	Adversary Moves				
Operator	No Threat	Threat			
Moves	Size.0	Size.1	Size.2	...	Size.N
Pass	\$0	$\$10^8 X_1$	$\$10^8 X_2$	...	$\$10^8 X_N$
Detain	$\$10^3$	$-\$10^5$	$-\$10^5$	...	$-\$10^5$

The detector's signal distribution also depends upon the amount of radiation contained in the bomb, resulting in a distribution that can be considered to be a function of X (i.e. $f(S|config.i) = f(S|X_i)$). In fact, a reasonable assumption is that the mean of the signal distribution is proportional to the bomb size, resulting in detector distributions similar to those illustrated in Figure 2.1. A specific example of this type of distributional model that will be used in this report is;

$$f(S|X_i) = \phi(S - \frac{4.66}{0.05}X_i) \tag{2.3}$$

where;

- S is the detector signal,
- X_i is the size of the dirty bomb (KG), and
- $\phi(z)$ is a standard normal density function.

It should be noted that the game theoretic formulation can help the checkpoint operator focus his search for a dirty bomb. The adversary *could* make a dirty bomb as small as a cell-phone or as large as the truck. Intuition suggests that a very small bomb should not be selected by the adversary; it would not be worthwhile to him. On the other hand, a very large bomb should also the unreasonable because it would be too easy to detect. One would expect that the adversary would select bombs from a certain size-range. If there is such a size-range, this would be important information for the checkpoint operator, affecting the way he would conduct searches in his secondary.

3.0 General Solution to the Threat Detection Game

In this section, we outline the decision theory solution to the Threat Detection Game and discuss some of the solution's more important properties. The solution to this type of problem is called the *minimax* solution Ferguson (1967); it is the solution that simultaneously minimizes the operator's risk while maximizing the adversary's gain. The operator's/adversary's minimax strategies produce a stable solution to this game; if either player deviates from the minimax strategy, the player who deviates will lose.

We would argue that the expected loss of the game under the minimax strategy produces the best description of the detector capability; evaluation of the detectors with any other adversary/operator strategy will produce an unrealistic description of detector performance. If one does not use the minimax strategy for the adversary, and one tries to find a 'good" strategy for the operator, the result will be too optimistic. If, on the other hand, one does not employ the minimax strategy for the operator, and then employs a red team analysis, the result will be too pessimistic.

3.1 Mathematical Nomenclature

The basic components of the threat detection game consist of a *loss matrix* and a set of *detector signal distributions*. Mathematically, these components can be described by the following notation;

Loss Matrix is represented by L_{ij} or $L(i,j)$, the loss when the checkpoint operator chooses i and the adversary chooses j. The index i is defined by;

$i = 0$ Pass the container
$i = 1$ Detain and Inspect the container

For the adversary, j identifies the choices;

$j = 0$ Let a clean container pass through.
$j > 1$ Send a threat container with configuration j through.

Signal Distribution is represented by $f(S|j)$ or $f(S|Config = j)$. This represents the distribution of the detector signal, S, when a container with configuration j is measured. Note that the zero-configuration distribution, $f(S|j = 0)$, is not under the control of the adversary, but the other distributions are.

Risk Matrix: The risk matrix is calculated for the loss matrix and signal distributions using the formula;

$$R(\delta, j) = \int (L(\delta(S), j) f(S|j)\, dS \tag{3.1}$$

where δ represents the operator's detection algorithm, and j the adversary's choice of container configuration.

3.2 Properties of the Game Theory Solution

The game theoretic solution to the threat detection problem consists of a detection algorithm, δ_{best}, and a *mixed strategy* P_{LFD} for the adversary that attains the minimax value of the game. A *mixed strategy* for the adversary is described by a vector of probabilities, $P = (P_0, P_1, ...P_N)$ which specify the frequency with which he chooses each configuration (i.e. Configuration j is chosen with probability P_j). The game theory solution for the adversary is also called the *least favorable distribution* and denoted by P_{LFD} because it is least favorable for the operator.

The risk associated with a mixed strategy, P is computed with the formula;

$$R(\delta, P) = \sum_j R(\delta, j) P_j \tag{3.2}$$

It is important to note that a best strategy will not generally exist for the adversary unless he is allowed to employ a mixed strategy. One would expect the same to hold for the detector operator, but this can be shown to be not the case Ferguson (1967); the operator's best solution will consist of a single decision function from the risk matrix, and not a mixed combination. However, this decision function might include a random component in its construction.

The game solution (δ_{best}, P_{LFD}) are defined to be those player strategies that achieve the minimax loss. In other words, we must find a δ and P that satisfy;

$$R(\delta_{best}, P_{LFD}) = \min_{\delta} \max_{P} R(\delta, P) \tag{3.3}$$

It can be shown that a solution always exists for the threat detection problem. Also, it can be shown that the solution also obeys;

$$R(\delta_{best}, P_{LFD}) = \max_{P} \min_{\delta} R(\delta, P) \tag{3.4}$$

These two equations are used to compute the solution and define its desirable properties. Verbally stated, the first equation shows that if the detector operator uses δ_{best}, he cannot lose more than $R(\delta_{best}, P_{LFD})$, and may lose even less if the adversary deviates from P_{LFD}. Conversely, the second equation shows that the adversary who uses P_{LFD}, will cause a loss of at least $R(\delta_{best}, P_{LFD})$ to the operator, and may cause an even greater loss if the operator deviates from his best strategy.

3.3 The Operator's Best Strategy

In this section, we present the form the operator's decision function must have to be considered as a possible solution to the problem. The form rests upon equation 3.4. From equation 3.4, one can show that the best decision function to use against the fixed adversary strategy, P, is the one that minimizes the risk function $R(\delta, P)$. This detector strategy is also called the "Bayes solution" to the problem. The Bayes decision functions form an *admissible* set of solutions to the problem and consequently the search for the best function can be limited to this set.

