Australian Government

Australian Safeguards and Non-Proliferation Office

Australian Safeguards Support Program

Randomized inspections at UDU storage facilities under Integrated Safeguards, Mean Time to Detection

Task AUL C 01208: **Re-Examination of Basic Safeguards Implementation Parameters**

Russell LESLIE, Annette BERRIMAN and John CARLSON

Australian Safeguards and Non-Proliferation Office (ASNO)
RG Casey Bldg, John McEwen Crescent, Barton ACT 0221 Australia

June 2005

AUL Report 2005-02

Page intentionally left blank

SYNOPSIS

Unirradiated Direct Use (UDU) material (Pu, including MOX, and HEU) is the most proliferation-sensitive nuclear material. Safeguards measures must take into account that access to significant quantities of UDU would allow a potential proliferator the quickest route to developing nuclear explosive devices.

In the move towards integrated safeguards (IS) it will be necessary to develop IS inspection schemes for facilities that contain UDU. The report will examine the possibility of IS inspection schemes that ensure necessary effectiveness while taking full advantage of the flexibility allowed under IS including increased reliance on randomised inspections (whether unannounced or allowing for shortened notice periods).

Issues addressed by the report include:

- Average inspection frequency;
- Calculation of mean time to detection under a variety of assumptions;
- The possibility of determining optimal inspection strategies

EXECUTIVE SUMMARY

The IAEA is faced with the task of finding the optimum combination of available measures in formulating its safeguards approach for any given state under Integrated Safeguards (IS). In making the necessary judgements as to what constitutes the optimum combination of measures, the IAEA has to work within certain constraints (including the need to be effective, efficient and non-discriminatory). In this context a method is needed to assist in evaluating which approach is best suited to the IAEA's needs (and which best fits within the above listed constraints). One possible method is to determine a "mean time to detection" (MTTD) for each approach. Another possible method is to calculate the probability of detecting a diversion at periods shorter than the traditional timeliness goal. Both methods can be useful guides to decision making by the IAEA.

It is important to emphasise that both methodologies have a limited range of applicability—they are useful for comparing different types of inspection schemes (e.g. fixed vs. SNRI, X inspections per year vs. Y inspections per year) but they would not be useful for comparing schemes based on widely different underlying approaches (e.g. inspection schemes vs. information analysis).

Current inspection approaches are mostly based on the use of inspections fixed in time—there has been only limited use of randomised inspection schemes for specific facility types. As a basis for comparison between traditional safeguards approaches and possible new IS safeguards approaches, it should be noted that traditional safeguards inspection schemes for facilities with UDU will have an expected mean time to detection[1] of 61 days—this figure is expected to have a standard deviation of 30 days. The probability of detecting a diversion within 61 days is approximately 78%. When examining the various inspection schemes detailed in the report it is important to keep in mind that it is not a comparison between a perfect current system with a less perfect future—no inspection system is perfect.

1 The meaning of the term "mean time to detection" is defined and discussed in full in this paper.

The MTTD methodology can be used to prove that it would be possible to maintain the same value for mean time to detection for different inspection schemes by conducting fewer, more stringent inspections per year (e.g. by changing the design detection probability for interim inspections from the current 50% to 75% in the case of an 8-inspection per year scheme).

The report recommends that the IAEA conduct a short notice random inspection (SNRI) trial for a UDU facility using the methodology developed in this report. The trial would make use of 1 fixed PIV inspection supplemented by an average of 5 SNRIs per year—the daily probability of inspection value is (approx.) 1.5%. Seven detailed recommendations on the performance of the proposed trial are contained in the recommendations section of the report.

Two results arising from the report are that:

Result 1: It is possible to achieve equivalent mean time to detection values for different numbers of randomised inspections per inspection year by varying the target values for the probability of detection for each inspection.

Result 2: The value of the mean time between inspections/detections parameter is not significantly changed by the imposition of a lower boundary condition provided an appropriate adjustment is made to the daily probability of inspection.

The overall conclusions of the report are that:

Conclusion 1: The use of SNRIs provides the possibility of detecting diversion of UDU at intervals shorter than the timeliness goal for the material in question.

Conclusion 2: It is possible to develop specific inspection schemes which utilise a mixture of fixed and SNRIs as a means of achieving specific "mean time to detection" targets. The numeric differences between a statistically ideal situation and the real-life situation are not significant in most practical cases.

Conclusion 3: When combined with policy decisions as to acceptable levels of the probability of achieving specific timeliness goals, mean time to detection calculations provide a useful guide to inspection decision making.

Conclusion 4: For the purposes of this discussion the exponential distribution is a useful analogue for the binomial distribution. The differences between results calculated via the two methods are not significantly different from each other for inspection opportunity groupings of 14 days or less.

1.1—PREAMBLE

The level of assurance that can be derived from a given level of inspection effort has been the subject of intensive academic study since the early seventies. The formulae used in the stratification, sampling and verification of nuclear material subject to safeguards have been the subject of examination and reconsideration since they were first derived. The approach taken by the Agency has gained wide acceptance.

The report assumes the basic theoretical underpinnings of the existing safeguards system. The aim is to show that the situations under consideration are entirely analogous to problems that have been solved in the literature and whose solutions have gained acceptance within the international community. The paper does not return to first principles in its argumentation or in the derivation of formulae. Instead use is made of work that has been done by a number of authors and this body of accepted work serves as the foundation for the conclusions reached in the paper.

The report is closely related to AUL Report 2005-01 "Minimum inspection frequencies under Integrated Safeguards"—it extends and expands upon the methodology contained in that report.

1.2—INTRODUCTION

With the establishment of Integrated Safeguards (IS) the IAEA is faced with the task of finding the optimum combination of available measures in formulating its safeguards approach for any given state. In its choice of methods the IAEA is bound by a series of important constraints, these include (*inter alia*):

- the need to be effective;
- the need to be efficient;
- the need to work within available resource constraints (including both human and financial resources);
- the need to be non-discriminatory;
- the need to maintain the confidence of the international community in the effectiveness of the safeguards system.

One proposal for the development of IS safeguards approaches to increase the use of Short Notice Random Inspection (SNRI). These are inspections that are timed in a way that is unpredictable to both the facility operator and the state. The use and potential value of SNRI will be examined in context of potential safeguards approaches.

In choosing any particular safeguards approach over others that are available it is helpful to be able to develop a method to assist in the evaluating which approach is best suited to the IAEA's needs (and which best fits within the above listed constraints). One useful methodology is to determine a "mean time to detection" (MTTD) for each approach and use the calculated figure as a guide to decision making. It is important to emphasise that such a methodology has a limited range of applicability—it is useful for comparing different types of inspection schemes (e.g. fixed vs. SNRI, X inspections per year vs. Y inspections per year) but it would not be useful for comparing schemes based on widely different underlying approaches (e.g. inspection schemes vs. information analysis).

It is also important to note that an inspection scheme that provides for a smaller number of inspections per facility per year for the same MTTD value may result in the expenditure of a greater amount of resources overall, because of the required increase in the value of probability of detection (p) and consequent decrease in the value of the non-detection probability (ß=1-p). If an inspection scheme requires that all items be verified (rather than a subset sampled) it may prove to be impractical in the field. Experience has shown that it is usually prohibitively costly to achieve p>90%.

There are many other possible methods of comparison between different inspection schemes that can be considered to assist in evaluating competing potential inspection approaches. After discussing the MTTD methodology the report will also explore issues relating to the proportion of inspections taking place at equal to or shorter than the stated timeliness goal and/or inspections and detections taking place at times considerably shorter than the stated timeliness goal.

Unirradiated Direct Use (UDU) material (Pu, including MOX, and HEU) is the most proliferation-sensitive nuclear material. Safeguards measures must take into account that access to significant quantities of UDU would allow a potential proliferator the quickest route to developing nuclear explosive devices. The report will look at the issue of the adequacy of SNRI schemes for addressing the most sensitive material subject to safeguards. The methodology used will be via the use of calculated MTTD and via the probability of achieving inspection intervals and detection intervals shorter than traditional timeliness goals.

Current inspection approaches are mostly based on the use of inspections fixed in time—there has been only limited use of randomised inspection schemes for specific facility types. As a basis for comparison for all of the discussion that follows it should be noted that traditional safeguards inspection schemes for facilities with UDU will have an expected mean time to detection of 61 days—this figure is expected to have a standard deviation of 30 days. The probability of detecting a diversion within 61 days is equal to (roughly) 78%[2]. When examining the various inspection schemes detailed in the report it is important to keep in mind the point that it is not a comparison between a perfect current system with a less perfect future—no inspection system is perfect.

1.3.1—IMPORTANT NOTES

> Under traditional safeguards approaches to UDU facilities the predicted mean time to the detection of a diversion is 61 days. This figure is an expected mean—both higher and lower values are expected to arise in real inspection situations.
>
> Due to the statistical nature of the process of detecting diversion, there is a significant probability of any diversion being undetected for periods considerably longer than 61 days under traditional safeguards approaches.

2 Assuming: a timeliness goal of 30 days; fixed inspection interval averaging to 30.4 days; a PIV detection probability of 90%; and a interim inspection detection probability of 50%.

