INSPECTOR LEADERSHIP WITH INCOMPLETE INFORMATION

by

Rudolf Avenhaus, Akira Okada and Shmuel Zamir

Abstract

Inspection problems arise when an inspector has to decide whether an inspectee has behaved legally, i.e., according to a formal agreement. For such a decision the inspector uses a well-defined procedure which is based on observations of random variables and which may or may not be announced to the inspectee.

Whereas in the latter case (no announcement) in equilibrium the inspectee will behave illegally with positive probability, in the former case he will act legally with certainty: the inspector deters the inspectee from illegal behavior by means of his "inspector leadership procedure".

In the case of complete information on both sides the leadership solution is instable. It is shown that in case the inspector has only incomplete information about the inspectee's payoff for undetected illegal action this instability disappears.

For the purpose of illustration the results are applied to material accountancy and data verification problems, and it is shown that the inspector's equilibrium strategies are the statistical tests commonly used in these contexts.

1. Introduction

Inspection problems arise when a person or organization (an *inspector*) has to decide whether a person, a group of persons, an organization or a State (*inspectee* in the following) has behaved legally, i.e., according to a law or a formal agreement. The inspector responsible for the surveillance of the behavior has two possibilities: He can declare the inspectee's behavior to be legal, or he can question it and induce a *second action level* in order to settle his suspicion.

Typically, the decision on these two alternatives has to be made with only uncertain information. Then stochastic methods have to be used. In terms of our theory, this means that the inspector is observing random variables whose distributions depend on the inspectee's behavior, and he uses these observations for deciding between two alternative hypotheses. In order to find the "best" decision procedure the inspector needs information about the utilities of the inspectee and about his possible actions. This means that inspection problems have to be modelled game theoretically in contrast to problems like process or quality control, where the adversary is inanimate nature or a technical process.

Apparently, the first analyses of this kind are due to DRESHER (1962) and KUHN (1963) who developed their ideas in the context of arms control and disarmament. MASCHLER (1967/68) makes decisive progress[1]: He abandons the zero-sum assumption and introduces the *price-leadership principle* into the field which he takes from STACKELBERG's work (1934) on market equilibria.

Without knowledge of this work, BIERLEIN (1969) analyzed zero-sum game theoretic models for nuclear material safeguards in the framework of the Non-Proliferation Treaty. In later papers (e.g., 1970,1983) he kept the zero-sum assumption but avoided the explicit use of payoffs to the inspector: BIERLEIN interpreted the zero-sum assumption such that ideally the inspector should consider the gain resp. loss of the inspectee as his loss resp. gain, and he was primarily interested in

[1] For the inspectee he uses the term *violator*. We do not use it, since it is not certain a priori (and, in fact, the inspector seeks to prevent this) that the inspectee violates the law or agreement. In the nuclear material safeguards application to be discussed later, the term *operator* of a nuclear plant under safeguards is sometimes used. We prefer the neutral term *inspectee* in order to indicate the wide range of applications of the theory.

reliable inspection strategies. HOPFINGER (1974) also determined reliable inspection strategies for a special inspection problem, where the total inspection effort is determined stochastically.

Except for a few attempts (MASCHLER (1967), BRAMS (1985)) the papers mentioned so far, and many more, are purely game theoretical in their scope; they do not deal with the evaluation of measurements and their errors which are typical features of many inspection problems. One of the first approaches in this direction was made by AVENHAUS and FRICK (1977) who emphasized that, in general, false alarms cannot be avoided and that, therefore, the interests of both the inspector and the inspectee are only partially opposing: Both parties want to avoid false alarms which cause losses to both. Inspection problems of this kind have to be modelled by non-zero-sum games.

The solution of concrete problems of this kind may become very complicated even in rather simple applications (see, e.g., FRICK (1976), BATTENBERG (1982)). Therefore, it turned out to be very convenient that the solution of games of this kind can be performed in two steps. In the first, the equilibrium false alarm probability is determined. In the second (and, for applications more important) step, a zero-sum game is solved. The solution is found with the help of the Lemma of Neyman and Pearson, well-known in statistical theory. In this way, a series of applied problems was solved (see, e.g., AVENHAUS (1986)).

A disadvantage of models where both players act simultaneously is that in equilibrium the inspectee behaves illegally with positive probability, whereas the primary objective of any inspection is the total deterrence of such behavior. FICHTNER (1985) was the first to show that the inspector leadership method (as mentioned, already introduced by MASCHLER) solves this problem. It has to be emphasized, however, that this method means that the inspector commits himself in a credible way to stay to the announced strategy. If one doubts that this is always possible for private or even national "inspectors", it can be safely assumed for the case of international inspection problems because of the attention paid to these procedures by the States involved; thus, here this method has a satisfactory justification.

This paper shows how the leadership game differs from the "simultaneous" game considered earlier and how in equilibrium the inspectee behaves legally. (FICHTNER, who in line with MASCHLER and STACKELBERG used the term leadership *solution*, did not formulate this explicitly.) Nevertheless, the result is not yet fully satisfying, since, if the inspector plays the equilibrium strategy, the inspectee's best reply is not exclusively legal behavior and, furthermore, because of what is called the "knife-edge" property of this equilibrium: a slight mistake by the inspector in determining his equilibrium strategy can lead to a major loss. It turns out, however, that the treatment of problems of this kind with incomplete information models solves this problem effectively. This approach is the more justified if one admits that the knowledge of the inspector about the gain of the inspectee in case of undetected illegal action is indeed incomplete.

Clearly, the above outlined problems are very closely related to a variety of economic models known as *principal agent problems* extensively studied in the economic literature (see, e.g., KANODIA (1985), DYE (1986)). The model and the applications in this paper are not primarily economic but rather what one may call "political". The main issues here are: violation, deterrence, detection, false alarms, etc. The traditional approach to such problems has been basically that of statistical analysis. The statistical tests were occasionally incorporated as strategies in a game theoretical model.

This paper which is based on two earlier versions (AVENHAUS and OKADA, and AVENHAUS and ZAMIR (1988)) is organized as follows:

In the next section, the nuclear material safeguards example, with its basic tools material accountancy and data verification, is introduced since it represents an important application of the models to be discussed subsequently; in fact, it stimulated the development of these models. In the third section, quite a general game theoretical model is proposed which generalizes the statistical procedure such that by appropriate *calibration* of the model - i.e., restrictions on payoffs and strategies - the well-known Neyman-Pearson test emerges as the unique Nash equilibrium of the game. The model is sequential with the potential of capturing the importance of the time in situations in which *time is valuable*. This is not done in existing statistical and game theoretical models which are basically static.

In the fourth section, the differences between games without and with a leadership of the inspector are analyzed with the help of simplified models which lead, as outlined above, to the formulation and treatment of inspector leadership games with incomplete information; this is achieved in the fifth chapter. In the concluding section, the results are applied to the problems outlined in the second one. Among others, it is shown that practitioners' questions like those for best decision procedures can indeed be answered without precise knowledge of payoff parameter values which especially in the application considered here are hard to estimate.

