EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

ELSEVIER European Journal of Operational Research 90 (1996) 383-394

Invited Review

Inspection games in arms control

Rudolf Avenhaus [a,*], Morton Canty [b], D. Marc Kilgour [c], Bernhard von Stengel [a], Shmuel Zamir [d]

[a] *University of the Federal Armed Forces Munich, D-85577 Neubiberg, Germany*
[b] *Research Center Jülich, D-52425 Jülich, Germany*
[c] *Wilfrid Laurier University, Waterloo, Canada*
[d] *The Hebrew University, Jerusalem, Israel*

Received August 1995

Abstract

An inspection game is a mathematical model of a situation in which an *inspector* verifies the adherence of an *inspectee* to some legal obligation, such as an arms control treaty, where the inspectee may have an interest in violating that obligation. The mathematical analysis seeks to determine an optimal inspection scheme, ideally one which will induce legal behavior, under the assumption that the potential illegal action is carried out strategically; thus a non-cooperative game with two players, inspector and inspectee, is defined. Three phases of development in the application of such models to arms control and disarmament may be identified. In the first of these, roughly from 1961 through 1968, studies that focused on inspecting a nuclear test ban treaty emphasized game theory, with less consideration given to statistical aspects associated with data acquisition and measurement uncertainty. The second phase, from 1968 to about 1985, involves work stimulated by the Treaty on the Non-Proliferation of Nuclear Weapons (NPT). Here, the verification principle of material accountancy came to the fore, along with the need to include the formalism of statistical decision theory within the inspection models. The third phase, 1985 to the present, has been dominated by challenges posed by such far-reaching verification agreements as the Intermediate Range Nuclear Forces Agreement (INF), the Treaty on Conventional Forces in Europe (CFE) and the Chemical Weapons Convention (CWC), as well as perceived failures of the NPT system in Iraq and North Korea. In this connection, the interface between the political and technical aspects of verification is being examined from the game-theoretic viewpoint.

Keywords: Arms control; Game theory; Inspection; Attribute sampling; Variable sampling

1. Introduction

In the context of arms control, inspections are procedures designed to provide data with which an agent's compliance to an agreement (or other set of rules) can be assessed. There is always, potentially at least, a conflict between the inspection authority and the agent (state, organization, or person) required to comply. Of course, if the agent were not tempted to violate the agreement, then inspections would be unnecessary. It is thus natural that quantitative models of inspections should be non-cooperative games with two players, *inspector* and *inspectee*.

Inspection games should be distinguished from two related topics: Inspections for *quality control*, or for

* Corresponding author. E-mail: avenhaus@informatik.unibw-muenchen.de

SSDI 0377-2217(95)00261-8

prevention of other kinds of random accidents, for which there is no adversary who acts strategically; and inspections that are *search problems*, where an adversary attempts to escape a searcher with well-defined and legitimate strategies, like a submarine escaping a destroyer in war. Neither situation is described by an inspection game in our sense; for us, the salient feature is that the inspector tries to prevent the inspectee from behaving illegally in terms of the agreement. In other words, our inspectee might decide not to violate, so that there is nothing to search for; such *deterrence* is generally a high priority goal for the inspector.

Inspections cause conflicts in many real world situations. In economics, this a central problem of *principal-agent* relationships where the principal (e.g. employer) delegates work or responsibility to the agent (employee) and chooses a payment schedule that best exploits the agent's self-interested behavior. The agent, of course, chooses his action so as to maximize his own utility given the fee schedule proposed by the principal. *Environmental* agreements obviously give rise to inspection problems, but these have not yet received as much attention from modelers as one might have expected. To date, most methodological analyses of inspection games have been performed in the context of *arms control and disarmament* (ACD). This review therefore focuses on ACD inspection games.

Immediately after von Neumann and Morgenstern's pioneering book (1944) *Theory of Games and Economic Behavior*, ACD inspections may have been analyzed game-theoretically as classified military research; this is not known for sure but may be inferred from papers published later. Non-classified work started vigorously in the early 1960s with analyses for ACDA, the United States Arms Control and Disarmament Agency. These dealt with very general ACD problems, and also with concrete problems of test ban treaty verification, as surveyed in detail below. A second phase of inspection game development started around 1968 in connection with the verification of the Treaty on the Non-Proliferation of Nuclear Weapons (NPT). Finally, the advent of a new series of ACD treaties in the mid-1980s, including the treaties on Intermediate Nuclear Forces (INF) and on Conventional Forces in Europe (CFE), along with the Chemical Weapons Convention (CWC), opened a new era of research which continues to flourish today.

The exposition below, covering more than thirty years of ACD inspection analyses, generally follows this historical arrangement. Our survey is naturally partial. To convey the flavor of the field, we present a characteristic model in some detail in each section. Different models of this kind are presented by Avenhaus, von Stengel and Zamir (1996).

2. The beginnings: test ban treaties

Probably the first genuine inspection game in the open literature was the recursive game developed by Dresher (1962). We present it in some detail since it was seminal for later work. The inspector has to distribute a limited number m of inspections over n stages. At each stage, the inspector may or may not use an inspection. The inspectee may decide at a stage to act legally or illegally, and will not perform more than one illegal act throughout the game. Illegal action is detected if and only if there is an inspection at the same stage. An inspection that has taken place is observed by the inspectee.