1.3.2—KEY ASSUMPTION

Key Assumption 1 (KA1): It is assumed throughout this discussion that diversion is not a random act, the diverter is prepared to take all necessary steps to prevent detection (in effect to maximise the time to detection). One key element of any strategy to maximise the time to detection is to divert material immediately after the completion of an inspection.

Key Assumption 1 (KA1) holds throughout the discussion that follows.

1.3.3—SIMPLIFYING ASSUMPTIONS

In order to simplify some of the calculations that are used in the report a series of assumptions are made that have the net result of making the calculations less mathematically complex. The effect of varying these assumptions is discussed, as appropriate, throughout the paper. The assumptions are as follows:

Simplifying Assumption 1 (SA1): Assume all inspections are perfect (if p is the probability of detection and ß is the non-detection probability then p=1 and ß =0). Under SA1, if there is a diversion it will be detected at the next inspection[3].

Simplifying Assumption 2 (SA2): Assume that the same p and the same ß apply to all inspections.

Simplifying Assumption 3 (SA3): There is an average of 365.25 inspection opportunities per year.

Simplifying Assumption 4 (SA4): The timeliness goal is 30 days.

Simplifying Assumption 5 (SA5): There is no diversion scenario for the facility in question that can be concealed in less than 24 hours.

Simplifying Assumption 6 (SA6): The PIV takes place at the first inspection opportunity of every inspection year.

1.3.4—CAVEAT

Caveat: The statistical equations used in the report are only strictly valid when all of the inspections taking place at the facility, including the PIV, are randomly scheduled. Fixing the PIV in time as per SA6 (above), introduces very minor differences to the various percentage values calculated for all inspection schemes.

The scale of these numeric differences is discussed in the Annex to this paper.

3 The values of p and ß are crucial parameters used by the IAEA to determine, *inter alia*, the number of samples and the level of effort expended during an inspection. SA1 assumes p=1 and ß=0.

2.1—Calculating the Mean Time to Detection (MTTD)

The calculation of a "mean time between inspections" parameter μ_i for either fixed or randomised inspections is trivial. It is dependent only upon the length of the inspection year[4] and the target number of inspections. If there is an average of only one inspection per facility per year, then the long term average interval between inspections at each facility must tend towards 365.25 days. Shorter or longer average inspection intervals per facility may be observed over relatively small collections of inspection years (1 to 100 inspection years), but the long term average must tend towards 365.25.

Demonstrating the logic behind this simple result is straightforward. Assuming, for example, that the long term average inspection interval per facility is any number smaller than 365.25 (for example 300 days) then over 1000 inspection years[5] there will be approximately 1220 inspections. In this example, the shorter average inspection interval would result in 22% more inspections than required over the period. The average would not be 1 inspection per facility per year but instead would be 1.22 inspections per facility per year. Similarly if the average inspection interval is any number larger than 365.25 (for example 460 days) then over one thousand inspection years there will be approximately 800 inspections The longer average inspection interval would result in 20% fewer inspections than required over the period.

The calculation of an average inspection interval is not a matter that requires recourse to complex mathematical algorithms, it is a matter of simple arithmetic. Any methodology that that does not result in a long term average interval between inspections for a 1 inspection per facility per year case of 365.25 days is not appropriate to the case under consideration. The same logic applies in cases in which the target number of inspections per facility per year is greater than one.

$$\mu_i = \psi / \acute{\iota}$$

$$= \textbf{mean time between inspections}$$

where

ψ = the number days in the inspection year

and

$\acute{\iota}$ = the target number of inspections per facility per year.

Once a value for the mean time between inspections, μ_i, has been determined, there are several different methodologies that can be used to calculate a mean time to detection (MTTD) parameter μ_d.

Under SA1 (i.e. perfect inspections) in combination with KA1[6] it is clear that the mean time to detection is exactly the same as mean time between inspections (i.e. $\mu_i = \mu_d$). Restating this point for clarity—assuming perfect inspections leads logically to the point that any diversion will be detected at the next inspection. Under KA1 it is assumed that the diverter will seek to

4 The term "inspection year" is used to mean one year of Agency inspection effort at one facility with the year starting at the facility Physical Inventory Verification (PIV).

5 1000 inspection years can occur as any combination of numbers of facilities and inspection years that sum to 1000 inspection years (for example, 5 inspection years at each of 200 facilities, 10 inspection years at each of 100 facilities or 1000 inspection years at one facility).

6 The key assumption of this paper (KA1) is that diversion is not a random act, it is a deliberate attempt by a diverter to defeat the safeguards system and that the diverter will seek to maximise the time to detection.

maximise μ_d hence it is conservatively assumed that the diverter will divert immediately after an inspection and $\mu_i = \mu_d$.

Varying SA1 (i.e. more realistic, imperfect inspections) represents the more reasonable case in which $p<1$ and $ß>0$. In that case the calculation of μ_d has to be re-examined ($\mu_i \neq \mu_d$). If $p<1$ then μ_d is no longer simply equal to ψ / í. The new formula becomes:

$$\mu_d = (\psi / í) / p$$

If p=0.5 with 12 inspections per facility per year out of 365.25 inspection opportunities then

$$\begin{aligned}\mu_d &= (\psi / í) / p \\ &= (365.25/12)/0.5 \\ &= 60.9 \text{ (cf 30.4 for p=1).}\end{aligned}$$

There are many methods that can be used to determine the distribution of inspections and detection intervals across the inspection year. One of the simplest, mathematically, is by the use of the exponential cumulative probability distribution and the closely related exponential distribution[7].

The exponential cumulative probability distribution equation can be stated as:

$$F(x) = 1- e^{(-x/\mu)} \quad \text{for } 0 \leq x < \infty, \mu > 0$$

The exponential distribution function can be stated as:

$$f(x) = (1/\mu)*e^{-(1/\mu)*x} \quad \text{for } 0 \leq x < \infty, \mu > 0$$

If the population mean[8] (μ) of the times between events is known, it can be used to define the exponential function parameter λ, in general the value of λ represents the probability per time period of the event of interest happening, λ is crucial in doing calculations with the exponential distribution. Once again we will use two parameters, λ_i is the daily probability of an inspection happening and λ_d is the daily probability of a detection happening (the numeric relationship between λ_i and λ_d involves the average value of the probability of detection).

$$\begin{aligned}\lambda_i &= 1/\mu \\ &= í / \psi\end{aligned}$$

Having introduced λ, the basic equations can be restated in terms of the new variable.

---

7 The use of the binomial distribution for this purpose is discussed in Section 2.4.

8 It is important to note that there are two separate μ variables that will be used in this discussion -
μ_i = the "mean time between inspections" and
μ_d = the "mean time between detections"
If the generic term μ is used, the details given apply equally to μ_i and μ_d. Similarly if the generic term λ is used the details apply equally to λ_i and λ_d.

$$\mathbf{F(x) \quad = \quad 1- e^{(-x*\lambda)} \qquad for\ 0 \le x < \infty, \lambda > 0}$$

and

$$\mathbf{f(x) \quad = \quad \lambda*e^{-\lambda*x} \qquad for\ 0 \le x < \infty, \lambda > 0}$$

Using SA1 (perfect inspections) the mean time to detection for a target of an average of 12 inspections per facility per year is 365.25/12 = 30.4[9] days, then

$$\mathbf{\lambda_i \quad = \quad (12/365.25) = 0.03285}$$

If the value of λ is known there are several other factors that can be calculated in straightforward fashion:

first quartile	=	**$\ln(^4/_3)/\lambda$**
median	=	**$\ln(2)/\lambda$**
third quartile	=	**$\ln(4)/\lambda$**
variance (σ^2)	=	**$1/\lambda^2$**
standard deviation (σ)	=	**$1/\lambda$**
	=	**μ**

and if we assume that p=0.5

$$\mathbf{\lambda_d \quad = \quad 0.06571} \text{ (cf 0.03285 for p=1)}$$

Varying SA2 (assuming that p and ß are not constant across all inspections) the calculations take a further step closer to realistic inspection schemes—allowing for the usual IAEA practice of using one set of quite restrictive values for p and ß for PIV inspections (p_{piv}) and a less restrictive set of values for p and ß in the case of interim inspections for timeliness purposes ($p_{interim}$). Assuming that there is never more than one PIV inspection per facility per year, the average detection probability can be approximated[10] over all inspection years as:

$$\mathbf{p_{av} \quad = \quad (p_{piv}+(í-1)*p_{Interim})/í}$$

The variable p_{av} is intended to approximate the average detection probability for each inspection across all inspection years—it is **<u>not</u>** intended to represent the cumulative detection probability across the inspection year p_{cum}. The value for p_{cum} is given by the following formula:

---

9 The calculation of μ assumes many years of inspections (the limiting case). The number defined as μ is unlikely to be the observed value for any given calendar year or periods of time that might commonly be considered long in an inspection context (e.g. 10 inspection years or more). It is the limiting value that would be approached by averaging observed values over large numbers of inspection years (e.g. 1000 inspection years or more).