Since in this paper statistical and game theoretical methods are matched, sometimes there arises the problem of having to use symbols which are common in the one, but not in the other field, or even worse which are used in both with different meanings. Thus, compromises have to be found which occasionally are not too statisfying both for the statistician and the game theorist.

2. The Nuclear Material Safeguards Example

There are many important applications of the general framework which will be developed subsequently (see, e.g., AVENHAUS (1986)). The case of nuclear material safeguards, however, is of primary importance; in fact, this special application stimulated the development of the formalism presented here as already mentioned before, therefore, it will be discussed in some detail in the following.

In partial fulfillment of the *Treaty for the Non-Proliferation of Nuclear Weapons* (NPT), the International Atomic Energy Agency (IAEA) in Vienna controls the nuclear material of the peaceful sectors[2] of those States' nuclear fuel cycles who have ratified this Treaty (by now more than 100 States) in order to achieve the objective of safeguards which is according to the model agreement (IAEA 1971)

[2] There are five States who officially have non-peaceful nuclear fuel cycles; three of them signed and ratified the NPT. According to the NPT provisions, these States are not obliged to have IAEA safeguards applied to their peaceful nuclear fuel cycles; they accept it, however, for good will reasons.

> "...the timely detection of diversion of significant quantities of nuclear material from peaceful nuclear activities to the manufacture of nuclear weapons or of other nuclear explosive devices or for purposed unknown and deterrence of such diversion by the risk of early detection."

The basic principle of IAEA safeguards is *material accountancy*, i.e., the comparison of book with physical inventory at the end of a given inventory period. The procedures for the performance of safeguards were also laid down in the model agreement: The operator of a nuclear plant generates all source data for the establishment of a material balance and reports these data to a national or regional authority; the international authority *verifies these reported data with the help of independent measurements* made on a random sampling basis. If there are no significant differences between the operator's reported data and the international authority's own findings, then all the operator's data are accepted and the material balance is closed on the basis of the operator's data alone.

According to this procedure, the operator of a nuclear plant, who intends to divert material, has two basic possibilities: Either he can simply divert material without data falsification, hoping that the measurement uncertainty of the whole balance covers this diversion, or he can falsify data such that the balance appears to be correct, hoping that the measurement and sampling uncertainty of the data verification procedure covers this falsification. Naturally, he can use both 'strategies'; we will come back to this point.

There are model agreement statements about the *maximum routine inspection effort* which provides the basis for all random sampling procedures. In order to make meaningful statements about this effort, the importance of the nuclear material in the sense of the NPT processed in the different plants of the nuclear fuel cycle had to be defined (the "critical" mass of about 8 kg plutonium or highly enriched uranium is necessary for the construction of one nuclear device). Thus, the concept of *effective kilogram* was introduced. Accordingly 1 kg of plutonium corresponds to one effective kilogram, whereas 1 kg of uranium with an enrichment of 0.01 and above corresponds to a quantity in effective kilograms that is obtained by taking the square of enrichment. The maximum routine inspection effort, given in inspection man-hours spent in a nuclear plant, is determined on the basis of the annual throughput or inventory of nuclear material expressed in effective kilograms.

Now, let us look at material accountancy more closely: Consider a so-called material balance area (a nuclear plant or a part of it). At time t_{i-1} the beginning inventory I_{i-1} is measured; during the inventory period $[t_{i-1}, t_i]$ receipts R_i and shipments S_i are measured which together with the beginning inventory add up to the book inventory B_i at t_i. This is compared to the ending inventory I_i of this period such that the difference between B_i and I_i is used (which for traditional reasons is called Material Unaccounted For, or shortly MUF),

$$MUF_i = X_i := B_i - I_i = I_{i-1} + R_i - S_i - I_i, \; i = 1, \ldots, n. \qquad (2\text{-}1)$$

Since in observing these quantities measurement errors cannot be avoided, the variances of which are known to all parties involved, the vector $X' = (X_1, \ldots, X_n)$ is a random vector with covariance matrix Σ. If there are no systematic errors, then the only correlation between the differences (2-1) for two subsequent inventory periods is given by the variance of the common intermediate inventory, i.e., the elements of the covariance matrix Σ are given by

$$(\Sigma)_{ij} = \mathrm{cov}(X_i, X_j) = \begin{cases} \mathrm{var}(X_i) & \text{for} \quad i = j \\ -\mathrm{var}(I_i) & \text{for} \quad i \neq j \\ 0 & \text{otherwise.} \end{cases} \qquad (2\text{-}2)$$

In case the operator of the plant, being the inspectee, does not divert any material, the expectation of the random vector X is zero due to the conservation of matter. If, according to general experience, we assume the measurement errors to be normally distributed, then the random vector X is multivariate normally distributed with the expectation zero and covariance matrix Σ.

$$X \sim N(0, \Sigma) \quad \text{for legal behavior.} \qquad (2\text{-}3a)$$

In case the operator diverts the amount μ_i of material in the i-th period, the expectation vector of X is $\mu' = (\mu_1, \ldots, \mu_n)$, and we get

$$X \sim N(\mu, \Sigma) \quad \text{for illegal behavior.} \qquad (2\text{-}3b)$$

Let us assume that the illegal behavior of the operator consists in choosing a vector μ such that

$$\sum_{i=1}^{n} \mu_i = e' \cdot \mu = M, \; e' = (1, \ldots, 1). \qquad (2\text{-}4)$$

Then the problem of the inspector is to decide with the help of an observation of the random vector X if the operator behaved legally, or on the contrary, he diverted the amount M of material. The value of M will be discussed later.

Let us consider now a simplified data verification problem: A nuclear plant is shut down for some time, and there is an inventory of nuclear material which is collected in n items. The operator of the plant presents via his national authority his estimates Y_i, $i = 1,\ldots,n$ either correctly, or he may falsify the ith-estimate by the amount μ_i, $i = 1,\ldots,n$; if he does this and is not detected, he can divert the corresponding amount of material and the material balance still appears to be correct. The inspector estimates k data independently; his estimates are Z_i, $i = 1,\ldots,k$ without loss of generality. Since he is only interested in deviations, he uses the differences

$$X_i := Y_i - Z_i,\ i = 1,\ldots,k \qquad (2\text{-}5)$$

of reported estimates and independently generated ones in order to solve his decision problem which is very similar to the one described above: In case the inspector takes the maximal sample size (k=n), X is multivariate normally distributed with expectation zero and some known covariance matrix Σ. In case the operator falsifies the data Y_i by the amounts μ_i, $i = 1,\ldots,n$, the expectation of X is $\mu' = (\mu_1,\ldots,\mu_n)$. If, again, it is assumed that the total falsification is M, then we have the same problem as before, only the covariance matrix being different.

