INSPECTION GAMES

January 30, 1998

RUDOLF AVENHAUS*

Universität der Bundeswehr München

BERNHARD VON STENGEL*

Eidgenössische Technische Hochschule Zürich

SHMUEL ZAMIR*

The Hebrew University of Jerusalem

*This work was supported by the German-Israeli Foundation (G.I.F.) and by the Volkswagen Foundation.

1. Introduction

Inspection games form an applied field of game theory. An inspection game is a mathematical model of a situation where an *inspector* verifies that another party, called *inspectee*, adheres to certain legal rules. This legal behavior may be defined by an arms control treaty, for example, and the inspectee has a potential interest in violating these rules. Typically, the inspector's resources are limited so that verification can only be partial. A mathematical analysis should help in designing an optimal inspection scheme, where it must be assumed that an illegal action is executed strategically. This defines a game-theoretic problem, usually with two players, inspector and inspectee. In some cases, several inspectees are considered as individual players.

Game theory is not only adequate to describe an inspection situation, it also produces results which may be used in practical applications. The first serious attempts were made in the early 1960's where game-theoretic studies of arms control inspections were commissioned by the United States Arms Control and Disarmament Agency (ACDA). The International Atomic Energy Agency (IAEA) performs inspections under the Nuclear Non-Proliferation Treaty. The decision rules of IAEA inspectors for detecting a deviation of nuclear material can be interpreted as equilibrium strategies in a zero-sum game with the detection probability as payoff function. In these applications, game theory has proved itself useful as a technical tool in line with statistics and other methods of operations research. Since the underlying models should be accessible to practitioners, traditional concepts of game theory are used, like zero-sum games or games in extensive form. Nevertheless, as we want to demonstrate, the solution of these games is mathematically challenging, and leads to interesting conceptual questions as well.

We will not consider inspection problems that are exclusively statistical, since our emphasis is on games. We also exclude industrial inspections for quality control and maintenance, except for an interesting worst-case analysis of timely inspections. Similarly, we do not consider *search games* (Gal, 1983; O'Neill, 1994) modeling pursuit and evasion, for example in war. Search games are distinguished from inspection games by a symmetry between the two players whose actions are usually equally legitimate. In contrast, inspection games as we understand them are fundamentally asymmetrical: Their salient feature is that the inspector tries to prevent the inspectee from behaving illegally in terms of an agreement.

In Section 2, we survey applications of inspection games to arms control, auditing and accounting and economics, and other areas like environmental regulatory enforcement or crime control.

In Section 3, we provide a general game-theoretic framework to inspections which extends the statistical approach used in practice. In these statistical hypothesis testing problems, the distribution of the observed random variable is strategically manipulated by the inspectee. The equilibrium of the general non-zero-sum game is found using an auxiliary zero-sum game in which the inspectee chooses a violation procedure and the inspector chooses a statistical test with a given false alarm probability. We illustrate this by two specific important models, namely material accountancy and data verification.

Inspections over time, so far primarily of methodological interest, are the subject of Section 4. In these games, the information of the players about the actions of their respective opponent is very important, which is best understood if the game is represented in extensive form. If the payoffs have a simple structure, then the games can sometimes be represented recursively and solved analytically. In another timeliness game, the optimal inspection times, continuously chosen from an interval, are determined by differential equations.

In Section 5 we discuss the inspector leadership principle. It says that the inspector may commit himself to his inspection strategy in advance, and thereby gains an advantage compared to the symmetrical situation where both players choose their actions simultaneously. Obviously, this concept is particularly applicable to inspection games.

2. Applications

The majority of publications on inspection games concerns arms control and disarmament, usually relating to one or the other arms control treaty that has been signed. We survey these models first. Some of them are described in mathematical detail in later sections. Inspection games have also been applied to problems in economics, particularly in accountancy and auditing, in enforcement of environmental regulations, and in crime control and related areas. Some papers treat inspection games in an abstract setting, rather than modeling a particular application.

2.1. Arms Control and Disarmament

Inspection games have been applied to arms control and disarmament in three phases. Studies in the first phase, from about 1961 to 1967, analyze inspections for a nuclear test ban treaty, which was then under negotiation. The second phase, about 1968 to 1985, comprises work stimulated by the Non-Proliferation Treaty for Nuclear Weapons. Under that treaty, proper use of nuclear material is verified by the International Atomic Energy Agency (IAEA) in Vienna. The third phase lasts from about 1986 until today. The end of the Cold War brought about new disarmament treaties. Verification of these treaties has also been analyzed using game theory.

In the 1960's, a *test ban treaty* was under negotiation between the United States and the Soviet Union, and verification procedures were discussed. Tests of nuclear weapons above ground can be detected by satellites. Underground tests can be sensed by seismic methods, but in order to discriminate them safely from earthquakes, on-site inspections are necessary. The two sides never agreed on the number of such inspections that they would allow to the other side, so eventually they decided not to ban underground tests in the treaty.

While the test ban treaty was being discussed, the problem arose how an inspector should use a certain limited number of on-site inspections, as provided by the treaty, for verifying a much larger number of suspicious seismic events. Dresher (1962) modeled this problem as a recursive game. The estimated number of events and the number of inspections that can be used are fixed parameters. We explain this game in Section 4.1. It formed the basis for much subsequent work.

With political backing for scientific disarmament studies, the United States Arms Control and Disarmament Agency (ACDA) commissioned game-theoretic analyses of inspections to the Mathematica company in the years 1963 to 1968. Game-theoretic researchers involved in this or related work were Aumann, Dresher, Kuhn, Maschler, Selten, among others. Publications on inspection games of that group are Kuhn (1963) and, in general, the reports to ACDA edited by Anscombe et al. (1963; 1965), as well as Maschler (1966; 1967). In a related study, Saaty (1968) presents some of the existing developments in this area. Many of these papers extend Dresher's game in various ways, for example by generalizing the payoffs, or by assuming an uncertainty in the signals given by detectors.

The *Non-Proliferation Treaty* (NPT) for Nuclear Weapons was inaugurated in 1968. This treaty divided the world into weapons states, who promised to reduce or even eliminate their nuclear arsenals, and non-weapon states who promised never to acquire such weapons. All states which became parties to the treaty agreed that the IAEA in Vienna verifies the nuclear material contained in the peaceful nuclear fuel cycles of *all* states.

The verification principle of the IAEA is *material accountancy*, that is, the comparison of book and physical inventories for a given material balance area at the end of an inventory period. The plant operators report their balance data via their national organizations to the IAEA, whose inspectors verify these reported data with the help of independent measurements on a random sampling basis.

Two kinds of sampling procedures are considered in all situations, depending on the nature of the problem. *Attribute sampling* is used to test or to estimate the percentage of items in the population containing some characteristic or attribute of interest. In inspections, the attribute of interest is usually if a safeguards measure has been violated. This may be a broken seal of a container, or an unquestioned decision that a datum has been falsified. The inspector uses the rate of tampered items in the sample to estimate the population rate or to test a hypothesis.

The second kind of procedure is *variable sampling*. This is designed to provide an estimate of or a test on an average or total value of material. Each observation, instead of being counted as falling in a given category, provides a value which is totaled or averaged for the sample. This is described by a certain test statistic, like the total of Material Unaccounted For (MUF). Based on this statistic, the inspector has to decide if nuclear material has been diverted or if the result is due to measurement errors. This decision depends on the probability of a false alarm chosen by the inspector.

Game-theoretic work in this area was started by Bierlein (1968; 1969), who emphasized that payoffs should be expressed by detection probabilities only. In contrast, inspection costs are parameters that are fixed externally. This is an adequate model for the IAEA which has a certain budget limiting its overall inspection effort. The agency has no intent to minimize that effort further, but instead wants to use it most efficiently.

Since 1969, international conferences on nuclear material safeguards are regularly held by the following institutions. The IAEA organizes about every five years conferences on Nuclear Safeguards Technology and publishes proceedings under that title. The European Safeguards Research and Development Association (ESARDA) as well as the Institute for Nuclear Material Management (INMM) in the United States meet every year and publish proceedings on that subject as well. Here, most publications concern practical matters, for example measurement technology, data processing, and safety. However, decision theoretical approaches, including game-theoretic methods, were presented throughout the years. Monographs which emphasize the theoretical aspects are Jaech (1973), Avenhaus (1986), Bowen and Bennett (1988), and Avenhaus and Canty (1996).

Some studies on nuclear material safeguards are not related to the NPT. The U.S. Nuclear Regulatory Commission (NUREG) is in charge of safeguarding nuclear plants against theft and sabotage to guarantee their safe operation, in fulfillment of domestic regulations. In a study for that commission, Goldman (1984) investigates the possible use of game theory and its potential role in safeguards.

In the mid-eighties, when the Cold War ended, new disarmament treaties were signed, like the treaty on Intermediate Nuclear Forces in 1987, or the treaty on Conventional Forces in Europe in 1990 (see Altmann et al. 1992 on verification issues). The verification of these new treaties was investigated in game-theoretic terms by Brams and Davis (1987), by Brams and Kilgour (1988), and by Kilgour (1992). Variants of recursive games are described by Ruckle (1992). In part, these papers extend the work done before, in particular Dresher (1962).

2.2. Accounting and Auditing in Economics

In economic theory, inspection games have been studied for the auditing of accounts. In insurance, inspections are used against the 'moral hazard' that the client may abuse his insurance by fraud or negligence. Our summary of these topics is based on Borch (1990). We will also consider models of tax inspections.

Accounting is usually understood as a system for keeping track of the circulation of money. The monetary transactions are recorded in accounts. These may be checked in full detail by an inspector, which is called an *audit*. Auditing of accounts is often based on sampling inspection, simply because it is unnecessarily costly to check every voucher and entry in the books. The possibility that any transaction may be checked in full detail is believed to have a deterring effect likely to prevent irregularities.

A theoretical analysis of problems in this field naturally leads to a search for suitable sampling methods (for an outline see Kaplan, 1973). The concepts of attribute and variable sampling described above for material accountancy apply to the inspection of monetary accounts as well. In particular, there are tests to validate the reasonableness of account balances, without classifying a particular observation as correct or falsified. These may be considered as variable measurement tests and are sometimes termed 'dollar value' samples in the literature.

These theoretical investigations include the use of game-theoretic methods. An early contribution that employs both noncooperative and cooperative game theory is given by Klages (1968), who describes quite detailed models of the practical problems in accountancy and discusses the merits of a game-theoretic approach.

Borch (1982) formulates a zero-sum inspection game between an accountant and his employer. The accountant, who is the inspectee, may either record transactions faithfully or cheat to embezzle some of the company's profits, and the employer may either trust the accountant or audit his accounts. If the inspectee is honest, he receives payoff zero irrespective of the actions of the inspector. In that case, the inspector gets payoff zero if he trusts and has a certain cost if he audits. If the accountant steals while being trusted, he has a gain that is the employer's loss. If the inspector catches an illegal action, he has the same auditing cost as before but the inspectee must pay a penalty. Borch interprets mixed strategies as 'fractions of opportunities' in a repeated situation, but does not formalize this further.

The employer may buy 'fidelity guarantee insurance' to cover losses caused by dishonest accountants. The insurance may require a strict auditing system that is costly to the employer. Borch (1982) considers a three-person game where the insurance company inspects or alternatively trusts the employer if he audits properly. The employer is both inspectee, of the insurance company, and inspector, of his accountant. No interaction is assumed between insurance company and accountant, so this game is fully described by two bimatrix games. For general 'polymatrix' games of this kind see Howson (1972).

Borch (1990) sees many potential applications of games to the economics of insurance: In any situation where the insured has undertaken to take special measures to prevent accidents and reduce losses, there is a risk – usually called *moral hazard* – that he may neglect his obligations. The insurance company reserves the right to inspect that the insured really carries out the safety measures foreseen in the insurance contract. This inspection costs money, which the insured must pay for by an addition to the premium.

Moral hazard has its price. Therefore, inspections of economic transactions raise the question of the most efficient design of such a system, for example so that surveillance is minimized or even unnecessary. These problems are closely related to a variety of economic models known as *principal–agent problems*. They have been extensively studied in the economic literature, where we refer to the surveys by Baiman (1982), Kanodia (1985), and Dye (1986). Agency theory focuses on the optimal contracted relationships between two individuals whose roles are asymmetric. One, the principal, delegates work or responsibility to the other, the agent. The principal chooses, based on his own interest, the payment schedule that best exploits the agent's self-interested behavior. The agent chooses an optimal level of action contingent on the fee schedule proposed by the principal. One important issue in agency theory is the asymmetry of information available to the principal and the agent. In inspection games, a similar asymmetry exists with respect to defining the rules of the game, see Section 5.

