EVALUATING LATE DETECTION CAPABILITY AGAINST DIVERSE INSIDER ADVERSARIES

UCRL--97740

DE88 008678

ALAN SICHERMAN, Lawrence Livermore National Laboratory*
Livermore, California 94550
415/423-8854

ABSTRACT

This paper describes a model for evaluating the late (after-the-fact) detection capability of material control and accountability (MC&A) systems against insider theft or diversion of special nuclear material. Potential insider cover-up strategies to defeat activities providing detection (e.g., inventories) are addressed by the model in a tractable manner. For each potential adversary and detection activity, two probabilities are assessed and used to fit the model. The model then computes the probability of detection for activities occurring periodically over time. The model provides insight into MC&A effectiveness and helps identify areas for safeguards improvement.

INTRODUCTION

The threat of theft or diversion of special nuclear material (SNM) by insiders is a key concern for safeguards planners. Different types of employees having varying degrees of access to both SNM and safeguards systems pose a difficult challenge for theft detection. Safeguards planners rely on physical security, material control, and material accountability to provide detection of a theft attempt. When detection occurs too late to prevent a theft, it is called a late or after-the-fact detection.

After-the-fact indication that material may be missing is usually provided by a material control and accountability (MC&A) system. MC&A activities include maintaining records for tracking nuclear material, verifying that all material is in its authorized location and conducting periodic inventories and audits to establish material balance. Inventory differences exceeding acceptable limits or material not in its authorized location when requested for processing are examples of MC&A system alarms. Late detection includes both an alarm and resolution of its cause, e.g., an error, falsification of records, or an actual theft or diversion.

Late detection is beneficial if it is timely enough to: 1) improve the ability to determine the cause of an alarm, 2) prevent an incorrect response to a threat demand (e.g., a hoax), 3) speed recovery of SNM, or 4) promote assurance that no theft has occurred in the absence of an alarm. We have developed a model for quantifying late detection capability. The model computes the probability of after-the-fact detection as a function of time elapsed since a theft or diversion. Effective MC&A should provide a high probability of material loss detection within a short time.

If a theft has occurred, the probability that the MC&A system will provide an indication of theft depends on many factors. These factors include MC&A measurement uncertainties, process holdup, and tampering or falsification (e.g., theft cover-up) by an adversary. Many detection events are repeated at certain time intervals, e.g., daily, weekly, or monthly inventories. An inventory taken shortly after a theft may be more likely to trigger an alarm than subsequent inventories. These factors complicate the task of estimating the probability of late detection versus time. Our simple model to evaluate late detection capability addresses these and other complications in a tractable manner.

Overview of Model for Late Detection Capability

We begin modeling MC&A late detection capability by identifying events that may provide an indication of material loss. Examples of such events are a daily check, physical inventory and DOE audit occurring daily, monthly, and yearly, respectively. Next, we characterize potential insider adversaries according to their MC&A access and authority (e.g., material custodian, manager, operator, guard). Such access may be used by insiders to hamper discovery of a theft. For each adversary and detection event two probabilities are assessed: the probability of detection for the first occurrence of the event and, if applicable, that of the second occurrence if no alarm resulted from the first occurrence. These

*Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

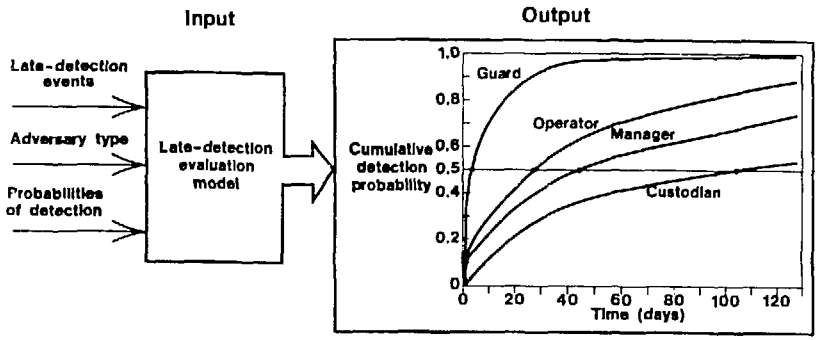

Fig. 1. Schematic of the inputs and outputs of the Weibull late detection model

probabilities must take into account potential falsification or tampering strategies. As described in the paper, these two numbers are sufficient to fit a Weibull[1] probability distribution to estimate the probability of detection on subsequent occurrences of the event. We will discuss the reason behind the choice of the Weibull distribution.