We assume the total amount of material or, equivalently, the total falsification M to be in the order of magnitude of the critical mass described above. Naturally, one may assume that the total diversion consists of severall small diversions using different *paths* (inventories at different points of time, diversion without data falsification by using the measurement uncertainty of the material balance and others). There are, however, both practical and theoretical reasons to assume that in case of diversion the operator will use only one path if he really wants to divert material.

3. Simple Inspection Games

There are two players, namely the inspector and the inspectee. $M_1, M_2, \ldots, M_n$ are the (pure) action sets of the inspectee in stages

1,2,...,n, respectively. The action set of the inspector is the same for all stages and consists of two elements A (alarm) and A (clear or continue).

$E_1,\dots,E_n$ are the signals (or observations) sets in stages 1,2,... . E_t is typically a finite dimensional Eucledian space or a subset of it. $f_1,\dots,f_n$ are transition probabilities:

$$f_t : M_1 \otimes E_1 \otimes \dots \otimes M_{t-1} \otimes E_{t-1} \otimes M_t \to E_t,$$

here, $M_0 = E_0 = \phi$ by convention. That is, for any actions $\mu = (\mu_1,\dots,\mu_t)$, $\mu_i \in M_i$ of the inspectee in stages 1,...,t and observations $x_1,\dots,x_{t-1}$, $x_i \in E_i$ at previous stages 1,...,t-1, $f_t(\mu_1,x_1,\dots,\mu_{t-1},x_{t-1},\mu_t)$ is the probability density of the random variable X_t observed at stage t.

For t = 1,...,n, I_t and O_t are real functions on $M_1 \times \dots \times M_t$. These are the payoff functions to the inspector and to the inspectee, respectively, if the game is stopped by alarm after stage t < n. Finally, I and O are two real functions on $M_1 \times \dots \times M_n$, the payoff functions when the game ends after n stages with no alarm called.

The game is played as follows:

- At stage 1, the inspectee chooses an action $\mu_1 \in M_1$ (not observable by the inspector). An observation x_1 is drawn from a random variable X_1 according to the density function $f_1(.|\mu_1)$. This observation becomes common knowledge to both players.

- The inspector chosses either $\bar{A}$ in which case the game continues to stage 2 or A in which case the game terminates with payoffs $(I_1(\mu_1), O_1(\mu_1))$.

- Inductively: At stage t, t = 1,...,n, the inspectee chooses $\mu_t \in M_t$, an observation x_t is made from a random variable,

  $X_t \quad f_t(.|\mu_1,x_1,\dots,\mu_{t-1},x_{t-1},\mu_t)$.

  The inspector chooses either $\bar{A}$ in which case the game continues to stage t + 1 if t < n, or terminates with payoffs $(I(\mu_1,\dots,\mu_n), O(\mu_1,\dots,\mu_n))$ if t = n, or A in which case the game ends with payoffs $(I_t(\mu_1,\dots,\mu_t), O_t(\mu_1,\dots,\mu_t))$. Note that by this notation we assume that the payoffs are determined only by the actions of the inspectee and not by the observations. This is obvious if no alarm is called. In case of alarm this requires the implicit assumption that the actual diversions are checked and found out after the alarm.

The extensive form of this game is scetched in Figure 1.

The first step in analyzing our model is to consider a special case corresponding to situations in which *time is not important* - inspite of the sequentiality of the problem; the only issues of interest are whether or not there was illegal behavior of the operator, and whether or not an alarm was called.

In addition to being a natural simple case to start with, this first step may be considered as *calibration* of the model, namely relating it to the existing, mainly statistical analyses of the problem. To do this, we choose a special, simple payoff structure capturing the underlying assumptions of most existing analyses.

The formal reduction to the non-sequential case is obtained by the following set of assumptions.

Assumption 1

- The inspector decides on whether to call an alarm only once at the end of the n-th stage, i.e., after having observed the whole random vector $X' = (X_1, \dots, X_n)$. Note that this does not exclude the possibility that the inspector decides on calling an alarm already after $t < n$ observations. However, since *time is not important*, he may as well postpone his announcement to the end of the game. In other words, this assumption still leaves room to some inherent sequentiality and intrinsic order of the test procedure.

- The inspectee's diversion strategy is completely determined at the beginning of the game, i.e., he decides either not to divert ($\mu_t = 0, t = 1, \dots, n$) or to divert according to the plan $\mu' = (\mu_1, \dots, \mu_n)$ where $\mu_t \geq 0$, $\Sigma \mu_t = M$. The constant $M > 0$ is the "critical diversion" as motivated in the previous section. ■

Under these assumptions the game in extensive form can now be described as given by Figure 2. We denote this game by Γ_0.