We conclude this section with applications to tax inspections. Schleicher (1971) describes an interesting recursive game for detecting tax law violations, which extends Maschler (1966). Rubinstein (1979) analyzes the problem that it may be unjust to penalize illegal actions too hard since the inspectee might have committed them unintentionally. In a one-shot game, there is no alternative to the inspector but to use a high penalty, although its potential injustice has a disutility. Rubinstein shows that if this game is repeated, a more lenient policy also induces the inspectee to legal behavior.

Another model of tax inspections, also using repeated games, is discussed by Greenberg (1984). A tax function defines the tax to be paid by an individual with a certain income. An audited individual who did not report its income properly must pay a penalty, and the tax authorities can audit only a limited percentage of individuals. Under reasonably weak assumptions about these functions and the individuals' utility functions on income, Greenberg proposes an auditing scheme that achieves an arbitrary small percentage of tax evaders. In that scheme, the individuals are partitioned into three groups that are audited with different probabilities, and individuals are moved among these groups after an audit depending on whether they cheated or not. A similar scheme of auditing individuals with different probabilities depending on their compliance history is proposed by Landsberger and Meilijson (1982). However, their analysis does not use game theory explicitly. Reinganum and Wilde (1986) describe a model of tax compliance where they apply the sequential equilibrium concept. In that model, the income reporting process is considered explicitly as a signaling round. The tax inspector is only aware of the overall income distribution and reacts to the reported income.

2.3. Environmental Control

Environmental control problems call for a game-theoretic treatment. One player is a firm which produces some pollution of air, water or ground, and which can save abatement costs by illegal emission beyond some agreed level. The other player is a monitoring agent whose responsibility is to detect or better to prevent such illegal pollution. Both agents are assumed to act strategically. Various problems of this kind have been analyzed. However, contrary to arms control and disarmament, these papers do not yet address specific practical cases.

Several papers present game-theoretic analyses of pollution problems but deal only marginally with monitoring problems. Bird and Kortanek (1974) explore various concepts in order to aid the formulation of regulations of sources of pollutant in the atmosphere related to given least cost solutions. Höpfinger (1979) models the problem of how to determine and adapt global emission standards for carbon dioxide as an infinite stage game with three players: regulator, producer, and population. Kilgour, Okada, and Nishikori (1988) describe the load control system for regulating chemical oxygen demand in water bodies. They formulate a cost sharing game and solve it in some illustrative cases.

In the last years, pollution control problems have been analyzed with game theoretic methods. Russell (1990) characterizes current enforcement of U.S. environmental laws as very likely inadequate, while admitting that proving this proposition would be extremely difficult, exactly because there is so little information about the actual behavior of regulated firms and government activities. As a remedy, a one-stage game between a polluter and an environmental protection agency is used as a benchmark for discussing a multiple-stage game in which the source's past record of discovered violations determines its future probabilities of being monitored. It is shown that this approach can yield significant savings in limiting the extent of violations to a particular frequency in the population of polluters.

Weissing and Ostrom (1991) examine how irrigation institutions affect equilibrium rates of stealing and enforcement. Irrigators come periodically into the position of turntakers. A turntaker chooses between taking a legal amount of water and taking more water than authorized. The other irrigators are turnwaiters who must decide whether to expand resources for monitoring the behavior of the turntaker or not. For no combination of parameters the rate of stealing by the turntaker drops to zero, so in equilibrium some stealing is always going on.

Güth and Pethig (1992) consider a polluting firm that can save abatement costs by illegal waste emission, and a monitoring agent whose job it is to prevent such pollution. When deciding on whether to dispose of its waste legally or illegally the firm does not know for sure whether the controller is sufficiently qualified or motivated to detect the firm's illegal releases of pollutant. The firm has the option of undertaking a small-scale deliberate 'exploratory pollution accident' to get a hint about the controller's qualification before deciding on how to dispose of its waste. The controller may or may not respond to that 'accident' by a thorough investigation, thus perhaps revealing his type to the firm. This sequential decision process along with the asymmetric distribution of information constitutes a signaling game whose equilibrium points may signal the type of the controller to the firm.

Avenhaus (1994) considers a decision theoretic problem. The management of an industrial plant may be authorized to release some amount of pollutant per unit time into the environment. An environmental agency may decide with the help of randomly sampled measurements whether or not the real releases are larger than the permitted ones. The 'best' inspection procedure can be determined by the use of the Neyman Pearson Lemma; see Section 3 below.

2.4. Miscellaneous Models

A number of papers do not belong to the above categories. Some of these deal – at least theoretically – with smuggling or crime control, other papers treat inspection games in an abstract setting.

Thomas and Nisgav (1976) consider a game where a smuggler tries to cross a strait in one of M nights. The inspecting border police has a speedboat for patrolling in k of these nights. In a patrolled night, the smuggler has some risk of being caught. The game is described recursively in the same way as the game by Dresher (1962). The only difference is that the smuggler *must* traverse the strait even if there will a patrol at every night, or otherwise receive the same worst payoff as if he is caught. The resulting recurrence equation for the game has a very simple solution where the game value is a linear function of k/M.

Baston and Bostock (1991) generalize the paper by Thomas and Nisgav (1976). They clarify some implicit assumptions made by these authors, in particular the full information of the players about past events, and the detection probability associated with a night patrol. Then, they study the case of two boats and derive explicit solutions that depend on the detection probabilities if one boat or both are on patrol. The case of three boats is solved by Garnaev (1994). Sequential games with three parameters, namely number of nights, patrols, and smuggling acts, are solved by Sakaguchi (1994) and Ferguson and Melolidakis (1996). Among other things, Baston and Bostok (1991) and Sakaguchi (1994) reproduce the result by Dresher (1962), which they are not aware of.

Goldman and Pearl (1976) study inspections in an abstract setting, where the inspector has to select among several sites where an inspectee can cheat, and only a limited number of sites can be inspected in total. A simple model is successively refined to study the effects of penalty levels and inspection resources.

Feichtinger (1983) considers a differential game with a suggested application to crime control. The dynamic control variables of police and thief are the 'rate of law enforcement' and the 'pilfering rate', respectively. Avenhaus (1997) analyzes inspections in local public transportation systems and shows that inspection rates used in practice coincide with the game-theoretic equilibrium.

Filar (1985) applies the theory of stochastic games to a generic 'traveling inspector model'. There are a number of sites, each with an inspectee that acts as an individual player. Each site is a state of the stochastic game. That state is deterministically controlled by the inspector who chooses the site to be inspected at the next time period. The inspector has unspecified costs associated with inspection levels, travel, and if he fails to detect a violation. The players' payoffs are either sums up to a finite stage, or limiting averages. In this model, all inspectees can be aggregated into a single player without changing equilibria. For finite horizon payoffs, the game has equilibria in Markov strategies, which depend on the time and the state but not on the history of the game. For limiting average payoffs, stationary strategies suffice, which only depend on the state.

3. Statistical Decisions

Many practical inspection problems are based on random sampling procedures. Furthermore, measurement techniques are often used which inevitably produce random and systematic errors. Then, it is appropriate to think of the inspection problem as being an extension of a statistical decision problem. The extension is game-theoretic since the inspector has to decide whether the inspectee has behaved illegally, and that action is strategic and not random.

In this section, we first present a general framework were we extend a classical statistical testing problem to an inspection game. The 'illegal part' of that game, where the inspectee has decided to violate, is equivalent to a two-person zero-sum game with the non-detection probability as payoff to the violator. If that game has a value, then the Neyman Pearson Lemma can be used to determine an optimal inspection strategy. In this framework, we then discuss two important inspection models that have emerged from statistics: material accountancy and data verification.

3.1. General Game and Analysis

In a classical hypothesis testing problem, the statistician has to decide, based on an observation of a random variable, between two alternatives (H_0 or H_1) regarding the distribution of the random variable. To make this an inspection game, assume that the distribution of the random variable is strategically controlled by another 'player' called the *inspectee*. More specifically, the inspectee can behave either *legally* or *illegally*. If he behaves legally, the distribution is according to the null hypothesis H_0. If he chooses to act illegally, he also decides on a *violation procedure* which we denote by ω. Thus, the distribution of the random variable Z under the alternative hypothesis H_1 depends on the procedure ω which is also a strategic variable of the inspectee. The statistician, called the *inspector*, has to decide between two actions, based on the observation z of the random variable Z: calling an *alarm* (rejecting H_0) or *no alarm* (accepting H_0). The random variable Z can be a vector, for instance in a multi-stage inspection. This rather general inspection game is described in Figure 3.1.

The pairs of payoffs to inspector and inspectee, respectively, are also shown in Figure 3.1. The status quo is legal action and no alarm, represented by the payoff 0 to both players. The payoffs for undetected illegal behavior are -1 for the inspector and 1 for the inspectee. In case of a detected violation, the inspector receives $-a$ and the inspectee $-b$. Finally, if a false alarm is raised, the inspector gets $-e$ and the inspectee $-h$. These parameters are subject to the restrictions

$$0 < e < 1, \qquad 0 < a < 1, \qquad 0 < h < b, \tag{3.1}$$

since the worst event for the inspector is an undetected violation, an alarm is undesirable for everyone, and the worst event for a violator is to be caught. Sometimes it is also assumed that $e < a$, which means that for the inspector a detected viola-

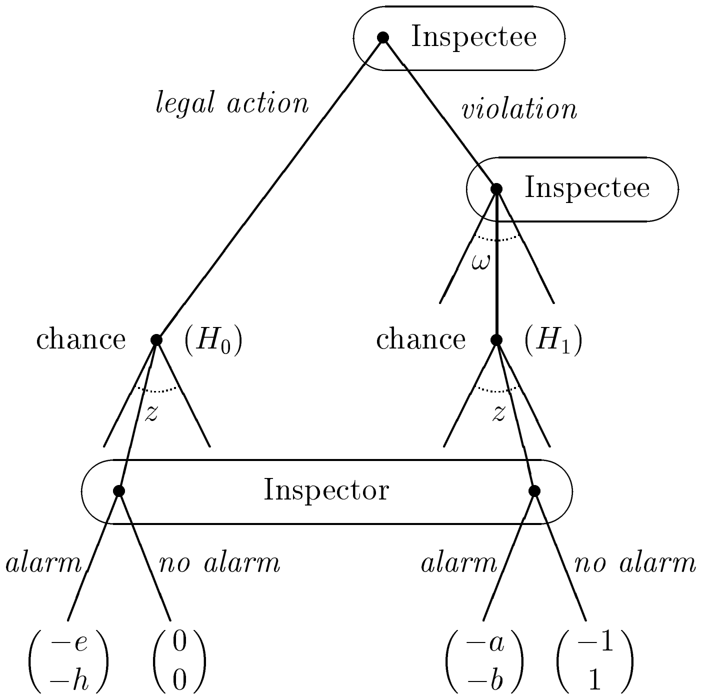

Figure 3.1. Inspection game extending the statistical decision problem of the inspector, who has to decide between the null hypothesis H_0 and alternative hypothesis H_1 about the distribution of the random variable Z. H_0 means legal action and H_1 means illegal action of the inspectee with violation procedure ω. The inspector is informed about the observation z of Z, but not if H_0 or H_1 is true. There is a separate information set for each z. The inspector receives the top payoffs, the inspectee the bottom payoffs. The parameters a, b, e, h are subject to (3.1).

tion, representing a 'failure of safeguards', is worse than the inconvenience of a false alarm.

A pure strategy of the inspectee consists of a choice between legal and illegal behavior, and a violation procedure ω. A mixed strategy is a probability distribution on these pure strategies. Since the game under consideration is of perfect recall, such a mixed strategy is equivalent to a behavior strategy given by the probability q for acting illegally and a probability distribution on violation procedures ω given that the inspectee acts illegally. Since the set Ω of violation procedures may be infinite, we assume that it includes, if necessary, randomized violations as well, which are therefore also denoted by ω. That is, a behavior strategy of the inspectee is represented by a pair (q, ω).

A pure strategy of the inspector is an *alarm set*, that is, a subset of the range of Z, with the interpretation that the inspector calls an alarm if and only if the observation z is in that set. A mixed strategy of the inspector is a probability distribution on pure strategies. A strategy of the inspector (pure or mixed) is also called a *statistical test* and will be denoted by δ. The alarm set or sets used in such a test are usually determined by first considering the *error probabilities* of falsely rejecting or accepting the null hypothesis, as follows.

A statistical test δ and a violation strategy ω determine two conditional probabilities. The probability of an error of the first kind, that is, of a *false alarm*, is the probability $\alpha(\delta)$ of an alarm given that the inspectee acts legally, which is independent of ω. The probability of an error of the second kind, that is, of *non-detection*, is the probability $\beta(\delta,\omega)$ that no alarm is raised given that the inspectee acts illegally.