It can be shown that the best δ against strategy P can be defined in terms of the risk achieved

when a signal of S observed. If the operator detains the container, his risk is;

$$R(S, i=1) = \sum_j L(1,j) f(S|j) P_j \tag{3.5}$$

and if he passes the container the risk is;

$$R(S, i=0) = \sum_j L(0,j) f(S|j) P_j \tag{3.6}$$

It would make sense for the operator to minimize his risk by selecting the option $i = 0, 1$, (Pass, Detain) so that $R(S,i)$ is minimized. This, in fact, can be shown to produce the Bayes decision function for the operator. Thus, the decision function, $\delta(S)$, is defined by;

$$\delta(S) = \begin{bmatrix} 1 & R(S,i=1) < R(S,i=0) \\ 0 & R(S,i=1) > R(S,i=0) \\ Z & R(S,i=1) = R(S,i=0) \end{bmatrix} \tag{3.7}$$

The case when the two risks are equal merits some comment. For many problems, the chance that this will occur has probability zero, and one can arbitrarily choose $Z = 0$. However, for an important set of cases, equality will occur with finite probability, and in these cases, the outcome, as represented by Z *must be randomly chosen*. For these cases, Z represents a binomial variable, chosen to equal 1 at some probability.

To emphasize that the decision function, $\delta(S)$, defined by Equation 3.7 is best against the adversary's strategy P, we will write it as $\delta(S,P)$ to emphasize that fact. This formula produces δ_{best} for us, but to do this, we first need to have found the adversary's best strategy. If the adversary's best strategy, P_{LFD}, is known one can show that;

$$\delta_{best}(S) = \delta(S, P_{LFD}) \tag{3.8}$$

Figure 3.1 illustrates the Bayes decision functions for a simple case. For this case, both the threat and no-threat distributions are assumed to be bivariate normal, with the threat signal showing a lower variance than the no-threat signal. This example also assumes that there is only one threat configuration being used by the adversary. The two distributions are indicated by colored contour plots in the Figure and specific decision functions, $\delta(S,P)$ are represented by lines, which bound the detain-region ($\delta(S,P) = 1$) from the pass-region ($\delta(S,P) = 0$). Three decision functions are plotted in the Figure; those Bayes decision functions associated with $P_0 = 0.10, 0.50, 0.90$, adversary strategies.

When the signal distribution is multivariate normal, one can show that the optimal decision function is a quadratic form involving the distribution means and covariances. This is exactly what we see, in Figure 3.1; in this case, the quadratic form produces circular boundaries.

3.4 The adversary's best strategy; P_{LFD}

Since all the adversary's possible strategies are represented by a finite vector of real values (i.e. the vector P), it is possible to employ a numerical optimization routine to find the vector P that

Figure 3.1. Example of Best Decision Functions when the Signal Distribution is Bivariate Normal

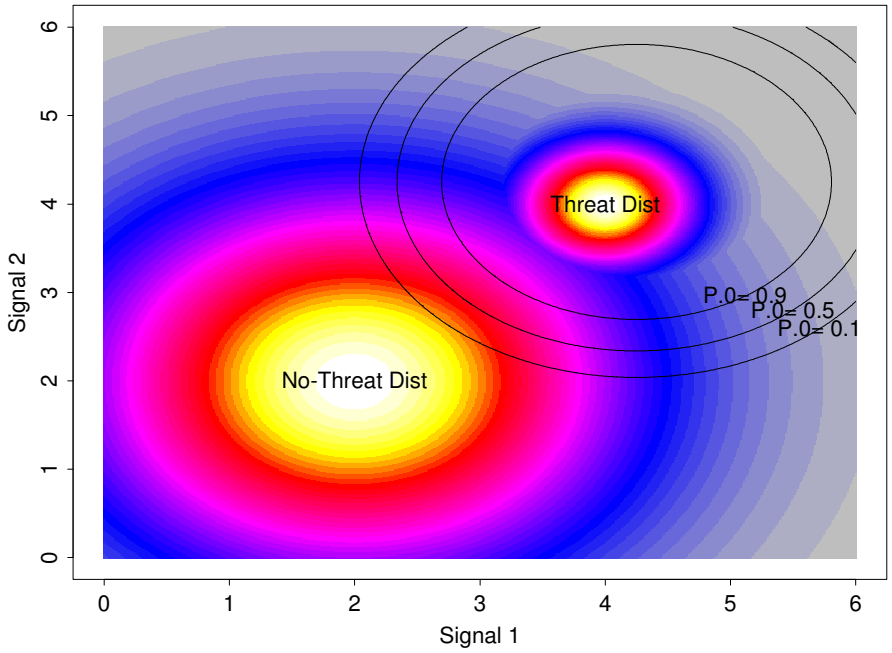

maximizes the risk. By substituting in the operator's solution found in Equation 3.7, one reduces the maximization problem to;

$$\max_P \min_\delta R(\delta, P) = \max_P R(\delta(\cdot, P), P) \tag{3.9}$$

So the least favorable distribution, P_{LFD} is the probability vector P that produces the maximum for the function;

$$R_0(P) = R(\delta(\cdot, P), P) = \sum_{j=0}^{N} H_j(P) P_j \tag{3.10}$$

where $H_j(P)$ is defined as;

$$H_j(P) = \int L(\delta(S, P), j) f(S|j)\, dS \tag{3.11}$$

To solve Equation 3.10, we employ an iterative solution to this problem that rests on playing the game McKinsey (2003). For this solution, we start by assuming some value for P, say $P = (1/n, 1/n, \ldots 1/n)$, and the operator calculates his detection function assuming this P. Each player makes his first play using this P. At play k of the game, one constructs an estimate $\hat{P}$ for the adversary's strategy by averaging together all the previous moves the adversary has made. The detector operator uses this estimate to calculate the decision function he will use for play k.

The adversary makes his choice by calculating the functions $H_j(\hat{P})$, and choosing the largest one. In other words, his choice for play k will be the configuration associated with the largest risk. One can show that as the number of plays gets large $\hat{P}$ will converge to (the *least favorable distribution*), P_{LFD}.

Although this algorithm is simple to implement and understand, it may not be the most efficient algorithm for obtaining a solution. However, it does illustrate an evolutionary property of the minimax solution; If two intelligent players play against each other over a long period of time, and carefully analyze past results to improve their play, they should arrive at the minimax solution.

Figure 3.2 displays the LFD solution for the example described by Figure 3.1. In this example, the adversary only has two alternatives to choose from, so his mixed strategy can be described by a single probability, P_0 (The probability he chooses the no-threat alternative). For this example, it is easy to find the maximum for $R_0(P_0)$ by an exhaustive search.

Figure 3.2 plots $R_0(P_0)$ versus P_0. The maximum occurs at $P_0 = 80\%$, so the adversary's best strategy is to choose the no-threat alternative 80% of the time and the threat alternative 20% of the time. From this Figure, we also see that the payoff is \$10 per move. This is the loss the operator can expect to see when both players use their optimal strategies.