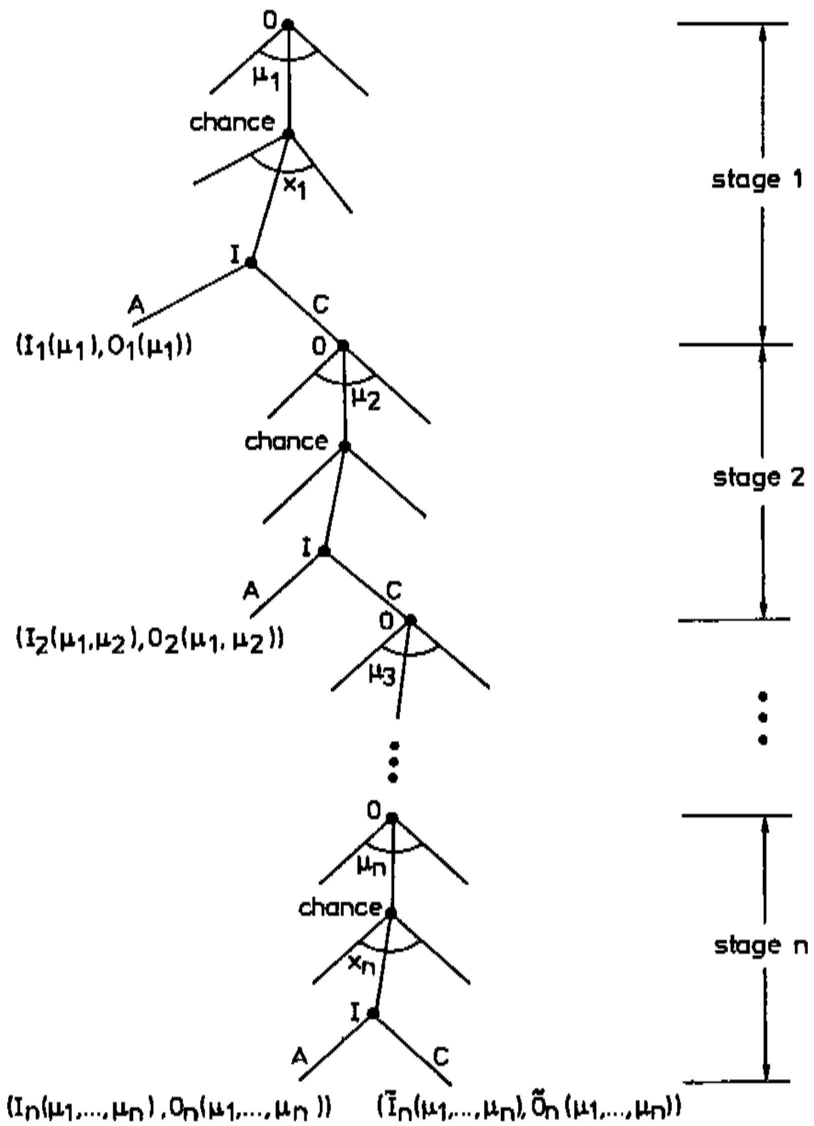

Figure 1: The sequential game in extensive form

The notation in Figure 2 is self-explaining: 1 and $\bar{1}$ are legal and illegal actions, respectively. $\mu' = (\mu_1, \ldots, \mu_n) : \Sigma \mu_i = M$ is a diversion plan. $x' = (x_1, \ldots, x_n)$ is the vector of observations. The payoffs at the endpoints of Figure 1 are listed in Table 1 which is to be read as follows: $I(1,\bar{A}) = O(\bar{1},\bar{A}) = 0$, $I(1,A) = -e$, $O(\bar{1},A) = -b$ etc.

Here, we normalized the utilities so that the status quo $(A,1)$ has zero utility, and undetected diversion is worth $d > 0$ for the inspectee and $-c$ for the inspector. The fact that all entries in line A are negative is to reflect the idea that an alarm (i.e., an open conflict between the two players) is bad for both players, compared to the status quo. $a < c$ is the obvious requirement that, if diversion occurred, then the inspector is better off detecting it than not detecting it. $h < b$ means that detected diversion is worse for the inspectee than a false alarm (reflecting some element of punishment), and finally $e < a$ means that for the inspector a detection of a diversion which means "a failure" of safeguards is worse than the inconvenience of a false alarm.

A general pure strategy of the inspector is an alarm set $s \subset E_1 \otimes \ldots \otimes E_n$, that is: According to s the inspector calls an alarm at the end of period n if and only if $(x_1, \ldots, x_n) \in s$. A mixed strategy σ is a probability distribution on pure strategies.

A pure strategy of the inspectee is a choice between 1 and $\bar{1}$ and a diversion plan $\mu' = (\mu_1, \ldots, \mu_n)$ satisfying $\Sigma \mu_i > M$. Thus, a mixed strategy is a joint distribution on such pairs. This is equivalent to a behavioral strategy $\tau = (q,p)$ where q is the probability of $\bar{1}$, and p is the probability distribution over diversion plans μ if he chooses $\bar{1}$.

Any pair of strategies (σ,τ) determines two conditional probabilities, namely the false alarm probability

$$\alpha(\sigma,\tau) = \text{prob}(A|1,\sigma) \tag{3-1}$$

and the non-detection probability

$$\beta(\sigma,\tau) = \text{prob}(\bar{A}|\bar{1},\sigma,\tau). \tag{3-2}$$

Clearly, $\alpha(\sigma,\tau)$ does not depend on τ while $\beta(\sigma,\tau)$ depends on τ only through p, and it is linear in p; hence, we shall write $\alpha(\sigma)$ and $\beta(\sigma,p) = \int \beta(\sigma,\mu)\, dp(\mu)$.

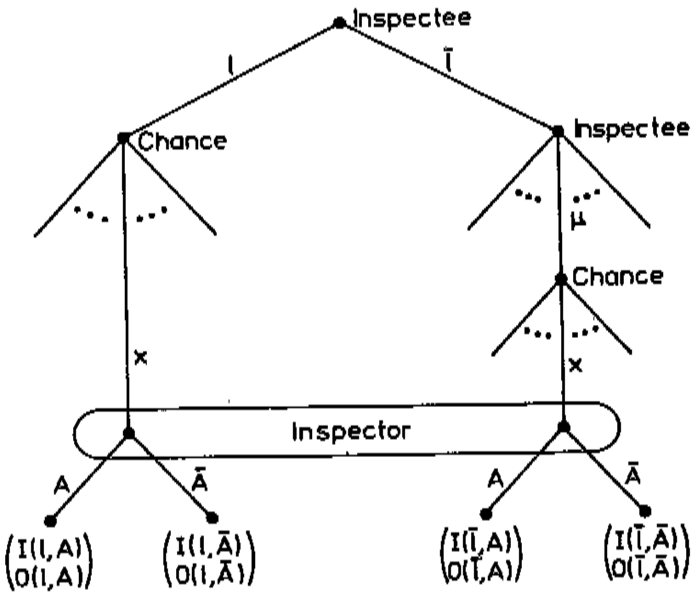

Figure 2: The "essentially" non-sequential game

Inspectee / Inspector	Legal behavior l	Illegal behavior $\bar{l}$
No alarm $\bar{A}$	0 / 0	d / -c
Alarm A	-h / -e	-b / -a

able 1: The simplified payoff structure for the essentially non-sequential inspection game as given by Figure 1.
$0 < e < a < c$, $0 < h < b$, $d > 0$.

Any σ and τ determine expected payoff functions which we denote by $I(\sigma,\tau)$ and $O(\sigma,\tau)$ for the inspector and for the inspectee:

$$I(\sigma,\tau) = -q\cdot(a+(c-a)\cdot\beta(\sigma,p)) - (1-q)\cdot e\cdot\alpha(\sigma) \tag{3-3}$$

$$O(\sigma,\tau) = -q\cdot(b+(-d-b)\cdot\beta(\sigma,p)) - (1-q)\cdot h\cdot\alpha(\sigma). \tag{3-4}$$

First, we observe that the game has no extreme Nash equilibrium (in the following simply called equilibrium), i.e., no equilibrium in which the inspectee deviates with probability 0 or 1 from legal behavior and no equilibrium in which the inspector's total probability for calling an alarm at some stage is 0 or 1. This is easily seen by looking at the payoff structure as given by Table 1. In particular, this means that in any equilibrium one must have $0 < q < 1$. In the next section we will see that in a variant of this game not only that $q = 0$ can be the case in equilibrium, but it must be: The only equilibrium of that model will be with $q = 0$.