Dresher modeled this situation as a two-person zero-sum recursive game. The inspector's payoff is $+1$ unit for a detected violation, zero for legal action throughout the game, and -1 unit for an undetected violation. The (minmax) value of the game, the equilibrium payoff to the inspector, is denoted by $I(n,m)$ for the parameters $0 \leqslant m \leqslant n$. For $m = n$, the inspector will inspect at every stage and the inspectee will act legally, and similarly the decision is unique for $m = 0$ (and $n \geqslant 1$) where the inspectee can safely violate, so that

$$I(n,n) = 0 \quad \text{and} \quad I(n,0) = -1 \quad \text{for } n > 0. \tag{1}$$

For $0 < m < n$, the game is represented by the recursive payoff matrix shown in Table 1. The rows denote the inspector's possible actions *at the first stage*, and the columns the inspectee's. If the inspectee violates, then he is either inspected and caught, so the game terminates and the inspector receives $+1$, or not, in which case he will act legally throughout, so that the game eventually terminates with payoff -1 to the inspector. After a legal action of the inspectee, the game continues as before, with $n-1$ instead of n stages

Table 1
The Dresher game, showing the decisions at the first of n stages, with at most one intended violation and m inspections, for $0 < m < n$. The game has value $I(n,m)$. The recursively defined entries denote the payoffs to the inspector

Inspector	Inspectee	
	legal action	violation
inspection	$I(n-1,m-1)$	$+1$
no inspection	$I(n-1,m)$	-1

and $m-1$ or m inspections left (the underlying assumption being that inspections are observable by the inspectee).

Since the inspector prefers to have more inspections (i.e. $I(n,m) > I(n,m')$ for any $n \geqslant m > m' \geqslant 0$, a fact that will be confirmed by the solution of the game), Table 1 implies a circular structure of the players' preferences. That is, the inspector prefers to use his inspection if and only if the inspectee violates, while the inspectee in turn prefers to violate if and only if the inspector does not inspect. This means that the game has a unique mixed Nash equilibrium. Furthermore, at this equilibrium, both players choose both of their actions with positive probability. This requires that their expected payoffs for both actions be the same. If the inspector's probability for inspecting at the first stage is p, then the inspectee is indifferent between legal action and violation if and only if

$$p \cdot I(n-1,m-1) + (1-p) \cdot I(n-1,m) = p + (1-p) \cdot (-1),$$

and both sides of this equation are equal to the game value $I(n,m)$. Solving for p and substituting yields

$$I(n,m) = \frac{I(n-1,m) + I(n-1,m-1)}{I(n-1,m) + 2 - I(n-1,m-1)}.$$

With this recurrence equation for $0 < m < n$ and the initial conditions (1), the game value is determined for all parameters. Dresher showed that these equations have an explicit solution, namely

$$I(n,m) = -\binom{n-1}{m} \Big/ \sum_{i=1}^{m} \binom{n}{i}.$$

Dresher suggested several possible arms control problems as applications for his model, in particular verification of a test ban treaty. It thus became an important tool when, at the beginning of the 1960s, ACDA started sponsoring systems analysis research on ACD in general and verification in particular. Even though test ban treaties were the central application of these models, their scope became much more general.

From 1963 to 1968, prominent mathematicians and game theorists analyzed ACD problems as members of Mathematica, Incorporated of Princeton, NJ. According to the first contract (Mathematica, 1963), the objectives were "to identify and explore potential applications of statistical methodology to the inspection aspects of arms control and disarmament; to develop and analyze techniques from the disciplines of sampling, decision theory and the theory of games for application to inspection in connection with arms control and disarmament planning and negotiation; and to evaluate the adequacy of the statistical methodology and the techniques developed as bases for determining if and how desired levels of verification can be achieved in connection with various arms control and disarmament measures."

The Mathematica publications are quite substantial, but are unfortunately rather inaccessible. One report, Mathematica (1965), is completely anonymous, and one can only infer the authors of its papers. After a few general remarks, we will survey this work in detail. First, central to the Mathematica research program was the antagonistic situation between the United States and the Soviet Union at the height of the Cold War. Second, the Mathematica contributions became progressively more mathematical and abstract. Nevertheless, at least in the initial phase, the political problems were approached from both practical and the mathematical points of view. Finally, no specific applications were given, perhaps because negotiations were under way, so that possible applications could not be cleared for publication. For the sake of abstractness in this sense, or for simplicity, false alarms caused by measurement errors were not taken into account explicitly. False alarms later became extremely important for nuclear material safeguards.

In the following, we discuss only Mathematica papers pertaining to inspection games. There are three categories of papers (not in chronological order), which describe specific inspection problems, analyze general features of ACD verification, and deal with extensions of Dresher's game, respectively.

Initially, special non-sequential games were developed by Anscombe and Davis (Papers 3 and 4 in Mathematica, 1963) to model inspection against clandestine rearmament: The controlled unit (plant, region) is subdivided into K subunits. The inspectee, assumed to undertake clandestine rearmament, selects subunits for that purpose, trying to attain a certain global arms potential. The inspector randomly selects certain subunits for inspection, subject to a limited total inspection effort. With these strategies, a zero-sum game is considered with the probability of detecting at least one illegally rearmed subunit as payoff to the inspector. One paper presents the model and discusses assumptions and certain qualitative aspects of the solution. A second paper gives approximate quantitative solutions of the game.

It may be questioned whether these games have been used for the intended real world inspection problems. Later, however, similar games were developed independently (Avenhaus, 1986) in connection with nuclear safeguards and the NPT (see Section 3). There, classes (*strata* in statistical terminology) instead of subunits are to be verified with the help of attribute sampling techniques; similar formulae for the distribution of inspection effort were obtained.