10 The calculation of p_{av} will be used throughout the paper, though it is acknowledged by the authors that it is not a true analytical value, it is a useful working value (a simple approximation) rather than a definitive calculation. The numeric differences between the approximate value and the analytical value are discussed in the Annex to this paper. As noted for the value for μ, the calculated p_{av} may not be the observed value for any given calendar year.

$$p_{cum} = 1 - (1-p_{piv})*(1- p_{Interim})^{(i-1)}$$

It is possible to approximate[11] the value for p_{cum} in terms of p_{av} as follows

$$p_{cum} = 1 - (1-p_{av})^{i}$$

The approximate value for p_{av} can be used to calculate the first directly useful result of this discussion. It is possible to demonstrate that inspection schemes with different numbers of inspections per facility per year can have broadly equivalent MTTD parameters (provided p and ß are varied to accommodate the changing numbers of inspections).

For example—the calculated MTTD parameter for 12 fixed inspections per facility per year (p_{piv}=0.9 and $p_{interim}$=0.5) is approximately 61 days. Table 1 lays out SNRI inspection schemes with MTTD parameters roughly equal to 61 days using varying target numbers of inspections per facility per year and varying values for $p_{interim}$.

Table 1 – Maintaining MTTD by varying p and ß for varying numbers of SNRI

Target number of inspections/year	Avg no. of inspection opp. /year	p_{piv}	$p_{interim}$	p_{av}	MTTD (days)
12	365.25	90%	50%	53%	57
11	365.25	90%	55%	58%	57
10	365.25	90%	60%	63%	58
9	365.25	90%	66%	69%	59
8	365.25	90%	75%	77%	59
7	365.25	90%	88%	88%	59
6[12]	365.25	99%	99%	99%	61

Result 1 (R1): It is possible to achieve equivalent mean time to detection values for different numbers of randomised inspections per inspection year by varying the target values for the probability of detection for each inspection.

For example 12 inspections per facility per year with a PIV detection probability of 90% and an interim inspection detection probability of 50% has a similar mean time to detection as conducting 8 inspections per facility per year with a PIV detection probability of 90% and an interim inspection detection probability of 75%.

11 The value for p_{cum} expressed in terms of p_{av} is a population average across many trials across many inspection years—it does not represent the value that would be observed across randomly chosen inspection years—except for the trivial case of large numbers of exact multiples of one inspection year.

12 Any inspection scheme that calls for a value of p>0.9 and a corresponding value for ß<0.1 is going to prove impractical to implement under most circumstances. The value for 6 inspections is included for completeness only—it does not represent a practical inspection scheme under the set of assumptions used in this example. Under a different set of assumptions a 6-inspection case may be both realistic and acceptable.

Caveat on R1: Increasing the value of $p_{interim}$ as a trade-off may not result in decreased costs overall due to the increased costs associated with performing interim inspections to a higher detection probability (an example follows which illustrates this point).

2.2 – MTTD AT A MODEL FACILITY

Assume the following hypothetical example for a model facility:

- Facility has only one stratum;
- Stratum to be verified contains 1000 batches;
- Each batch consists of one item;
- Each batch consists of PuO_2 powder of broadly equivalent, known isotopics;
- Each batch contain 1 kg (element weight) of Pu;
- Stratum is under a form of containment or surveillance (C/S) that qualifies as acceptable single C/S under IAEA criteria and bias defect measurements are not required;
- An instrument is available which allows each measurement to be made in a way that satisfies IAEA requirements for a partial defect test;
- Each measurement takes 15 minutes;
- Samples can be held ready for immediate measurement (i.e. selection and retrieval times are either zero or can be included in the measurement time for previous samples);
- All measurements are done at the partial defect level, as there is no significant difference in time between the performance of a partial defect test or a gross defect test;
- One inspector performs each inspection;
- 8 hours represents 1 person day of inspection (PDI)—presented in integer multiples (e.g. 7 hours 55 mins represents 1 PDI—8 hours 5 mins represents 2 PDI).

Using the simplified form of the standard sampling formula

$$N_0 = N*(1- \beta^{(x/M)})$$

Where

N = the number of items in the stratum to be verified—which in this case is 1000;

N_0 = the subset of the population that will need to be verified;

β = the non-detection probability for the inspection (1-p);

x = the average amount of material in each item in the stratum, in the example this is 1 kg;

M = the defined significant quantity of the material (e.g. 8 kg for Pu, assuming the entire stratum is available for verification).

Table 2 - Examination of PDI and sample numbers for different inspection schemes

12-Inspection Case			
PIV (1 per year)		Interim Inspections (11 per year)	
p	= 90%	p	= 50%
Samples required	= 126	Samples required	= 42
PDI required	= 4	PDI required	= 2
Totals for year			
PDI required	= 26	Samples required	= 588
8-Inspection Case			
PIV (1 per year)		SNRI (7 per year)	
p	= 90%	p	= 75%
Samples required	= 126	Samples required	= 80
PDI required	= 4	PDI required	= 3
Totals for year			
PDI required	= 25	Samples required	= 686
6-Inspection Case[13]			
PIV (1 per year)		SNRI (5 per year)	
p	= 99%	p	= 99%
Samples required	= 219	Samples required	= 219
PDI required	= 7	PDI required	= 7
Totals for year			
PDI required	= 42	Samples required	= 1314

As can be seen from the figures in Table 2—saving on numbers of inspections in any given year may not result in any savings in overall resource expenditure. Such changes would need to be examined on a case-by-case basis.

In the example the 8-inspection case requires 98 more samples to be measured than the 12-inspection case (requiring 24.5 hours of additional measurement time). As noted above, PDI figures are always rounded-up to integer periods of 8 hours. In this case the rounding-up has

13 The 6 inspection case, under the assumptions given, is not realistic—it has been included as extreme to illustrate that saving on numbers of inspections does not necessarily represent an overall saving—even when MTTD is equivalent. Under a different set of assumptions a 6-inspection case may be both realistic and acceptable.

resulted in seemingly anomalous PDI numbers when compared to the required total number of samples that are to be measured under the different inspection approaches. PDI figures are a very blunt tool to use when comparing inspection schemes. In this example the “samples required” figure gives a more realistic basis for the comparison of different schemes.

2.3–FEWER THAN 365.25 INSPECTION OPPORTUNITIES PER FACILITY PER YEAR

Varying SA3 allows consideration of cases in which there are fewer than 365.25 inspection opportunities per facility per year. Three separate cases of these restrictions upon available inspection opportunities will be presented.

1. The imposition of a boundary condition—in order to avoid imposing an excessive inspection burden upon the operator it may be agreed to accept a defined minimum interval between inspections I_{min} (e.g. this could be set at 7 days[14]).

2. Inspections restricted to facility working days only[15]—in order to ensure that required facility staff are available during inspections it may be agreed to restrict inspections to the normal working week of the facility.

3. Both a boundary condition and inspection restrictions—this approximates the actual SNRI scheme most likely to be acceptable to the broad range of stakeholders[16] in the introduction of such an inspection scheme.

Looking at each of these cases in turn

2.3.1—THE IMPOSITION OF A BOUNDARY CONDITION

The exponential cumulative probability distribution can be used to determine what proportion of the probability curve is being removed by imposing an I_{min} value (where I_{min} represents the lower boundary condition).

$$F(x) = 1- e^{(-x*\lambda)}$$

$$F(I_{min}) = 1- e^{(-I_{min}*\lambda)}$$ excluding below the I_{min} value[17]

Assuming 12 inspections per facility per year, and using the appropriate numbers in the equations given above, allows the calculation of the following value

$$F(I_7) = 21.79\%$$

The value of $f(I_{min})$ allows the examination of the overall effect of introducing a boundary condition upon the cumulative distribution function. For the purposes of further discussion a

14 For a 30-day timeliness goal, setting I_{min} much longer than 7 days would remove many of the benefits from a random inspection regime.

15 No inspections on weekends or public holidays.

16 Stakeholders include (*inter alia*): facility operators; IAEA Operations Divisions and IAEA SGCP

17 Some of the reviewers of this paper had the expectation that this equation was formally and explicitly linked to the timeliness goal, though the equation does have a far wider range of validity. It should be clear from the foregoing discussion that there are analytical solutions to this equation for all positive values of x ($0 \leq x < \infty$).

new parameter is required, $\bar{\lambda}$ (referred to as "lambda bar" or "lambda adjusted"). The new variable $\bar{\lambda}$ represents the adjusted value of the daily probability of inspection that is necessary to maintain the target number of inspections per facility per year.

The exponential function parameter λ, which was introduced earlier, has a simple numeric relationship to the population mean μ.

$$\lambda_i = 1/\mu_i$$

Imposing I_{min} on an inspection scheme removes a minimum of $\acute{\iota}*(I_{min}-1)$ inspection opportunities from the inspection year (as no inspection can take place within I_{min} days of any other inspection). This adjustment removes the simple numeric relationship between $\bar{\lambda}$ and μ, with $\bar{\lambda}$ defined as:

$$\bar{\lambda}_i = \acute{\iota}/(\psi - \acute{\iota}*(I_{min}-1))$$

Table 3 and Figure 1 give values for a variety of the parameters that have been defined so far.

Table 3 – Variable values for a range of inspection schemes featuring a fixed PIV, SNRIs and a lower boundary condition (minimum inspection interval)

Inspections /year	Inspection[18] opportunities	$f(I_7)$	λ	$\bar{\lambda}$
12	293	21.79%	3.29%	4.10%
11	299	19.81%	3.01%	3.68%
10	305	17.86%	2.74%	3.28%
9	311	15.94%	2.47%	2.89%
8	317	14.05%	2.19%	2.52%
7	323	12.19%	1.92%	2.17%
6	329	10.36%	1.64%	1.82%

18 Noting that a boundary condition removes minimum of $\acute{\iota}*(I_{min}-1)$ inspection opportunities from the inspection year.