Figure 3.2. Example of the LFD for a Bivariate Normal Signal Distribution

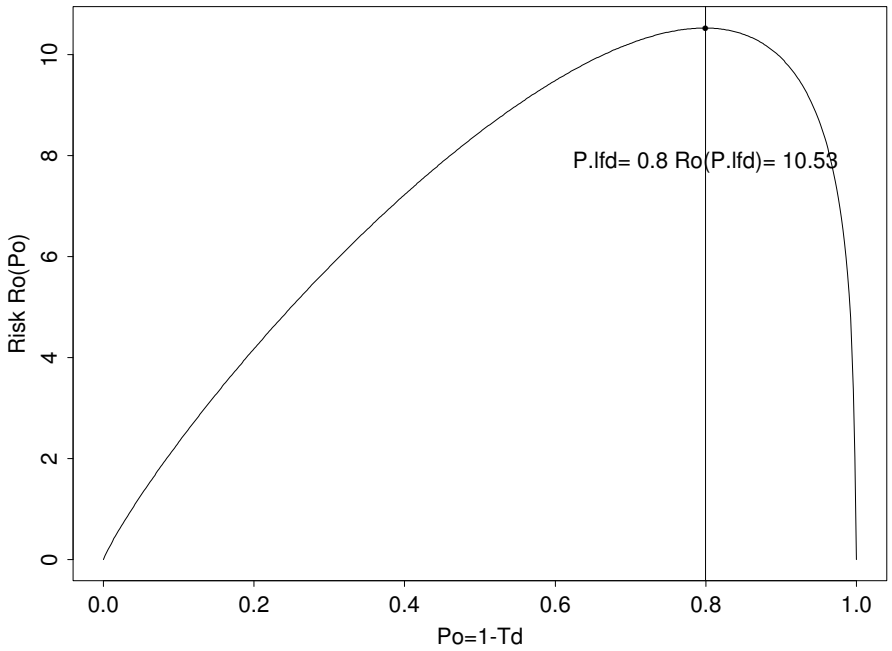

4.0 Solution to Nuclear Weapons Smuggling Example

In this section, we discuss interesting characteristics that the game-theory solution has in the case of the Nuclear Weapons Smuggling problem outlined in Section 2.1.

For simplicity, we assume the detector produces a single signal, such as an average gamma-count, so that the signal distributions, $f(S|j)$, are univariate. When the adversary chooses configuration j, the resulting signal, S_j is produced according to the regression model;

$$S_j = X_j + E \tag{4.1}$$

where the random variable E represents Gaussian measurement error with mean zero and a standard deviation of 1 (In other words, the raw signals have been divided by the measurement standard deviation to produce S_j). The term X_j represents the scaled gross count produced by a container of configuration j.

This model allows us to clearly distinguish those portions of the signal distribution that are determined by the adversary and those that are determined by nature. The measurement error, E, and also the counts originating from a clean container, X_0 are determined by nature. X_0 represents the gross count from a "clean" container and it's magnitude is determined by the natural radioactivity of the container's cargo. X_0 is therefore a random variable, which is defined by the clean container stream. This distribution can be estimated from measurements on the clean traffic stream and Figure 4.1 presents the distribution for X_0 that has been used for this example.

Figure 4.1. An Example of the Distributions Produced by a Clean Container Stream

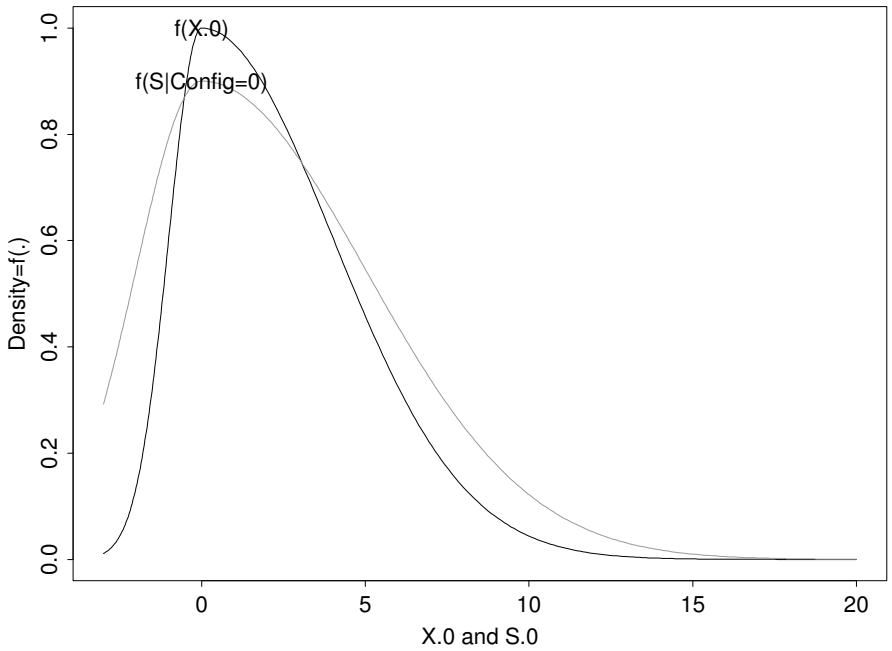

On the other hand, the count-rates associated with any bomb configuration, as represented by X_j, $j > 0$, are controlled by the adversary. The adversary can alter the magnitude of X_j by the amount of shielding he uses in configuration j. There is an important constraint on the adversary's choices; weight restrictions would place a limit on the maximum amount of shielding he could use, so that any X associated with a configuration he could produce would be greater than a certain minimum count-rate. There would also be an upper bound on the count rate (as long as he didn't hide his bomb in radioactive cargo); the upper bound would be the count rate associated with an unshielded nuclear bomb. So the adversary is free to construct a configuration that produces a count-rate of X, under the restriction;

$$X_{max.shield} < X < X_{no.shield} \tag{4.2}$$

We formulate a discrete version of this set of adversary alternatives by assuming he will consider $j = 1, 2, ...11$ configurations with values X_j equally spaced between $X_{max.shield}$ and $X_{no.shield}$. For this specific example, we assume that $X_{max.shield} = 7$ and $X_{no.shield} = 17$, as illustrated in Figure 2.1. The distributional model discussed above is applicable to any detector based on the measurement of a single magnitude, and has broader application than to the nuclear problem. One could argue that the left-hand bound, $X_{no.shield}$ should not be included for this problem. However, from the solution, it will be easy for the reader to see what the solution would be if it were omitted.