As mentioned, these models are zero-sum games with the probability of detection as payoff to the inspector. This is intuitive for specific situations but does not always meet the intent of ACD verification, since the inspector's highest priority is usually legal behavior of the inspectee (that is, successful deterrence) rather than uncovering illegal behavior. A second category of papers addresses these questions. At an early stage, Maschler (Papers 9 and 10 in Mathematica, 1963), and later Harsanyi (Paper 1 in Mathematica, 1966) represented the preferences of inspector and inspectee by individual utility functions that were not always zero-sum. If one normalizes the payoffs to both players to be zero for legal behavior of the inspectee and its recognition by the inspector, then a detected illegal action will yield negative payoffs to both since this is not desired by either player. Only if the inspectee acts illegally and detection is the matter of concern are the players' interests antagonistic. In that case the zero-sum game with the probability of detection as payoff to the inspector is adequate.

In the first of two papers by Maschler (Paper 9 in Mathematica 1963), a non-constant-sum game is discussed as a model for deciding whether a nuclear test ban treaty should be signed. This seems to be the first investigation of bargaining in the context of ACD. Inspections matter only for the different possible treaties (signed treaty with or without inspections, etc.). The approach demonstrates impressively the quantification of verbally described situations and the possibility of drawing conclusions even in the absense of specific utility functions.

Harsanyi (Papers 1 and 4 in Mathematica, 1966) later extended these considerations. In his first contribution, he discusses and explains aspects of utility theory, such as aggregated preferences, risk and uncertainty, bounded rationality, and preference elicitation. Similarly, game theory is discussed in general terms, such as zero-sum, non-cooperative and cooperative games. The conclusions for ACD problems are rather general, recommending the use of non-cooperative games with incomplete information. Harsanyi's second contribution continues the discussion but without reference to inspection games.

The third category of contributions, by Kuhn and Maschler among others, represents extensions and applications of Dresher's (1962) model. Kuhn's first paper (Paper 5 in Mathematica, 1963) generalizes Dresher's model for monitoring a test ban treaty. An assumed random number n of seismic events (referred to above as stages) are considered. Each one might be due to an earthquake (E) or a nuclear test (T). An earthquake produces the (correct) signal E with probability $1 - p$ and a test generates the signal T with probability $1 - q$. With probabilities p and q, respectively, a doubtful signal (D) is generated. The inspector may use m inspections in total. The zero-sum payoffs are the same as Dresher's in Table 1. Kuhn presented analytical solutions of this model and some considerations for the case of more than one test.

Following up on his more general study mentioned above, Maschler (Paper 9 in Mathematica, 1963) extended Dresher's model by introducing non-zero-sum payoffs. Papers 9 and 10 of Mathematica (1965) introduce a probability q of failing to detect a test in an inspection. (However, false alarms are not included, so the model is not genuinely statistical.) An important concept introduced in these papers (later published in Maschler, 1966, 1967) is *inspector leadership*, which gives the inspector the commitment power to announce his (randomized) inspection strategy. This is appropriate for the asymmetric situation of inspections, in which the inspectee cannot announce his intentions to behave illegally. The announced optimal inspection strategy induces the inspectee to *legal behavior* for certain (as long as there are inspections left), which cannot be achieved without such an announcement.

These models were extended in the mid-1980s in the context of new ACD developments. A recursive zero-sum game with the number of intended violations as an additional parameter was solved explicitly by von Stengel (1991). Rinderle (1995) showed formally that the solution of the leadership game obtained by Maschler (1966) is indeed a Nash equilibrium, thus removing some unnecessary assumptions, and introduced false alarms, with the false alarm probability as a strategic variable of the inspector.

As mentioned, the Mathematica papers dealt only partially with inspection games. Later Mathematica work, by Aumann, Harsanyi, Maschler, Selten, and Stearns among others, concerned problems of information in bargaining and repeated games, which became seminal for later game-theoretic developments. These works were recently published as a book (Aumann and Maschler, 1995).

The monograph by Saaty (1968), who was then scientific representative of ACDA, complements the Mathematica reports, with emphasis on modeling global ACD development rather than inspections.

3. Coming to maturity: nuclear safeguards

The Treaty on the Non-Proliferation of Nuclear Weapons (NPT) was inaugurated in 1968. It was intended to freeze the status quo of nuclear weapons and non-nuclear weapons states, the former pledging themselves to a long-term reduction and ultimate elimination of their nuclear arsenals, while the latter agreed not to acquire such arsenals. The International Atomic Energy Agency (IAEA) in Vienna was given responsibility to verify compliance to the treaty, in particular to inspect the peaceful nuclear activities of the member states.

In the course of subsequent negotiations of the practical implementation of IAEA inspections, *material accountancy* emerged as the basic verification principle: Through periodic comparisons of book and physical inventories at nuclear installations, a quantitative statement regarding the continued presence of nuclear material was to be made. This system requires that plant operators, via their national control authorities, report all relevant material balance data to the IAEA, while the international inspectors verify those data by making independent measurements on a random sampling basis. Extensive use of automatic surveillance equipment and seals is also made in order to minimize the number of measurements.

Inspection sampling procedures used in NPT safeguards are conventionally of two kinds, depending on the nature of the verified material: *Attribute sampling* is used to test or estimate the number of items in a population having some qualitative characteristic or attribute of interest, usually referred to as a gross defect, such as a broken seal, illegally substituted material, or a large falsification of content. *Variable sampling* on the other hand involves quantitative measurements with known precision. Each observation is totaled or averaged for the population to form a *test statistic*, for example the book-physical inventory difference, better known as MUF (Material Unaccounted For). Based on this statistic, the inspector has to decide if material has been diverted or if the result is due to measurement errors, a decision which will in part depend on his own choice of false alarm probability.

Since 1969 international conferences on nuclear material safeguards have been held regularly. Conferences on Nuclear Safeguards Technology are organized by the IAEA roughly every four years and published under that title. Annual Symposia on safeguards and nuclear materials management are also sponsored and published by the European Safeguards Research and Development Association (ESARDA) and by the American Institute for Nuclear Materials Management (INMM).