Figure 1 - Comparison of the values of lambda and lambda bar.

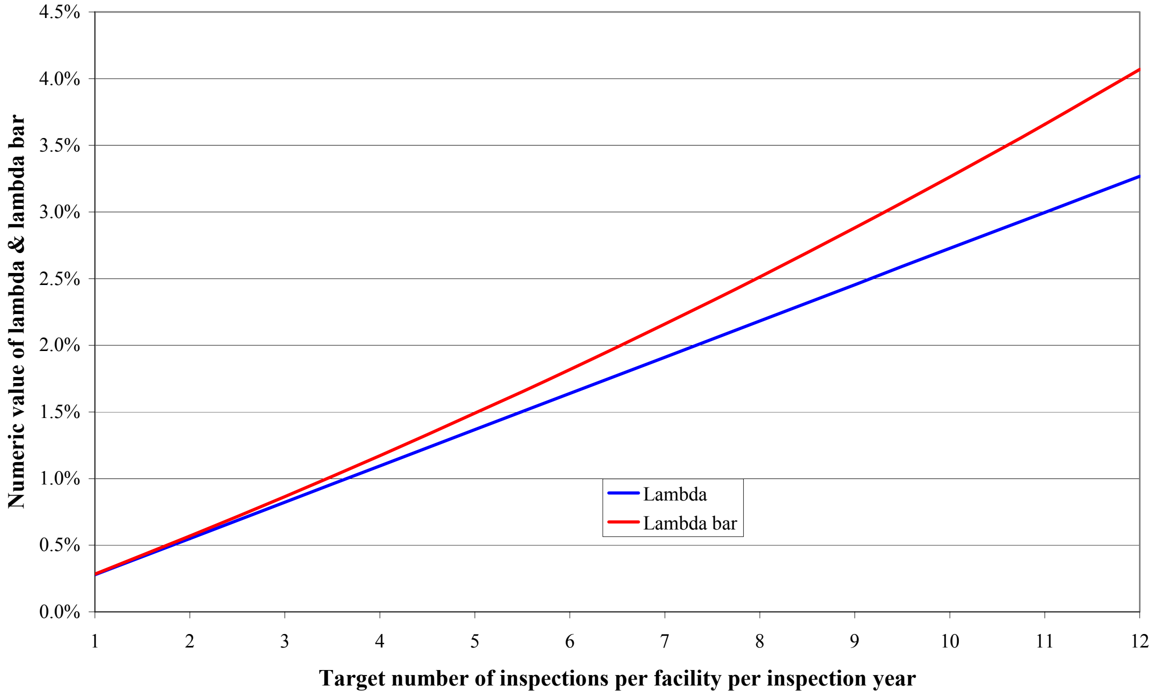

As the number of inspections per facility per year does not change with the introduction of ƛ and the mean time between detections (μ_d) is dependent only upon:

- the unadjusted number of inspection opportunities per facility per year;
- the target number of inspections per facility per year; and
- the average value of the probability of detection (p_{av}),

this gives rise to the second important result of the calculations:

Result 2 (R2): The value of μ (the mean time between inspections/detections parameter) is <u>not significantly changed</u> by the imposition of a lower boundary condition[19].

This means that:
The <u>mean time between inspections</u> remains at its noted value ($\mu_i=\psi/\acute{\iota}$).
The <u>mean time between detections</u> remains at its noted value ($\mu_d=(\psi/\acute{\iota})/p_{av}$).

19 Introduction of an upper boundary condition (i.e. a maximum interval between inspections) while maintaining a target number of inspections is mathematically much more complex than the simple introduction of the ƛ parameter and is beyond the scope of this paper.

2.3.2—INSPECTIONS RESTRICTED TO FACILITY WORKING DAYS ONLY

There are at least three options that could be used to accommodate restrictions upon the available inspection opportunities. Choosing a favoured option is a question of policy rather than a matter that can be decided on its technical merits.

Option 1: Adjust the number of inspection opportunities down (subtracting 104 days for weekends and additional days for the number of declared facility holidays) and calculate a new value for λ_i (referred to as λ_{i1} or "lambda eye one") based on the smaller number of inspection opportunities. Inspections are assigned to inspection opportunities by a random process with the daily probability value being λ_{i1}, only facility working days would be eligible for selection (for example with 250 working days per facility per year λ_{i1}=((365.25/250)*λ)=í/250). Option 1 can therefore be treated as entirely analogous to the ƛ formalism that was discussed under the boundary condition topic above.

Option 2: Keep the value of λ_i unchanged, but conduct any inspections that are randomly scheduled for weekends or facility holidays on the first working day after the day ineligible for inspection[20]. This option would result in a tripling of the probability of inspections taking place on Mondays and, similarly, result in a proportionate increase in the probability of inspections immediately after prolonged periods of holiday (eg perhaps a 9-fold increase on the first working day after Golden Week in Japan or a 10-fold increase on the first working day after the Christmas-New Year periods in many Western countries). The λ_{i2} value is unchanged from the unadjusted value for λ_i ($\lambda_{i2}=\lambda_i$).

Option 3: Count inspection opportunities in a different way (e.g. making each week an inspection opportunity rather than seven inspection opportunities). This may be acceptable if the safeguards approach for the facility includes the judgement that there are no diversion scenarios that can be successfully concealed in less than one week. This resolves the issue of weekend inspections and all but the longest of facility holiday breaks. Prolonged facility unavailability would still need to be addressed using either Options 1 or 2. A λ_{i3} value for weeklong inspection opportunities is simply seven times λ_i ($\lambda_{i3}=7*\lambda_i$). The use of Option 3 takes the discussion beyond the formal range of validity of the exponential distribution and requires a recalculation of all of the previous discussion in terms of the binomial distribution (see section 2.4 below).

20 Deviating from the principle of uniform daily probability of inspection would be considered undesirable in a general sense but may be attractive in a policy sense. The increased probability of inspection after facility unavailability may be seen as an additional measure to address/deter the possibility of diversion/misuse during facility downtime.

2.3.3—BOTH A BOUNDARY CONDITION AND INSPECTION RESTRICTIONS

The format for dealing with both a boundary condition and inspection restrictions would be heavily dependant upon the option chosen for dealing with inspection restrictions.

Option 1—dealt with by calculating λ based on the actual number of inspection opportunities/year (taking into account both days lost to the lower boundary condition and days lost to weekends and facility holidays).

Option 2—no change to the way in which λ is calculated (taking into account only those days lost to the lower boundary condition).

Option 3—if the lower boundary condition and the time period chosen for inspection opportunity consolidation are the same then it is possible to perform calculations based solely on λ_3—prolonged facility unavailability can be dealt with under either Option 1 or Option 2.

2.4–The Binomial Distribution

2.4.1—Poisson Approximation

The foregoing discussion relied on the large number of inspection opportunities per facility per year (and corresponding small daily probability of inspection) in using the exponential distribution and the exponential cumulative distribution. The distribution of inspections across the inspection year approximates the Poisson distribution in cases such as those discussed above which were characterised by a long inspection year (ψ) and a small daily probability of inspection (π) and where the product ($\psi * \pi$) is approximately equal to 1.

Figure 2 - Plot of the Poisson distribution for the probability of a variety of SNRI inspection schema.