Figure 4.2 illustrates the signal distributions associated with this problem. The broad distribution in the figure represents the signal distribution from clean containers, and is the convolution of the clean container distribution in Figure 4.1 and the measurement distribution. The distribution of the signal produced by configuration $i > 0$ is just a standard normal distribution centered at X_j. An adversary's mixed strategy produces a "bomb distribution" of;

$$\sum_{j=1}^{N} \phi(S - X_j) P_j \tag{4.3}$$

which is the discrete form of a convolution.

A critical restriction on the adversary's choices is the fact that he cannot produce a bomb with a gross count *lower than* X_{shield}. If he could produce any gross count, one can show that his best strategy would be to use the mixed strategy $P_j \approx f_{X_0}(X_j)$ for selecting the container configurations and this strategy would render the detector completely ineffective. In other words, the detector could do no better than random selection of detained containers.

This observation provides us with some simple guesses for the adversary's best strategy. We might try to select a probability vector P so that the clean distribution $f(S|0)P_0$ is as close as possible to the threat distribution $\sum_j f(S|j)P_j$. If, for example, we interpreted "as close as possible" in the least squares sense, one would solve for P using regression. Figure 4.3 presents such a regression solution for P. One can see that the resulting "threat distribution" is almost identical in shape to the clean container distribution between X_{shield} and $X_{no.shield}$, indicating that any signal in this region would have no information about the true state of the container. If the adversary were to employ the strategy produced by least squares, he would at least defeat the detector whenever it produced a value in this region.

Figure 4.2. Signal Distributions Associated With the Bomb-Detection Problem

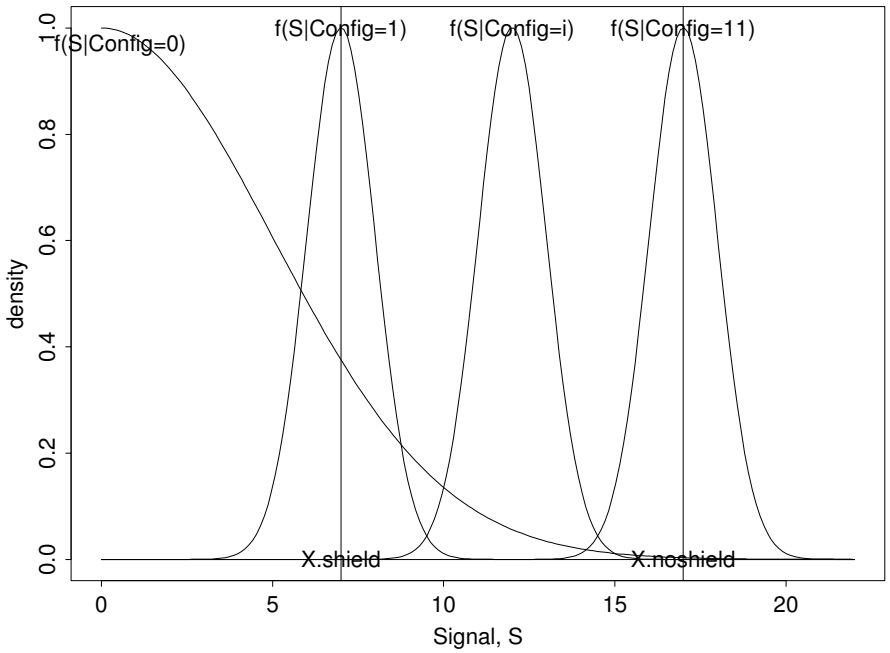

Figure 4.3. Solution from Least Squares fit of the Threat Distribution to the Clean Distribution

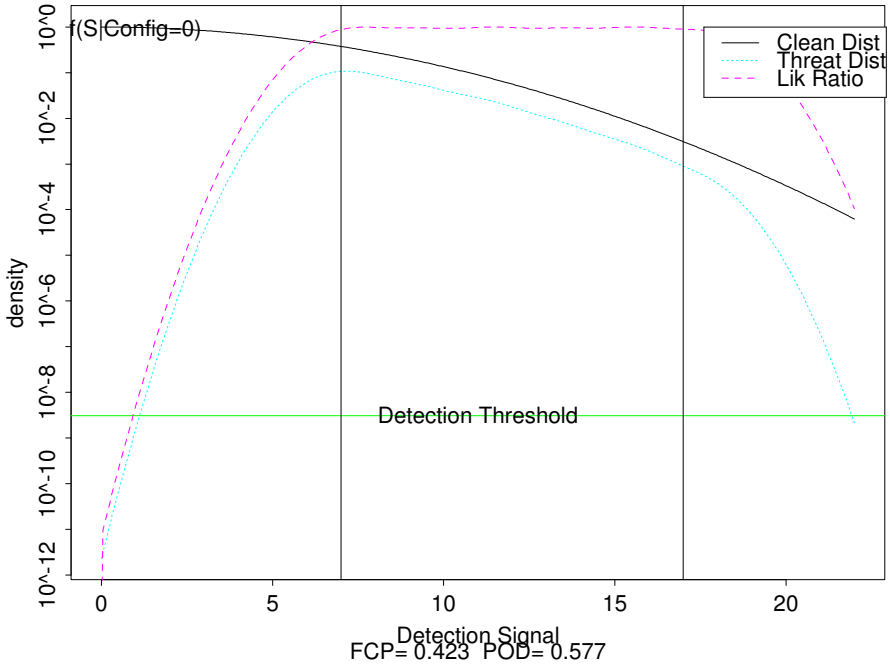

On the other hand, one might argue that the adversary should select the threat configurations that are closest to the null distribution, and this would lead one to select threat distributions defined by the region boundaries, This sort of reasoning would lead one to a mixed strategy that included only configurations 0,1 and 11.

4.1 Minimax Solution

One can calculate the LFD for this example by maximizing the function in Equation 3.10. When the loss matrix presented in Table 2.1 is used, this results in a least favorable distribution that assigns probability to only three configurations: $P_{LFD}(0) = 0.944$, $P_{LFD}(1) = 0.056$, and $P_{LFD}(11) = 10^{-5}$. In other words, one of the simple guesses described in the last section describes the solution; the adversary's solution only uses the boundary configurations.

