MASTER

DISCLAIMER

# The Goals of Measurement Systems for International Safeguards*

by

J. M. de Montmollin
Sandia Laboratories, Albuquerque, NM

E. V. Weinstock
Brookhaven National Laboratory, Upton, NY

## ABSTRACT

The safeguards applied by the International Atomic Energy Agency are based on technical performance goals and criteria that have been developed, but not officially adopted by the Agency. The goals derive in part from the external consequences that safeguards are intended to prevent and in some cases on internal considerations of feasibility. To the extent that these goals may not be attainable, as may be the case with large-throughput bulk reprocessing plants, the Agency is placed in a difficult position. In this paper safeguards goals and criteria and their underlying rationales are critically examined. Suggestions for a more rational and workable structure of performance goals are offered.

KEYWORDS: International Atomic Energy Agency; international safeguards; safeguards performance criteria; safeguards objectives

International Atomic Energy Agency safeguards have been the subject of much attention in recent years, since concerns over nuclear-weapon proliferation have come to have a predominant influence on US policies affecting nuclear power. The performance criteria and goals that are being developed for the technical safeguards system are reflected in the perceptions of the policymakers and the reliance they are willing to place on IAEA safeguards. Technical performance criteria therefore have an impact on policy objectives, and they must be considered in that broader context. In this paper we address the problem of the specification of goals and criteria that are both practical from the technical standpoint and that will support the attainment of non-proliferation policy objectives.

## IAEA SAFEGUARDS

Let us begin with a brief description of IAEA safeguards, and how they relate to domestic safeguards such as those administered by NRC. They differ in important respects, and form complementary but distinct elements in the overall structure.

As shown in Figure 1, the two systems include both common and separate elements. The IAEA system relies in an essential way on the State System of Accounting and Control, or SSAC. The Agency provides general guidelines for the SSAC that the State is expected to operate, and it avoids duplication of SSAC data, making independent measurements only for verification of the State's records.[1] The Agency cannot impede or obstruct the movement of materials in any way, and it relies on the containment -- walls, shielding, barriers, containers -- that is available. The Agency performs independent audits of records required by the State, as well as independent measurements, observations, and surveillance of containment and certain operations.

*The views expressed herein are those of the authors, and do not necessarily reflect those of Brookhaven or Sandia Laboratories or the Department of Energy.

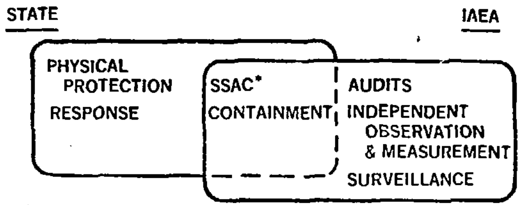


**STATE SYSTEMS**

MONITOR AND CONTROL
RESPONSE ACTION

--

STATE REGULATION
AND ENFORCEMENT

**AUTHORITY**

NATIONAL LAWS

**IAEA SAFEGUARDS**

MONITOR BY MUTUAL AGREEMENT
VERIFICATION
GUIDELINES FOR PHYSICAL
PROTECTION

**AUTHORITY**

IAEA STATUTE
- ART III
  - NPT[2]
  - TLATELOLCO[3]
  - BILATERAL AGREEMENTS
- ART XII
  - IAEA PROJECTS & MATERIALS

*STATE SYSTEM OF ACCOUNTING AND CONTROL

FIGURE 1. RELATIONSHIP OF IAEA SAFEGUARDS TO DOMESTIC SYSTEMS

While the Agency system provides only a capability to detect unauthorized activities, the State system covers control and response functions as well. It includes physical protection -- the monitoring and control of personnel and the exclusion of those not authorized, and means to respond to any attempted violation while it is in progress. It protects against sabotage, as well as misuse or diversion of materials and facilities.

The two systems differ further in that the State system is backed by the legal authority of the State, and is directed against threats to State interests by subnational groups or individuals. The IAEA system is intended to promote that protection by the State, but principally to provide assurance that the State itself does not violate its commitments pledging not to misuse materials or facilities. IAEA safeguards are applied only with the consent of the State, with the technical details subject to negotiation. While the State is the nominal adversary, there is a common interest in providing assurance that legitimate activities are not a cover for weapon production. The credibility of the assurance provided by IAEA safeguards depends in large measure on the capability to detect a diversion, should it occur. Since an actual attempted diversion of nuclear material would be a rare and unique event, the perception of the IAEA capability, rather than a history of successful detections, is the basis for the reliance on IAEA safeguards. Performance criteria and the ability of the IAEA to meet them are the basis for those perceptions.

## THE OBJECTIVE OF SAFEGUARDS

The agreement establishing the IAEA states that the purpose of safeguards is

> ". . . to ensure that special fissionable and other materials, services, equipment, facilities, and information . . . are not used to further any military purpose . . ."[4]

The authority of the Agency to prevent misuse is limited to monitoring, and not direct control or sanctions.* The first document describing the IAEA safeguards system was more specific:

> "The purpose of safeguards inspections shall be to verify compliance with safeguards agreements and to assist States in complying with such agreements . . . ."[5]

The NPT[2] obligated non-weapon-State signatories to accept IAEA safeguards. A group of technical consultants was convened by the IAEA for the purpose of drafting a new document covering safeguards under the NPT.[6] The group considered the problems of how a system could operate so as to provide assurance with positive measures of performance that diversion did not occur.[7,8] The solution that they arrived at was to state the objective of safeguards in positive terms:

> ". . . the objective of safeguards is the timely detection of diversion . . . of nuclear material . . . and the deterrence of such diversion by the risk of early detection."[9]

Although that definition of safeguards objective is widely cited and accepted, it differs from that stated in the safeguards document covering non-signatories of the NPT[10] and in the NPT itself in that it states that the objective is the detection of diversion rather than verification of compliance with non-proliferation obligations.** The problem of providing assurance that diversion has not occurred is not really solved by stating safeguards objectives in terms of detection of diversion. Safeguards operations

*There is limited additional authority in connection with Agency-supplied materials or Agency-sponsored projects, but that is of little practical significance.