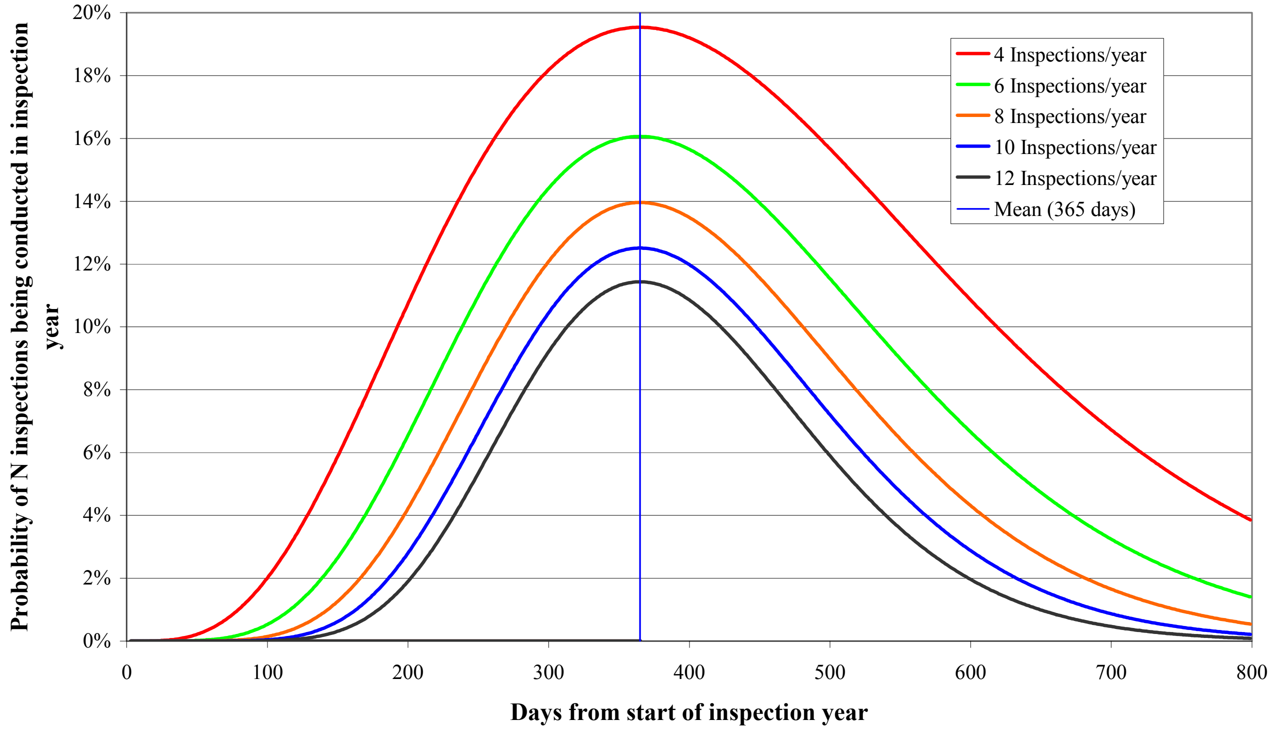

Cases in which inspection opportunities are grouped[21] move beyond the formal area of validity of the Poisson distribution and it becomes necessary to perform the same set of calculations making use of the binomial distribution. This can be stated as:

$$p(n,k) = \frac{n!}{k!(n-k)!}\pi^{k}(1-\pi)^{n-k}$$

With

n = the number of inspection opportunities;
k = the target number of inspections per facility per year;
π = probability of inspection at any opportunity

In the limiting case discussed above the Poisson Theorem provides that:

21 Whether they are grouped as weeks, fortnights, months or in some other way.

$$p(n,k)=\frac{n!}{k!(n-k)!}\pi^{k}(1-\pi)^{n-k}\approx e^{-n\pi}\frac{(n\pi)^{k}}{k!}$$

Using the above formulae it is simple to plot the probable distribution of inspections using the Poisson approximation and binomial distributions for a variety of groupings of inspection opportunities for a given inspection scheme (Figure 3 below).

Figure 3 - Plot of probability of fixed PIV + 5 SNRI occurring within a particular time period when inspections are grouped in time (plot of probability vs days from start of inspection year).

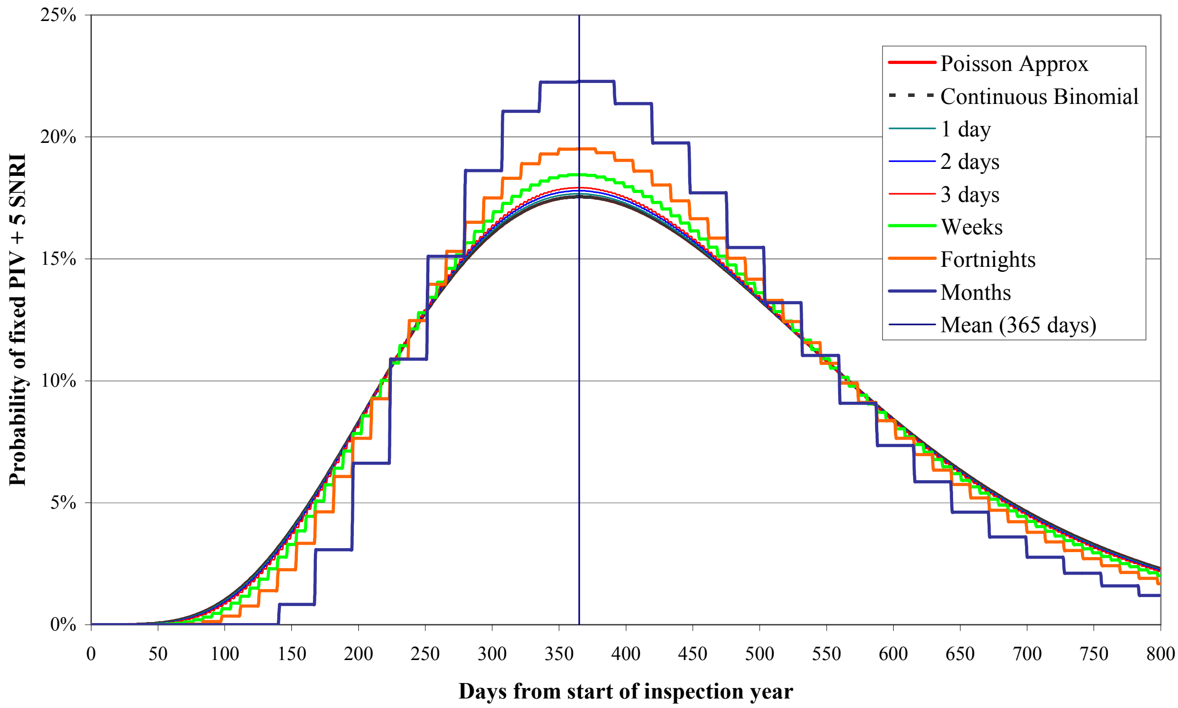

The major effect of the Poisson approximation is to overstate the probability that all of the SNRI will occur at a period shorter than the product of the target number of SNRI and the grouping period (other over- and understatements are present but less significant).

Taking into account the problem with over- and understatements, it is clear from Figure 3 that the Poisson theorem provides an extremely good approximation to a continuous binomial distribution and that it continues to provide a good approximation to binomial distributions based on various inspection opportunity groupings up until inspection opportunities are grouped in months.

2.4.2— Binomial Cumulative Probability

The binomial equivalent of the exponential cumulative probability distribution is given by

$$P(n,k)=\sum_{k}^{n}\pi(n_i,k)$$

Where n, k and π are as given above.

While it is possible to define $\pi(n_i,k)$ for all values of n from 1 to n, its value is zero for values of n < k.

Figure 4 - Plot of cumulative probability of fixed PIV + 5 SNRI occurring within a particular time period when inspections are grouped in time (plot of cumulative probability vs days from start of inspection year).

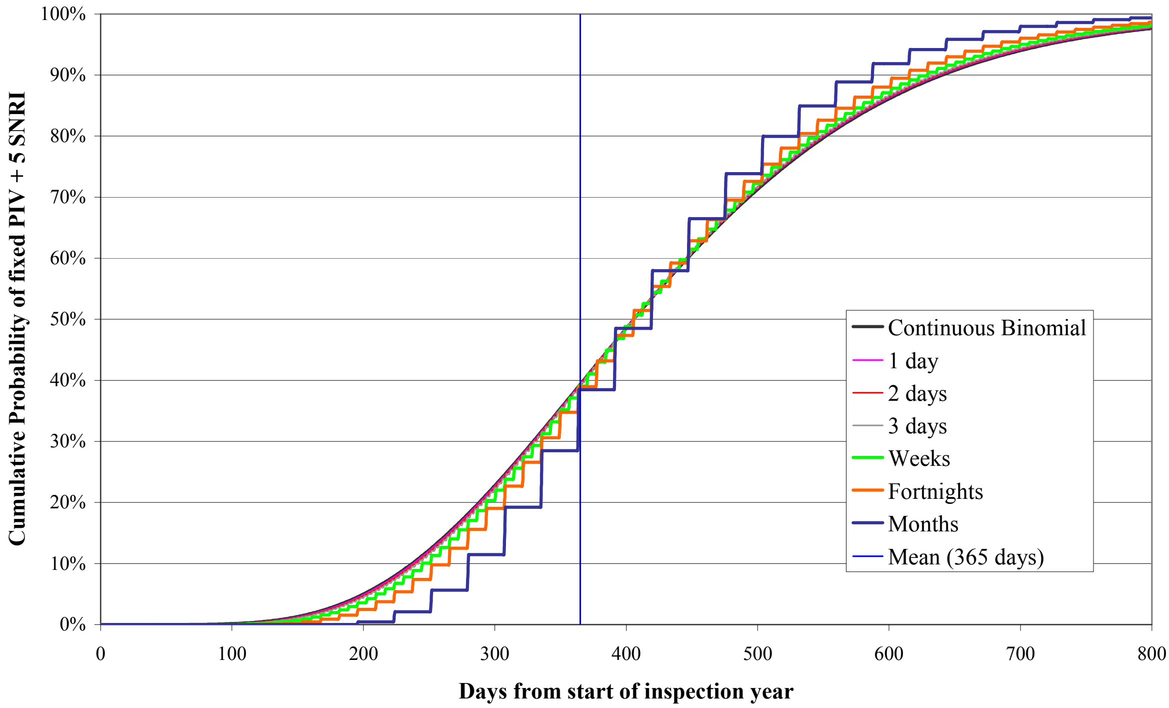

Figure 5 - Replot of Figure 4 in terms of differences between Poisson approximation and various binomial plots.