Table 4.1. Minimax Solution for Nuclear Weapon Smuggling Example

Minimax Loss	\$425
Operator Strategy $\delta(S) =$ detain	 $3.5 < S < 21$
Adversary Strategy $P_{LFD}(0)$ $P_{LFD}(1)$ $P_{LFD}(0)$	 0.994 0.056 10^{-5}

The general form of the adversary's strategy is more complicated than this. There are some conditions where the least-squares solution for the threat distribution is close to the adversary's optimal strategy and to illustrate this, we varied the costs in the loss matrix. Since the solution only depends upon the ratio of the two costs in the loss matrix (i.e. L_{miss}/L_{fcp}), we only had to vary one of the costs to examine the complete set of solutions.

We found that as the L_{miss} cost is lowered, the adversary's solution becomes more like the least-squares solution. When L_{miss}/L_{fcp} is near 1, the adversary's threat distribution closely resembles the least squares solution presented previously. But when this ratio much greater than 1, the adversary should choose a threat distribution that favors the two extreme configurations.

The adversary's general strategy is a combination of these two cases and Figure 4.4 illustrates this. The top graph in this figure plots the threat distribution as L_{miss} varies from \$10 to $\$10^{12}$. With $L_{miss} = 10$, the threat distribution resembles the least squares solution. As L_{miss} is increased, the configurations associated with the two end-points are preferentially selected, resulting in a smaller region that resembles the least-squares solution. For extremely large values of L_{miss}, there is no region that resembles the clean distribution and the adversary's best strategy only involves the end points.

Figure 4.4. Solutions for various loss matricies: $L_{fcp} = \$1000$ while L_{miss} varies from \$10 to \$$10^12$

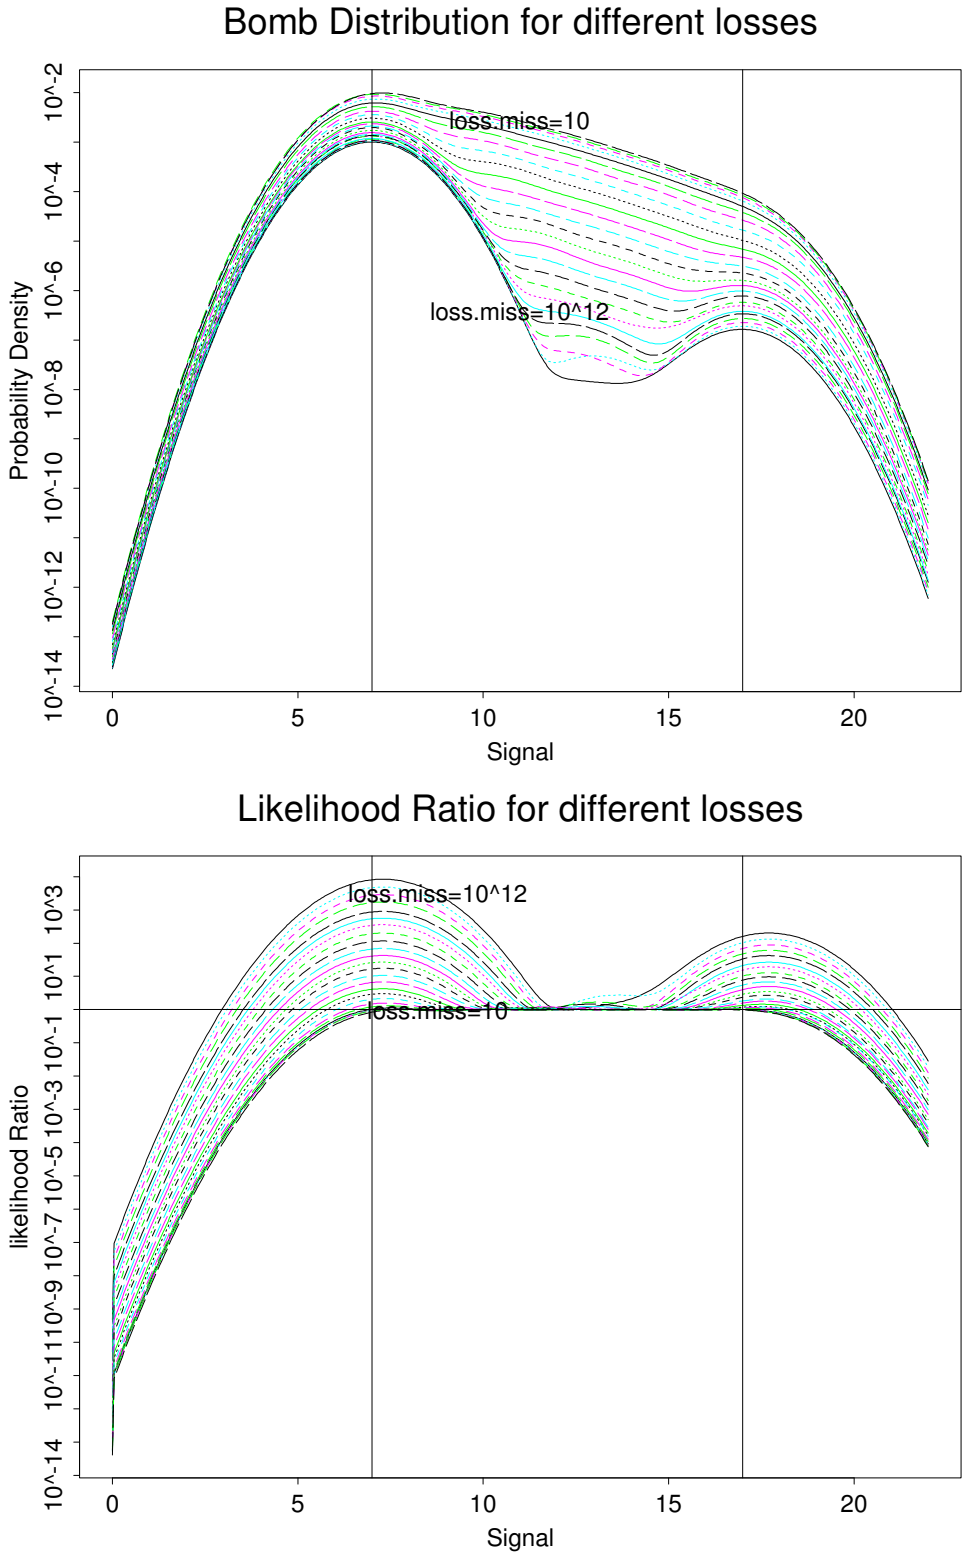

4.2 Form of Operator's detection Algorithm, $\delta(S)$

The loss matrix presented in Table 2.1 has a special structure that allows the decision function to be simplified. In this loss matrix, (1) the losses are the same for each threat configuration, and (2) the diagonals in the matrix are zero. This structure will allow the optimal decision function, $\delta(S|P)$ to be expressed in terms of a well known statistic employed in hypothesis testing, the *likelihood ratio statistic.* For this loss matrix, the decision function can also expressed as;