It follows from (3-3) that because of $0 < q < 1$ in any equilibrium p must be a maximizer of $\beta(\sigma,p)$ that means that the support of p has to contain only μ maximizing $\beta(\sigma,\mu)$. Denote for any strategy σ of the inspector

$$\beta(\sigma) := \sup_p \beta(\sigma,p) = \sup_\mu \beta(\sigma,\mu). \tag{3-5}$$

If (σ,μ) is an equilibrium, we can rewrite (3-3) and (3-4) as

$$I(\sigma,\tau) = -q\cdot(a+(c-a)\cdot\beta(\sigma)) - (1-q)\cdot e\cdot\alpha(\sigma) \tag{3-6}$$

$$O(\sigma,\tau) = -q\cdot(b+(-d-b)\cdot\beta(\sigma)) - (1-q)\cdot h\cdot\alpha(\sigma). \tag{3-7}$$

Note that we actually eliminated the diversion strategy plan μ from the equilibrium strategy of the inspectee, since it is automatically determined by (3-5). It will reappear in the auxiliary game to be discussed later.

Let us proceed to determine q and σ in equilibrium. The inspectee's payoff is $(d+b)\cdot\beta(\sigma) - b$ if he chooses $\bar{1}$ and $-h\cdot\alpha(\sigma)$ if he chooses 1. Therefore, in the (α,β)-space the inspectee's best reply is to choose $\bar{1}$ or 1 or both according to whether

$$\beta(\sigma) \gtreqless \frac{b}{d+b} - \frac{h}{d+b}\cdot\alpha(\sigma), \tag{3-8}$$

see Figure 3. In this figure, also the corresponding zones in the (α,β)-plane and the resulting payoff functions for the inspector are given.

Note that the inspector's payoff is discontinuous along the line L where it is also not defined since the inspectee, being indifferent between $\bar{1}$ and 1, can use any mixture of them as a best reply to the strategy σ.

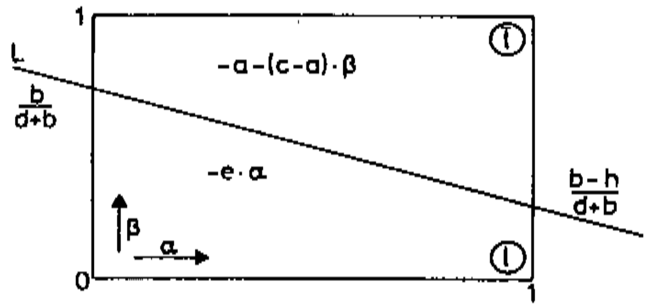

Figure 3: The payoff for the inspector when using the strategy σ, and the inspectee is doing his best reply

In order to maximize the inspector's payoff function, we first notice that its domain in the $(\alpha(\sigma),\beta(\sigma))$-space is not the whole unit square, but is restricted by the line $\beta = \beta(\alpha)$ of the relation between α and β for most powerful tests, that is

$$\beta(\alpha) = \inf_{\sigma}\{\beta(\sigma)\,|\,\alpha(\sigma)=\alpha\}$$

or, with (3-5),

$$\beta(\alpha) = \inf_{\sigma:\alpha(\sigma)=\alpha}\ \sup_{\mu}\ \beta(\sigma,\mu). \qquad (3\text{-}9)$$

Obviously, the function $\beta(\alpha)$ is determined by parameters of the problem, in particular, by the distributions of the random vector $(x_1,\ldots,x_n)$. We make the usual assumptions on $\beta(\alpha)$, satisfied, for instance, for the multinormal distribution:

Assumption 2

The line $\beta = \beta(\alpha)$ is well-defined for all $\alpha \in [0,1]$. It is differentiable, convex, strictly decreasing and satisfies $\beta(0) = 1$, $\beta(1) = 0$. ■

Consider now the following auxiliary game played by the inspector and the inspectee which we denote by G_α and call the *axiliary game with parameter* α:

- The strategy set of the inspector is $\Sigma_\alpha := \{\sigma | \alpha(\sigma) \leq \alpha\}$.
- The pure strategy set of the inspectee is the set of all diversion plans $M := \{\mu' = (\mu_1, \ldots, \mu_n) : \mu_i \geq 0, i=1,\ldots,n, \sum \mu_i \geq M\}$.
- The payoff from the inspectee to the inspector, when (σ, μ) are used, is the detection probability $1 - \beta(\sigma,\mu)$.

Assumption 2 can now be written as follows:

Assumption 2'

For any α, $0 < \alpha < 1$, the game G_α has a value $v(\alpha) = 1 - \beta(\alpha)$, and both players have optimal strategies (i.e., sup and inf may be replaced by max and min, respectively). Furthermore, $v(\alpha)$ is differentiable, concave, strictly increasing and satisfies $v(0) = 0$, $v(1) = 1$. ■

Note that G_α depends on α, but not on the parameters of the original game. Of course, the value of α, for which Γ_α plays a role, will depend on the data of the original game. Later on, we show that Assumption 2' is satisfied under the usual assumptions of multinormally distributed signals.

Now, we add the line $\beta = \beta(\alpha)$ to Figure 3 and maximize the inspector's payoff, see Figure 4. In region 1, the payoff decreases in α for any fixed β, whereas in region 1 it decreases in β for fixed α. It follows that the supremum in σ is reached on the line $\beta = \beta(\alpha)$.

On that part of the line, in region 1 the payoff decreases in σ, similarly on the part in region 1 it decreases in β. It follows that the supremum is at the intersection of the line $\beta = \beta(\alpha)$ with the straight line L. Thus, we conclude, if (σ^*, τ^*) is an equilibrium, that the corresponding (α^*, β^*) is the unique solution of

$$\beta(\alpha^*) = \frac{b}{b+d} - \frac{h}{b+d} \cdot \alpha^* . \qquad (3\text{-}10)$$

Since (α^*,β^*) lies on L, by its definition σ^* makes the inspectee indifferent between legal and illegal behavior, so, an illegal behavior with any probability q will be a best reply against σ^*. However, the payoff to the inspector at (α^*,β^*) is a function of q,

$$I(\sigma^*,q) = q\cdot(-a-(c-a)\cdot\beta(\alpha^*)) - (1-q)\cdot e\cdot\alpha^* . \qquad (3\text{-}11)$$

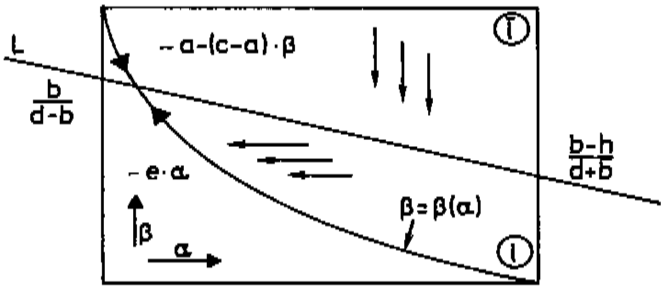

Figure 4: The Nash equilibrium values (α^*,β^*)

The value q^* of q now has to be determined so as to satisfy the second condition for an equilibrium which is that α^* is a best reply to q, that is, it maximizes

$$q^*\cdot(-a-(c-a)\cdot\beta(\alpha)) - (1-q^*)\cdot e\cdot\alpha$$

as a function of α. By the assumption of differentiability of $\beta(\alpha)$, this function has one local maximum at α^* satisfying the first derivate condition which is

$$\frac{e}{q^*} = e - (e-a)\cdot - \frac{d}{d\alpha}\,\beta(\alpha)\Big|_{\alpha^*} . \qquad (3\text{-}12)$$

Thus, q^* is determined since α^* is fixed by (3-10).