The bulk of the work published in these proceedings is concerned with practical matters, for example measurement and surveillance technology, data processing, and plant safety and security. However decision-theoretic approaches, including game theory, have been presented through the years. Monographs emphasizing theoretical aspects are Avenhaus (1986), Bowen and Bennett (1988), and Avenhaus and Canty (1996).

Game-theoretic work in this area was started by Bierlein (1968, 1969) and continued by Höpfinger (1974). Bierlein emphasized that payoffs should be represented by detection probabilities only, with inspection costs as external boundary conditions. This is an adequate model for IAEA verification, where inspection effort is limited by a fixed budget: The IAEA has no intention to minimize it further, but rather tries to make the most efficient use of available resources.

Perhaps the easiest class of inspection problems to tackle from a game-theoretic point of view is that involving random *attribute sampling*. The strategic aspect first crops up if the sampled population is stratified. The inspectee then has the freedom to distribute his gross defects over the strata, while the inspector must find an optimal number of samples for each stratum subject to his effort restrictions. One treats the detection probability as the inspector's payoff in a zero-sum game and seeks a saddle point in the combined strategy space of the protagonists. The zero-sum assumption can be justified as being part of the Nash equilibrium of a supergame in which the inspectee may decide to behave legally or illegally and the utilities of the protagonists are included explicitly. Closed solutions have been found under rather general conditions (Avenhaus, 1986). A heuristic formula widely used by the IAEA for calculating plans for attribute sampling without replacement was also derived formally by Avenhaus and Canty (1989) as the equilibrium of a leadership game.

As an illustration of attribute sampling, related to the work by Anscombe et al. (Mathematica, 1963) mentioned above, we consider an inspection of K classes of material. The ith class contains N_i items, whose data are reported to the inspector. Different classes are characterized by their batch numbers, by the measurement techniques and – related – by the efforts ε_i of the inspector for verifying one datum; the inspector has total inspection effort ε at his disposal. If he verifies n_i data in class i, these inspections are thus constrained by

$$\sum_{i=1}^{K} \varepsilon_i n_i = \varepsilon.$$

Correspondingly, we assume that the inspectee falsifies r_i data items in the ith class by the amount μ_i such that his total falsification is μ, that is,

$$\sum_{i=1}^{K} \mu_i r_i = \mu.$$

Note that the maximal total falsification $\mu_{\max}$ is given by

$$\mu_{\max} = \sum_{i=1}^{K} \mu_i N_i.$$

Based on sampling with replacement, we get the probability of detecting at least one falsified datum as

$$1 - \beta(\boldsymbol{n}, \boldsymbol{r}) = 1 - \prod_{i=1}^{K} \left(1 - r_i/N_i\right)^{n_i}.$$

Thus, when the values of N_i, ε_i, μ_i, $i = 1, \dots, K$, μ, and ε are known, we have defined a zero-sum game with the sets of strategies given by the sample sizes $\boldsymbol{n} = (n_1, \dots, n_K)$ and the falsification plan $\boldsymbol{r} = (r_1, \dots, r_K)$, and with the probability of detection $1 - \beta(\boldsymbol{n}, \boldsymbol{r})$ as payoff to the inspector.

If the sample sizes are treated as continuous variables, then the solution of the game is given by

$$n_i^* = \frac{\varepsilon}{\sum_j \mu_j \varepsilon_j N_j \exp(-\kappa \varepsilon_j)} \cdot \mu_i N_i \exp(-\kappa \varepsilon_i),$$

$$r_i^* = N_i \cdot (1 - \exp(-\kappa \varepsilon_i)), \quad i = 1, \dots, K,$$

$$1 - \beta^* = 1 - \exp(-\kappa \varepsilon),$$

where the parameter κ is uniquely determined by the equation

$$\sum_{i=1}^{K} \mu_i N_i \exp(-\kappa \varepsilon_i) = \sum_{i=1}^{K} \mu_i N_i - \mu.$$

If μ is small compared to the maximal falsification then we get the very simple expressions

$$n_i^* \approx \frac{\varepsilon}{\sum_j \mu_j \varepsilon_j N_j} \cdot \mu_i N_i,$$

$$r_i^* \approx \frac{\mu}{\sum_j \mu_j \varepsilon_j N_j} \cdot \varepsilon_i N_i, \quad i = 1, \dots, K,$$

$$1 - \beta^* \approx \frac{\mu \cdot \varepsilon}{\sum_i \mu_i \varepsilon_i N_i}.$$

These simplified expressions, already found by Anscombe, Davis, and Kuhn (Mathematica, 1963), allow an intuitive interpretation: $N_i \varepsilon_i$ is the effort for verifying *all* batch data in class i, and $N_i \mu_i$ is the maximal falsification in this class. Thus, the inspector's optimal sample sizes are proportional to the maximal possible falsifications in the respective classes, and conversely, the inspectee's optimal level of falsification in a given class is proportional to the inspector's efforts for verifying all data in that class.

Methodologically, the most important innovation in the treatment of NPT safeguards problems has been the marriage of statistical and game-theoretic modeling. This was necessary to deal with the concrete problems posed by the IAEA's material accountancy and variable sampling verification procedures. In particular, the use of the Neyman Pearson Lemma (see e.g. Lehmann, 1959) for the determination of equilibrium test procedures has turned out to be exceedingly fruitful.