The inspection game has then the following normal form. The set of strategies δ of the inspector is Δ. The set of (behavior) strategies (q,ω) of the inspectee is given by $[0,1]\times\Omega$. The payoffs to inspector and inspectee are denoted $I\big(\delta,(q,\omega)\big)$ and $V\big(\delta,(q,\omega)\big)$, respectively, where the letter 'V' indicates that the inspectee may potentially violate. In terms of the payoffs in Figure 3.1, these payoff functions are

$$\begin{aligned} I\big(\delta,(q,\omega)\big) &= (1-q)\,\big(-e\,\alpha(\delta)\big) \;+\; q\,\big(-a-(1-a)\,\beta(\delta,\omega)\big), \\ V\big(\delta,(q,\omega)\big) &= (1-q)\,\big(-h\,\alpha(\delta)\big) \;+\; q\,\big(-b+(1+b)\,\beta(\delta,\omega)\big). \end{aligned} \tag{3.2}$$

We are looking for an *equilibrium* of this noncooperative game. This is a strategy pair $\delta^*,(q^*,\omega^*)$ so that

$$\begin{aligned} I\big(\delta^*,(q^*,\omega^*)\big) &\geq I\big(\delta,(q^*,\omega^*)\big) && \text{for all } \delta\in\Delta, \\ V\big(\delta^*,(q^*,\omega^*)\big) &\geq V\big(\delta^*,(q,\omega)\big) && \text{for all } q\in[0,1],\quad \omega\in\Omega. \end{aligned} \tag{3.3}$$

Usually, there is no equilibrium in which the inspectee acts with certainty legally ($q^*=0$) or illegally ($q^*=1$). Namely, if $q^*=0$, then by (3.2), the inspector would choose a test δ^* with $\alpha(\delta^*)=0$ excluding a false alarm. However, then the equilibrium condition for the inspectee in (3.3) requires $-b+(1+b)\,\beta(\delta^*,\omega^*)\leq 0$, which means that the non-detection probability $\beta(\delta^*,\omega^*)$ has to be sufficiently low. This is usually not possible with a test δ^* that has false alarm probability zero. On the contrary, such a test usually has a non-detection probability of one. Similarly, if $q^*=1$, then the inspector could always raise an alarm irrespective of his observation to maximize his payoff, so that $\alpha(\delta^*)=1$ and $\beta(\delta^*,\omega)=0$. However, then the equilibrium choice of the inspectee would not be $q^*=1$ since $h<b$ by (3.1). Thus, in equilibrium,

$$0<q^*<1\,, \tag{3.4}$$

and the inspectee is indifferent between legal and illegal behavior:

$$-h\,\alpha(\delta^*) \;=\; -b+(1+b)\,\beta(\delta^*,\omega^*)\,. \tag{3.5}$$

With respect to the non-detection probability β, the payoff function in (3.2) is monotonically decreasing for the inspector and increasing for the inspectee. Thus, given any test δ, the inspectee will choose a violation mechanism ω so as to maximize

$\beta(\delta, \omega)$. In equilibrium, the inspectee will therefore determine ω^* (which depends on δ) so that $\beta(\delta, \omega^*) = \max_{\omega \in \Omega} \beta(\delta, \omega)$. Typically, β is continuous and ω ranges in a compact domain, so the maximum is achieved.

On the side of the inspector, it is useful to choose first a fixed false alarm probability α and then consider only those tests that result in this error probability. Denote this set of tests $\{\delta \in \Delta \mid \alpha(\delta) = \alpha\}$ by Δ_α. For these tests, the inspector's payoff in (3.2) depends only on the non-detection probability. Thus, he will choose that test δ^* in Δ_α that minimizes the worst-case non-detection probability $\beta(\delta^*, \omega^*)$. It follows that in equilibrium, the non-detection probability is represented by

$$\beta(\alpha) := \min_{\delta \in \Delta_\alpha} \max_{\omega \in \Omega} \beta(\delta, \omega). \tag{3.6}$$

That is, the inspector's strategy is a minmax strategy with respect to the non-detection probability. This suggests that, given a false alarm probability α, the equilibrium non-detection probability $\beta(\alpha)$ is the minmax value of a certain auxiliary *zero-sum* game G_α which we call the 'non-detection' game:

Definition. *Given the inspection game in Figure 3.1 and $\alpha \in [0,1]$, the* non-detection game *G_α is the zero-sum game $(\Delta_\alpha, \Omega, \beta)$, where: The set of strategies δ of the inspector is Δ_α, the set of all statistical tests with false alarm probability α. The set of strategies Ω of the inspectee consists of all violation strategies ω. The payoff to the inspectee is the non-detection probability $\beta(\delta, \omega)$, which he tries to maximize and the inspector tries to minimize.*

In other words, the non-detection game G_α captures one part of our original game, namely the situation resulting after the inspectee has decided to behave illegally and the inspector has decided to use tests with false alarm probability α. Thus, G_α does not depend directly on the parameters a, b, e, h of the original game, but only through the fact that these parameters determine the value of α in equilibrium.

In order to find an equilibrium as in (3.3), we assume that the minmax value $\beta(\alpha)$ in (3.6) of the game G_α is a function $\beta : [0,1] \to [0,1]$ that fulfills

$$\beta(0) = 1, \qquad \beta(1) = 0, \qquad \beta \text{ is convex, and continuous at } 0. \tag{3.7}$$

These conditions imply that β is continuous and monotonically decreasing. An extreme example is the function $\beta(\alpha) = 1 - \alpha$ which applies if the inspector ignores his observation z and calls an alarm with probability α, independently of the data. The properties (3.7) are standard for statistical tests. In particular, β is convex since the inspector may randomize.

Theorem 3.1. *Assume that for any false alarm probability α, the non-detection game G_α has a value $\beta(\alpha)$ fulfilling (3.7). Then the inspection game in Figure 3.1 has the equilibrium $\big(\delta^*, (q^*, \omega^*)\big)$, where: The false alarm probability $\alpha^* = \alpha(\delta^*)$ is the solution of*

$$-h\,\alpha^* \;=\; -b + (1+b)\,\beta(\alpha^*)\,. \tag{3.8}$$

The probability q^* of illegal behavior is given by

$$q^* = \frac{e}{e - (1-a)\beta'(\alpha^*)} \tag{3.9}$$

where $\beta'(\alpha^*)$ is a sub-derivative of the function β at α^*. The inspection strategy δ^* and the violation procedure ω^* are optimal strategies in G_{α^*}.

Proof. Equation (3.8) is implied by (3.5). The probability q^* of illegal behavior is determined as follows. By (3.7), the inspector's payoff in (3.2) is a convex combination of two concave functions of α, one decreasing, the other increasing. For $q = q^*$, the resulting concave function must have its maximum at α^* (see also Figure 5.6 in Section 5 below). This implies (3.9), if the derivative $\beta'(\alpha^*)$ of the function β at α^* exists. Otherwise, one may take any sub-derivative $\beta'(\alpha^*)$, defined as the derivative of any tangent supporting the convex function β at its argument α^*. Since the zero-sum game G_{α^*} has a value, its solution δ^*, ω^* is part of an equilibrium, as argued above. □

Usually, this equilibrium is *unique*. It is easy to see that by (3.7), equation (3.8) has a unique solution α^* in $(0, 1)$. If β is differentiable at α^*, then q^* is also a unique solution to (3.9). The optimal strategies in the game G_{α^*} are often unique as well. At any rate, they are equivalent since the value $\beta(\alpha^*)$ of the game is unique.

Theorem 3.1 shows that the inspection game is in effect 'decomposed' into two games which can be treated separately, such that the solution of one game becomes a parameter of the second game. Technically, this decomposition works as follows: Solve the non-zero sum game in which the strategies are $\alpha \in [0, 1]$ (false alarm probability) for the inspector, $q \in [0, 1]$ (probability for acting illegally) for the operator and in which the payoff functions are given by (3.2). The equilibrium (α^*, q^*) is determined by equations (3.8) and (3.9) and α^* becomes a parameter of the zero-sum game G_{α^*}. The solution of G_{α^*} determines the optimal test (of significance α^*) and the optimal diversion strategy.

The non-zero sum game which determines the false alarm and violation probabilities is of a 'political' nature. The significance of the numerical values of these probabilities, thus, may be debated by practitioners since the latter depend on payoffs which are highly questionable: they are hard to measure since they are supposed to quantify political effects. The zero-sum game G_{α^*}, which we may refer to as the non-detection game, is of a technical nature, devoid of political considerations: Given that the operator has decided to act illegally and the inspector has decided to restrict himself to false alarm probability α^*, what is the optimal statistical test for the inspector and what is the optimal illegal behavior (for the operator)? This game is less debatable 'politically' but usually more challenging mathematically. In fact, many inspection games in the literature deal only with zero-sum models of this kind. Operationally, the most important tool for solving G_α is the *Neyman Pearson Lemma* (see, for example, Lehmann, 1959). We *assume* first that the game G_α has a value, which can therefore be written as

$$\beta(\alpha) = \max_{\omega \in \Omega} \min_{\delta \in \Delta_\alpha} \beta(\delta, \omega).$$

Hereby, the non-detection probability $\beta(\delta, \omega)$ for a given false alarm probability α is minimized, assuming that ω (and thus the alternative hypothesis) is fully known. The corresponding *best test* δ^* is described by the Neyman Pearson Lemma. This lemma provides a test procedure to decide between two *simple hypotheses*, that is, hypotheses which determine the respective probability distributions completely. With δ^* thus found, it is often possible to describe an optimal violation procedure ω^* against this test. Then, it remains to show that these strategies form an equilibrium of the game G_α. In the following, we will demonstrate the use of the Neyman Pearson Lemma with two important applications.

3.2. Material Accountancy

Material accountancy procedures are designed to control rare, dangerous, or precious materials. In particular, the material balance concept is fundamental to IAEA safeguards, where the inspector watches for a possible diversion of nuclear material and the inspectee is the operator of a nuclear plant. We consider the case of periodic inventory measurements. Further details of this model are described in Avenhaus (1986, Section 3.3).

Consider a certain physical material balance area and a sequence of n *inventory periods*. At the beginning of the first period, the amount I_0 of material under control is measured in the area. During the ith period, $1 \leq i \leq n$, the known net amount S_i of material enters the area, and at the end of that period, the amount I_i is measured in the area. The quantity

$$Z_i = I_{i-1} + S_i - I_i, \qquad 1 \leq i \leq n,$$

is called the material balance test statistic for the ith inventory period. Under H_0, its expected value is zero (because of the preservation of matter),

$$E_0(Z_i) = 0, \qquad 1 \leq i \leq n. \tag{3.10}$$

Under H_1, its expected value is μ_i,

$$E_1(Z_i) = \mu_i, \qquad 1 \leq i \leq n, \qquad \sum_i \mu_i = \mu, \tag{3.11}$$

where μ_i is the amount of material diverted in the ith period, and where μ is the *total* amount of material to be diverted throughout all periods.

Given normally distributed measurement errors, the random (column) vector $\boldsymbol{Z} = (Z_1, \ldots, Z_n)^\top$ is multivariate normally distributed with covariance matrix $\boldsymbol{\Sigma}$, which is the same for H_0 and H_1. The elements of this matrix are given by

$$\operatorname{cov}(Z_i, Z_j) = \begin{cases} \operatorname{var}(Z_i) & \text{for } i = j \\ -\operatorname{var}(I_i) & \text{for } i = j-1 \\ -\operatorname{var}(I_j) & \text{for } j = i-1 \\ 0 & \text{otherwise.} \end{cases}$$

The two hypotheses to be tested by the inspector are defined by the expectation vectors $E_0(\boldsymbol{Z})$ and $E_1(\boldsymbol{Z})$ as given by (3.10) and (3.11).

According to our general model, the 'best' test procedure is an equilibrium strategy of the inspector. By Theorem 3.1, this is an optimal strategy of the non-detection game $G_\alpha = (\Delta_\alpha, \Omega, \beta)$. In this zero-sum game, the inspector's set of strategies Δ_α is, as above, the set of test procedures with a fixed false alarm probability α. The inspectee's set of strategies Ω is the set of all diversion vectors $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)^\top$ that have fixed total diversion μ, that is, fulfill $\mathbf{1}^\top \boldsymbol{\mu} = \mu$ with $\mathbf{1} = (1, \ldots, 1)^\top$. We assume – and will show later – that this game has an equilibrium. Furthermore, this equilibrium involves only pure strategies of the players.