Note that $0 < q^* < 1$ as expected from our assumptions. By straightforward verification it can be shown that this local maximum α^* is also the global maximum, i.e., it is the unique best reply to q^*.

As soon as (α^*,q^*) are uniquely determined in equilibrium, the rest of the strategies are determined as follows: the inspector chooses σ^* as a maxmin strategy in the game Γ_{α^*}, and the inspectee uses p^* which is any probability distribution over best replies μ^* to σ^* in that game. We summarize this discussion in

Theorem 1

The game Γ_0 as given by Figure 2 has a unique equilibrium $(\sigma^*,(q^*,\mu^*))$ in which

(i) the false alarm probability $\alpha^* = \alpha(\sigma^*)$ is the solution of (3-10),

(ii) the probability q^* for illegal behavior is given by (3-12),

(iii) the test strategy σ^* is a strategy guaranteeing the lower value $v(\alpha^*) = \max_{\sigma} \min_{\mu}(1-\beta(\sigma,\mu))$ in the auxiliary game G_{α^*}; μ^* is a best reply to σ^* in that game. ■

Note that the uniqueness claimed in Theorem 1 holds only for (α^*,β^*,q^*). There may well be more than one σ^*, or μ^*, although as we will see later, may also be unique, e.g., in the multinormal case.

A special case worth noticing is when $\beta(\alpha) = 1 - \alpha$. In this case, equations (3-10) and (3-12) yield

$$\alpha^* = \frac{d}{d+b-h}, \quad q^* = \frac{e}{c+e-a},$$

which is the (unique) mixed equilibrium of the matrix game in Table 1. This describes a situation in which the inspector has no specific test procedure for detecting illegal behavior and, therefore, $(\alpha^*,1-\alpha^*)$ appears as mixed equilibrium in that game.

The equilibrium of the inspection game Γ_0 has two undesirable and related features: First, as was already mentioned, the unique equilibrium has a positive probability for illegal behavior ($q^* > 0$). Second, the inspector's payoff at the equilibrium point is drastically discontinuous. This is particularly disturbing in view of the fact that the inspectee is indifferent between legal and illegal behavior at that point. Therefore, by "doing as well as in the equilibrium point" he may inflict a severe loss to the inspector by behaving illegally with probability 1.

In the following sections, we address these two problems: First, we propose a variant of the model which takes into account the natural leadership role of the inspector. We show that in the leadership game there is a unique equilibrium in which the inspector uses the same rategy σ^* as in the equilibrium of the game Γ_0, but the inspectee es $q^* = 0$, i.e., he behaves legally in equilibrium.

then turn to the discontinuity problem and introduce some certainty of the inspector about the utility d of the inspectee in se of undetected illegal action. We show that the result of such certainty removes the discontinuity of the inspector's payoff at uilibrium.

Deterrence and Commitment

important feature of safeguards games is the fundamental asymmetry the players' roles. The inspector, generally representative of an ficial authority, can, and usually does, announce his inspection rategy while it is inconceivable for the inspectee to make public s deviation strategy.

rthermore, the inspector has the *power of commitment*: not only can make public his inspection strategy, but he can credibly commit mself to follow the announced strategy (e.g., to sign a contract plemented and enforced by court). In game theoretical terminology is means that the inspector can become a *Stackelberg leader* in the ne if he chooses so.

mentioned in the introduction, the idea of granting a leadership le to the inspector was already brought into the literature on feguards and arms control (see, e.g., MASCHLER (1967/68), FICHTNER 985)). However, to our best knowledge, the precise distinction tween the equilibrium with and without inspector leadership was not plicitly formulated.

s this power of commitment give an advantage to the inspector, and nce, does it induce a significant change in the analysis? Let us st look at a very simple example. Before, we recall the notion of ckelberg leadership which we would like to emphasize as being a sion of the game rather than a version of the solution.

Definition 1

Given any game Γ (in extensive or in strategic form) with players set $N = \{1,\ldots,n\}$ and corresponding (mixed) strategy sets $\Sigma_1,\ldots,\Sigma_n$. The version of the game with commitment power to player i is the game Γ_i played as follows: Player i chooses $\sigma_i \in \Sigma_i$ which is made common knowledge to all players. Then, the game Γ is played by the players $(N\backslash\{i\})$ where the moves of player i are mady by a "machine" executing the instructions of σ_i. ■

At a first glance, one may suspect that such a modification should not be very important, since after all in Nash equilibrium reasoning each player assumes the strategies of the others to be given. In order to see this more carefully, consider the two-person game G in strategic form as given by Table 2.

The game G has a unique equilibrium which is (1/2,1/2) for player 1 and (3/4,1/4) for player 2, yielding expected payoffs (3/4,1/2).

2 1	Left L	Right R
Top T	0 1	1 0
Bottom B	1 0	0 3

Table 2: Strategic or normal form game G

The game $\tilde{G}$ with commitment power to player 1 is a game the extensive form of which is sketched in Figure 5.

In equilibrium, at any information set U_x, player 2 will choose L if $x < 1/2$, R if $x > 1/2$ and be indifferent if $x = 1/2$. The only parameter needed to complete the specification of the second player's strategy in equilibrium is $0 \le q \le 1$, the probability of playing L at

/2. Denote the resulting strategy by τ_q. The first player's payoff gainst τ_q is given by, see Figure 6,

$$f_q(x) = \begin{cases} x & 0 \le x < \frac{1}{2} \\ q \cdot x + 3 \cdot (1-q) \cdot (1-x) \quad \text{for} & x = \frac{1}{2} \\ 3 \cdot (1-x) & \frac{1}{2} < x \le 1. \end{cases}$$

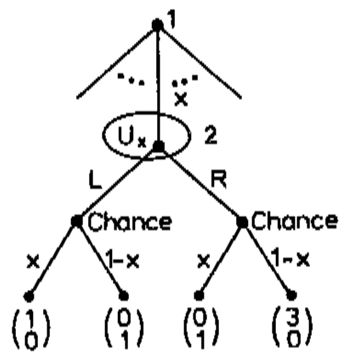

Figure 5: The game G with commitment power

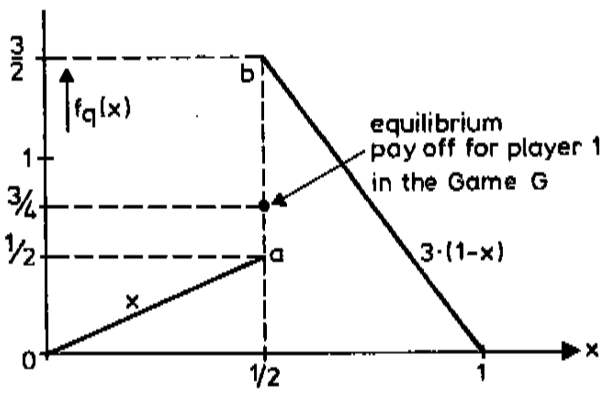

Figure 6: The function $f_q(x)$

ow, an equilibrium point is a pair (x^*, τ_{q^*}) such that τ_{q^*} is best eply to x^* which is satisfied automatically for all (x^*, q^*) by the onstruction of τ_{q^*}, and $f_{q^*}(x)$ has a global maximum over $[0,1]$ at x^*. ut $f_{q^*}(x)$ *has no maximum* unless $q^* = 0$, i.e., player 2 plays R. It follows that the game G has a unique equilibrium point at $x^* = 1/2$, $q^* = 0$ with expected outcome (3/2,1/2).