**Article III of the NPT states in part that safeguards are ". . . for the exclusive purpose of verification of the fulfillment of obligations assumed under this treaty . . ."

lead to periodic findings by the IAEA that compliance is verified or not verified, and that is the objective as stated in INFCIRC/66 Rev. 2 and the NPT. The credibility of that verification depends on a capability to detect any diversion that might be attempted. A capability to detect diversion is a means to the end objective of verification. In the general case where no diversion occurs the objective is verification, and the problem of proving non-diversion remains.

## VERIFICATION, NOT DETECTION

If the findings generated by the safeguards system could be so conclusive and unambiguous that we could always expect either hard evidence of diversion or convincing assurance of no diversion, the definition of safeguards objective would have less significance. That is, if the two cases were mutually exclusive, so that the absence of one condition necessitated the other, the objectives would be equivalent. However, in the real world there is a third state: a condition of uncertainity in which it can neither be stated that compliance with obligations is verified, nor that diversion has been detected. That situation may be due in part to other factors besides the imperfect performance of the technical safeguards system. Negotiated safeguards arrangements are neither perfect nor all-embracing; the safeguards operations can go no further than to verify that the agreements were complied with and that no indications of possible diversion were detected in the process. Failure to verify compliance may or may not be related to an actual diversion. It makes considerable difference whether the safeguards operation is intended to meet some required probability of detecting any diversion that might be attempted or whether it is expected only to detect indications of apparent violations and hence to reach a finding of inability to verify compliance.

Figure 2 diagrams the relationships between detection of diversion and verification of compliance as they are generally idealized and as they actually are in safeguards operations. The horizontal axis represents a scale of varying degrees of assurance or confidence. On the basis of the safeguards data that is collected and analyzed, a finding is reached with some degree of assurance represented by point A or B. A, for example, might be based on a more intensive sampling plan, smaller MUF, or fewer anomalous indications. As we move to the left, assurance diminishes to the point where we are unable to verify compliance with a degree of assurance deemed to be adequate. Current safeguards criteria recognize some small band of uncertainty, represented by the 95 percent confidence limits, wherein there is a risk of false alarm or failure to detect. No data are collected that would establish a point on the left side (except in the rare case of actual diversion), and therefore no statement could be made on the basis of available data other than verification or inability to verify. Inability to verify, unless it could be dismissed as being within the allowable margin for error represented by the statistical confidence limits, is often assumed by many in the technical safeguards community to be tantamount to detection of diversion.

The uncertainty is not limited only by the precision of the technical safeguards system. As Figure 2(b) indicates, detection of diversion and verification of compliance do not form a continuum, since agreements are not perfect or all-encompassing, and non-compliance is not neccssarily associated with diversion. Agreements must accomodate practical considerations, often after hard negotiations, including such things as intrusiveness and operational impact, proprietary rights, costs, and political constraints. Material accounting, the principal safeguards measure, can only be implicitly assumed to approximate reality because of IAEA resource limitations and the possibilities of undeclared streams and the falsification of records. Finally, the safeguards systems' technical performance is limited by coverage and reliability as well as measurement precision.

The relatively large uncertainty band that separates assured detection from assured verification in Figure 2(b) is the reason why detection and verification cannot be accepted as being effectively mutually exclusive. A finding of verification can be made only if there is supporting data at some point such as A or B. If such a point cannot be established, and if there is no direct evidence to substantiate an actual diversion, the only finding that can be made is that compliance cannot be verified. If there is an

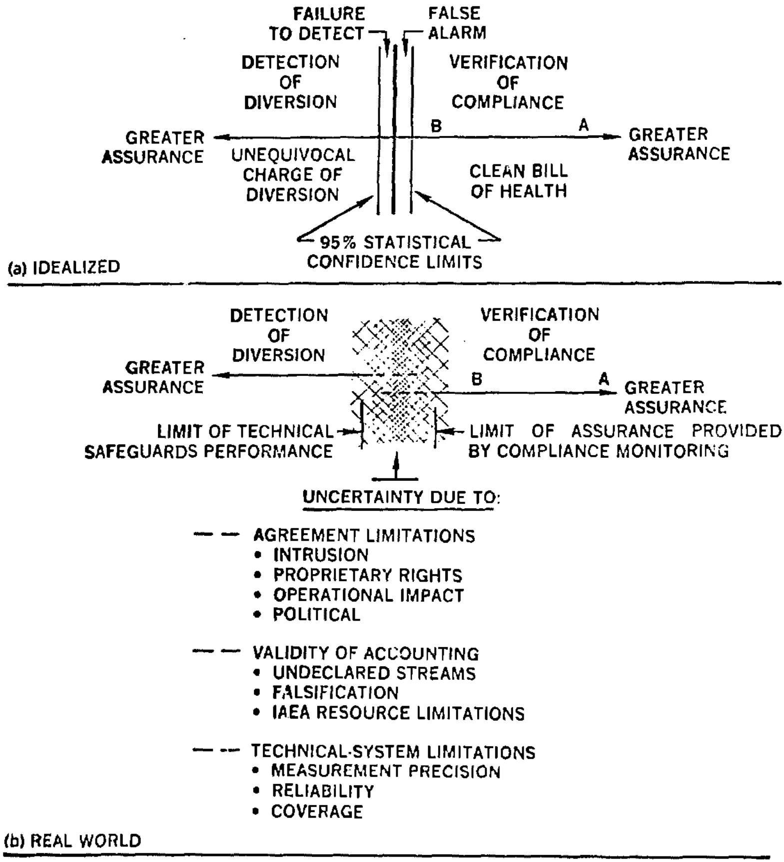


FIGURE 2. RELATIONSHIP OF DETECTION OF DIVERSION TO VERIFICATION OF COMPLIANCE

actual diversion and an effort is made to conceal it, the diverter is likely to be able to make it sufficiently ambiguous so that the information falls into the region of uncertainty with inadequate assurance to support a direct confrontation.