For fixed diversion $\boldsymbol{\mu}$, the alarm set of the best test (the set of observations where the null hypothesis is rejected) is described by the Neyman Pearson Lemma. It is given by the set of observations $\boldsymbol{z} = (z_1, \ldots, z_n)^\top$

$$\{ \, \boldsymbol{z} \mid \frac{f_1(\boldsymbol{z})}{f_0(\boldsymbol{z})} > \lambda \, \}, \tag{3.12}$$

where f_1 and f_0 are the densities of the random vector $\boldsymbol{Z} = (Z_1, \ldots, Z_n)^\top$ under H_1 respectively H_0. These are

$$f_1(\boldsymbol{z}) = \frac{1}{\sqrt{(2\pi)^n \cdot |\boldsymbol{\Sigma}|}} \exp\Bigl(-\frac{1}{2} \cdot (\boldsymbol{z} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{z} - \boldsymbol{\mu})\Bigr),$$

and the density f_0 has the same form with $\boldsymbol{\mu}$ replaced by the zero vector. Finally, λ is a parameter that is determined according to the false alarm probability α. This alarm set is equivalent to the set

$$\{ \, \boldsymbol{z} \mid \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{z} > \lambda' \} \tag{3.13}$$

for some λ', which means that the test statistic is the linear expression $\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}$.

Since $\boldsymbol{Z}$ is multivariate normally distributed, a linear combination $\boldsymbol{a}^\top \boldsymbol{Z}$ of its components, for any n-vector $\boldsymbol{a}$, is univariate normally distributed with expectation 0 under H_0 and $\boldsymbol{a}^\top \boldsymbol{\mu}$ under H_1, and variance $\boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a}$. Thus, the expected value of $\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}$ is

$$E(\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}) = \begin{cases} 0 & \text{under } H_0 \\ \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} & \text{under } H_1 \end{cases}$$

and its variance

$$\operatorname{var}(\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{Z}) = \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \, .$$

Surprisingly, mean and variance are equal under H_1. Therefore, the probability of detection as a function of the false alarm probability is given by

$$1 - \beta = \Phi\Bigl(\sqrt{\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}} - q_{1-\alpha}\Bigr) \tag{3.14}$$

where Φ denotes the Gaussian or normal distribution function and $q_{1-\alpha}$ the corresponding quantile (given by the inverse of Φ).

Theorem 3.2. *The optimal strategy of the inspectee in the non-detection game G_α for material accountancy is*

$$\boldsymbol{\mu}^* = \boldsymbol{\Sigma 1} \cdot \frac{\mu}{\mathbf{1}^\top \boldsymbol{\Sigma 1}}\,. \tag{3.15}$$

The value β^* of the game is given by the equation

$$1 - \beta^* = \Phi\Big(\frac{\mu}{\sqrt{\mathbf{1}^\top \boldsymbol{\Sigma 1}}} - q_{1-\alpha}\Big), \tag{3.16}$$

which describes the probability of detection in equilibrium. The optimal strategy δ^* of the inspector is based on the test statistic $\mathbf{1}^\top \boldsymbol{Z}$ with the alarm set $\{\, \boldsymbol{z} \mid \mathbf{1}^\top \boldsymbol{z} > \sqrt{\mathbf{1}^\top \boldsymbol{\Sigma 1}} \cdot q_{1-\alpha}\}$. The strategy pair $\delta^*, \boldsymbol{\mu}^*$ is an equilibrium of G_α.

Proof. By (3.14), the optimal diversion vector $\boldsymbol{\mu}^*$ minimizes $\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$ for all $\boldsymbol{\mu}$ subject to $\mathbf{1}^\top \boldsymbol{\mu} = \mu$. Since $\boldsymbol{\Sigma}^{-1}$ is symmetric, the gradient of this objective function is $2\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1}$, whereas the gradient of the constraint equation is $\mathbf{1}^\top$. For the optimum $\boldsymbol{\mu} = \boldsymbol{\mu}^*$, these gradients are collinear with a suitable Lagrange multiplier, which means that $\boldsymbol{\mu}^*$ is a multiple of $\boldsymbol{\Sigma 1}$ as in (3.15). With $\boldsymbol{\mu} = \boldsymbol{\mu}^*$ in (3.14) one obtains (3.16). The test statistic used in the alarm set (3.13), suitably scaled, is $\mathbf{1}^\top \boldsymbol{Z}$. Since it has variance $\mathbf{1}^\top \boldsymbol{\Sigma 1}$, the alarm set corresponding to the false alarm probability α is $\{\, \boldsymbol{z} \mid \mathbf{1}^\top \boldsymbol{z} > \sqrt{\mathbf{1}^\top \boldsymbol{\Sigma 1}} \cdot q_{1-\alpha}\}$.

Finally, we show that these optimal strategies fulfill the equilibrium criterion

$$\beta(\delta^*, \boldsymbol{\mu}) \le \beta(\delta^*, \boldsymbol{\mu}^*) \le \beta(\delta, \boldsymbol{\mu}^*) \qquad \text{for all } \delta \in \Delta_\alpha,\ \boldsymbol{\mu} \in \Omega.$$

Since we have constructed the Neyman Pearson test for any diversion strategy $\boldsymbol{\mu}$, the right inequality is just the Neyman Pearson Lemma applied to $\boldsymbol{\mu}^*$. The left inequality holds for all $\boldsymbol{\mu}$ as equality since the optimal test δ^* is based on the statistic $\mathbf{1}^\top \boldsymbol{Z}$ which is normally distributed with mean $\mathbf{1}^\top \boldsymbol{\mu}$ and variance $\mathbf{1}^\top \boldsymbol{\Sigma 1}$, that is, this distribution, and hence the detection probability, does not depend on the diversion vector $\boldsymbol{\mu}$, as long as $\mathbf{1}^\top \boldsymbol{\mu} = \mu$. □

This solution deserves several remarks. Both players use *pure* strategies. In particular, the inspectee does not randomize over diversion plans. By proving the equilibrium property with the Neyman Pearson Lemma, we did not have to consider all possible tests δ of the inspector explicitly. Even defining the set of these strategies would be complicated. The optimal test procedure δ^* does not depend on the total diversion μ. This adds to the appeal of the model since the total diversion μ might not be known. It is a parameter of the non-detection game which is only relevant for the violation strategy in (3.15). The optimal test statistic of the inspector is

$$\mathbf{1}^\top \boldsymbol{Z} = \sum_{i=1}^n Z_i = I_0 \;+\; \sum_{i=1}^n S_i \;-\; I_n$$

which is the *overall material balance for all periods.* This means that the intermediate inventories are *ignored.* In fact, the same result is obtained if one considers the problem of subdividing a plant into several material balance areas: For the general payoff structure used so far, the equilibrium strategy of the inspector does *not* require subdividing the plant.

This example demonstrates the power of the Neyman Pearson Lemma of statistics in connection with the game-theoretic concept of an equilibrium. We should emphasize that, as in many game-theoretic problems, finding a pair of equilibrium strategies can be rather difficult, while it is usually easy to verify its properties.

3.3. Data Verification

As already mentioned, IAEA safeguards is based on material accountancy, which derives from the physical principle of the preservation of matter. A second, operational principle of the IAEA is *data verification*. The inspector has to compare the material balance data reported by the plant operators (via their national authorities) with his own findings in order to verify that these data are not falsified for the purpose of concealing diversion of nuclear material. Since both sets of data, those of the operators and of the inspector, are based on measurements, statistical errors cannot be avoided. Furthermore, since one has to assume that the operators will – if at all – falsify their reported data strategically, decision problems arise which are very similar to those analyzed before. Again, the Neyman Pearson Lemma has to be applied, and again very interesting solutions, both from a theoretical and practical point of view, are obtained. We present an analytical model taken from Avenhaus (1986). It has a number of variations which have been solved in the literature or which are a topic of current research, as outlined below.

Assume that a nuclear plant is shut down and that the plant operator, the inspectee, has collected the nuclear material in K classes with N_i batches in the ith class, $i = 1, \ldots, K$. Typically, each class represents a particular type of measurement method. The operator has measured all contents of all batches and has reported the results to an inspector sent by the IAEA. The inspector, in turn, measures some of the material contents in order to compare them with the reported ones. His decision problem will be how to allocate his limited inspection effort. Before we can describe this, we have to make assumptions about the measurement statistics, and, as we will see, about the possible falsification of the data.

Let ν_{ij} be the true value of the content of the jth batch of the ith class, and let e_{ij} be the operator's random measurement error. The measurement result x_{ij} can therefore be written as

$$x_{ij} = \nu_{ij} + e_{ij} \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, N_i\,. \tag{3.17}$$

We assume that the random variables E_{ij} associated to the errors are independent and normally distributed with zero expectation and variances $\sigma^2_{O_i}$ that are specific to the class i and to the operator's instruments, written as

$$E_{ij} \sim N(0, \sigma^2_{O_i}) \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, N_i\,.$$

Under the null hypothesis H_0, no data are falsified. Then, the data reported by the operator are given by (3.17). Under the alternative hypothesis H_1, the operator is assumed to falsify all batch data of one class by a class specific positive amount μ_i, so that each batch in that class is reported to contain μ_i more material than it actually does. In that case, the reported data are given by

$$x_{ij} = \nu_{ij} + \mu_i + e_{ij} \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, N_i\,.$$

Of course, this is a very restrictive assumption about the falsification that takes place. It may be justified with practical arguments, for example manipulation of the calibration of the measuring instruments. More general assumptions are very difficult to handle. We will mention approaches in that direction below.

The inspector verifies n_i randomly chosen items ith class with the help of own measurements. Without loss of generality, let these be the first n_i ones. Then, the inspector's findings y_{ij} are

$$y_{ij} = \nu_{ij} + d_{ij} \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, n_i\,.$$

The random variables D_{ij} associated to the inspector's errors d_{ij} are independent and normally distributed with zero expectation and variances $\sigma^2_{I_i}$,

$$D_{ij} \sim N(0, \sigma^2_{I_i}) \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, n_i\,.$$

The inspector is *not* interested in estimating the true values ν_{ij} but only in verifying the reported data. Therefore, he uses the differences

$$z_{ij} = x_{ij} - y_{ij} \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, n_i$$

for constructing his verification test. According to our assumptions, we have

$$Z_{ij} \sim N(\mu_{it}, \sigma_i^2) \qquad \text{for } i = 1, \ldots, K, \quad j = 1, \ldots, n_i, \quad t = 0, 1,$$

where $\sigma_i^2 = \sigma^2_{O_i} + \sigma^2_{I_i}$, and where

$$\mu_{it} = \begin{cases} 0 & \text{for } t = 0 \ (H_0), \\ \mu_i & \text{for } t = 1 \ (H_1). \end{cases}$$

The inspector uses the Neyman Pearson test. As above in (3.12), its alarm set is with $\boldsymbol{z} = (z_{11}, \ldots, z_{Kn_K})$ given by

$$\{\, \boldsymbol{z} \mid \frac{f_1(\boldsymbol{z})}{f_0(\boldsymbol{z})} > \lambda \,\},$$

where for $t = 0, 1$

$$f_t(\boldsymbol{z}) = \Big(\prod_{i=1}^{K} \frac{1}{(\sqrt{2\pi}\sigma_i)^{n_i}}\Big) \cdot \exp\Big(-\frac{1}{2} \sum_{i=1}^{K} \sum_{j=1}^{n_i} \frac{(z_{ij} - \mu_{it})^2}{\sigma_i^2}\Big).$$

Therefore, the alarm set is explicitly given by

$$\{ \boldsymbol{z} \mid \sum_{i=1}^{K} \frac{\mu_i}{\sigma_i^2} \sum_{j=1}^{n_i} z_{ij} \ > \ \lambda' \}.$$

Now we have under H_t, $t = 0, 1$,

$$\sum_{i=1}^{K} \frac{\mu_i}{\sigma_i^2} \sum_{j=1}^{n_i} Z_{ij} \ \sim \ N\Big(\sum_{i=1}^{K} n_i \frac{\mu_{it}^2}{\sigma_i^2}, \ \sum_{i=1}^{K} n_i \frac{\mu_i^2}{\sigma_i^2}\Big).$$

Thus, the probabilities of error of the first and second kind are given by

$$\alpha = 1 - \Phi\Big(\frac{\lambda'}{\sqrt{\sum_{i=1}^{K} n_i \frac{\mu_i^2}{\sigma_i^2}}}\Big) \qquad \text{and} \qquad \beta = \Phi\Big(\frac{\lambda' - \sum_{i=1}^{K} n_i \frac{\mu_i^2}{\sigma_i^2}}{\sqrt{\sum_{i=1}^{K} n_i \frac{\mu_i^2}{\sigma_i^2}}}\Big).$$

If we eliminate λ' with the help of α, then we get the non-detection probability

$$\beta = 1 - \Phi\Big(\sqrt{\sum_{i=1}^{K} n_i \frac{\mu_i^2}{\sigma_i^2}} - q_{1-\alpha}\Big). \tag{3.18}$$

In order to describe the strategy sets of the players, let us assume that the operator wants to falsify all data by the fixed total amount μ. Furthermore, let one measurement by the inspector in the ith class, $i = 1, \ldots, K$ cost the effort ε_i, where the total effort available, measured in monetary terms or inspection hours, is ε. This means, according to Theorem 3.1, that we consider the following non-detection game G_α.