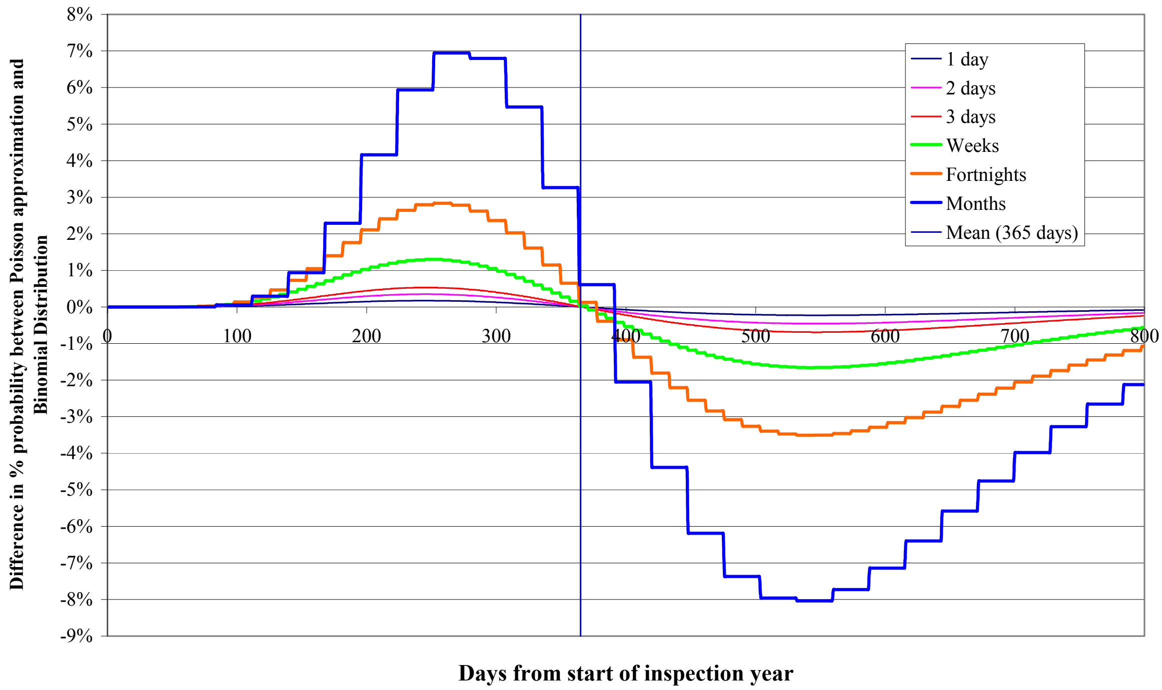

It is clear from Figure 4 and Figure 5 that the effect of grouping inspection opportunities upon these calculations is small for groupings of between 2 and 14 days (maximum of +/-

3%). For groupings of months (28 days) the difference is more significant but the differences are not large enough to affect the conclusions reached so far in this discussion. It is clear that all such differences are readily quantifiable and, on the basis of the foregoing, calculations throughout the paper will make use of the exponential functions already noted with any differences between the exponential and binomial distributions highlighted only where they are significant or effect the results or conclusions.

Please note that Figure 3, Figure 4 and Figure 5 extend 800 days from the start of the inspection year. This is a simplification of the real situation as it ignores the fact that each inspection year is both followed and preceded by another inspection year. Under SA6 a PIV is scheduled to occur at the first inspection opportunity of each inspection year—the effect of the preceding inspection years upon the current inspection year is additive.

3.1—Detection at Shorter Intervals than current timeliness goals

3.1.1—Traditional Safeguards Approaches

As a basis for comparison for the discussion that follows it should be noted that traditional safeguards inspection schemes for facilities with UDU have an expected mean time to detection of 61 days—this figure has a standard deviation of 30 days. The probability of detecting a diversion within 61 days is equal to (roughly) 78%[22]. Figure 6 provides a plot of the cumulative probability of detection and detection probability for idealised inspection years under traditional safeguards (the inspection year has been rounded to 360 days to ensure integer results).

Figure 6 - Plot of cumulative probability and % detection probability at a given inspection for an idealised traditional safeguards inspection scheme for a UDU facility.

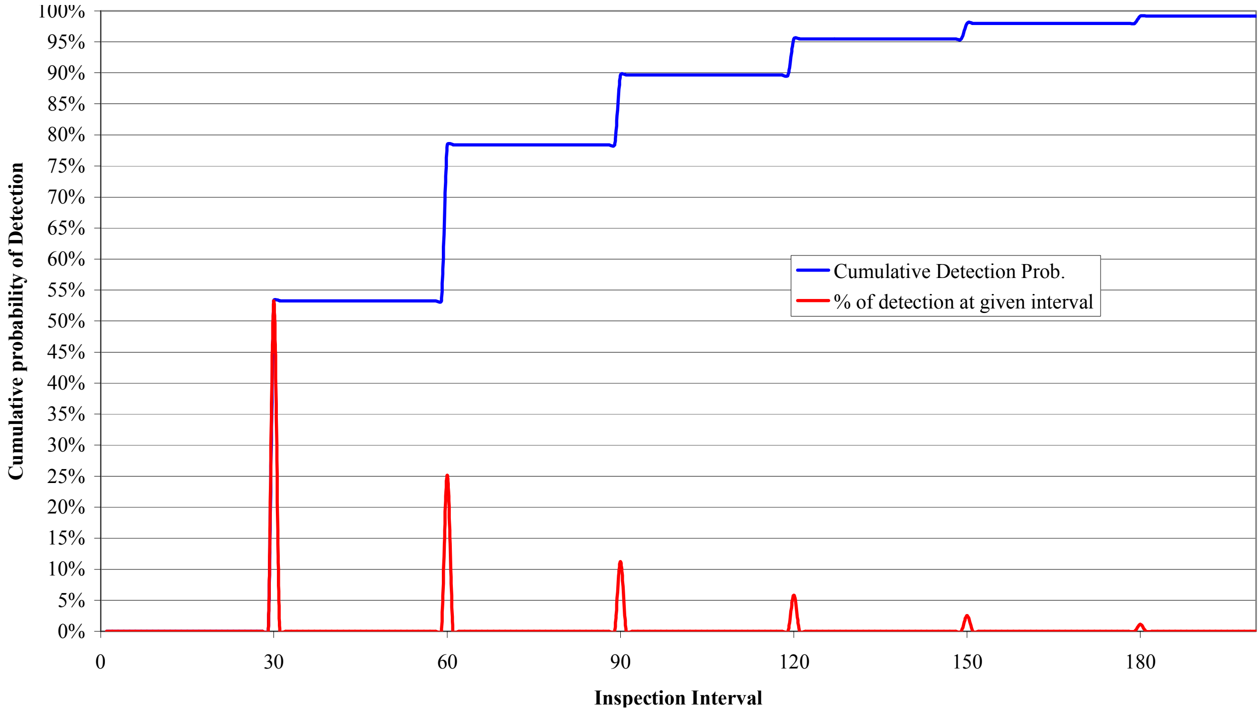

It can be seen from Figure 6 that, under the idealised case, detections are possible only at the fixed inspection interval of 30 days and that there is an average 47% probability that the detection interval will exceed the 30-day timeliness goal. There is also an average 22% probability that detection interval will exceed the design detection goal as expressed in the defined detection probability for the inspection (61 days).

When comparing SNRI schemes with traditional inspection approaches it is important to acknowledge that there are limitations on the effectiveness of traditional safeguards approaches. Under traditional safeguards approaches to UDU facilities, there is a significant probability (approximately 22%) of any diversion being undetected for periods longer than 61 days.

---

22 The idealised inspection year assumes: a timeliness goal of 30 days; inspections intervals of exactly 30 days, a PIV detection probability of 90%; and an interim inspection detection probability of 50% on an idealised 360-day inspection year.

3.1.2—SAFEGUARDS APPROACHES USING SNRI

In all of the discussion that follows it is assumed that: the timeliness goal is 30 days; inspection intervals are randomly timed; a PIV has a design detection probability of 90%; and an interim inspection has a design detection probability of 50%. The figures that follow include (*inter alia*) markers to indicate: the lower boundary condition; one half the inspection goal; the timeliness goal and twice the timeliness goal for various inspection schemes.

Figure 7 - Plot of cumulative probability of inspection in an ideal SNRI approach with no fixed inspections and no boundary conditions.