By assuming probabilistic independence among the various detection events, we can derive the cumulative detection probability as a function of time based on all MC&A late detection events. Figure 1 illustrates the input and output for the Weibull model. This model is implemented as part of a PC-based computer program called ET.[2] Figure 1 illustrates the use of the model for evaluating a hypothetical facility. The hypothetical results shown in the figure demonstrate the possible variation in safeguards performance among adversaries.

Despite its simplifications, the Weibull model provides useful insight about how the timing and efficacy of late detection events affect the overall probability of late detection as a function of time. Potential upgrades such as a computerized accounting system, additional independent inventories, rotating of inventory personnel, or more frequent inventories can be analyzed by this model for their impact on the late detection effectiveness against a spectrum of insider adversaries.

COMPUTING LATE DETECTION PROBABILITY AS A FUNCTION OF TIME

In this section we describe the series of steps followed in our modeling and evaluation of late detection capability. For each step we describe assessments and assumptions needed to address complicating factors encountered.

Identifying Late Detection Events

To evaluate the probability of late detection as a function of time, we focus on MC&A activities or events that could reveal loss of SNM. Many of the late detection events are periodic. It is possible that if the first occurrence of a late detection event does not indicate missing SNM, a subsequent occurrence of the event will. Table I illustrates this first step of defining late detection events and their periods of recurrence if appropriate.

TABLE I
Illustrative Identification of Late Detection Events

Late Detection Events	Time Lag (Units = Days)	Repeated
Daily check	1	Yes
Forms reconciliation	3	No
Process call	15	Yes
Physical inventory	30	Yes
DOE audit	365	Yes

Assessing Probabilities of Detection

After an occurrence of a detection event, there is a probability of a loss indication and proper alarm resolution resulting from that event, given material is missing. We can estimate such a probability based on knowledge of measurement uncertainties, statistical methods (e.g., if inventories or checks are based on random sampling techniques) or error propagation models. When adversary tampering or falsifi-

cation is a possibility, subjective judgment by experts at the facility may be required to provide an estimate of the probability. For example, an inventory may detect a loss with a certain probability absent any tampering. But if a particular adversary participated in the inventory taking, falsification of data could significantly degrade the detection probability. Ideally, accountability exercises or tests would be conducted to "validate" subjective probability asssessments. But such judgments are required to address the insider threat and are strongly related to each adversary's knowledge, access, and authority. Thus we need to review each insider's potential strategies for delaying or defeating late detection from an event and incorporate this knowledge into the assessment of the detection probability.

As mentioned previously, detection events are often repeated at periodic intervals. The probability of detection for the second occurrence given no detection on the first occurrence of an event will, in general, be different than that for the first occurrence of the event. We can imagine that if an adversary's falsification strategy for an event worked the first time, it is more likely to work again the second and subsequent times (e.g., falsifying an inventory). It is also possible to imagine a falsification strategy that is more likely to fail as it is repeated. Of course, the probability could remain the same for subsequent events as well. To address repeat occurrences of events, we must at least assess the probability of detection for the second occurrence of the event given no detection on the first occurrence. This second probability assessment gives some indication of the trend of effectiveness (e.g., decreasing) for subsequent events.

Table II shows some illustrative probability assessments for different adversaries and detection events. In this illustration, detection probabilities are different for each adversary because of special access or strategies available to each adversary. For example, the probability of the daily check detecting a material theft by a custodian is effectively zero if the custodian, as assumed here, is the sole individual performing the daily check. Table II illustrates events all of which repeat. To address a nonrepeating event, the probability of detection for the second occurrence can be set to zero. A nonrepeating event is just like a repeating event with no chance of detection except on the first occurrence.