Theorem 3.3. *Let G_α be the zero-sum game in which the sets of strategies for inspector and inspectee are, respectively,*

$$\{ \boldsymbol{n} = (n_1, \ldots, n_K) \mid \sum_{i=1}^{K} \varepsilon_i n_i = \varepsilon \} \qquad \textit{and} \qquad \{ \boldsymbol{\mu} = (\mu_1, \ldots, \mu_K) \mid \sum_{i=1}^{K} N_i \mu_i = \mu \}.$$

The payoff to the inspectee is $\beta(\boldsymbol{n}, \boldsymbol{\mu})$ defined by (3.18). Considering $\boldsymbol{n}$ as a vector of continuous variables, this game has an equilibrium $\boldsymbol{n}^*, \boldsymbol{\mu}^*$, given by

$$\begin{aligned} n_i^* &= \frac{\varepsilon}{\sum_{j=1}^{K} N_j \sigma_j \sqrt{\varepsilon_j}} \cdot \frac{N_i \sigma_i}{\sqrt{\varepsilon_i}} \\ \mu_i^* &= \frac{\mu}{\sum_{j=1}^{K} N_j \sigma_j \sqrt{\varepsilon_j}} \cdot \sigma_i \sqrt{\varepsilon_i} \end{aligned} \qquad \textit{for } i = 1, \ldots, K. \tag{3.19}$$

The value of the game is

$$\beta(\boldsymbol{n}^*, \boldsymbol{\mu}^*) = 1 - \Phi\Big(\frac{\mu\sqrt{\varepsilon}}{\sum_{j=1}^{K} N_j \sigma_j \sqrt{\varepsilon_j}} - q_{1-\alpha}\Big).$$

Proof. The equilibrium conditions say that for all $\boldsymbol{n}, \boldsymbol{\mu}$ in the strategy sets

$$\beta(\boldsymbol{n}^*, \boldsymbol{\mu}) \ \leq \ \beta(\boldsymbol{n}^*, \boldsymbol{\mu}^*) \ \leq \ \beta(\boldsymbol{n}, \boldsymbol{\mu}^*)$$

holds. By (3.18), this is equivalent to

$$\sum_{i=1}^{K} n_i^* \frac{\mu_i^2}{\sigma_i^2} \ \geq \ \sum_{i=1}^{K} n_i^* \frac{\mu_i^{*2}}{\sigma_i^2} \ \geq \ \sum_{i=1}^{K} n_i \frac{\mu_i^{*2}}{\sigma_i^2} \qquad \text{for all } \boldsymbol{n}, \boldsymbol{\mu}.$$

The right inequality is identically fulfilled. Since $\sum_{i=1}^{K} N_i \mu_i = \mu$, the left inequality is just a special case of the Cauchy-Schwarz inequality $(\sum a_i^2)(\sum b_i^2) \geq (\sum a_i b_i)^2$ with $a_i^2 = N_i \mu_i^2 / \sigma_i \sqrt{\varepsilon_i}$ and $b_i^2 = N_i \sigma_i \sqrt{\varepsilon_i}$. $\square$

The sampling design defined by (3.19) is known in the statistical literature as *Neyman-Chuprov* sampling (see, for example, Cochran, 1963). The above model provides a game-theoretic justification for this sampling design.

This data verification problem has a lot of ramifications if one assumes more general falsification strategies. It may be unrealistic to assume that all batches in class i are falsified by the same amount μ_i. In particular, this is the case for a high total diversion μ that can no longer be hidden in measurement errors for the individual batches. In that case, it turns out that the inspectee will concentrate his falsification on as few batches as possible, in the hope that these will not be measured by the inspector. A theoretical analysis of the general case is very difficult. Some approaches in this direction have been made for unstratified problems, that is, only a single class of batches; see Avenhaus, Battenberg and Falkowski (1991), Mitzrotsky (1993), Avenhaus and Piehlmeier (1994), Piehlmeier (1996), and Battenberg and Falkowski (1997).

4. Sequential Inspections

Certain inspections evolve over a number of stages. The number of inspections may be limited, and statistical errors may occur. The models presented in this section vary with respect to the information of the players about the course of the game, the rules for detecting a violation, and the resulting payoffs.

4.1. Recursive Inspection Games

We consider sequential inspections with a finite number of inspection stages. Their models can be defined recursively and their solution is given by a recurrence equation, which in some cases is solved explicitly. The recursive description implies that each player learns about the action of the other player after each stage, which is not always justified. This is not problematic if the inspectee can violate at most once, so we consider this case first.

The inspection game introduced by Dresher (1962) has n stages. At each stage, the inspector may or may not use an inspection, with a total of m inspections permitted for all stages. 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.

Dresher modeled this game recursively, assuming zero-sum payoffs. For the inspector, this payoff is one unit for a detected violation, zero for legal action throughout the game, and a loss of one unit for an undetected violation. The value of the game, the equilibrium payoff to the inspector, is denoted by $I(n,m)$ for the parameters $0 \le m \le n$. For $m = n$, the inspector will inspect at every stage and the inspectee will act legally, and similarly the decision is unique for $m = 0$ where the inspectee can safely violate, so that

$$I(n,n) = 0 \quad \text{for } n \ge 0 \qquad \text{and} \qquad I(n,0) = -1 \quad \text{for } n > 0. \tag{4.1}$$

For $0 < m < n$, the game is represented by the recursive payoff matrix as shown in Figure 4.1. The rows denote the possible actions at the first stage for the inspector and the columns those for the inspectee. If the inspectee violates, then he is either caught, where the game terminates and the inspector receives 1, or not, after which 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 and $m-1$ or m inspections left.

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

Figure 4.1. 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.

In this game, it is reasonable to assume a circular structure of the players' preferences. That is, the inspector prefers to use his inspection if and only if the inspectee violates, who in turn prefers to violate if and only if the inspector does not inspect. This means that the game has a unique mixed equilibrium point where both players choose both of their actions with positive probability. Hence, if the inspector's probability for inspecting at the first stage is p, then the inspectee must be indifferent between legal action and violation, that is,

$$p \cdot I(n-1,m-1) \;+\; (1-p) \cdot I(n-1,m) = p \;+\; (1-p) \cdot (-1). \tag{4.2}$$

Both sides of this equation denote the game value $I(n,m)$. Solving for p and substituting yields

$$I(n,m) = \frac{I(n-1,m) + I(n-1,m-1)}{I(n-1,m) + 2 - I(n-1,m-1)}\,. \tag{4.3}$$

With this recurrence equation for $0 < m < n$ and the initial conditions (4.1), the game value is determined for all parameters. Dresher showed that equations (4.1) and (4.3) have an explicit solution, namely

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

The *payoffs* in Dresher's model have been generalized by Höpfinger (1971). He assumed that the inspectee's gain for a successful violation need not equal his loss if he is caught, and also solved the resulting recurrence equation explicitly. Furthermore, zero-sum payoffs are not fully adequate since a caught violation, compared to legal action throughout, is usually undesirable for *both* players since for the inspector this demonstrates a failure of his surveillance system. A non-zero-sum game that takes account of this is shown in Figure 4.2. There, $I(n,m)$ and $V(n,m)$ denote the equilibrium payoff to the inspector and inspectee, respectively ('V' indicating a potential violator), and a and b are positive parameters denoting the losses to both players for a caught violation, where $a < 1$. Analogous to (4.1), the cases $m = n$ and $m = 0$ are described by

$$\begin{matrix} I(n,n) = 0 \\ V(n,n) = 0 \end{matrix} \quad \text{for } n \geq 0 \qquad \text{and} \qquad \begin{matrix} I(n,0) = -1 \\ V(n,0) = 1 \end{matrix} \quad \text{for } n > 0. \tag{4.4}$$

Inspector \ Inspectee	legal action	violation
inspection	$V(n-1,m-1)$ $I(n-1,m-1)$	$-b$ $-a$
no inspection	$V(n-1,m)$ $I(n-1,m)$	1 -1

Figure 4.2. Non-zero-sum game at the first of n stages with m inspections, $0 < m < n$, and equilibrium payoff $I(n,m)$ to the inspector and $V(n,m)$ to the inspectee. Legal action throughout the game has reference payoff zero to both players. A caught violation yields negative payoffs to both, where $0 < a < 1$ and $0 < b$.

Theorem 4.1. *The recursive inspection game in Figure 4.2 with initial conditions (4.4) has a unique equilibrium. The inspectee's equilibrium payoff is given by*

$$V(n,m) = \binom{n-1}{m} \Big/ \sum_{i=0}^{m} \binom{n}{i} b^{m-i} \tag{4.5}$$

and the inspector's payoff by

$$I(n,m) = -\binom{n-1}{m} \Big/ \sum_{i=0}^{m} \binom{n}{i} (-a)^{m-i}. \tag{4.6}$$

The equilibrium strategies are determined inductively, using (4.5) and (4.6), by solving the game in Figure 4.2.

Proof. As an inductive hypothesis, assume that the player's payoffs in Figure 4.2 are such that the game has a unique mixed equilibrium; this is true for $n = 2$, $m = 1$ by (4.4). The inspectee is then indifferent between legal action and violation. As in (4.2) and (4.3), this gives the recurrence equation

$$V(n,m) = \frac{b \cdot V(n-1,m) \; + \; V(n-1,m-1)}{V(n-1,m-1) + b + 1 - V(n-1,m)}$$

for $V(n,m)$, and a similar expression for $I(n,m)$. The verification of (4.5) and (4.6) is shown by Avenhaus and von Stengel (1992), which generalizes and simplifies the derivations by Dresher (1962) and Höpfinger (1971). The assumed inequalities about the mixed equilibrium can be proved using the explicit representations. □

The same payoffs, although with a different normalization, have already been considered by Maschler (1966). He derived an expression for the payoff to the inspector that is equivalent to (4.6), considering the appropriate linear transformations for the payoffs. Beyond this, Maschler introduced the *leadership* concept where the inspector announces and commits himself to his strategy in advance. Interestingly, this improves his payoff in the non-zero-sum game. We will discuss this in detail in Section 5.

The recursive description in Figure 4.1 and 4.2 assumes that all players know about the game state after each stage. This is usually not true for the inspector after a stage where he did not inspect, since then he will not learn about a violation. The constant payoffs after a successful violation indicate that the game terminates. In fact, the game continues, but any further action of the inspector is irrelevant since then the inspectee will only behave legally.

As soon as the the players' information is such that the recursive structure is not valid, the problem becomes difficult, as already noted by Kuhn (1963). Therefore, even natural and simple looking extensions of Dresher's model are still open problems. Von Stengel (1991) shows how the recursive approach is still possible with special assumptions. Similar sequential games with three parameters, where the game continues even after a detected violation, are described by Sakaguchi (1977;

1994) and Ferguson and Melolidakis (1996). The inspector's full information after each stage is not questioned in these models.

In the situation studied by von Stengel (1991) the inspectee intends at most k violations and can violate at most once per stage. The inspectee collects one unit for each successful violation. As before, n and m denote the number of stages and inspections. Like Dresher's game, the game is zero-sum, with value $I(n, m, k)$ denoting the payoff to the inspector. That value is determined if all remaining stages either will or will not be inspected, or if the inspectee does not intend to violate:

$$I(n,n,k) = 0, \qquad I(n,0,k) = -\min\{n,k\}, \qquad I(n,m,0) = 0. \tag{4.7}$$

If the inspector detects a violation, the game terminates and the inspectee has to pay a positive penalty b to the inspector for being caught, but the inspectee does not have to return the payoff he got for previous successful violations. This means that at a stage the inspector does not inspect, a violation reduces k to $k-1$ and a payoff of 1 is 'credited' to the inspectee, and legal action leaves k unchanged. Furthermore, it is *assumed* that even after a stage with no inspection, it becomes common knowledge whether there was a violation or not. This leads to a game with recursive structure, shown in Figure 4.3.

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

Figure 4.3. First stage of a zero-sum game with n stages, m inspections, and up to k intended violations, for $0 < m < n$. The inspectee collects one unit for each successful violation, which he can keep even if he is caught later and has to pay the nonnegative amount b. The inspector is informed after each stage, even with no inspection, if there was a violation or not.