## THE CONSTRAINT OF ACCEPTABILITY TO THE SAFEGUARDED

The IAEA Statute does not obligate any State to accept safeguards.* A commitment to accept safeguards must be otherwise motivated, usually from participation in some regional or international agreement such as the NPT or the Treaty of Tlatelolco or because it was a condition imposed by an exporting State. Typically the commitment is only to submit to "IAEA safeguards," with the actual arrangements to be subject to negotiations within general guidelines proposed by the Agency. While the State is committed to accept some kind of IAEA safeguards, the specific arrangements are confidential between the IAEA and the State. The guidelines for the agreements are provided in INFCIRC/66 Rev. 2 and INFCIRC/153. They have no legal force in themselves; for example, INFCIRC/66 Rev. 2, states specifically that its provisions will become legally binding only when, and to the extent that, they become incorporated into an individual agreement[11,12]. Similarly, the title page of INFCIRC/153 contains the following statement:

> "The Board of Governors has requested the Director General to use the material reproduced in this booklet as the basis for negotiating safeguards agreements between the Agency and . . . States . . . "[6]

Since what constitutes IAEA safeguards is not specified in the basic agreement such as the Statute or the NPT, in which the State agrees to accept them, the IAEA cannot unilaterally impose its own performance criteria. The agreement will in the end be based principally on a commonality of interests, bounded by the Agency's need to maintain credibility and the State's willingness to accept some intrusion on operations and national sovereignty. In setting performance goals and criteria, the Agency is motivated to press for greater effectiveness, striving for further improvements in the future.[13] At the same time, the impact on the State's nuclear activities must be reasonable.[14]

## EXTERNAL AND INTERNAL GOALS

The distinction between external goals -- wherein performance goals are specified in terms of quantities related to prohibited uses of the diverted material -- and internal goals based on current or reasonable extensions of state-of-the-art capabilities, is an important one. Current goals are ambivalent in this respect: threshold quantities and timeliness goals are based at least in principle on the time and amounts required for a diverter to fabricate a single explosive device, while detection and false-alarm probabilities are based on statistical limits of measurement precision and not on deterrent value or consequences of failure.

If goals could be set which are both reasonably achievable, now or in the near future, and which would be effective as deterrents and in limiting the consequences of diversion, the distinction between external and internal goals would be unimportant. That is the case with many safeguards applications: external goals can be met with measures that are reasonable and acceptable. However, that is not true for large-throughput bulk processing plants. Acceptance of external goals as requirements that cannot be met will in the end diminish confidence in the entire safeguards institution, as the inability to meet stated goals becomes more widely apparent.

The IAEA has the mission of applying the most effective safeguards that it can, subject to the constraints of impact on operations, limited resources, and acceptability to the State. It is therefore bound by internal constraints, however desirable it may be to meet more difficult external goals.

*With the unimportant exception of Agency-sponsored materials or projects.

## GOALS VS. REQUIREMENTS

If goals are not immediately attainable, we must distinguish between goals and requirements, or what might more accurately be called current standards of performance. Standards that are currently achievable with reasonable effort and impact form a reference point for negotiating ageements. Requirements,* in the sense that an IAEA system that does not meet them is considered inadequate, should be based on what is currently achievable, and hence on internal goals.

Since the basic agreements obligate the Agency to apply safeguards without specifying what performance constitutes adequate safeguards, there can be no requirements based on external goals that are not immediately attainable. From the Agency's standpoint, there is no question of the "adequacy" of safeguards beyond its internal performance standards --"adequate" is doing the best it can with what it has.

On the other hand, external goals are all that matter to the outside world, and to States whose actions are influenced by safeguards. They can certainly set performance levels in terms of the function they hope IAEA safeguards will accomplish for them, and they are free to judge whether safeguards are "adequate" to them or not on the basis of those requirements. However, if the Agency adopts external requirements which it cannot meet, it is caught in a serious dilemma. It would be in a position of having to apply safeguards which by its own criteria it must concede to be inadequate, or to apply no safeguards at all on the grounds that adequate safeguards were infeasible. Either choice would be detrimental to the IAEA and the non-proliferation effort.

## US REQUIREMENTS

The example of the US Nuclear Non-Proliferation Act of 1978[15] is a good illustration of the different forces that influence national and IAEA safeguards policies. We will concentrate here on the important single issue of timeliness.

There is a view held by some in the US that there should be a rigid requirement that timeliness of detection should be enough shorter than conversion time** so as to allow preventive action to be taken before the first explosive device could be fabricated. The preventive action is understood to be some form of diplomatic pressure sufficient to cause immediate cessation of the conversion operation, but the time that is to be allowed for the action is not specified. In any case, since the conversion time for plutonium in any form that is free of fission products is stated by the IAEA to be on the order of one to three weeks (days for metal), there is little time for the IAEA to reach a finding that is unambiguous enough to generate any kind of immediate action that could be expected to stop the conversion operation in its tracks. Furthermore, the available time is stated by advocates such as Wohlstetter to be a matter of hours,[16] or even "practically zero" (Gilinsky).[17]

The official US position on the timeliness requirement is written into the Nuclear Non-Proliferation Act. With regard to whether reprocessing of US-exported spent fuel is to be permitted, the Act states:

"Sec 303(b)(2) . . . foremost consideration will be given to whether or not the reprocessing or retransfer will take place under conditions that will ensure timely warning to the United States of any diversion well in advance

---

*Requirements can only be self-imposed by the Agency, and there is no authority to require States to accept safeguards measures merely because they are claimed to be necessary to meet the IAEA's stated performance goals.