A good example of the power of this synthesis is provided by the investigation of a special variant of material accountancy, known as *near real time accountancy* (NRTA). This technique was studied extensively in the 1970s and 1980s in connection with safeguarding large plutonium reprocessing facilities. It involves frequent material balance closings in a running plant and the sequential analysis of some appropriate test statistic, such as cumulative MUF. It was initially hoped that the method would improve the sensitivity of IAEA accountancy procedures for detection of protracted diversions over some given reference time period, and many sophisticated test procedures were proposed. In a game-theoretic treatment utilizing the Neyman Pearson Lemma, Avenhaus and Jaech (1981) showed that, given the freedom of the inspectee to distribute his diversion any way he wished, the optimal test in the sense of maximal overall detection probability makes no use of the intermediate balance data at all. NRTA was therefore shown to be of no advantage in improving the sensitivity of material accountancy for safeguards.

With regard to the relevance of NRTA to timely detection, Avenhaus and Okada (1992) constructed a general two-person non-zero-sum sequential game in which the utilities were dependent on the period in which detection occurs. They showed that if the players' utilities are discounted exponentially, being higher for the inspector at earlier detection times and higher for the inspectee at later detection times, then the inspector should minimize the average run length to detection under the diversion hypothesis, for a given value of the average run length under the null hypothesis. This may be interpreted as a game-theoretic justification for average run length as an optimization criterion for timely detection, a criterion which had been used intuitively in NRTA investigations for some time (see also Canty and Avenhaus, 1991).

Independently, without reference to any specific arms control problem, Diamond (1982) presented an elegant method for choosing the times of unobservable inspections so as to minimize the expected time to detection. The time horizon is finite and the number of inspections is fixed. The optimal inspections are randomized over a one-parameter family of strategies.

In nuclear safeguards, the optimal use of both reported and independently verified data in closing material balances is a rather involved problem for the inspector because of the many illegal strategies available to the inspectee; the latter can falsify any portion of his reported flow and inventory data and, independently of this, divert material. If a series of balance periods is to be considered, the problem is compounded still. A fundamental question is the following: Given the a priori untrustworthiness of the reported data, should they be included in the final test procedure at all? A game-theoretic model of an idealized material balance area (Avenhaus and Canty, 1996) showed that, unless the inspector has independent data on all material flows and inventories, he should indeed make use of the inspectee's reported data in testing the diversion/falsification hypothesis. Another fundamental game-theoretic conclusion in this connection, again formally justifying IAEA practice, is that the optimal test statistic combining reported and verified data, given that the data are initially aggregated as MUF and D, is MUF–D (Avenhaus, 1986). Here D is the sum of the inspectee-inspector measurement differences extrapolated to the material balance as a whole.

Variable sampling, like NRTA, can be modeled as a zero-sum game with infinitely many strategies: The inspectee distributes some total falsification, treated as an external parameter, continuously across a finite population of reported data. The inspector's strategy set is the set of all statistical decision procedures having a given false alarm probability, and the payoff to the inspector for a given test and a falsification strategy, is the detection probability. Again the zero-sum assumption can be justified in terms of a non-cooperative game involving both legal and illegal strategies, with only the inspector's choice of the false alarm probability depending upon the subjective utilities of the protagonists (see Avenhaus and Canty, 1996).

Solutions for the special cases in which just one datum or in which the entire population is sampled have been obtained (Avenhaus, Battenberg, and Falkowski, 1991). However the case of real practical interest in which n samples of a population of N are taken has not been solved for all values of the total falsification. The simple D-test, involving the sum of the differences of reported and verified data, is generally applied for all magnitudes of falsification, albeit without formal justification. Avenhaus and Piehlmeier (1994) reviewed the state-of-the-art of single stratum variable sampling problems, while Avenhaus and Canty (1996) also derive equilibria for stratified variable sampling.

In concluding this section, we should also mention studies that are not related to the NPT. The U.S. Nuclear Regulatory Commission (NUREG), responsible for national safeguards and security of nuclear installations, sponsored a study (Goldman 1984), investigating the use of game theory, or "strategic analysis" as it was referred to, for safeguards. This work presents a large number of references. It discusses some basic issues in the application of game theory for modeling and implementing inspections, such as its understandability for practitioners, difficulties in defining payoffs, and the use of mixed strategies.

4. New challenges: INF, CFE and the CWC

Since the mid-1980s, the evolution of inspection games in the arms control context has been driven by the inspection problems arising in new and qualitatively different arms control regimes, and by the need to understand the impact of political and cost parameters on optimal inspection strategies and on compliance behavior. The end of the Cold War brought a willingness to use arms control to address a broader range of international problems. Agreements to limit nuclear and other weapons and forces, and to destroy existing weapons, have posed significant new verification challenges. At the same time, attention has focused on cost control and on the "political" factors, including treaty characteristics, that affect arms control success.

The first of the new agreements was the Intermediate Range Nuclear Forces (INF) Treaty, a bilateral agreement between the United States and the Soviet Union that eliminated forever an entire category of weapons - nuclear-armed missiles with ranges between 500 and 5,500 km. Its ratification in 1988 marked the first time that the superpowers had agreed to mutual on-site inspections - in this case, of missile storage and launch facilities, and of destruction operations.

Verification of the INF Treaty depended on each party's National Technical Means, or unilateral intelligence and monitoring capability, and on on-site inspections - both routine and short-notice. Because only a limited number of short-notice inspections were permitted, and because the possession of even one prohibited system constituted a violation, INF verification possessed many strategic features captured in the earlier Mathematica models. For instance, some aspects of INF verification were analyzed by Brams, Davis, and Kilgour (1991) using a zero-sum Dresher-type model in which the inspectee begins with a finite amount of "cheating resources" (missiles that he prefers not to destroy, for example) and allocates these resources over time slots (or sites); the inspector chooses which slots to inspect. Among the authors' conclusions is the observation that, if the ratio of inspections to slots is held fixed, then optimal inspection strategies detect violations more effectively as the numbers of both inspections and slots become larger. For example, allowing 24 inspections per year makes for a more effective treaty than allowing two per month.