So, each of the two games G and $\tilde{G}$ has a unique equilibrium (which is also a perfect equilibrium), but *they are* not the same. While they yield the same payoff for player 2 the outcome is higher for player 1 in the equilibrium of G. The difference 3/2 - 3/4 may be viewed as the *value of his commitment power*.

Let us conclude this example with two remarks.

1) In order to see how the commitment power changes the equilibrium of the original game, observe that, if in the game tree of G the moves at the nodes following U_x would be made by player 1 instead of chance, then (x^*=1/2,R) is not an equilibrium, since it is not equilibrium in the subgame starting at $O_{1/2}$: The best reply to R is not $x_* = 1/2$, but $x = 0$. So, player 1 commits himself not to *change his mind* and play B even though this would be more profitable against R.

2) Equilibrium of $\tilde{G}$ can be viewed as a limit of points (x=1/2+ε,R) which are not equilibria, but from the point of view of player 1, all he can improve his payoff is by no more than ε, and in reward to this "inefficiency" R becomes the *only best reply*, unlike the situation in the equilibrium. Hence, these points are more stable than the equilibria against deviations of player 2.

Our observations on the commitment power were probably noticed, though not spelled out explicitly, in the works on refinements of equilibria, see, e.g., VAN DAMME (1986). However, it is especially relevant for our context: Commitment power is a natural thing for inspection authorities and is usually practiced in exactly the same way we modified our game from G to $\tilde{G}$, namely the inspection strategy is officially declared and is executed by "machines" or agents without decision authority.

To say this more precisely, let us first give a second representation of the game Γ_0 which is equivalent to that in Figure 2, see Figure 7. In this figure, we used the same notation as before: μ is the inspectee's diversion plan, σ is the test strategy of the inspector,

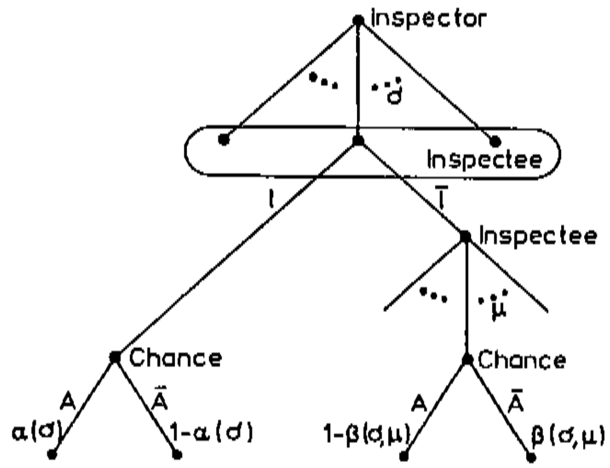

Figure 7: An equivalent representation of the game Γ_0

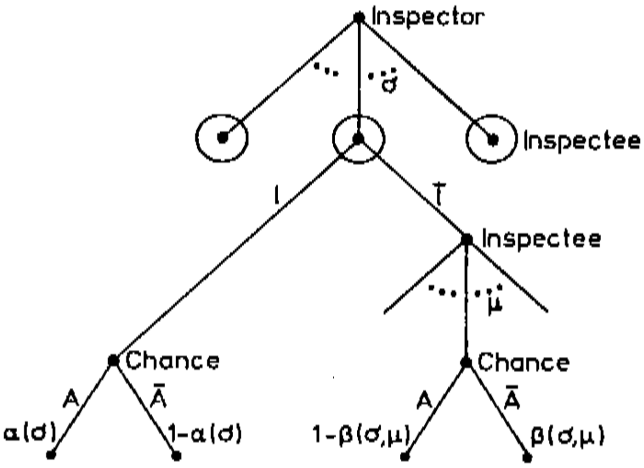

Figure 8: The inspector leadership game Γ_1

and $\alpha(\sigma)$ and $\beta(\sigma,\mu)$ are the false alarm and not detection probabilities.

If the inspector announces his test strategy σ which is then executed by a neutral mechanism, we obtain the following "leadership game" which is denoted by Γ_1 and represented in extensive form in Figure 8.

A "pure" strategy of the inspector in the game Γ_1 is a test procedure σ. For the inspectee, a pure strategy is a map τ which assigns to each σ a strategy $\tau(\sigma) = (q,\mu)$.

The analysis of the game Γ_1 is very similar to that of Γ_0 (modified as done in our example). For any fixed α, the corresponding μ and σ is given by the auxiliary game G_α. The values α^*, β^* and μ^* are therefore identical to those in Γ_0.

The inspectee's best reply as a function of α is (fixing $\beta = \beta^*$, see Figure 4)

$$\text{play} \quad \begin{matrix} \bar{1} \\ \\ 1 \end{matrix} \quad \text{if} \quad \begin{matrix} 0 \le \alpha < \alpha^* \\ \\ \alpha^* < \alpha \le 1, \end{matrix} ,$$

and if $\alpha = \alpha^*$, play any mixed strategy $(q,1-q)$. The corresponding payoff to the inspector is, see Figure 9,

$$f_q(\alpha) = \begin{cases} -a - (c-a)\cdot\beta(\alpha) & 0 \le \alpha < \alpha^* \\ q\cdot(-a-(c-a)\cdot\beta^*) - (1-q)\cdot e\cdot\alpha^* \quad \text{for} & \alpha = \alpha^* \\ -e\cdot\alpha & \alpha^* < \alpha \le 1. \end{cases}$$

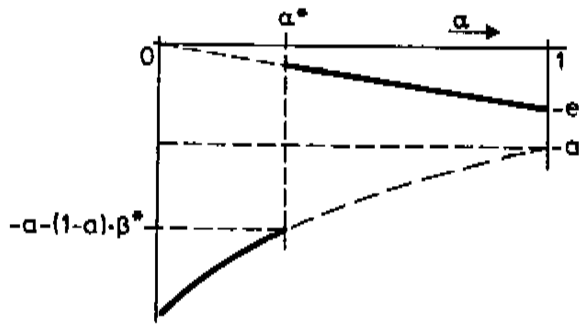

Figure 9: Payoff to the inspector in the leadership game Γ_1

Now, in equilibrium, $f_q(\alpha)$ has to have its maximum at α^*, but $f_q(\alpha)$ does not have a maximum unless $q = 0$. We thus get