**The time required to convert material to a weapons-usable form.

of the time at which the non-nuclear state could transform the diverted material into a nuclear explosive device."

In the floor debate on the bill, Senator Glenn explained that, while there might be circumstances in which the timely warning requirement might be relaxed somewhat, those circumstances would be very unusual:

". . . there is no part of this bill that is of more significance for the prevention of nuclear proliferation than the elevation of the 'timely warning' standard to statutory force.

"I think it is important to note that the concept of timely warning, the basic concept upon which the entire international safeguards program rests, is strictly a measure of whether the warning of a diversion will be received far enough in advance of the time when the recipient could transform the diverted material into an explosive device to permit an adequate diplomatic response."

". . . the bill requires that foremost consideration be given to the question of timely warning. This implies that the latter will receive the greatest weight among all factors. Although this does not require denial of a request when timely warning is not clearly determinable, the language suggests that, in the absence of a clear determination that timely warning will indeed be provided, a strong combination of other factors would be necessary to compensate for this weakness in safeguards. It then follows that a decision to approve reprocessing in the absence of such a clear determination would be an unusual event . . ."[18] (underlining added).

Senator Glenn subscribes to a view of the function of IAEA safeguards, as indicated in the underlined passages, that Gilinsky and others have held since as early as 1970.[19] Commissioner Gilinsky, whose views closely parallel those of the Carter Administration, reiterated the point in his statement to the Senate Energy Committee:

"Safeguards . . . are an alarm that warns of illicit activity. To be effective the warning has to come in time for us to do something about it."[17]

It is important to recognize the difference between the US position and IAEA objectives. The US position is that reprocessing should not be permitted if safeguards cannot provide "timely warning" of diversion, as defined by Senator Glenn. Where the US has control, such as in the approval of the reprocessing of US-origin spent fuel, the consequence of inability to meet the US requirement for timely warning is to forbid reprocessing. The consequences for the IAEA are entirely different. It is obligated to apply safeguards to the best of its ability under the current state of the art. It may adopt the "timely-warning" criterion as an ultimate goal, but it can hardly adopt it as a requirement that must be satisfied before safeguards can be considered adequate, if the requirement cannot be met. Unlike the US, the IAEA cannot prohibit reprocessing, and it must safeguard those operations as it is best able to do so. The views as to what constitutes "adequate" safeguards that are reflected in the Non-Proliferation Act can form a basis for action by the US, if it chooses to adopt them as criteria for approving exports. They are not valid, and are in fact counterproductive for the IAEA in discharging the responsibilities assigned to it, if the requirements cannot be met.

## THE NEED FOR A MORE RATIONAL BASIS FOR IAEA SAFEGUARDS GOALS

The performance criteria currently used by the IAEA in the design of safeguards arrangements are rationalized in part by external objectives and in other instances by internal considerations. The external objectives have a superficial logic which upon closer examination is much less compelling. As the safeguards tasks become more difficult, with the advent of large-throughput bulk processing plants involving plutonium, a critical reexamination of the present criteria seems appropriate.

Let us preface the discussion by recalling the point made earlier that the objective of safeguards is to verify compliance with safeguards agreements, and that the assurance provided by that verification depends in part on the perceived capability to detect a diversion if it should be attempted. The perceived capability also has value as a deterrent. Finally, the capability defines an alarm level upon which States can plan their response, which may be to seek confirmatory information from other sources or to take more direct action to counter the diversion. The criteria we are discussing here apply to the capability for detection, which is a function that supports the primary objective of verification.

Before specific criteria are considered, the roles of external and internal goals for the IAEA should be clarified. Internal performance goals are necessarily related to feasibility; they are technical, and they govern the performance of the technical system in providing the capability for detection. External goals are political; they are determined by the broader purposes that the safeguards operation is intended to serve: providing assurance that non-proliferation commitments are honored and deterring violations of them.

We suggest the following structure for technical-system performance criteria:

1. Current and near-term performance standards should be based on what it is feasible to do. The current standard is the reference point for negotiation of an agreement, and the IAEA is obligated to apply whatever measures current technology will permit. It can do no more.

2. The standards should not be static. To stimulate improved capabilities and to enhance future safeguards performance, standards could be scheduled for higher levels in the same way that progressively-tighter air quality and automobile economy standards are scheduled by law to become effective at future dates.

3. External goals are needed to guide the evolving internally-driven performance standards and to drive R&D. Quantitative external goals need not be universal, fixed values. There are no clear limits for external goals that are based on external considerations alone;* the ultimate goal is perfection. The significance of a diversion sufficient for one explosive depends to a large extent on who the diverter is. It must be recognized that any threshold of detection capability would in principle allow a sufficient amount of material for the fabrication of one or more explosive devices to escape detection, if long-protracted or multiple diversions are assumed. External goals can, however, provide a logical basis for balancing performance goals in the planning of R&D programs.

## INTERACTIONS AMONG QUANTITATIVE PERFORMANCE GOALS

Various measures of performance interact so that improvement of one is often at the expense of another that is of possibly greater importance. While we have suggested that performance standards should be based on what is feasible, external consequences should be considered in balancing tradeoffs.

Consider the example of allowable probabilities of failure to detect and false alarm. These are the uncertainties shown in Figure 2a. There is no mention of allowable uncertainty in the safeguards documents, INFCIRC/66 Rev. 2 or INFCIRC/153. The Safeguards Manual[20] addresses uncertainty in the statistical terms of "probability of detection" and "confidence level of detection", cites recommendations of experts from member States

---

*For example, if the goal is to detect a single diversion of one significant quantity, what about two diversions of half that threshold? Or, a protracted diversion of 13 months rather than 12 months?

of values between 90% and 99% for each, and states that in most cases 95% will be used. As shown in Figure 3, those goals require a capability to detect a diversion of 3.3 times the combined standard deviation of the State's and the IAEA's measurements. For a given diversion (expressed in units of the standard deviation--in the present case, $3.3\sigma$), a point on the curve represents a particular combination of false alarm and detection probability. Tradeoffs between them can be made by sliding along the curve. The actual size, in kilograms, of the diversion that can be detected with a given probability and false alarm rate depends on the internal capability of the safeguards system--i.e., on the precision of the measurement system. The selection of a particular operating point on a curve, however, should be based on external consequences; that is, on the relative consequences of a false alarm and of a failure to detect a diversion of a given magnitude.