Another landmark in the history of arms control was the Conventional Forces in Europe (CFE) Treaty, signed November 19, 1990. The CFE Treaty bound the states of the Warsaw Pact and the North Atlantic Treaty Organization (NATO) to reduce conventional military hardware to agreed limits within an extremely large area - from the Atlantic to the Urals. It required the withdrawal and destruction of tens of thousands of pieces of military equipment - including tanks, fixed-wing aircraft, and helicopters - and the monitoring for compliance of thousands of military bases. Likewise, the CFE Treaty included many important verification innovations, including detailed data exchange verification, declared-site inspections with very limited refusal rights, and challenge inspections of undeclared sites.

The CFE Treaty occasioned some rethinking of inspection strategies. Perhaps the most important new

Table 2
A simple two-site inspection game. Note that when a violation site i is not inspected, the inspector's loss $-w_i$ and the inspectee's gain v_i depend on the site. When a violation site is inspected, the inspector and inspectee have losses $-F$ and $-P$, respectively, that do not depend on the site

Inspector	Inspectee			
	Comply	Violate 1	Violate 2	Violate 1 & 2
Inspect 1	0, 0	$-F$, $-P$	$-w_2$, v_2	$-F$, $-P$
Inspect 2	0, 0	$-w_1$, v_1	$-F$, $-P$	$-F$, $-P$

problem was the allocation of limited numbers of inspections across sites of different values. For instance, illegal activity at some declared sites may be considerably more threatening than at others, as a consequence of the size, nature, location, etc., of each site. Even when inspection is perfect (a violation is detected if, and only if, the site is in violation), the simple strategy of inspecting only the most valuable sites is not very appealing, for, once known, it leaves every uninspected site vulnerable. Improvements can be accomplished by utilizing uncertainty – that is, by choosing a subset of sites to inspect according to an appropriate probability distribution.

It is generally possible to calculate the optimal inspection distribution over sites of different values but, regretably, little is known in general about how the optimal random inspection pattern depends on the pattern of values of the sites, and on other parameters of the problem (see e.g. Canty and Avenhaus, 1994). Described next is a simple two-person non-zero-sum game, based on Kilgour (1992), that illustrates this dependence. Suppose that the inspectee can choose to violate or to behave legally at each of sites 1 and 2, and that the inspector is committed to inspect at exactly one of the two sites. Take both players' utilities to be 0 if there is legal behavior, and $-F$ (to the inspector) and $-P$ (to the inspectee) if there is a violation at the inspected site, which is assumed to be detected with certainty. If there is a violation at site i but site i is not inspected, let the utilities be $-w_i$ to the inspector and v_i to the inspectee. This produces the bimatrix game shown in Table 2.

The bimatrix game in Table 2 can be solved by standard methods and illustrates well not only the complex dependence of optimal inspection strategies on values, but also the dependence of behavior on political parameters. First, note that one of the inspectee's strategies, "Violate 1 & 2", is strictly dominated, so it is never selected at equilibrium. Assume $w_i > F > 0$ for all i (the inspector most prefers to deter violations; but, if a violation occurs, he prefers to detect it). Define $P_0 = \sqrt{v_1 \cdot v_2}$. If $P < P_0$, the game of Table 2 has a unique Nash equilibrium at which violations always occur. Specifically, the inspectee chooses to violate at the ith site with probability q_i as given by

$$q_i = \frac{w_j - F}{w_1 + w_2 - 2F}, \quad i = 1, 2, \quad j \neq i.$$

One sees that $q_1 + q_2 = 1$, i.e. the inspectee never complies. Meanwhile the inspector chooses to inspect the ith site with probability p_i as given by

$$p_i = \frac{v_i + P}{v_1 + v_2 + 2P}, \quad i = 1, 2.$$

For instance if $v_1 = 4v$ and $v_2 = v$, then site 1, which is four times as important to the inspectee, should be inspected four times as frequently when the punishment parameter P is 0, but only twice as frequently when the punishment parameter equals its threshold value $P_0 = 2v$. The equilibrium payoff for the inspectee is

$$\frac{v_1 v_2 - P^2}{v_1 + v_2 + 2P}.$$

When $P > P_0$, the situation changes dramatically. There are infinitely many Nash equilibria involving certain legal behavior on the part of the inspectee; the inspector must simply choose an inspection probability for site 1 that lies between $v_1/(P + v_1)$ and $P/(P+v_2)$. The situation is illustrated in Fig. 1, which shows how larger values of P allow the inspector greater flexibility in deterring violations.

It can be shown that there are no other equilibria, in particular none which mix legal and illegal inspectee strategies (except for the case $P = P_0$, which can be neglected).

Thus this simple site selection game shows how "political" parameters, such as the level of punishment (sanctions) for a detected violation, can affect behavior. When the situation is favorable, there can be considerable flexibility in the "technical" choice of where to inspect; but when the situation is unfavorable, the inspector cannot deter violations, but only minimize their impact.