TABLE II
Illustrative Assessed Probabilities of Detection for the First and Second* Occurrences of Detection Events

	Detection Events							
	Daily Check 1 day		Process Call 15 days		Physical Inventory 30 days		DOE Audit 365 days	
Adversary	1st	2nd	1st	2nd	1st	2nd	1st	2nd
Operator	1	.01	1	1	.3	.2	.2	0
Custodian	0.0	0.0	.01	.01	.3	.1	.2	0
Manager	1	.01	.01	.01	.3	.2	.2	0
Guard	.3	.1	.1	.1	.8	.3	.2	0

*The probability of detection for the second occurrence given no detection from the first occurrence.

Computing Late Detection Probability as a Function of Time for Each Event

From the previous discussion, we need at least two probabilities to characterize the trend for late detection for each adversary and detection event. We need to estimate MC&A effectiveness, however, for other times--not simply for the first and second occurrences of an event. Assessing additional input probabilities for this purpose is not desirable. Such assessments are impractical to make in a consistent fashion. We require a model to extrapolate in a reasonable manner for subsequent occurrences of a detection event. The model must be tractable enough to calibrate using two input probabilities.

In addition to such extrapolation, the model must also address another issue--that of the possible timing of malevolent acts (such as theft or diversion) in relation to the MC&A detection events. Theft attempts could occur in random relation to detection event schedules. Even if we assumed conservatively that a theft attempt occurred to maximize the length of time before one type of detection event took place, it is difficult to assume rigid timing with respect to all event types; e.g., theft always occurs exactly 365 days before the DOE audit. Avoiding late detection is but one concern of an adversary and avoiding detection during the theft act may dictate when the theft attempt occurs. Also, the timing of detection events can vary. In fact, for safeguards reasons, it is not unusual for certain events like inventories to be scheduled somewhat randomly. Thus, events like the process call and inventory occur on average, say, at 15- and 30-day intervals respectively, rather than exactly at those intervals. Because of these issues, we must consider how smoothly late detection probabilities should vary as time elapses.

To address these concerns, we selected a "Weibull distribution" calibrated using the probability assessments for the first two occurrences of an event in order to model late detection capability. The equation for the Weibull-based model is shown in Table III. The proper-

TABLE III
Definition of the Weibull Late Detection Model

$$PLD_i(t) = \left(t/T_i\right) P_{i1} \quad \text{for } 0 \le t < T_i$$

$$= 1 - \text{Exp}\left[-\left(B_i \quad t/T_i\right)^{C_i}\right] \quad \text{for } t \ge T_i$$

with

$$C_i = \text{Ln}\{\text{Ln}[(1 - P_{i1})(1 \quad P_{i2})]/\text{Ln}(1 \quad P_{i1})\}/\text{Ln}2$$

and

$$B_i = [\ \text{Ln}(1 \quad P_{i1})]^{1/C_i}$$

where

$PLD_i(t)$ = probability of late detection from event i as a function of time

T_i = average period of recurrence for detection event i

P_{i1} = probability of detection for the first occurrence of event i

P_{i2} = probability of detection for the second occurrence of event i, given no detection on the first occurrence.*

*Though technically P_{i1} and P_{i2} must lie between 0 and 1, when made very close to either boundary they produce the effect expected at the boundary

ties of the Weibull model are discussed in the following paragraphs.

Eq. (1) in Table III may seem complex but the important features of the model can be described in simpler terms. The model is Weibull-based because of the functional form for $PLD_i(t)$ for $t \ J \ T_i$, which is that of a Weibull cumulative distribution function. The Weibull distribution is used in a variety of applications where the rate of "failure" (e.g., failure to detect) is not constant but changes over time. The Weibull model offers the flexibility of several key extrapolations given basic input on the probability of detection for the first and second occurrences of the event.