It is not apparent that the 95%/95% criteria have any relation to external consequences. Such cutoff limits are appropriate, for example, for quality-control limits in a manufacturing operation, where the consequences of a "false alarm"--rejection of an acceptable item--can be readily assessed and accepted in terms of economic cost. In the case of international safeguards, however, if a false conclusion of diversion is reached and immediate action is taken in response, driven by US concerns over timeliness, for example, the resulting political impact on the IAEA could be very serious indeed. The potential magnitude of the problem is illustrated by projections for the international nuclear industry, according to which between 50 and 100 large reprocessing, fabrication, and enrichment plants may be in existence by the year 2000. A material accountability system with a positive alarm limit of $2\sigma$ would then produce one to two false alarms per year, purely on the basis of statistics. From the past history of US plants handling strategic special nuclear material, the actual alarm rate can be expected to be much higher, due to human failures, blunders, unanticipated processing conditions, unsuspected sources of error, etc. Even the purely statistical rate, however, would probably be intolerable.

Another consideration is that the need to avoid the consequences of what may be a false alarm would increase the uncertainty in concluding that an actual diversion had occurred. The resulting uncertainty would necessarily raise the threshold of detection, which would tend to reduce the deterrent value. At the same time, insistence on too-stringent goals for detection probability is likely to result in compromise of other goals on the basis of subjective judgment.

For these reasons, it would probably be desirable eventually to reduce the false alarm probability substantially below the 1 to 10% rates cited in the Safeguards manual.[20] From Figure 3, reducing the false alarm probability from 5% to, say, 0.1% would reduce the detection probability for a $3.3\sigma$ diversion from 95% to 48%. A further reduction of the false alarm rate to 0.01% would reduce the detection probability to 34%, but at a slightly higher diversion level of $4\sigma$ the detection probability at this same very low false alarm rate would be as high as 60%.

Some have suggested that the statistical false-alarm rate resulting from measurement uncertainties could be reduced by containment and surveillance. That is, an apparent MUF would not be assessed as a diversion unless it were confirmed by indications from surveillance. While a <u>positive</u> indication might reinforce a conclusion from the MUF that a diversion <u>had</u> occurred, it appears unlikely that the absence of such an indication could provide sufficient basis to reject a conclusion based on MUF. Such confidence in the reliability and completeness of containment and surveillance, to detect <u>any</u> diversion, is not justified. Apparent diversions can be identified as false alarms only by a determination that the material is actually present, which leaves us with the statistical false-alarm rate of the measurement system as the irreducible minimum, short of additional sampling and measurement.

Sanborn[21] has pointed out another problem with low false alarm and high detection probabilities, namely that they are hard to demonstrate in a rigorous mathematical manner. There are two reasons. First, the central limit theorem, on which the assumption of normality of the sample distribution is based, cannot be relied on in the region of the tails of the distribution, for small sample sizes. Second, the data for the distribution are being supplied by the diverter himself, and hence any assumption of normality may be invalid.

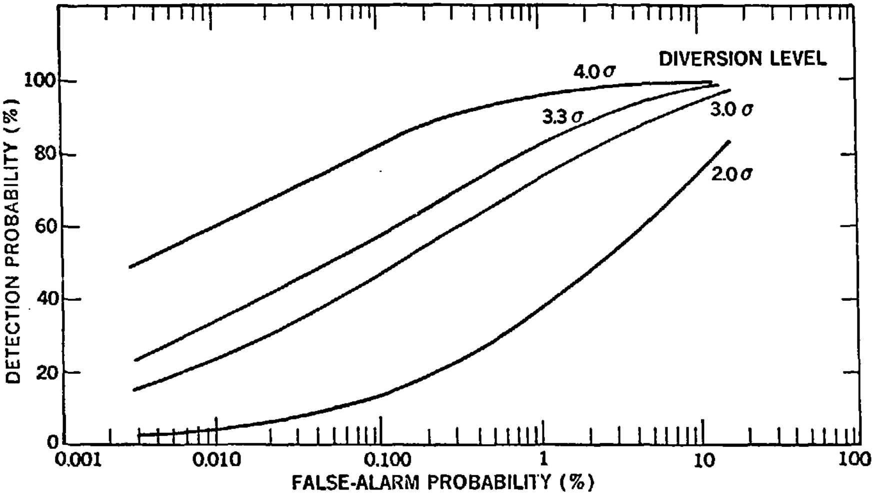


FIGURE 3. DETECTION PROBABILITY VS ONE-SIDED FALSE-ALARM PROBABILITY FOR VARIOUS DIVERSION LEVELS

If the measurement precision cannot be improved, lower false-alarm rates can be achieved only at the expense of some other, externally-derived, goal: the threshold quantity to be detected or the probability of detection.* What is the consequence of accepting a lower probability of detection for diversion of a threshold quantity?

Let us first consider probability of detection for a specified minimum quantity to be detected. What are the consequences of a too-low probability of detection? The perceived capability of detection affects the deterrent value of safeguards, and less directly, the degree of assurance that compliance with non-proliferation agreements is verified. Deterrent value depends on the disincentives that are the consequence of being caught, as well as the probability of detection. What sanctions might be imposed as a consequence of an IAEA finding of diversion is beyond the scope of safeguards. We may assume, however, that the consequences of a State being detected in a deliberate attempt to secretly violate its non-proliferation pledges would be sufficiently grave that the State would not take any significant risk of detection. For that reason, it seems very unlikely that a State that would be deterred by a 95% probability of detection would accept a risk on the order of 50% or even less. The external consequences of detection and false alarm are entirely different, and much lower probability of detection than probability of no false alarm could be accepted.