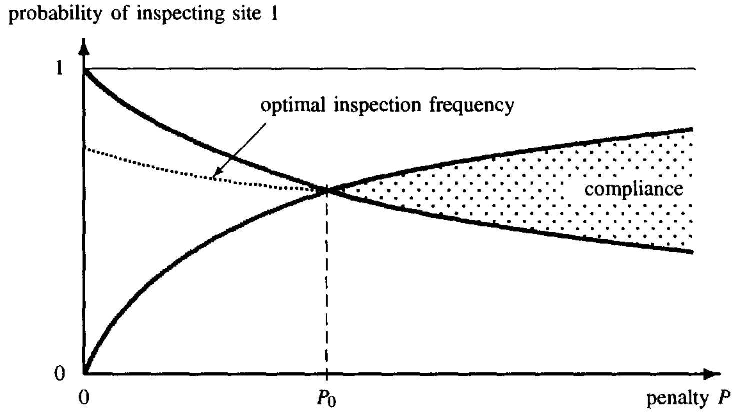

Fig. 1. Optimal inspection frequency for site 1 in the game of Table 2, versus punishment parameter P. Along each bold line, the inspectee is indifferent between violating and complying at one site. All violation is deterred in the dotted area (where $P > P_0$)

The site selection game of Table 2 illustrates several difficulties with the general problem of selecting an inspection strategy when the objects of inspection have different values. Other examples, some given by Kilgour (1992), show even more pathological behavior. For instance, when there are more than two sites, there can be Nash equilibria at which equally valuable sites are not inspected equally often. Also, a Nash equilibrium strategy for an inspectee may involve sometimes violating at different numbers of sites. In fact, even when all sites are of equal value, it is not clear at how many sites violations should occur in equilibrium. Ruckle (1983, p. 25) showed that if there are n sites, all with value v to the inspectee, if the inspector has m inspections, and if $P = -v$, then the number of sites at which to violate (in equilibrium) is $-(n-m)/(m+1)$ that is, slightly less than n/m.

Another significant event in the history of arms control was the Chemical Weapons Convention (CWC), which was signed in 1993 but will not enter into force until approximately 1996, when sufficient ratifications have been accumulated. The CWC is the first comprehensively verifiable multilateral treaty to eliminate completely an entire class of weapons, and to regulate activities that may contribute to the production of such weapons. The Organization for the Prohibition of Chemical Weapons (OPCW), which will administer the CWC, will catalogue the compulsory national declarations, carry out routine inspections of declared facilities, and, at the request of any state party, conduct a short-notice challenge inspection of any site.

While it is clear that CWC will represent an enormous verification problem, many aspects of that problem are not yet clear, as the OPCW operating rules are still being spelled out. One early study (Kilgour, 1990) determined optimal inspection strategies in a zero-sum model with variable violation levels and quantity-dependent attribute sampling; that is, both the value of a violation, and the probability of detecting it in an inspection, depend linearly on the violation amount.

Cost considerations will certainly be important in the CWC due to the sheer volume of the undertaking, and these are known to have significant strategic implications. For instance, imagine the game of Table 2 when the inspector must pay a small amount for each inspection, and where he also has a third strategy – "Do Not Inspect" – available. Then the inspectee's strategy "Violate 1 & 2" is no longer dominated, and in fact can be selected with positive probability at equilibrium. In general, one-period ("simultaneous") analysis predicts that there will always be a low probability of violation whenever an inspecting side must take its own costs of inspection into account (Kilgour and Brams, 1992).

While it may be especially vulnerable to this cost-of-inspection problem, the CWC shares with most other arms control regimes a dependence on other, so-called "political" parameters such as the utility loss to a violator as a result of sanctions. Efforts to separate the effects of political and technical considerations continue, but most recent models have focused on the interaction of these factors. Downs and Rocke (1990), for instance, suggest treaty-maintenance strategies involving violation triggers that depend on (possibly noisy) signals, and other forms of "tacit bargaining." Kilgour and Brams (1992) argue that the inspector leadership principle has quite general applicability. Brams and Kilgour (1988) and Kilgour and Avenhaus (1994) model this interaction in other contexts.

5. Inspection games in the future

Inspection games continue to be an active area of research, and it is virtually certain that many of the conclusions reached to date will be sharpened, deepened, or revised by future models. In addition, the time is now ripe to consolidate knowledge over many treaty types, and to develop general theories of inspection which have now begun to appear, such as O'Neill (1994), Avenhaus and Canty (1996), and Avenhaus, von Stengel, and Zamir (1996).

Knowledge of arms control inspection will certainly grow, in response to new treaties and to increased experience with existing treaties. On the horizon now are the "93 + 2" strengthening of the NPT verification procedures, a possible Comprehensive Test Ban Treaty, a fissile materials "cutoff" agreement, verification provisions for the Biological and Toxin Weapons Convention, and various measures to constrain the proliferation of conventional weapons, especially small arms. Much has been learned from the less-than-satisfactory experiences of NPT verification in Iraq and North Korea, and also from the operations of the United Nations Special Commission in Iraq since 1991. Future inspection games will no doubt incorporate, and elaborate on, these lessons.

One noteworthy feature, common to many of these practical inspection problems but not yet incorporated in game models of arms control, is the use of combinations of inspection modalities. For example, limitations on conventional weapons might be based first on overhead surveillance (satellite and aircraft) combined with fixed tamperproof cameras and/or remote perimeter and portal monitors; only after a suspect event or area has been identified by one or several of these techniques would an on-site inspection be called for. McFate et al. (1992) argue that such combinations would synergistically achieve dramatic improvements in verification cost-effectiveness.

References

Aumann, R.J., and Maschler, M. (1995), *Repeated Games with Incomplete Information*, MIT Press, Cambridge, MA.

Avenhaus, R. (1986), *Safeguards Systems Analysis*, Plenum, New York.

Avenhaus, R., Battenberg, H.P., and Falkowski, B.J. (1991), "Optimal tests for data verification", *Operations Research* 39, 341–348.

Avenhaus, R., and Canty, M. (1989), "Re-examination of the IAEA formula for stratified attributes sampling", *Proceedings of the 11th ESARDA Symposium*, Luxembourg (JRC Ispra), 351–356.

Avenhaus, R., and Canty, M. (1996), *Compliance Quantified – An Introduction to Data Verification*, Cambridge University Press, Cambridge, UK.

Avenhaus, R., and Jaech, J.L. (1981), "On subdividing material balances in time and/or space", *Journal of the INMM* X/3, 24–33.