$$\delta(S|P) = \begin{bmatrix} 1 & \gamma(S) > C_v \\ 0 & \gamma(S) < C_v \\ Z & \gamma(S) = C_v \end{bmatrix} \tag{4.4}$$

where C_v is the pass/detain decision threshold and $\gamma(S)$ represents the likelihood ratio statistic;

$$\gamma(S) = \frac{f(S|Threat)}{f(S|Clean)} = \frac{\sum_{j=1}^{N} f(S|j)\lambda_j}{f(S|0)} \tag{4.5}$$

This statistic is the ratio of the "Clean" distribution, $f(S|0)$ and the "Threat" distribution, $\sum_j f(S|j)\lambda_j$. This likelihood ratio statistic is best against an adversary's mixed strategy, P. with the assignments;

$$\lambda_j = \frac{P_j}{1-P_0} \tag{4.6}$$

and

$$C_v = \frac{L_{fcp}P_0}{L_{miss}(1-P_0)} = \frac{10^3 P_0}{10^{12}(1-P_0)} \tag{4.7}$$

This formulation for the decision function has important consequences. First, we can see that the decision function is only affected by the ratio of the costs (i.e. L_{fcp}/L_{miss}) in the loss matrix, and this ratio only appears in Equation 4.7, which determines the pass/detect threshold. Since the equation for the likelihood ratio contains no costs, one might conclude that it is independent of values in the cost matrix. This conclusion is not true, because the likelihood ratio is defined in terms of the λ_j, which do indirectly depend upon this cost ratio.

Figure 4.5 illustrates the shape of the likelihood ratio solution for two different loss matricies. The top graph describes the solution when $L_{miss} = 10^{12}$, (or $L_{miss}/L_{fcp} = 10^9$), while the bottom represents the case when $L_{miss} = 10^5$ ($L_{miss}/L_{fcp} = 10^2$). As one can see, the two likelihood ratio curves (indicated by the red curves) have substantially different shapes, and produce different decision regions. In the case $L_{miss} = 10^{12}$, the detection region corresponds to a simple threshold on the signal: any signal greater than 0.77 results in a detection.

However, the likelihood ratio results in a more complex decision algorithm for the $L_{miss} = 10^5$ case:

Pass: Pass the container if the signal is less than 4 or greater than 20.

Detain: Detain the container if the signal is in the interval (4,11) or in the interval (15,20).