The consequences for the IAEA of requiring high probabilities of detection must also be considered, in view of the increased sampling requirements. In Figure 4, from Sanborn[21], the fraction of the total population required to be sampled by the inspector is plotted as a function of detection probability and the number of items F falsified by the diverter. Consider, for example, the curve for F=8 falsified items. For a detection probability of 0.5, approximately 8% of the total population must be sampled; for a detection probability of 75% the sample size must be 15%, and for detection probabilities of 95 and 99%, the sample sizes must be approximately 30 and 43%, respectively. Increasing the detection probability from 50 to 95% therefore quadruples the sample size, while going to 99% increases it, overall, by a factor of more than five. These are very considerable increases in inspection effort, and one may well question whether they are worth it, in view of the cost in limited IAEA resources and the dubious value of high probability of detection.

As Figure 3 shows, the probability of detection is related to the size of the diversion, for a fixed probability of false alarm. Let us examine the question of the minimum quantity that must be detected in terms of external consequences of failure to detect and internal constraints on the ability to detect.

IAEA safeguards are based on the "significant quantity" of material necessary to fabricate a single explosive device, which they define as 8 kg of plutonium or 25 kg of highly-enriched uranium. Goals are based tentatively on the detection of one "significant quantity" in a single, abrupt diversion or in a protracted diversion spread over one year. As we have suggested, there is no rational basis for a detection capability based on a minimum quantity corresponding to a single explosive device, insofar as diversion by the State is concerned. In the case of subnational diversion it has more significance, and it is possible that thinking based on subnational diversion has been carried over. The external consequences of national diversion that are important arise from the detection of <u>any</u> diversion that indicates covert violation of non-proliferation agreements, and the actual amount has less significance than is generally supposed.

From the practical standpoint, it is necessary to establish some threshold quantity of detection, not on the basis of external consequences, but as a measure of system sensitivity. The threshold should therefore reflect system capability rather than any external consequence of diversion. The IAEA's "significant quantity" may be satisfactory, provided that it is in the range of feasibility. Feasibility limits will be exceeded, however, in the case of large-throughput bulk processing plants. A significant quantity represents a very small fraction of the quantities that must be measured, as indicated in

*These two goals can be balanced against each other. They are not independent, and both are measures of detection limit.

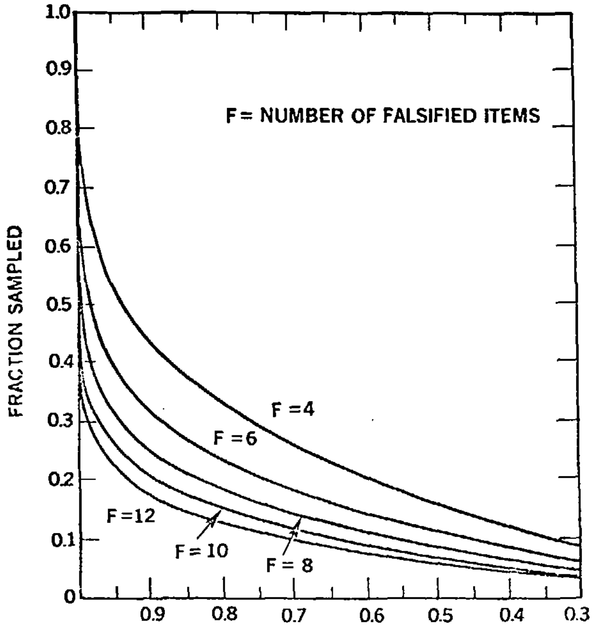


FIGURE 4. PROBABILITY OF DETECTION OF A FALSIFIED ITEM AS A FUNCTION OF SAMPLE SIZE

Table I. It is assumed here that present US NRC regulatory requirements for commercial reprocessing plants, namely that LEMUF(Pu) $\leq$ 1% of throughput for material balances every 6 months, can be met. It is also assumed that the uncertainties in the successive material balances add randomly, so that the uncertainty for a year is just $\sqrt{2}$ times the uncertainty for 6 months, which understates the actual error. The plutonium fraction is assumed to be 1% and 8%, respectively, for LWR and LMFBR spent fuel. A 200-te/yr LMFBR reprocessing plant is comparable to a 1500 te/yr LWR plant in terms of plutonium flows, and each is a likely size for economic operation.

As can be seen, for LWR reprocessing plants of $\geq$ 100 te/yr throughput the annual uncertainty is about equal to a significant quantity, while for a plant the size of the NFS West Valley plant it amounts to two or three significant quantities. These may be regarded as in the pilot-plant range, or typical of what a small industrialized or fairly large semi-industrialized country might want to support recycle. On the other hand, a 1500-te/yr LWR reprocessing plant would have a material balance uncertainty at least an order of magnitude larger than a significant quantity. It is therefore these large economically-sized plants that are the problem, as far as material accountancy is concerned, not the smaller ones in which some developing countries have expressed an interest.

## ABSOLUTE VS. RELATIVE GOALS

The preceding discussion raises the important question of whether, in plants like the ones considered, safeguards measurement goals should be expressed in terms of absolute or of relative quantities. There are two difficulties with absolute goals, one of which is the inability to meet them with large plants. The second difficulty is a related one. A rigid insistence on the ability to detect the diversion of one significant quantity over a period of one year would favor the building of small, relatively-uneconomical reprocessing plants. Plants of this kind could be ordered by developing or semi-industrialized nations with a weapons program in mind, and justified on the grounds that they meet safeguards measurement requirements. On the other hand, a large industrialized country, if it wanted to reprocess, would be forced to build a number of small plants, each capable of meeting the requirements but in the aggregate providing essentially the same opportunities for diversion as a single economical plant of the same total size. So in the first case the policy of absolute requirements would favor the very situation the US is most concerned about; that is, the spread of small reprocessing plants to developing countries, while in the second case it would favor an economically-irrational development without appreciably reducing the risk of undetected diversion. Insistence on absolute goals could be regarded, in fact, as discriminatory against large nations, since it imposes much more stringent measurement requirements on them--requirements that cannot even be met by present technology--than on small nations.