Theorem 4.2. *The recursive inspection game in Figure 4.3 with initial conditions (4.7) has value*

$$I(n,m,k) = \frac{\binom{n-k}{m+1} - \binom{n}{m+1}}{s(n,m)} \qquad \textit{where} \qquad s(n,m) = \sum_{i=0}^{m} \binom{n}{i} b^{m-i}.$$

The probability p of inspection at the first stage is $p = s(n-1, m-1)/s(n,m)$.

Proof. Analogously to (4.3), on obtains a recurrence equation for $I(n, m, k)$, which has this solution shown by von Stengel (1991). □

In this game, the inspector's equilibrium payoff $I(n,m,k)$ decreases with the number k of intended violations. However, that payoff is constant for all $k \geq n-m$ since then $\binom{n-k}{m+1} = 0$. Indeed, the inspectee will not violate more than $n-m$ times since otherwise he would be caught with certainty.

Theorem 4.2 asserts that the the probability p of inspection at the first stage, determined analogously to (4.2), is given by $p = s(n-1,m-1)/s(n,m)$ and thus *does not depend on* k. This means that the inspector *need not know* the value of k in order to play optimally. Hence, the assumption in this game about the knowledge of the inspector after a period without inspection can be removed: The solution of the recursive game is also the solution of the game *without* recursive structure, where – more realistically – the inspector does not know what happened at a stage without inspection. This is analyzed by von Stengel (1991) using the extensive form of the games.

Some sequential models take into account statistical errors of the first and second kind (Maschler, 1967; Thomas and Nisgav, 1976; Baston and Bostock, 1991; Rinderle, 1996). Explicit solutions have been found only for special cases.

4.2. Timeliness Games

In certain situations, the inspector uses a number of inspections during some time interval in order to detect a single violation. This violation is detected at the earliest *subsequent* inspection, and the *time* that has elapsed in between is the payoff to the inspectee that the inspector wants to minimize. The inspectee may or may not observe the inspections.

In reliability theory, variants of this game have been studied by Derman (1961) and Diamond (1982). An operating unit may fail which creates a cost that increases with the time until the failure is detected. The overall time interval represents the time between normal replacements of the unit. A minmax analysis leads to a zero-sum game with the operating unit as inspectee which cannot observe any inspections. A more common approach in reliability theory, which is not our topic, is to assume some knowledge about the distribution of the failure time.

Another application are interim inspections of direct-use nuclear material (see, for example, Avenhaus, Canty, and von Stengel, 1991). Nuclear plants are regularly inspected, say at the end of the year. If nuclear material is diverted, one may wish to discover this not after a year but earlier, which is the purpose of interim inspections. One may assume either that the inspectee can violate just after he has observed an interim inspection, or that such observations or alterations of his diversion plan are not possible.

The players choose their actions from the unit time interval. The inspectee violates at some time s in $[0,1)$. The inspector has m inspections ($m \geq 1$) that he uses at times $t_1, \ldots, t_m$, where his last inspection must take place at the end of the interval, $t_m = 1$. He can choose the other inspection times freely, with $0 \leq t_1 \leq \cdots \leq t_{m-1} \leq t_m = 1$. The violation is discovered at the earliest inspection following the violation time s. Then, the time elapsed is the payoff to the inspectee given by the function

$$V(s,t_1,\ldots,t_{m-1}) = t_k - s \qquad \text{for } \ t_{k-1} \le s < t_k, \quad k=1,\ldots,m, \tag{4.8}$$

where $t_0 = 0$ and $t_m = 1$. It is useful to assume that if the times of inspection and violation coincide, $t_{k-1} = s$, then the violation is not detected until the next inspection at time t_k. The inspector's payoff is the negative of the payoff to the inspectee.

If the inspector deterministically inspects at times $t_i = i/m$ for $i = 1,\ldots,m$, a violation is discovered at most after time $1/m$ has elapsed. This is indeed an optimal inspection strategy if the inspectee can *observe* the inspections. In that case, the inspectee will violate *immediately* after the kth inspection assuming he can react without delay, where k is uniformly chosen from $\{0,\ldots,m-1\}$. That is, the violation time s is uniformly chosen from $\{t_0,t_1,\ldots,t_{m-1}\}$. Note that s is conditional upon the inspector's action. Then by (4.8) the expected time after the violation is discovered is ${}^1\!/_m \cdot (t_1 - t_0) + {}^1\!/_m \cdot (t_2 - t_1) + \cdots + {}^1\!/_m \cdot (t_m - t_{n-1}) = {}^1\!/_m \cdot (t_m - t_0) = 1/m$.

A more involved model of such observable inspections is studied by Derman (1961). Each interim inspection detects a violation only with a certain fixed probability. In addition to the time until detection, there is a cost to be paid for each inspection. Derman describes a deterministic minmax schedule for the inspection times where t_i is not proportional to i as before with $t_i = i/m$, but a quadratic function of i, so that the inspections accumulate towards the end of the time interval because of their cost.

If the inspections *cannot be observed*, then the violation time s and the free inspection times $t_1,\ldots,t_{m-1}$ are chosen simultaneously. This is an appropriate assumption for inspecting an unreliable unit if the inspections do not interfere with the operation of the system. This game is completely solved by Diamond (1982). We will demonstrate the solution of the first nontrivial case $m = 2$.

For $m = 2$, the inspector has one free inspection at time t_1, for short t, chosen from $[0,1]$. The inspectee may select his violation at time $s \in [0,1)$. This is a zero-sum game over the unit square with payoff $V(s,t)$ to the inspectee where, by (4.8),

$$V(s,t) = \begin{cases} t - s & \text{if } s < t, \\ 1 - s & \text{if } s \ge t. \end{cases} \tag{4.9}$$

This describes a kind of *duel* with time reversed where both players have an incentive to act early but after the other. By (4.9), the inspectee's payoff is too small if he violates too late, so that he will select s with a certain probability distribution from an interval $[0,b]$ where $b < 1$. Consequently, the inspector will not inspect later than b.

Theorem 4.3. *The zero-sum game over the unit square with payoff $V(s,t)$ in (4.9) has the following solution. The inspector chooses the inspection time t from $[0,b]$ according to the density function $p(t) = 1/(1-t)$, where $b = 1 - 1/e$. The inspectee* *chooses his violation time s according to the distribution function $Q(s) = 1/e(1-s)$ for $s \in [0, b]$, and $Q(s) = 1$ for $s > b$. The value of the game is $1/e$.*

Proof. The inspector chooses $t \in [0, b]$ according to a density function p, where

$$\int_0^b p(t)dt = 1. \tag{4.10}$$

The *expected* payoff $V(s)$ to the inspectee for $s \in [0, b]$ is then given as follows, taking into account that the inspection may take place before or after the violation:

$$\begin{aligned} V(s) &= \int_0^s (1-s)p(t)dt \;+\; \int_s^b (t-s)p(t)dt \\ &= \int_0^s p(t)dt \;+\; \int_s^b t\,p(t)dt \;-\; s\int_0^b p(t)dt. \end{aligned}$$

If the inspectee randomizes as described, then this payoff must be constant (see Karlin, 1959, Lemma 2.2.1, p. 27). That is, its derivative with respect to s, which by (4.10) is given by $p(s) - s\,p(s) - 1$, should be zero. The density for the inspection time is therefore given by

$$p(t) = \frac{1}{1-t},$$

with b in (4.10) given by $b = 1 - 1/e$. The constant expected payoff to the violator is then $V(s) = 1/e$. For $s > b$, the inspectee's payoff is $1 - s$ which is smaller than $1/e$ so that he will indeed not violate that late.

The optimal distribution of the violation time has an atom at 0 and a density q on the remaining interval $(0, b]$. We consider its distribution function $Q(s)$ denoting the probability that a violation takes place at time s or earlier. The mentioned atom is $Q(0)$, the derivative of Q is q. The resulting expected payoff $-V(t)$ for $t \in [0, b]$ to the inspector is given by

$$\begin{aligned} V(t) &= t\cdot Q(0) \;+\; \int_0^t (t-s)q(s)ds \;+\; \int_t^b (1-s)q(s)ds \\ &= t\cdot Q(0) \;+\; t\int_0^t q(s)ds \;+\; \int_t^b q(s)ds \;-\; \int_0^b s\,q(s)ds \\ &= t\cdot Q(t) \;+\; Q(b) \;-\; Q(t) \;-\; \int_0^b s\,q(s)ds\,. \end{aligned}$$

Again, the inspector randomizes only if this is a constant function of t, which means that $(t-1)\cdot Q(t)$ is a constant function of t. Thus, because $Q(b) = Q(1-1/e) = 1$, the distribution function of the violation time s is for $s \in [0, b]$ given by

$$Q(s) = \frac{1}{e}\cdot\frac{1}{1-s},$$

and for $s > b$ by $Q(s) = 1$. The nonzero atom is given by $Q(0) = 1/e$. □

Unsurprisingly, the value $1/e$ of the game with unobservable inspections is better for the inspector than the value $1/2$ of the game with observable inspections.

The solution the game for $m > 2$ is considerably more complex and has the following features. As before, the inspectee violates with a certain probability distribution on an interval $[0, b]$, leaving out the interval $(b, 1)$ at the end. The inspections take place in *disjoint intervals*. The inspection times are random, but it is a randomization over a *one parameter family* of pure strategies where the inspections are fully correlated. The optimal violation scheme of the inspectee is given by a distribution function which is piecewise defined on the $m - 1$ time intervals and has an atom at the beginning of the game. For large m, the inspection times are rather uniformly distributed in their intervals and the value of the game rapidly approaches $1/(2m - {}^4\!/_3)$ from above. Based on this complete analytic solution, Diamond (1982) also demonstrates computational solution methods for non-linear loss functions.

5. Inspector Leadership

The leadership principle says that it can be advantageous in a competitive situation to be the first player to select and stay with a strategy. It was suggested first by von Stackelberg (1934) for pricing policies. Maschler (1966) applied this idea to sequential inspections. The notion of leadership consists of two elements: The ability of to *announce* one's strategy first and make it known to the other player, and the ability of the player to *commit* himself to playing it. In a sense, it would be more appropriate to use the term *commitment power*. This concept is particularly suitable for inspection games since an inspector can credibly announce his strategy and stick to it, whereas the inspectee cannot do so if he intends to act illegally. Therefore, it is reasonable to assume that the inspector will take advantage of his leadership role.

5.1. Definition and Introductory Examples

A leadership game is constructed from a simultaneous game which is a game in normal form. Let the players be I and II with strategy sets Δ and Ω, respectively. For simplicity, we assume that Δ is closed under the formation of mixed strategies. In particular, Δ could be the set of mixed strategies over a finite set. The two players select simultaneously their respective strategies δ and ω and receive the corresponding payoffs, defined as expected payoffs if δ originally represents a mixed strategy.

In the leadership version of the game, one of the players, say player I, is called the *leader*. Player II is called the *follower*. The leader chooses a strategy δ in Δ and makes this choice known to the follower, who then chooses a strategy ω in Ω which will typically depend on δ and is denoted by $\omega(\delta)$. The payoffs to the players are those of the simultaneous game for the pair of strategies $(\delta, \omega(\delta))$. The strategy δ is executed without giving the leader an opportunity to reconsider it. This is the *commitment power* of the leader.

Simultaneous game

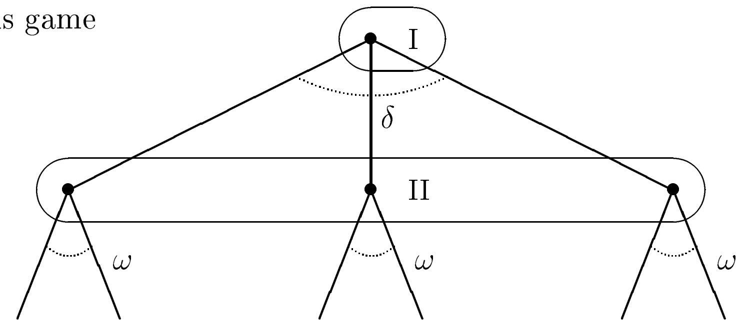

Leadership version

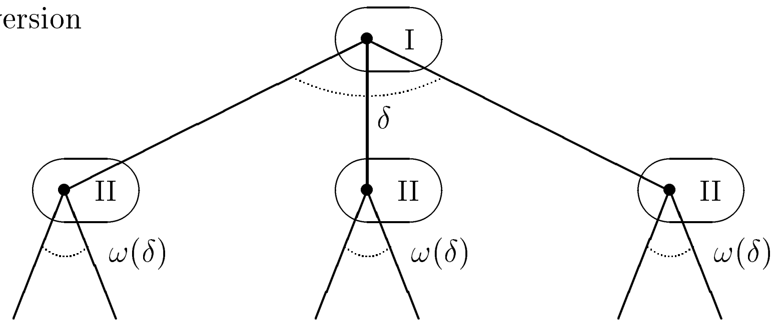

Figure 5.1. The simultaneous game and its leadership version in extensive form. The strategies δ of player I include all mixed strategies. In the simultaneous game, the choice ω of player II is the same irrespective of δ. In the leadership version, player I is the leader and moves first, and player II, the follower, may choose his strategy $\omega(\delta)$ depending on the announced leader strategy δ. The payoffs in both games are the same.