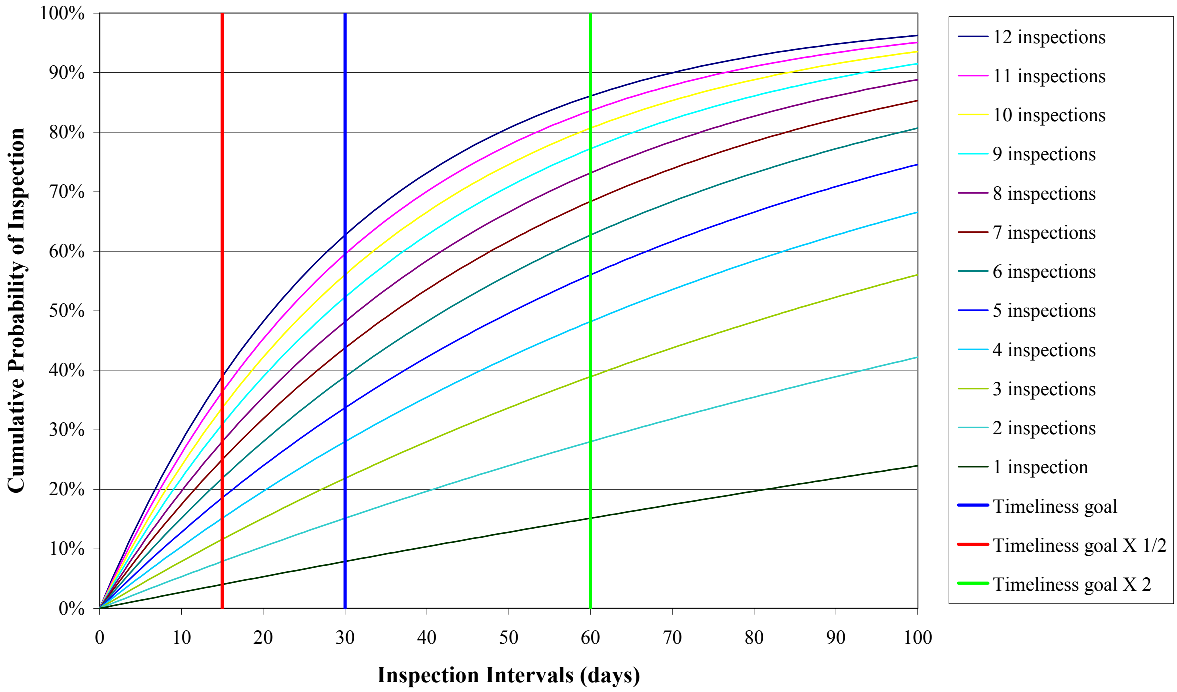

The attractions of randomised inspection schemes include that:
- they offer the possibility of <u>inspection</u> intervals shorter than traditional timeliness goals;
- they offer the possibility of <u>detection</u> intervals shorter than traditional timeliness goals; and
- they offer the possibility of significant resource savings when a reduced number of randomised inspections is used to replace traditional fixed inspection schemes.

As can be seen from Figure 7 and Figure 8, with SNRI there is a significant probability of <u>inspection intervals</u> of shorter than one half of the traditional timeliness goal, even in cases where there are, on average, half the traditional number of inspections conducted per facility per year (i.e. 6 inspections per year).

Figure 9 shows the cumulative probability of detection of a range of detection intervals and also shows that there is a significant probability of achieving detection intervals equal to or less than the timeliness goal with each of the inspection schemes shown.

Figure 8 - **Plot of inspection intervals for various inspection schemes with fixed date for PIV and SNRI for all other inspections, with (λ) and without (ƛ) a defined lower boundary condition of a 7 day minimum interval.**

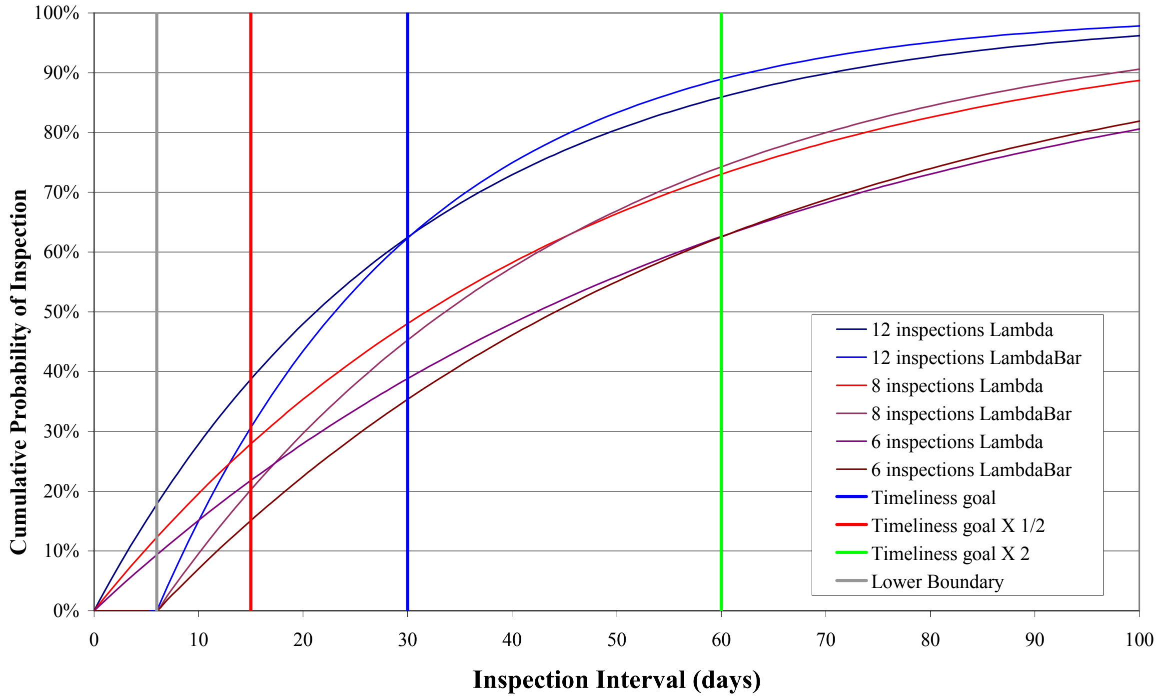

Figure 9 - **Plot of <u>detection intervals</u> (based on p_{av}) for various inspection schemes with fixed date for PIV and SNRI for all other inspections, with (λ) and without (ƛ) a defined lower boundary condition of a 7 day minimum interval (plot of % probability vs. <u>detection interval</u> (in days))**

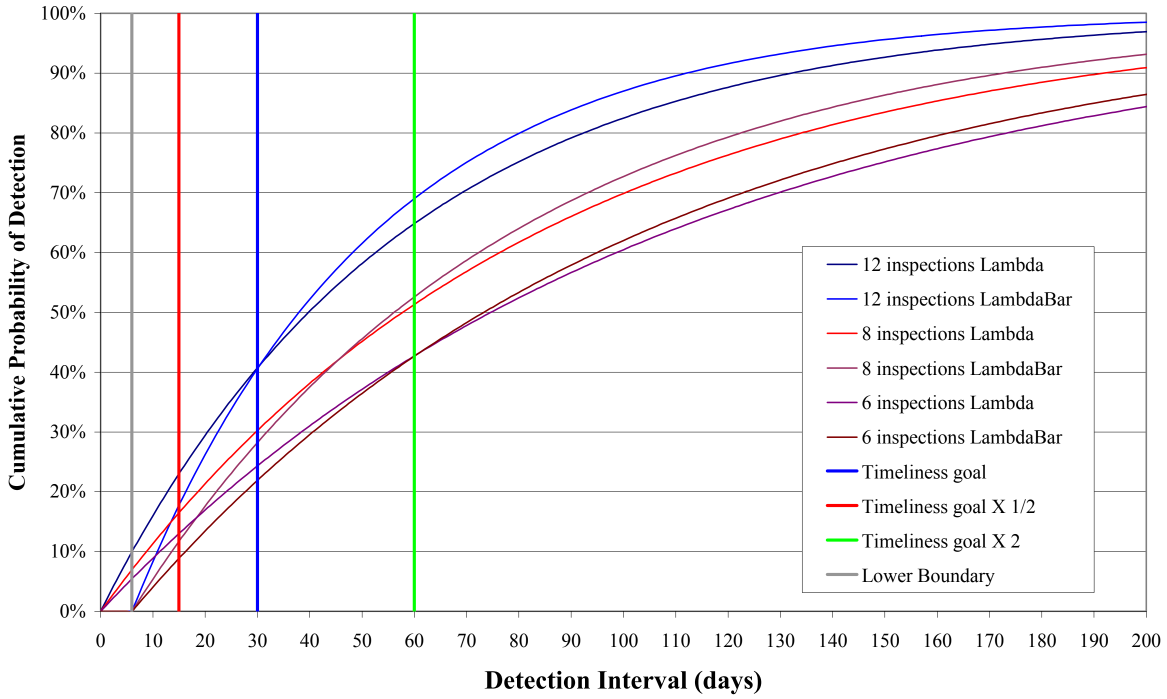

3.2—COMPARING TRADITIONAL APPROACHES AND VARIOUS SNRI SCHEMES

Table 4 - Table of detection probabilities and cumulative detection probabilities for an idealised traditional safeguards inspection scheme for a UDU facility.

Inspection Interval (days)	Detection Probability	Cumulative Detection Probability
15	N/A	N/A
30	53%	53%
60	25%	78%
90	11%	89%
120	6%	95%
150	3%	98%
180	1%	99%

Table 5 - Table of detection probabilities and cumulative detection probabilities for an inspection scheme with a minimum inspection interval of 7 days, utilising 8 inspections per facility per year (1 fixed PIV and 7 SNRI) for a UDU facility.

Inspection Interval (days)	Detection Probability (cumulative interval)	Cumulative Detection Probability
15	11%	11%
30	17%	28%
60	25%	53%
90	13%	69%
120	10%	79%
150	7%	86%
180	5%	91%

Table 6 - Table of detection probabilities and cumulative detection probabilities for an inspection scheme with a minimum inspection interval of 7 days utilising 6 inspections per facility per year (1 fixed PIV and 5 SNRI) for a UDU facility.