Table I

UNCERTAINTIES ($2\sigma$)* IN THE
MATERIAL BALANCES OF LWR AND LMFBR REPROCESSING PLANTS
AS A FUNCTION OF THROUGHPUT

| | Uncertainty (kg Pu per year) | |
|---|---|---|
| te HM/yr | LWR | LMFBR |
| 10 | 0.7 | 5.7 |
| 100 | 7 | 57 |
| 200 | 14 | 114 |
| 300 | 21 | |
| 1500 | 105 | |

*Assuming two material balances per year with random measurement errors only, and $2\sigma$ = 1% of throughput.

The consequences of the "significant-quantity" limit, when applied in the US export-approval conditions, will be contrary to US policy objectives. The large processing plants that fail to meet the measurement goals will, for the foreseeable future, all be located in large industrial countries. A credible nuclear weapons program in such a country would have to involve tens and probably hundreds of weapons. A mere half dozen or so weapons in the hands of a country like West Germany or Japan, for example, would not be very useful to it. Furthermore, such countries are quite capable, if they considered it necessary, of developing a large clandestine military program tailored to their specific needs. It is therefore highly unlikely that they would take the risk of diverting material from a safeguarded facility when the amounts to be gotten this way are so limited.

On the other hand, small semi-industrialized or developing countries are likely, for the foreseeable future, to operate only the smaller plants, in which the detection goal for material accountability can be met. The effect will therefore be to permit reprocessing of spent exported fuel in countries the US is most concerned about, while forbidding reprocessing by western-European States and Japan, where proliferation risks are minimal.

The proper balance of emphasis in reprocessing safeguards could be accomplished by defining limits of detection capability on the basis of a fraction of throughput rather then in absolute, "significant" quantities. In small plants, where it is feasible, the significant-quantity limit could be retained. A limit expressed as a small fraction of throughput would provide good assurance that the major countries having such plants did not divert, since, although the detection limit might not allow enough for several explosives to be detected, the quantity would be insignificant in military terms and would therefore provide credible assurance that no diversion at all had been attempted.

## TIMELINESS GOALS

Risks of false-alarm/detection failure can be balanced against timeliness as well as threshold-detection levels. Up to times on the order of a few months, improved timeliness in safeguards findings can only be achieved at the expense of uncertainty. A goal of extremely short "detection" time displays an oversimplified view of what "detection" really is. As discussed above, "detection" of diversion by the safeguards system is a final conclusion of a process that begins as an attempt to verify compliance with the safeguards agreement. A finding of diversion, if there is any serious attempt by the State to conceal the diversion, is arrived at by a process of elimination. Any abnormalities must be investigated to see if there are other causes than diversion. While strong indications might lead the inspectors to report directly, as provided for in the agreement, it seems unlikely that with covert diversion there would be sufficient basis for immediate and vigorous responsive action until all available information had been thoroughly assessed.

What is loosely referred to as "detection" is the conclusion of a process of analysis and investigation, beginning usually with routine inspection and collection of samples for analysis, and ending after a number of alternative explanations have been investigated and rejected. There is no single indication or event short of the final decision that could be called "detection". Rather, from the first indication there is the possibility of diversion, but with a high degree of uncertainty. As the assessment progresses, uncertainty diminishes but time passes. It is as if a kind of Heisenberg uncertainty principle were at work: as detection time is reduced the "detection" becomes more ambiguous, and the resolution of uncertainty is at the expense of timeliness. The provisions contained in INFCIRC/153[1] to resolve any uncertainties, which could require considerable time, illustrate the tradeoff between time and uncertainty.

Thus it appears that detection time can be reduced to the extreme degree considered necessary by some US policymakers only if it is done at the expense of increased false-alarm probability or if the diversion is very large. Since the purpose of very short detection times is to allow immediate diplomatic action to be taken before conversion to weapon components could be completed, a high degree of confidence in the finding of diversion would also be necessary. False-alarm limits would have to be more stringent, not less, if timely detection is to lead directly to responsive action.

The logical basis that rationalizes the need for extremely short detection times calls for closer examination. Implicit in the timeliness requirement, as stated in the Nuclear Non-Proliferation Act, is the assumption that when conversion time for one explosive is completed some kind of point of no return is reached, after which corrective diplomatic action is pointless or ineffective, and the diverting State must henceforth be accepted as a member of the nuclear club. It is not at all apparent or self-evident why that should be so. We are aware of no discussion of that fundamental assumption in all the voluminous literature on non-proliferation. It is accepted as a rule of some kind of game, which defines winning or losing in terms of who wins the race at the end of conversion time. A number of arguments against that point of view are offered:

(1) The non-proliferation objective is the prevention of the acquisition of nuclear armament by another State, and the threshold is a nuclear-weapon capability with effective delivery systems, not the final assembly of unit no. 1. This process affords considerably more time and opportunity to detect a clandestine weapons program, through both international safeguards and national intelligence activities.