From time zero until the first occurrence, the model in Eq. (1) linearly interpolates between probability zero and P_{i1}. This reflects the fact that at time T_i, we are sure exactly one detection event of this type occurred and thus the probability of detection is P_{i1} for that event. Before time T_i, it is still possible that the event occurred because of the relative timing of the theft vis-a-vis the detection event. For example, a theft may have occurred just a few days before a monthly inventory. The linear interpolation is a way of reflecting this possibility. Thus an event with period T_i can still be given some credit for earlier than time T_i detection. This type of extrapolation seems reasonable and also provides for a more smoothly varying probability of late detection as a function of time. (As a technical aside, the pure Weibull formula does not always extrapolate backwards to time 0 in a sensible manner for safeguards analysis; e.g., it indicates that a system with $P_{i1} = .5$, $P_{i2} = 0$ is better before time T_i than one with $P_{i1} = .5$, $P_{i2} = .5$, when there should be no difference before time T_i; hence the need for something like the linear interpolation.)

From time T_i onward, the Weibull model follows the trend established by the inputs for P_{i1} and P_{i2}. Between times T_i and $2T_i$, Eq. (1) reflects the possibility that due to the relative timing of the theft and detection event, more than one detection event may have occurred since the theft. (If in Eq. (1) we ignored any fractional part of the term (t/T_i) wherever it appears, we would have a stair-step function instead of a smooth function. Such a stair-step function is conservative and assumes rigid timing between theft and detection events.) To summarize, the Weibull cumulative detection probability increases smoothly as a function of time in a way that reasonably extrapolates the trend indicated by the user inputs P_{i1} and P_{i2}. We now explore the kinds of trends possible for repeated events and compare the Weibull to other schemes.

Modeling of Detection Trends for Repeated Occurrences of an Event

If P_{i1} and P_{i2} are both equal, say 0.5, and T_i is 1 day, then using Eq. (1), the $PLD_i(t)$ for times 1, 2, 3, and 4 are respectively, 0.5, 0.75, 0.875, and 0.9375. In essence, Eq. (1) reduces to $PLD_i(t) = 1 - (1 - P_{i1})^{t/T_i}$ for $t \ J \ T_i$. We can also arrive at this result using the following reasoning. The probability of detection after four occurrences of the event is one minus the probability that no detection resulted from any of the four occurrences, or $1 - (1 - 0.5)^4$. The Weibull model gives the results we would expect for these times and is hardly mysterious for this case. However, we need something else when P_{i1} and P_{i2} are different. Table IV shows three possibilities for how we might extrapolate given P_{i1} and P_{i2}.

In Table IV, trend (a) is an optimistic extrapolation of the assessed inputs. The probability of detection provided by a subsequent event occurrence given no detection on previous occurrences is no worse than that for the second event occurrence. For $T_i = 1$ day, the optimistic extrapolation implies a "certainty" of detection after 60 days (e.g., using the trend (a) equation in Table IV). In essence, even though each occurrence after the first has only 0.1 chance of detection, defeating 59 of them is just about impossible. When encountered by practitioners in the field, this behavior was judged overly optimistic. If detection is degrading markedly from the first

TABLE IV
Possible Trends for
Late Detection Probabilities

Illustrative assessed inputs:

	$P_{i1} = 0.5$, $P_{i2} = 0.1$		Extrapolated Probabilities	
Trend	P_{i1}	P_{i2}	P_{i3}	P_{i4}
a.	[a]0.5	0.1	0.1	0.1
b.	[b]0.5	0.1	0.07	0.05
c.	[c]0.5	0.1	0.0	0.0

[a] $PLD_i(t) = 1 - (1 - P_{i1})(1 - P_{i2})^{(t-T_i)/T_i}, t \geq 2T_i$.
[b] Corresponds to Weibull extrapolation.
[c] $PLD_i(t) = 1 - (1 - P_{i1})(1 - P_{i2}), t \geq 2T_i$.