Inspection Interval (days)	Detection Probability (cumulative interval)	Cumulative Detection Probability
15	9%	9%
30	13%	22%
60	21%	43%
90	15%	58%
120	11%	69%
150	8%	77%
180	6%	83%

Table 7 - Table of detection probabilities and cumulative detection probabilities for an inspection scheme with a minimum inspection interval of 7 days utilising 4 inspections per facility per year (1 fixed PIV and 3 SNRI) for a UDU facility.

Inspection Interval (days)	Detection Probability (cumulative interval)	Cumulative Detection Probability
15	6%	6%
30	10%	16%
60	16%	32%
90	13%	45%
120	10%	55%
150	9%	64%
180	7%	71%

Table 8 - Table of cumulative detection probabilities for: a traditional safeguards approach; 8 inspections per facility per year (1 fixed PIV and 7 SNRI); 6 inspections per facility per year (1 fixed PIV and 5 SNRI); and 4 inspections per facility per year (1 fixed PIV and 3 SNRI) for a UDU facility.

Inspection Interval (days)	Traditional Safeguards	1 fixed PIV + 7 SNRI	1 fixed PIV + 5 SNRI	1 fixed PIV + 3 SNRI
15	N/A	11%	9%	6%
30	53%	28%	22%	16%
60	78%	53%	43%	32%
90	89%	69%	58%	45%
120	95%	79%	69%	55%
150	98%	86%	77%	64%
180	99%	91%	83%	71%

Table 5, Table 6, Table 7 and Table 8 illustrate the different cumulative detection probabilities for the four different approaches considered. While the traditional safeguards case provides a higher probability of achieving a detection interval of 61 days or less, it does not provide any capability of achieving a detection interval of less than 30 days. In choosing among the options presented above, a policy decision would be needed as to the weight that should be given to the probability of detecting a diversion of UDU at intervals less than the traditional timeliness goals for the material in question.

Page intentionally left blank

SECTION 4.1—SPECIFIC RECOMMENDATIONS

Recommendation 1: That the IAEA should seek to arrange for a trial of a randomised inspection scheme for a UDU storage facility.

Recommendation 2: On the basis of Result 2, the IAEA should accept operator requests for a defined minimum interval between randomly timed inspections for the trial as there is expected to be no net effect upon the mean time to detection in this type of restriction upon inspection timing. If a policy decision is made to group inspection opportunities for inspection planning purposes it may be unnecessary to include a defined minimum inspection interval.

Recommendation 3: That the trial should make use of 6 inspections per facility per year (1 fixed PIV plus an average of 5 SNRI).

Recommendation 4: That the daily probability of an SNRI for the inspection trial should be set using the calculated λ value for this inspection scheme 1.4907%. If a policy decision is made to group inspection opportunities for inspection planning purposes, the product of the λ value and the number of inspection opportunities should be used for the purposes of setting inspection probabilities (e.g. for groupings of 7 inspection opportunities a value of 9.5825% should be used, and for groupings of 14 inspection opportunities a value of 19.1650% should be used).

Recommendation 5: That, if a policy decision is made to group inspection opportunities for inspection planning purposes, groupings of no more than two weeks (14 calendar days) should be used.

Recommendation 6: In making the calculation of the days on which to conduct inspections, the facility will be assumed to be available an average of 365.25 days a year. For any days that are selected for inspection which the facility operator has notified the IAEA in advance that the facility will not be available for inspection (including all weekends and facility holidays), the inspection will be carried out on the next available facility working day[23]. In the case of grouped inspection opportunities, the inspection should take place on a day within the planned inspection opportunity grouping.

Recommendation 7: If the operator declares the facility not available for inspection, without prior notice, on a day or within an inspection opportunity grouping which is randomly selected for inspection, this should be treated in accordance with SMO provisions for an Access Anomaly.

23 As noted in the body of this report—adopting this option without inspection opportunity grouping will result in an effective tripling of the probability of inspections being conducted on Mondays and proportionate increases in the probability of inspections being conducted after long periods of facility unavailability.

Page intentionally left blank

SECTION 4.2—RESULTS AND CONCLUSIONS

Result 1: It is possible to achieve equivalent mean time to detection values for different numbers of randomised inspections per inspection year by varying the target values for the probability of detection for each inspection.

Caveat on Result 1: Increasing the value of probability of detection as a trade-off may not result in decreased overall inspection costs across the inspection year, due to the increased costs associated with performing interim inspections to a higher detection probability[24].

Result 2: The value of μ (the mean time between inspections/detections parameter) is not significantly changed by the imposition of a lower boundary condition[25].

This means that the values for key parameters are as follows:
The mean time between inspections $\mu_i=\psi/\acute{\iota}$
The mean time between detections $\mu_d=(\psi/\acute{\iota})/p_{av}$

Conclusion 1: The use of SNRIs provides the possibility of detecting diversion of UDU at intervals shorter than the timeliness goal for the material in question.

Conclusion 2: It is possible to develop detailed inspection schemes which utilise a mixture of fixed and SNRIs as a means of achieving specific "mean time to detection" targets. The numeric differences between a statistically ideal situation and the real-life situation are not significant in most practical cases.

Conclusion 3: When combined with policy decisions as to acceptable levels of the probability of achieving specific timeliness goals, mean time to detection calculations provide a useful guide to inspection decision making.

Conclusion 4: For the purposes of this discussion the exponential distribution is a useful analogue for the binomial distribution. The differences between results calculated via the two methods are not significantly different from each other for inspection opportunity groupings of 14 days or less.

---

24 See section 2.2 for an illustration of this point.

25 Introduction of an upper boundary condition (i.e. a maximum interval between inspections) while maintaining a target number of inspections is mathematically much more complex than the simple introduction of the λ parameter and is beyond the scope of this paper.

BIBLIOGRAPHY:

[1] CRC Handbook of Chemistry and Physics, 82nd Edition.

[2] *"Algorithms to calculate samples sizes for Inspection Sampling Plans"*, STR-261, J.L.Jaech and M.Russell, IAEA Section for Statistical Analysis, March 1991.

[3] *"Schaum's Mathematical Handbook Of Formulas And Tables"*, M.R. Spiegel, J. Liu, McGraw-Hill, 1981.

Annex

A.1—NUMERIC APPROXIMATIONS USED IN THIS PAPER

The caveat section of the report noted that the calculations using the cumulative exponential probability function are strictly valid only in cases in which all inspections during the inspection year were randomly scheduled including the PIV.

The numeric differences between the ideal situation and the more practical situation with a fixed PIV are small and consistent. Table 9 lists these differences[26] in relative terms where λ_a refers to the ideal case (no fixed inspections) and λ_b refers to the case with the PIV fixed and all other inspections randomly scheduled. Except for the case of 1 inspection per facility per year, such differences are of the order of 0.5% of the relative value of λ (less than 0.02% of absolute value).

Table 9 - Numeric values for deviations from the statistical ideal when PIV is fixed in time

Inspections /year	λ_a (PIV not fixed in time)	λ_b (PIV fixed in time)	δ_1 $(\lambda_a - \lambda_b)/\lambda_a$	δ_2 $\delta_1/í$
12	3.288%	3.267%	0.6423%	0.0535%
11	3.014%	2.996%	0.5885%	0.0535%
10	2.740%	2.727%	0.4808%	0.0481%
9	2.466%	2.454%	0.4771%	0.0530%
8	2.192%	2.182%	0.4540%	0.0567%
7	1.918%	1.911%	0.3529%	0.0504%
6	1.644%	1.638%	0.3298%	0.0550%
5	1.370%	1.367%	0.1928%	0.0386%
4	1.096%	1.095%	0.0574%	0.0144%
3	0.822%	0.823%	-0.1344%	-0.0448%
2	0.548%	0.551%	-0.5985%	-0.2993%
1	0.274%	0.279%	-1.8715%	-1.8715%

A.2—AVERAGE DETECTION PROBABILITY—CALCULATION OF VALUES USED

The average detection probability for any given inspection year in which more than one detection probability applies will be dependent upon the number of inspections that are conducted in that year and their target detection probabilities.

Calculation of p_{av} for inspection schemes utilising fixed inspections is straightforward, as a simple weighted average of the PIV detection probability and interim inspection detection probability gives a good working value for p_{av}. The value is not valid for periods of less than one inspection year but is valid for all integer multiples of an inspection year.

Calculation of p_{av} for inspection schemes utilising random inspections only is more complex than the fixed inspection case, but if a long term weighted average (1000 inspection years or more) for the PIV detection probability and interim detection probability is used, the

26 The values for λ_b were calculated using a Monte-Carlo style simulation program.

uncertainties in the p_{av} will tend to be insignificant. The value for p_{av} is unlikely to be valid for any given inspection year (as the number of interim inspections will vary around the mean)—but the weighted average value will provide a useful year-to-year comparison.

Calculations of p_{av} for inspection schemes utilising a mixture of fixed and random inspections presents a number of problems that are related to the small but consistent deviation from the ideal form of the exponential cumulative probability distribution (see the discussion on numeric approximations at the beginning of this Annex). While p_{av} can be calculated using the same weighted average technique noted above for the random inspections only case, there will be a bias in the calculated result that will tend to overstate p_{av} for cases of 6 inspections or more per facility per year and understate p_{av} for cases of less than 6 inspections per facility per year. This problem was overcome for the report by making use of the same Monte-Carlo style simulations noted above to calculate a value inclusive of bias.