(2) An important and officially stated objective of the IAEA is to deter diversion, a function which depends on the perception of the risk of detection and the probable consequence of discovery and not on the time at which a violation is discovered. Since the nature of the response of the international community to a violation would probably not be any different whether or not conversion had occurred (even assuming that there was some way of knowing whether it had), the deterrent value should be the same in either case. This point can be illustrated by the example of India. There, knowledge of preparations before the test explosion did not lead to the application of effective response measures, and pressures to restrain them continue more than five years later. If the detection-time/conversion-time dependency were correct, such pressures after the test explosion would be pointless. In fact, it would be harmful to US objectives to convey the impression to a potential diverter that once he converts the material into a weapon he is "home free".[22]

(3) The objective of diplomatic pressure would presumably be to restore the status quo ante; that might include return of the diverted material and assurance that the weapon program itself had in fact been terminated. The accomplishment of neither of these objectives depends on whether the material has been fabricated into a weapon, since, clearly, it can be returned at any stage and, at least in principle, there is no clearly delineated point or threshold beyond which a weapons program cannot be terminated.

(4) The claimed shortness of conversion times robs the first weapon assembly of much of its significance, since it implies that the fabrication of a weapon is easy, even for a nation with no prior experience. If so, it cannot have great significance as a threshold.

(5) The concern over a single explosive is much more valid in the case of subnational diversion, where the threat is more immediate and the significance of additional numbers of weapons relatively less. The preoccupation with the minimum time to fabricate a single explosive seems to be the result of a blurring of the distinction between physical protection against subnational groups and safeguards against national diversion.

Thus, while timely warning is certainly of value, there is no apparent logical basis for requiring that detection time--the time from diversion to detection--not exceed conversion time, or that there is necessarily any direct, quantitative relationship.

## SUMMARY

Currently, the tentative goals for IAEA safeguards performance lack internal consistency and a sound logical basis. In particular, it is argued that there is no rational foundation for the claimed connection between detection time and conversion time; that present goals for false-alarm probability will most likely be unacceptable for the expanded nuclear industry of the future, while goals for detection probability are set unnecessarily high; and further, that rigid insistence on absolute thresholds for detection that are unattainable in high-throughput bulk processing facilities may be counterproductive.

A distinction must be made between external safeguards goals, based on the consequences of diversion, and internal goals that govern IAEA operations, if attainment of the external goals is not feasible. The IAEA should adopt quantitative performance goals that are clearly feasible, and upgrade them as the state of the art advances. This can be done in accordance with a prearranged schedule as is done in some US environmental regulations. External goals can provide balance in the establishment of IAEA goals, but they should not be adopted as quantitative requirements unless they are feasible.

IAEA safeguards are too often viewed in negative terms, a view that is more appropriately applied to domestic safeguards. Domestic safeguards agencies are regulatory agencies; the IAEA is not. The IAEA, like other international agencies, is based on respect for agreements that are accepted voluntarily. IAEA safeguards are a great step forward in the widespread acceptance of in-country monitoring of compliance with agreements. Expectations of what IAEA safeguards can accomplish, as reflected in stated performance goals, should be in terms of positive assurances that non-proliferation commitments are respected, and not in terms of a regulatory body that can compel compliance.

A capability to detect diversion is an essential element in the safeguards system; it establishes the level of credibility of assurance that safeguards provide. It is not reasonable, however, to expect that a detection capability that combines a high probability of detection with extremely short detection times can at the same time provide hard enough information to justify immediate and drastic diplomatic response. Such an expectation would appear to be based on a grossly-oversimplified view of what constitutes detection of diversion. Externally-derived performance goals should anticipate more realistic responsive actions. Those might include verification from other sources of early, ambiguous alarm signals, or more deliberate actions on the basis of carefully-assessed findings that require longer times.

In practice, the IAEA must be governed by internally-derived goals--what it is possible to do, now and in the future. The ability of the IAEA to accomplish what safeguards can do must not be compromised by striving for unattainable external goals, which can only lead to system distortions and less-than-optimum performance. External goals are needed to guide safeguards R&D, but not to define requirements that safeguards must meet in order to be considered adequate.

## References


1. The Structure and Content of Agreements Between the Agency and States Required in Connection with the Treaty on Nuclear Weapons, INFCIRC/153, IAEA, May 1971, par. 7, 31, 32.

2. Treaty on the Non-Proliferation of Nuclear Weapons, 1968.

3. Treaty for the Prohibition of Nuclear Weapons in Latin America, 1967.

4. Statute of the International Atomic Energy Agency, 1956, Article III.

5. The Agency's Safeguards System (1965, As Provisionally Revised and Extended in 1966 and 1968), INFCIRC/66, Rev. 2, IAEA, September 16, 1968, par. 46.

6. INFCIRC/153.

7. Rometsch, et. al., "Safeguards -- 1975-1985", IAEA-SM-201/103, Safeguarding Nuclear Materials, Proceeding of a Symposium, Vienna, 20-29 October, Vol. I, p. 5.

8. Paul C. Szasz, "The Law and Practices of The International Atomic Energy Agency", Legal Series No. 7, IAEA, 1970, Note 6, p. 637 and Note 67, p. 639.

9. INFCIRC/153, par. 28.

10. INFCIRC/66 Rev. 2, par. 46.

11. INFCIRC/66 Rev. 2, par. 14.

12. Szasz, p. 558.

13. INFCIRC/153, par. 6.

14. INFCIRC/153, par. 4, 5.

15. The Nuclear Non-Proliferation Act of 1978, Public Law 95-242.

16. Prof. Albert Wohlstetter, as quoted in the Congressional Record, February 2, 1978, p. S1077.

17. NRC Commissioner Gilinsky, Statement before the Senate Energy Committee, Congressional Record, February 21, 1978, p. S1077.

18. Senator John Glenn, Congressional Record, February 7, 1978, p. S1310.

19. Gilinsky and Hoehn, "Non-Proliferation Treaty Safeguards and the Spread of Nuclear Power, R-501, Rand Corporation, May 1970, p. 9.

20. International Atomic Energy Agency, IAEA Technical Safeguards Manual, IAEA-174, 1976, par. 5.1.3.

21. Jon Sanborn, Brookhaven National Laboratory. Personal communication.

22. Myron B. Kratzer, International Energy Associates Limited. Personal communication.