Avenhaus, R., and Okada, A. (1992), "Statistical criteria for sequential inspector-leadership games", *Journal of the Operations Research Society of Japan* 35, 134–151.

Avenhaus, R., and Piehlmeier, G. (1994), "Recent developments in and present state of variable sampling", IAEA-SM-333/7, *Proceedings of a Symposium on International Nuclear Safeguards 1994: Vision for the Future*, Vol. I. IAEA, Vienna, 307–316.

Avenhaus, R., von Stengel, B., and Zamir, S. (1996), "Inspection games", in: R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory, Vol. 3*, North-Holland, Amsterdam, to appear.

Bierlein, D. (1968), "Direkte Inspektionssysteme", *Operations Research Verfahren* 6, 57–68.

Bierlein, D. (1969), "Auf Bilanzen und Inventuren basierende Safeguards-Systeme", *Operations Research Verfahren* 8, 36–43.

Bowen, W.M., and Bennett, C.A. (eds.) (1988), "Statistical methodology for nuclear material management", Report NUREG/CR-4604 PNL-5849, prepared for the U.S. Nuclear Regulatory Commission, Washington, DC.

Brams, S., Davis, N.D., and Kilgour, D.M. (1991), "Optimal cheating and inspection strategies under INF", in: R. Avenhaus, H. Karkar, and M. Rudnianski (eds.), *Defense Decision Making*, Springer, Berlin, 318–335.

Brams, S., and Kilgour, D.M. (1988), *Game Theory and National Security*, Chapter 8: Verification. Basil Blackwell, New York.

Canty, M., and Avenhaus, R. (1991), "Inspection games over time", KFA-Report Jül-2472, KFA Jülich, Germany.

Canty, M., and Avenhaus, R. (1994), "Detection of undeclared activity: quantifying the unquantifiable", IAEA-SM-333/150P, *Proceedings of a Symposium on International Nuclear Safeguards 1994: Vision for the Future*, Vol. I, IAEA, Vienna, 505–511.

Diamond, H. (1982), "Minimax policies for unobservable inspections", *Mathematics of Operations Research* 7, 139–153.

Downs, G.W., and Rocke, D.M. (1990), *Tacit Bargaining, Arms Races, and Arms Control*, University of Michigan Press, Ann Arbor, MI.

Dresher, M. (1962), "A sampling inspection problem in arms control agreements: A game-theoretic analysis", Memorandum No. RM-2972-ARPA, The RAND Corporation, Santa Monica, CA.

Goldman, A.J. (1984), "Strategic analysis for safeguards systems: A feasibility study, Appendix", Report NUREG/CR-3926 Vol. 2, prepared for the U.S. Nuclear Regulatory Commission, Washington, DC.

Höpfinger, E. (1974), "Zuverlässige Inspektionsstrategien", *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 31, 35–46.

Kilgour, D.M. (1990), "Optimal cheating and inspection strategies under a chemical weapons treaty", *INFOR* 28, 27–39.

Kilgour, D.M. (1992), "Site selection for on-site inspection in arms control", *Arms Control* 13, 439–462.

Kilgour, D.M., and Avenhaus, R. (1994), "The optimal distribution of inspection effort", in: S. Mataija (ed.), *Non-Proliferation and Multilateral Verification: The Comprehensive Nuclear Test Ban Treaty*, Centre for International and Strategic Studies, York University, Toronto, Canada, 137–148.

Kilgour, D.M., and Brams, S. (1992), "Putting the other side 'on notice' can induce compliance in arms control", *Journal of Conflict Resolution* 36, 395–414.

Lehmann, E.L. (1959), *Testing Statistical Hypotheses*, Wiley, New York.

Maschler, M. (1966), "A price leadership method for solving the inspector's non-constant-sum game", *Naval Research Logistics Quarterly* 13, 11–33.

Maschler, M. (1967), "The inspector's non-constant-sum-game: its dependence on a system of detectors", *Naval Research Logistics Quarterly* 14, 275–290.

Mathematica (1963), "Applications of statistical methodology to arms control and disarmament", Final report, submitted to the U.S. Arms Control and Disarmament Agency under Contract No. ACDA/ST-3 by Mathematica, Inc., Princeton, NJ, F.J. Anscombe et al.

Mathematica (1965), "The application of statistical methodology to arms control and disarmament", Final report, submitted to the U.S. Arms Control and Disarmament Agency under Contract No. ACDA/ST-37 by Mathematica, Inc., Princeton, NJ.

Mathematica (1966), "Development of utility theory for arms control and disarmament", Final report, submitted to the U.S. Arms Control and Disarmament Agency under Contract No. ACDA/ST-80 by Mathematica, Inc., Princeton, NJ, R.J. Aumann et al.

McFate, P.B., Graybeal, S.N., Lindsey, G., and Kilgour, D.M. (1993), "Constraining Proliferation: The contribution of verification synergies", Arms Control Verification Studies No. 5, Department of Foreign Affairs and International Trade, Ottawa, Ontario, Canada.

O'Neill, B. (1994), "Game theory models on peace and war", in: R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory*, Vol. 2, North-Holland, Amsterdam, 995–1053.

Rinderle, K. (1995), "Mehrstufige sequentielle Inspektionsspiele mit Fehlern erster und zweiter Art", Doctoral dissertation, University of the Federal Armed Forces Munich, Germany.

Ruckle, W.H. (1983), *Geometric Games and their Applications*, Pitman, London.

Saaty, T.L. (1968), *Mathematical Models of Arms Control and Disarmament: Application of Mathematical Structures in Politics*, Wiley, New York.

von Neumann, J., and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, New Jersey.

von Stengel, B. (1991), "Recursive inspection games", Technical Report S-9106, University of the Federal Armed Forces Munich, Germany.