Figure 4.5. MiniMax Solution for the Nuclear Smuggling Example

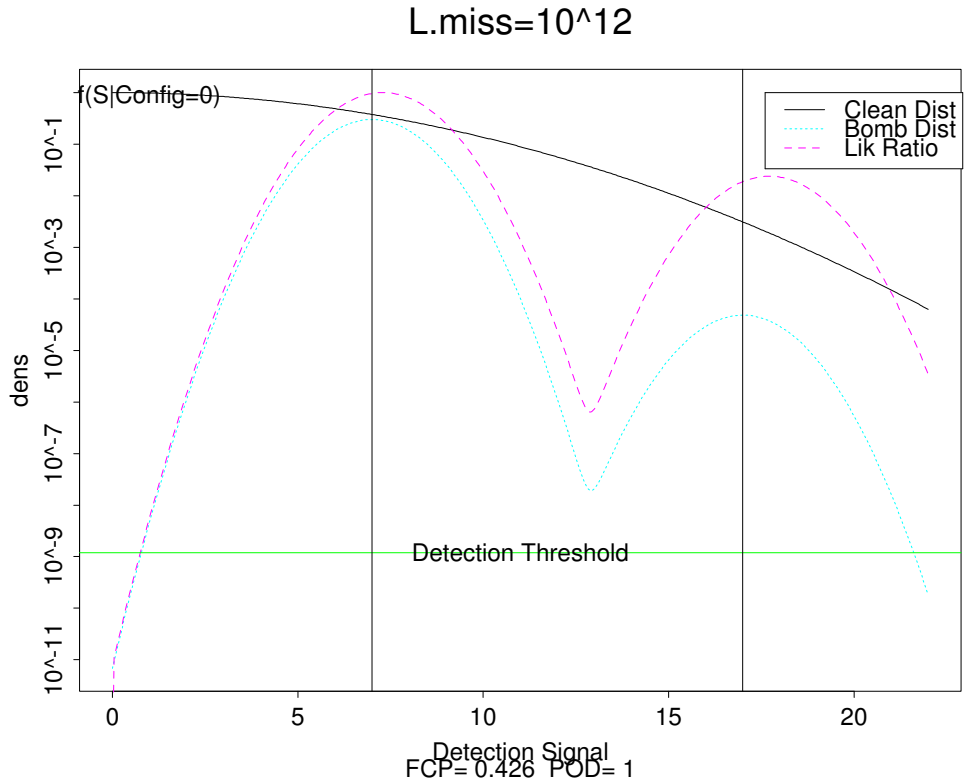

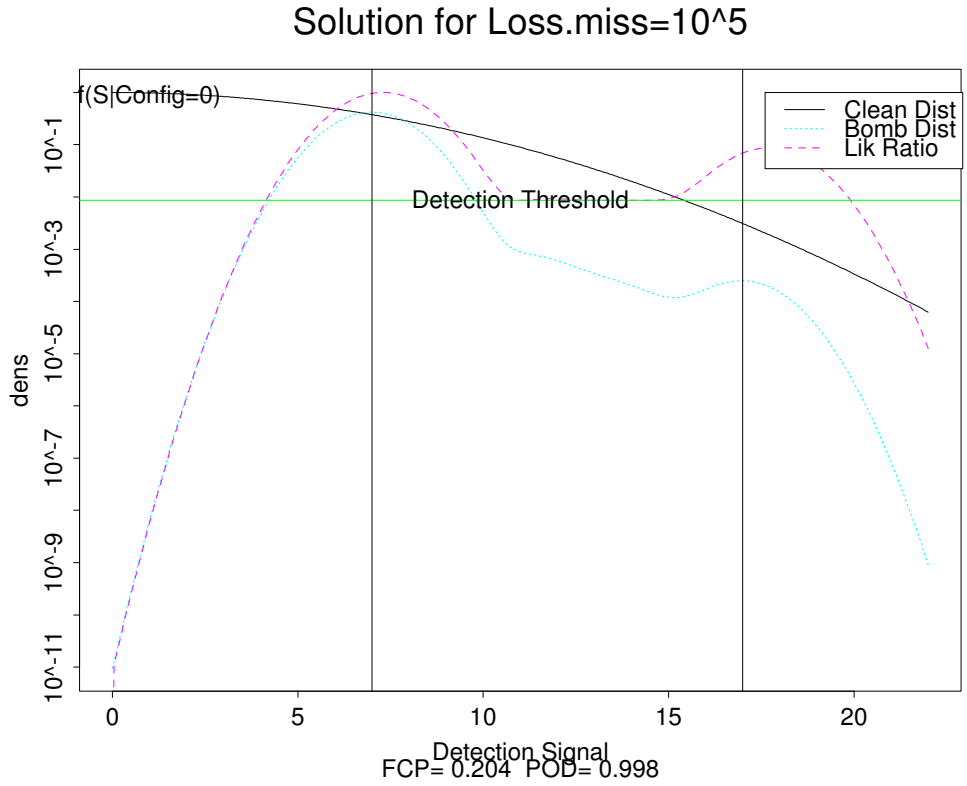

Detain with probability 5%: if the signal is in the interval (11,15).

In other words, for this case, random selection is a key portion of the optimal decision algorithm. In the region (11,15), the threat and clean distributions match each other, and produce a solution similar to the "least squares" solution illustrated in Figure 4.3.

The bottom graph in Figure 4.4 illustrates how the decision procedure is affected by the costs in the loss matrix. When the ratio L_{miss}/L_{fcp} is large, the decision procedure involves no randomization, and reduces to simple thresholding on the signal, S. But when the ratio is small, signals from the "middle" of the interval (7,17) are randomly passed/detained.

For any realistic nuclear threat detection problem, L_{miss} should represent a very large number, and one would expect the most relevant minimax solution to be the one associated with $L_{miss} = 10^{12}$. However, one must remember that most checkpoints operate with a false call constraint, as described in Section 1.5.

For example, the unconstrained solution for $L_{miss} = 10^{12}$ produces a FCP of 43%, a value that would be unrealistic at most checkpoints. If the FCP had to be reduced to 20% to be feasible, then the decision algorithm described in the bottom graph of Figure 4.5 would be the relevant solution. If the FCP had to be reduced to 1%, then the decision algorithm associated with $L_{miss} = \$10$ would be applicable.

4.3 ROC Curves associated with the Minimax Solution

One can use the minimax solutions to create an ROC Curve[(a)] as illustrated in Figure 4.6. This ROC curve is an upper-bound envelope on the (FCP,POD) performance of all decision procedures. The most striking feature of this ROC curve is the nearly linear region between $FCP = 0$ and $FCP = 0.05$ which is related to the randomization required when FCP is low.

To compare the least squares solution to the true minimax solution, we have also included its ROC curve in Figure 4.6. As one can see, the least-squares ROC matches almost perfectly with the MiniMax ROC when $FCP < 0.05$, but is slightly more optimistic for $0.05 < FCP < 0.20$. In this region, the least-squares solution isn't using the adversary's best strategy, so it produces POD values that are too large.

Another detection algorithm is presented in Figure 4.6, and it is included to illustrate why one should be interested in the MiniMax solution. The green ROC curve in Figure 4.6 represents the performance associated with a simple thresholding detection algorithm when attempting to detect a shielded threat. In other words, the detection algorithm detains a container when the signal S, is greater than a threshold, T. This is the most common detection algorithm employed on a univariate signal, and one can see that the Minimax detection algorithm performs *much* better at low FCP. In fact, the performance of simple thresholding in the region of $0 < FCP < .05$ is worse than using random container selection at the checkpoint, an indication of very poor performance!

(a) A 2-dimensional ROC curve can only be defined when the loss matrix has the form defined in Table 2.1

Figure 4.6. Comparison of ROC Curves for the Minimax, Least squares and Simple Thresholding Solutions

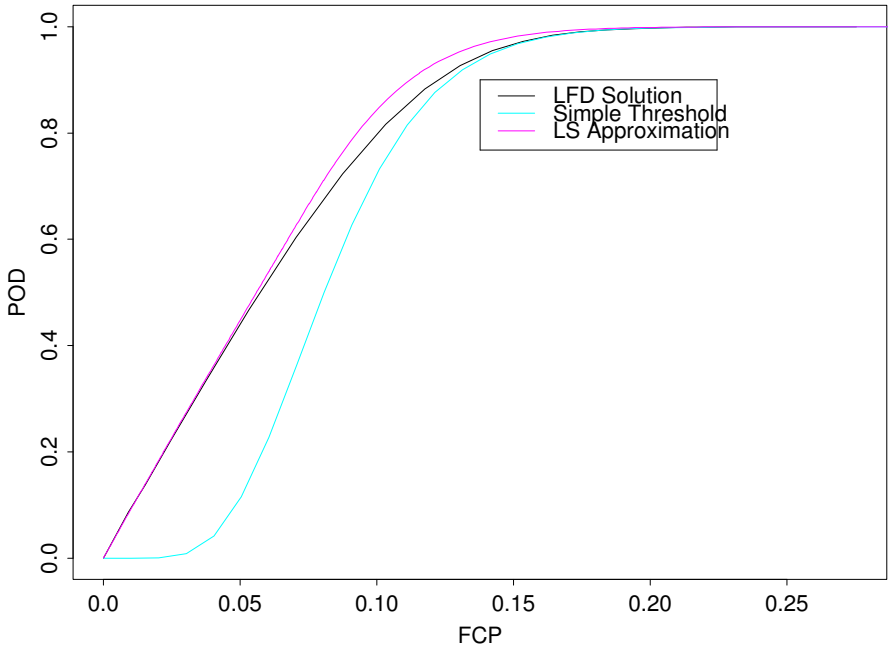

5.0 Solution to Dirty Bomb Example

The dirty bomb example described in Section 2.2 has a particularly simple solution for the adversary; his optimal strategy is to use a random mixture of two choices; the $X = 0$ (no-threat alternative and the $X = X_{best}$ threat alternative, with X_{best} identified by the minimax algorithm. For this example, the solution characteristics are summarized in Table 5.1, under the column labeled "Container."