The simplest way to construct the leadership version is to start with the simultaneous game in extensive form. Player I moves first, and player II has a single information set since he is not informed about that move, and moves second. In the leadership game, that information set is dissolved into singletons, and the rest of the game, including payoffs, stays the same, as indicated in Figure 5.1.

The leadership game has perfect information because the follower is precisely informed about the announced strategy δ. He can therefore choose a best response $\omega(\delta)$ to δ. We assume that he always does this, even if δ is not a strategy announced by the leader in equilibrium, that is, we only consider *subgame perfect* equilibria. Any randomized strategy of the follower can be described as a *behavior strategy* defined by a 'local' probability distribution on the set Ω of choices of the follower for each announced strategy δ.

In a zero-sum game which has a value, leadership has no effect. This is the essence of the minmax theorem: Each player can guarantee the value even if he announces his (mixed) strategy to his opponent and commits himself to playing it.

I \ II	L	R
T	8 9	0 0
B	0 0	9 8

Figure 5.2. Game where the unique equilibrium (T, L) of the leadership version is one of several equilibria in the simultaneous game.

In a non-zero-sum game, leadership can sometimes serve as a coordination device and as a method for equilibrium selection. The game in Figure 5.2, for example, has two equilibria in pure strategies. There is no rule to select either equilibrium if the players are in symmetric positions. In the leadership game, if player I is made a leader, he will select T, and player II will follow by choosing L. In fact, (T, L) is the *only* equilibrium of the leadership game. Player II may even consider it advantageous that player I is a leader and accept the payoff 8 in this equilibrium (T, L) instead of 9 in the other pure strategy equilibrium (B, R) as a price for avoiding the undesirable payoff 0.

I \ II	L	R
T	1 -1	-1 $-1/2$
B	0 0	1 -1

Figure 5.3. A simultaneous game with a unique mixed equilibrium. Its leadership version has a unique equilibrium where the leader announces the same mixed strategy, but the follower changes his strategy to the advantage of the leader.

However, a simultaneous game and its leadership version may have different equilibria. The simultaneous game in Figure 5.3 has a unique equilibrium in mixed strategies: Player I plays T with probability $1/3$ and B with probability $2/3$, and player II plays L with probability $1/3$ and R with probability $2/3$. The resulting payoffs are $-2/3, 1/3$.

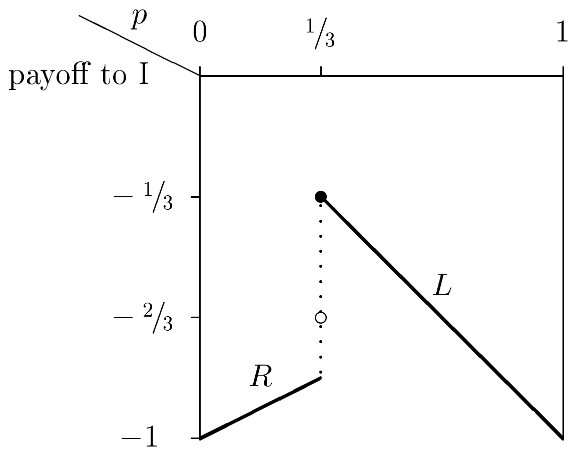

Figure 5.4. Payoff to player I in the game in Figure 5.3 as a function of the probability p that he plays T. It is assumed player II uses his best response (L or R), which is unique except for $p^* = 1/3$. Player I's equilibrium payoff in the simultaneous game is indicated by $\circ$, that in the leadership version by $\bullet$.

In the leadership version, player I as the leader uses the *same* mixed strategy as in the simultaneous game, and player II gets the *same* payoff, but as a follower he will act to the *advantage* of the leader, for the following reason. Consider the possible announced mixed strategies, given by the probability p that the leader plays T. If player II responds with L or R, the payoffs to player I are $-p$ and $p/2 - 1$, respectively. When player II makes these choices, his own payoffs are p and $1 - 2p$, respectively. He therefore chooses R for $p < 1/3$, L for $p > 1/3$, and is indifferent for $p = 1/3$. The resulting payoff to the leader as a function of p is shown in Figure 5.4. The leader tries to maximize this payoff, which he achieves only if the follower plays L. Thus, player I announces any p that is greater than or equal to $1/3$ but as small as possible. This means announcing exactly $1/3$, as pointed out by Avenhaus, Okada, and Zamir (1991): If the follower, who is then indifferent, would not choose L with certainty, then it is easy to see that the leader could announce a slightly larger p thus forcing the follower to play L and improve his payoff, in contradiction to the equilibrium property. Thus, the unique equilibrium of the leadership game is that player I announces $p^* = 1/3$ and player II responds with L and deterministically, as described, for all other announcements of p, which are not materialized.

5.2. Announced Inspection Strategy

Inspection problems are a natural case for leadership games since the inspector can make his strategies public. The inspectee cannot announce that he intends to violate with some positive probability. The example in Figure 5.3 demonstrates the effect of leadership in an inspection game, with the inspector as player I and the inspectee as player II. This game is a special case of the recursive game in Figure 4.2 for two stages ($n = 2$), where the inspector has one inspection ($m = 1$). Inspecting at the first stage is the strategy T, and not inspecting, that is, inspecting at the second stage, is represented by B. For the inspectee, L and R refer to legal action and violation at the first stage, respectively. The losses to the players in the case of a caught violation in Figure 4.2 have the special form $a = 1/2$ and $b = 1$.

In this leadership game, we have shown that the inspector announces his equilibrium strategy of the simultaneous game, but receives a better payoff. By the utility assumptions about the possible outcomes of an inspection, that better payoff is achieved by *legal* behavior of the inspectee. It can be shown that this applies not only to the simple game in Figure 5.3, but to the recursive game with general parameters n and m in Figure 4.2. However, the inspectee behaves legally only in the two-by-two games encountered in the recursive definition: If the inspections are used up ($n > 0$ but $m = 0$), then the inspectee can and will safely violate. Except for this case, the inspectee behaves legally if the inspector may announce his strategy.

The definition of this particular leadership game with general parameters n and m is due to Maschler (1966). However, he did not determine an equilibrium. In order to construct a solution, Maschler postulated that the inspectee would act to the advantage of the inspector if he is indifferent, and called this behavior 'pareto-optimal'. Then he argued, as we did with the help of Figure 5.4, that the inspector can announce an inspection probability p that is slightly higher than p^*. In that way, the inspector is on the safe side, which should also be recommended in practice. Maschler (1967) uses leadership in a more general model involving several detectors that can help the inspector determine under which conditions to inspect.

Despite the advantage of leadership, announcing such a precisely calibrated inspection strategy looks risky. Above, the equilibrium strategy $p^* = 1/3$ depends on the payoffs of the inspectee which might not be fully known to the inspector. Therefore, Avenhaus, Okada, and Zamir (1991) considered the leadership game for a simultaneous game with *incomplete information*. There, the gain to the inspectee for a successful violation is a payoff in some range with a certain probability distribution assumed by the inspector. In that leadership game, unlike in Figure 5.4, the inspector maximizes over a continuous payoff curve. He announces an inspection probability that just about forces the inspectee to act legally for *any* value of the unknown penalty b. This strategy has a higher equilibrium probability p^* and is on the safer side than in the simultaneous game.

The simple game in Figure 5.3 and the argument using Figure 5.4 is prototypical for many leadership games. In the same vein, we consider now the inspection game in Figure 3.1 as a simultaneous game, and construct its leadership version. Recall that in this game, the inspector has collected some data and, using this data, has to decide whether the inspectee acted illegally or not. He uses a statistical test procedure which is designed to detect an illegal action. Then, he either calls an alarm, rejecting the null hypothesis H_0 of legal behavior of the inspectee in favor of the alternative hypothesis H_1 of a violation, or not. The first and second error probabilities α and

β of a false rejection of H_0 or H_1, respectively, are the probabilities for a false alarm and an undetected violation.

As shown in Section 3.1, the only choice of the inspector is in effect the value for the false alarm probability α from $[0, 1]$. The non-detection probability β is then determined by the most powerful statistical test and defines the function $\beta(\alpha)$, which has the properties (3.7). Thereby, the convexity of β implies that the inspector has no advantage in choosing α randomly since this will not improve his non-detection probability. Hence, for constructing the leadership game as above in Section 5.1, the value of α can be announced deterministically. The possible actions of the inspectee are legal behavior H_0 and illegal behavior H_1. According to the payoff functions in (3.2), with $q = 0$ for H_0 and $q = 1$ for H_1, and the definition of $\beta(\alpha)$, this defines the game shown in Figure 5.5.

Inspector \ Inspectee	legal action H_0	violation H_1
$\alpha \in [0, 1]$	$-h\alpha$ (inspectee) $-e\alpha$ (inspector)	$-b + (1 + b)\beta(\alpha)$ (inspectee) $-a - (1 - a)\beta(\alpha)$ (inspector)

Figure 5.5. The game of Figure 3.1 with payoffs (3.2) depending on the false alarm probability α. The non-detection probability $\beta(\alpha)$ for the best test has the properties (3.7).

In an equilibrium of the simultaneous game, Theorem 3.1 shows that the inspector chooses α^* as defined by (3.8). Furthermore, the inspectee violates with positive probability q^* according to (3.4) and (3.9). This is different in the leadership version.

Theorem 5.1. *The leadership version of the simultaneous game in Figure 5.5 has a unique equilibrium where the inspector announces α^* as defined by (3.8). In response to α, the inspectee violates if $\alpha < \alpha^*$, and acts legally if $\alpha \geq \alpha^*$.*

Proof. The response of the inspectee is unique for any announcement of α except for α^*. Namely, the inspectee will violate if $\alpha < \alpha^*$, and act legally if $\alpha > \alpha^*$. The inspector's payoffs for these responses are shown in Figure 5.6. As argued with Figure 5.4, the only equilibrium of the leadership game is that in which the inspector announces α^*, and the inspectee acts legally in response. $\square$

Summarizing, we observe that in the simultaneous game, the inspectee violates with positive probability, whereas he acts legally in the leadership game. Inspector leadership serves as *deterrence*. The optimal false alarm probability α^* of the inspector stays the same. His payoff is larger in the leadership game.

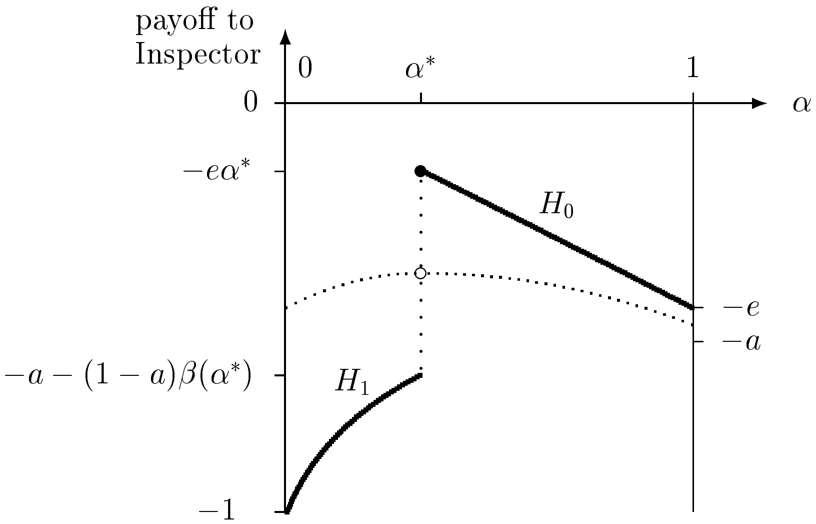

Figure 5.6. The inspector's payoff in Figure 5.5 as a function of α for the best responses of the inspectee. For $\alpha < \alpha^*$ the inspectee behaves illegally, hence the payoff is given by the curve H_1. For $\alpha > \alpha^*$ the inspectee behaves legally and the payoff is given by the line H_0. In the simultaneous game, the inspectee plays H_1 with probability q^* according to (3.9), so that the inspector's payoff, shown by the thin line, has its maximum for α^*; this is his equilibrium payoff, indicated by $\circ$. The inspector's equilibrium payoff in the leadership game is marked by $\bullet$.