to second occurrence, further degradation is also expected on subsequent occurrences. Trend (c) is the pessimistic extrapolation. The downward trend between the first two occurrences is sharply extrapolated to yield a probability of zero in subsequent occurrences beyond the second. Such an extrapolation essentially assumes the event has no detection value beyond the second occurrence. (This pessimistic model does not provide a normalized cumulative probability distribution that goes to 1 as time goes to infinity.) Trend (b) is the Weibull model extrapolation. The Weibull continues the trend of decreasing detection probability for a subsequent event given no detection on previous events. But the trend is not exaggerated and the Weibull provides a tractable cumulative probability distribution eventually going to 1 as time gets large. It represents a "happy medium" between the two extremes shown in Table IV.

Computing the Probability of Late Detection as a Function of Time for All MC&A Events

By assuming probabilistic independence among the various types of detection events, we can derive the cumulative detection probability as a function of time based on all MC&A late detection events. The resulting equation is

$$PLD(t) = 1 - \prod_i [1 - PLD_i(t)] \qquad (2)$$

where

PLD(t) = probability of late detection from all events as a function of time (t)
$\prod_i$ = product of i terms
$PLD_i(t)$ is as defined in Eq. (1), Table III for each detection event.

The reasoning behind Eq. (2) is the same as that encountered earlier. The probability of detection at any time t is just 1 minus the probability that no event has resulted in a detection by time t. Eq. (2), along with the input data in Table II, was used to generate the illustrative curves shown in Fig. 1 at the beginning of this paper.

USE OF MODEL IN ANALYZING LATE DETECTION CAPABILITY

The model described here is relatively easy to use while still addressing major concerns such as adversary falsification strategies and reasonable extrapolation of inputs. The model can be used to generate curves such as that shown in Fig. 1, or summary statistics such as the median (highlighted by the horizontal line in Fig. 1), or mean time for late detection for each adversary.

Because many inputs rely on subjective judgments of experts at a facility, the tractability of the model is important. The model's ease of use provides a simple way to test the sensitivity of results to alternative input assumptions and probability assessments. This sensitivity analysis can highlight those inputs requiring investigation and data collection and focus debate on the most important judgments.

The model helps identify those adversaries for which the late detection capability is weak and helps to pinpoint areas where corrective action or upgrades are required. The model can help one assess the change in late detection effectiveness prior to implementing upgrades so that the benefits of various alternatives can be highlighted.

CONCLUSION AND FUTURE MODELING DIRECTIONS

The model described in this paper characterizes late detection capability by the probability of late detection as a function of time. While an important indicator, this cumulative detection probability does not directly address key issues associated with quantifying the value of MC&A late detection. Specifically, each benefit provided by a good late detection capability described in the Introduction may place different emphasis on timeliness. The value of time is only considered informally when examining a cumulative detection probability. To address the value of time formally with respect to alarm resolution, alarm response, SNM recovery (or prevention of protracted theft), and assurance, value models [3] are needed. Such models quantify the relative importance of improved after-the-fact detection (e.g., how much better is late detection in 2 days rather than 3 days). However, these value models (like our probability model) require subjective input by decision makers at facilities and therefore must be tractable to allow for easy sensitivity

analysis. Value models combined with probability models, such as the one described here, will allow facility decision makers to characterize late detection capability more completely for purposes of resource allocation and comparing safeguards alternatives.

REFERENCES

1. R. E. WALPOLE, and R. H. MYERS, Probability and Statistics for Engineers and Scientists, MacMillan Publishing Company, Inc., New York (1978), pp. 133-136.

2. R. A. AL-AYAT, B. R. JUDD, and T. A. RENIS, "The Safeguards Method for Evaluating Vulnerability to Insider Threats," Proc. 27th Annual Meeting of Institute of Nuclear Materials Management, New Orleans, June 25, 1986, pp. 676-680.

3. R. L. KEENEY, and H. RAIFFA, Decisions with Multiple Objectives, John Wiley & Sons, Inc., New York (1976).