Table 5.1. Minimax Solution for RDD Problem

	Car	Container
Minimax Loss, L_{mm}	$928	$949
Fcp	.928	.949
Pnd	3.22e-2	1.96e-2
X_{best}	.3Kg	.5Kg
$P_{LFD,0}$	6.08e-4	4.30e-4

For this example, the bomb the adversary should use has a size of $X_{best} = 0.5Kg$. Bombs bigger or smaller than this will either be too easy to detect or not cause enough damage. Also, the adversary should attack with approximately 5 bombs in a traffic stream of 10,000. To counter the adversary, the checkpoint operator should run at a very high false call rate (95%), which produces a PND of 2% for X_{best}.

Also present in the table is a case for dirty bomb detection in hidden in a Car. The losses associated with the car are the same as those for the container, but it is assumed that it is easier to detect a dirty bomb in a car so the signal distribution for this case is;

$$f(S|X_i) = \Phi(S - \frac{6.44}{0.05}X_i) \tag{5.1}$$

Because of this improved detection capability, the adversary must use a smaller dirty bomb (0.3Kg), but this does not appreciably change the overall loss the detector operator suffers. For this configuration of losses and detection capability, the adversary is able to force the checkpoint operator to spend large amounts in secondary inspection.

6.0 Conclusions From the Game Theory Solution

Best Form of Operator's Detection Algorithm: The best decision algorithm will have the form as defined by Equation 3.7 for some choice of P. Any study focused on developing good decision algorithms for a detector should at least consider this class of detection algorithms as possible candidates. If the signal distributions, $f(S|j)$, are known and there is consensus regarding the costs in the loss matrix, one can argue that the decision algorithm produced by Equation 3.7 should be used at the checkpoint.

Red Team Evaluations: The typical attempt to account for the adversary's behavior involves some sort of "Red Team" activity. From the game theoretic solution, we would argue that a good red team evaluation should always consider mixed strategies, and in particular, the mixed strategy associated with the least favorable distribution. The ideal red-team scenario to use for detection algorithm design are the Adversary's LFD strategy.

Interaction between Red team activities and detection algorithm development: One important feature of the game theoretic formulation is that it does not separate red team evaluations (i.e. evaluation of adversary's strategy) and detector optimization. From the game theoretic solution, one can see that both activities must proceed simultaneously. Any investigation that attempts to separate red team evaluations from detector optimization will have limited success.

Relationship of game-theory to Risk Analysis: A standard probabilistic risk analysis model can be thought of as a component of the game-theory model; to compute elements in the Risk Matrix defined in Equation 1.1, one might construct a Monte-Carlo computer model, and the output of this model can also be described as a probabilistic risk analysis.

Results for a Nuclear Weapons Detection: For the magnitude detection problem described in Section 2.1, the following conclusions follow from the game-theory solution:

- A maximum detection threshold exists for a magnitude detector! If one is thresholding the raw detector signal, S, one should never choose a threshold above X_{shield} to achieve the desired false call rate. If the false call rate is not low enough with the use of X_{shield} as the threshold, achieve the desired false call rate through randomization.

- The LFD that produces a signal according to Equation 4.1 is a combination of two different distributions. The first distribution in the combination is the distribution that makes the bomb-distribution most closely resemble the clean distribution. The second distribution is a mixture of the two most extreme configurations available to the adversary.

- The best detection algorithm for a magnitude detector will use randomization to achieve low false call rates. When the L_{miss} cost is large, the detection algorithm will be equivalent to simple thresholding.

Dirty Bomb Size: There is a "best-size" dirty bomb for an adversary to use.

7.0 References

Avenhaus R. 1986. *Safeguards Systems Analysis with application to nuclear material safeguards and other inspection problems*. Plenum Press.

Bowen W. 1988. *Statistical Methods for Nuclear Material Management*. TK9152 .S72, Supt. of Docs., U.S. G.P.O.

Ferguson T. 1967. *Mathematical Statistics: A Decision Theoretic Approach*. Series in Probability and Mathematical Statistics. Academic Press.

McKinsey J. 2003. *Introduction to the Theory of Games*. Dover.

Morrow J. 1994. *Game Theory for Political Scientists*. Princeton University Press.

Myerson R. 1991. *Game Theory Analysis of Conflict*. Harvard University Press.

Osborne M and A Rubinstein. 1994. *A Course in Game Theory*. MIT Press.

Raiffa H and R Luce. 1957. *Games and Decisions*. Wiley.

Runkle R, K Anderson, and etal. 2005. *Advanced Decision Making for Low-count Spectra: Initial Algorithm Development*. PNNL-15442, PNNL.

Siciliano. 2004. *Comparison of PVT and NaI(TI) Scintillators for Vehicle Portal Monitor Applications*. PNNL-14487, PNNL.

Smith E and etal. 2004. *A Sensitivity Comparison of NaI and PVT Portal Monitors at a Land-Border Port of Entry*. PNNL-14974, PNNL.

Smith E and etal. 2005. *A Comparison of NaI and PVT Portal Monitors at a Land-Border Port of Entry, Part II: Representative Threats*. PNNL-15260, PNNL.

Zagare F and D Kilgour. 2000. *Perfect Deterrence*. Studies in International Relations:72. Cambridge Press.

Example Distribution List

No. of Copies

OFFSITE

Person∗∗
Organization
Address
City, State Zip

2 Person∗∗
Organization
Address
City, State Zip

3 Person∗∗
Organization
Address
City, State Zip
ATTN: R. Jim
W. Riggsbee

No. of Copies

ONSITE

DOE Office of River Protection

Person 1∗
Person 2∗
Person 3∗

5 DOE Richland Operations Office

Person 1∗
Person 2∗
Person 2 (5) A6-38∗∗∗

Distribution method:
∗ Email notification of availability at ERICA
∗∗ CD
∗∗∗ Hardcopy