6. Conclusions

One of our claims is that inspection games constitute a real field of applications of game theory. Is this justified? Have these games actually been used? Did game theory provide more than a general insight, did it have an operational impact?

The decision to implement a particular verification procedure is usually *not* based on a game-theoretic analysis. Practical questions overwhelm, and the allocation of inspection effort to various sites, for example, is usually based on rules of thumb. Most stratified sampling procedures of the IAEA are of this kind (IAEA, 1989). However, they can often be justified by game-theoretic means. We mentioned that the subdivision of a plant into small areas with intermediate balances has no effect on the overall detection probability if the violator acts strategically – such a physical separation may only improve localizing a diversion of material. With all caution with respect to the impact of a theoretical analysis, this observation may have influenced the design of some nuclear processing plants.

Another question concerns the proper scope of a game-theoretic model. For example, the course of second-level actions – after the inspector has raised an alarm – is often determined politically. In an inspection game, the effect of a detected violation is usually modeled by an unfavorable payoff to the inspectee. The particular magnitude of this penalty, as well as the inspectee's utility for a successful violation, is usually not known to the analyst. This is often used as an argument against game theory. As a counterargument, the signs of these payoffs often suffice, as we have illustrated with a rather general model in Theorem 3.1. Then, the part of the game where a violation is to be discovered can be reduced to a zero-sum game with the detection probability as payoff to the inspector, as first proposed by Bierlein (1968).

We believe that inspection models should be of this kind, where the merit of game theory as a technical tool becomes clearly visible. 'Political' parameters, like the choice of a false alarm probability, are exogenous to the model. Higher-level models describing the decisions of the states, like whether to cheat or not, and in the extreme whether 'to go to war' or not, should in our view not be blended with an inspection game. They are likely to be simplistic and will invalidate the analysis.

Which direction will or should research on inspection games take in the foreseeable future? The interesting concrete inspection games are stimulated by problems raised by practitioners. In that sense we expect continued progress from the application side, in particular in the area of environmental control where a fruitful interaction between theorists and environmental experts is still missing. Indeed there are open questions which shall be illustrated with the material accountancy example discussed in Section 3.2.

We have shown that intermediate inventories should be ignored in equilibrium. This means, however, that a decision about legal or illegal behavior can be made only at the end of the reference time. This may be too long if *detection time* becomes a criterion. If one models this appropriately, one enters the area of sequential games and sequential statistics. With a new 'non-detection game' similar to Theorem 3.1 and under reasonable assumptions one can show that then the probability of detection and the false alarm probability are replaced by the *average run lengths* under H_1 and H_0, respectively, that is, the expected times until an alarm is raised (see Avenhaus and Okada, 1992). However, they need not exist and, more importantly, there is no equivalent to the Neyman Pearson Lemma which gives us a constructive advice for the best sequential test. Thus, in general, sequential statistical inspection games represent a wide field for future research.

From a theoretical point of view, the leadership principle as applied to inspection games deserves more attention. In the case of sequential games, is legal behavior again an equilibrium strategy of the inspectee? How does the leadership principle work if more complicated payoff structures have to be considered? Does the inspector in general improve his equilibrium payoff by credibly announcing his inspection strategy? We think that for further research, a number of related ideas in economics, like principal–agent problems, should be compared more closely with leadership as presented here.

In this contribution, we have tried to demonstrate that inspection games have real applications and are useful tools for handling practical problems. Therefore, the major effort in the future, also in the interest of sound theoretical development, should be spent in deepening these applications, and – even more importantly – in trying to convince practitioners of the usefulness of appropriate game-theoretic models.

References

J. Altmann et al. (1992), eds., Verification at Vienna: Monitoring Reductions of Conventional Armed Forces. Gordon and Breach, Philadelphia.

F. J. Anscombe et al. (1963), eds., Applications of Statistical Methodology to Arms Control and Disarmament. Final report to the U.S. Arms Control and Disarmament Agency under contract No. ACDA/ST-3 by Mathematica, Inc., Princeton, New Jersey.

F. J. Anscombe et al. (1965), eds., Applications of Statistical Methodology to Arms Control and Disarmament. Final report to the U.S. Arms Control and Disarmament Agency under contract No. ACDA/ST-37 by Mathematica, Inc., Princeton, New Jersey.

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

R. Avenhaus (1994), Decision theoretic analysis of pollutant emission monitoring procedures. Annals of Operations Research 54, 23–38.

R. Avenhaus (1997), Entscheidungstheoretische Analyse der Fahrgast-Kontrollen. Der Nahverkehr 9/97, Alba Fachverlag, Düsseldorf, 27–30.

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

R. Avenhaus and M. J. Canty (1996), Compliance Quantified. Cambridge Univ. Press, Cambridge.

R. Avenhaus, M. J. Canty, and B. von Stengel (1991), Sequential aspects of nuclear safeguards: interim inspections of direct use material. Proceedings of the 4th International Conference on Facility Operations-Safeguards Interface, American Nuclear Society, Albuquerque, New Mexico, 104–110.

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

R. Avenhaus, A. Okada, and S. Zamir (1991), Inspector leadership with incomplete information. In: Game Equilibrium Models IV, ed. R. Selten, Springer, Berlin, 319–361.

R. Avenhaus and G. Piehlmeier (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.

R. Avenhaus and B. von Stengel (1992), Non-zero-sum Dresher inspection games. In: Operations Research '91, eds. P. Gritzmann et al., Physica-Verlag, Heidelberg, 376–379.

S. Baiman (1982), Agency research in managerial accounting: a survey. Journal of Accounting Literature 1, 154–213.

V. J. Baston and F. A. Bostock (1991), A generalized inspection game. Naval Research Logistics 38, 171–182.

H. P. Battenberg and B. J. Falkowski (1997), On saddlepoints of two-person zero-sum games with applications to data verification tests. International Journal of Game Theory, to appear.

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

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

C. G. Bird and K. O. Kortanek (1974), Game theoretic approaches to some air pollution regulation problems. Socio-Economic Planning Sciences 8, 141–147.

K. Borch (1982), Insuring and auditing the auditor. In: Games, Economic Dynamics, Time Series Analysis, eds. M. Deistler, E. Fürst, and G. Schwödiauer, Physica-Verlag, Würzburg, 117–126. Reprinted in: K. Borch (1990), Economics of Insurance, North-Holland, Amsterdam, 350–362.

K. Borch (1990), Economics of Insurance. Advanced Textbooks in Economics Vol. 29, North-Holland, Amsterdam.

W. M. Bowen and C. A. Bennett (1988), eds., Statistical Methodology for Nuclear Material Management. Report NUREG/CR-4604 PNL-5849, prepared for the U.S. Nuclear Regulatory Commission, Washington, D.C.

S. Brams and M. D. Davis (1987), The verification problem in arms control: a game theoretic analysis. In: Interaction and Communication in Global Politics, eds. C. Ciotti-Revilla, R. L. Merritt, and D. A. Zinnes, Sage, London, 141–161.

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

W. G. Cochran (1963), Sampling Techniques, 2nd edition. Wiley, New York.

C. Derman (1961), On minimax surveillance schedules. Naval Research Logistics Quarterly 8, 415–419.

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

M. Dresher (1962), A Sampling Inspection Problem in Arms Control Agreements: A Game-Theoretic Analysis. Memorandum No. RM–2972–ARPA, The RAND Corporation, Santa Monica, California.

R. A. Dye (1986), Optimal monitoring policies in agencies. RAND Journal of Economics 17, 339–350.

G. Feichtinger (1983), A differential games solution to a model of competition between a thief and the police. Management Science 29, 686–699.

T. S. Ferguson and C. Melolidakis (1996), On the inspection game. Manuscript, UCLA and Technical University of Crete.

J. A. Filar (1985), Player aggregation in the traveling inspector model. IEEE Transactions on Automatic Control AC-30, 723–729.

S. Gal (1983), Search Games. Academic Press, New York.

A. Yu. Garnaev (1994), A remark on the customs and smuggler game. Naval Research Logistics 41, 287–293.

A. J. Goldman (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, D.C.

A. J. Goldman and M. H. Pearl (1976), The dependence of inspection-system performance on levels of penalties and inspection resources. Journal of Research of the National Bureau of Standards 80B, 189–236.

J. Greenberg (1984), Avoiding tax avoidance: a (repeated) game-theoretic approach. Journal of Economic Theory 32, 1–13.

W. Güth and R. Pethig (1992), Illegal pollution and monitoring of unknown quality – a signaling game approach. In: Conflicts and Cooperation in Managing Environmental Resources, ed. R. Pethig, Springer, Berlin, 276–332.

E. Höpfinger (1971), A game-theoretic analysis of an inspection problem. Unpublished manuscript, University of Karlsruhe.

E. Höpfinger (1979), Dynamic standard setting for carbon dioxide. In: Applied Game Theory, eds. S. J. Brams, A. Schotter, and G. Schwödiauer, Physica-Verlag, Würzburg, 373–389.

J. T. Howson, Jr. (1972), Equilibria of polymatrix games. Management Science 18, 312–318.

IAEA (1989), IAEA Safeguards: Statistical Concepts and Techniques, Fourth Revised Edition. IAEA, Vienna.

J. L. Jaech (1973), Statistical Methods in Nuclear Material Control. Technical Information Center, United States Atomic Energy Commission TID-26298, Washington, D.C.

C. S. Kanodia (1985), Stochastic and moral hazard. Journal of Accounting Research 23, 175–193.

R. S. Kaplan (1973), Statistical sampling in auditing with auxiliary information estimates. Journal of Accounting Research 11, 238–258.

S. Karlin (1959), Mathematical Methods and Theory in Games, Programming, and Economics, Vol. II: The Theory of Infinite Games. Addison-Wesley, Reading.

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

D. M. Kilgour, N. Okada, A. Nishikori (1988), Load control regulation of water pollution: an analysis using game theory. Journal of Environmental Management 27, 179–194.

A. Klages (1968), Spieltheorie und Wirtschaftsprüfung: Anwendung spieltheoretischer Modelle in der Wirtschaftsprüfung. Schriften des Europa-Kollegs Hamburg, Vol. 6, Ludwig Appel Verlag, Hamburg.

H. W. Kuhn (1963), Recursive inspection games. In: Applications of Statistical Methodology to Arms Control and Disarmament, eds. F. J. Anscombe et al., Final report to the U.S. Arms Control and Disarmament Agency under contract No. ACDA/ST-3 by Mathematica, Inc., Princeton, New Jersey, Part III, 169–181.

M. Landsberger and I. Meilijson (1982), Incentive generating state dependent penalty system. Journal of Public Economics 19, 333–352.

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

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

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

E. Mitzrotsky (1993), Game-theoretical approach to data verification (in Hebrew). Master's Thesis, Department of Statistics, The Hebrew University at Jerusalem.

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

G. Piehlmeier (1996), Spieltheoretische Untersuchung von Problemen der Datenverifikation. Kovač, Hamburg.

J. F. Reinganum and L. L. Wilde (1986), Equilibrium verification and reporting policies in a model of tax compliance. International Economic Review 27, 739–760.

K. Rinderle (1996), Mehrstufige sequentielle Inspektionsspiele mit statistischen Fehlern erster und zweiter Art. Kovač, Hamburg.

A. Rubinstein (1979), An optimal conviction policy for offenses that may have been committed by accident. In: Applied Game Theory, eds. S. J. Brams, A. Schotter, and G. Schwödiauer, Physica-Verlag, Würzburg, 373–389.

W. H. Ruckle (1992), The upper risk of an inspection agreement. Operations Research 40, 877–884.

G. S. Russell (1990), Game models for structuring monitoring and enforcement systems. Natural Resource Modeling 4, 143–173.

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

M. Sakaguchi (1977), A sequential allocation game for targets with varying values. J. Operations Research Society of Japan 20, 182–193.

M. Sakaguchi (1994), A sequential game of multi-opportunity infiltration. Mathematica Japonicae 39, 157–166.

H. Schleicher (1971), A recursive game for detecting tax law violations. Economies et Sociétés 5, 1421–1440.

M. U. Thomas and Y. Nisgav (1976), An infiltration game with time dependent payoff. Naval Research Logistics Quarterly 23, 297–302.

H. von Stackelberg (1934), Marktform und Gleichgewicht. Springer, Berlin.

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

F. Weissing and E. Ostrom (1991), Irrigation institutions and the games irrigators play: rule enforcement without guards. In: Game Equilibrium Models II, ed. R. Selten, Springer, Berlin.