Theorem 2

The leadership game Γ_1 has a unique Nash equilibrium point (hence also perfect) in which the inspector commits himself to the most powerful test with false alarm probability α^* determined by (3-10), and the inspectee surely does not deviate from legal behavior. ∎

Note that, inspite of having $q = 0$ in equilibrium, the inspectee is still indifferent between $\bar{I}$ and 1. Nevertheless, this equilibrium may be called a *deterrence equilibrium* since it is the limit of points of the form $(\alpha^*+\varepsilon,1)$. These are not equilibria, but rather deterrence points in which the inspector gets an outcome ε-close to that in the equilibrium, and the inspectee strictly prefers legal to illegal behavior.

5. Inspector Leadership with Incomplete Information

In the last section, we assumed that the inspector and the inspectee know all payoff parameters in Table 1. There are, however, many real situations in which this assumption of complete information is too strong to be satisfied. In particular, it rarely happens that the inspector has complete knowledge of the inspectee's payoff. Therefore, in this section we will modify our rule of the inspector leadership game Γ_1 described in Figure 8 so that it can capture a situation in which the inspector is uncertain about the inspectee's payoff parameter. We will use HARSANYI's model (1967/68) of games with incomplete information and assume:

(i) The inspectee's payoff d in case the inspectee acts illegally ($\bar{I}$) and the inspector does not raise an alarm is a random variable on the interval $[0,\infty)$ with the (cumulative) distribution $F(.)$. This distribution $F(.)$ is known to both players. In what follows we will use the symbol t for the value of this random variable.

(ii) The inspectee knows the realized value t, but the inspector does not. All other payoff parameters in Table 1 are common knowledge. We may justify this in such a way that the losses of the inspectee in case of false alarms or illegal behavior are economic losses or well-defined sanctions which are publicly known, whereas his gain in case of undetected illegal action can only be guessed by the inspector.

Let $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$ denote a vector of observations. A general test procedure of the inspector is represented by a measurable function from $\mathbb{R}^n$ to $\{A, \bar{A}\}$, therefore, we denote by

$$\Sigma := \{\sigma \mid \sigma: \mathbb{R}^n \to \{A, \bar{A}\};\ \sigma \text{ is measurable}\} \qquad (5\text{-}1)$$

the set of all test procedures.

Remark that we defined here the set of pure strategies. We could have defined the set of mixed strategies (distributions on Σ) or the set of behavioral strategies

$$\Sigma = \{\sigma \mid \sigma: \mathbb{R}^n \to [0,1];\ \sigma \text{ is measurable}\}.$$

However, for our equilibrium analysis, this larger set will not be needed, and pure strategies will suffice.

Recall that a behavioral strategy of the inspectee in Γ_0 is $\tau = (q, \mu)$ where q is the probability for illegal behavior and $\mu \in \Delta_M$ is the diversion plan in case of illegal behavior. For $q = 1$ or $q = 0$ we shall also write $\tau = (\bar{1}, \mu)$ or $\tau = 1$, respectively. Any pair (σ, τ) determines a false alarm probability $\alpha(\sigma)$ which depends on σ only and non-detection probability $\beta(\sigma, \mu)$ which depends on τ only through μ. The rules of our game with incomplete information are described as follows:

- First, a chance move selects a value of the inspectee's payoff t in case of 'no detection' according to the distribution F(.). The inspectee is informed of this value (his own type), and the inspector is not.

- The inspector selects a test procedure $\sigma \in \Sigma$ and announces it to the inspectee.

- With complete knowledge of (t, σ), the inspectee selects $\tau = (q, \mu) \in \mathcal{T}$.

o Finally, a chance move selects "alarm" (A) or "no alarm" ($\bar{A}$) with the probability distribution

$(\alpha(\sigma), 1-\alpha(\sigma))$

in case the inspectee selects l or else

$(1-\beta(\sigma,\mu), \beta(\sigma,\mu))$

in case the inspectee selects $\tilde{l}$ with a diversion plan μ.

We call this game the inspector leadership game with incomplete information and denote it by Γ_2. The extensive form of this game is described in Figure 10.

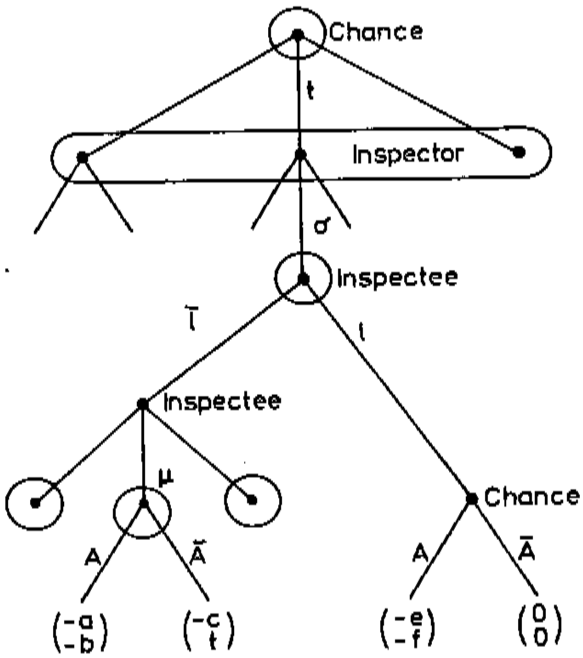

Figure 10: The extensive form of the inspector leadership game Γ_2 with incomplete information

Note that the set of pure strategies for the inspector is Σ as in Γ_0 or Γ_1 while for the inspectee the set of pure strategies is the set of maps

$[0,\infty) \otimes \Sigma \to \mathcal{T}$.

We shall denote this map also by τ and use for instance $\tau(t,\sigma)$ for the element (q,p) in chosen by the inspectee of type t when facing the test procedure $\sigma \in \Sigma$.

In the following analysis it turns out that, since the inspectee makes the last personal move, he may as well do without randomizing over 1 and $\bar{1}$ or over diversion plan μ, therefore, from now on we restrict the strategies τ to be pure, i.e.,

$$\tau \;:\; [0,\infty) \otimes \Sigma \rightarrow \mathcal{T}_0 \;:=\; \{1\} \cup \{(\bar{1},\mu) \mid \mu \in \Delta_M\}. \qquad (5\text{-}2)$$

Given a strategy combination (σ,τ), the conditional expected payoff for the inspector, given that t is selected, is

$$I_2(\sigma,\tau|t) = \begin{cases} -e\cdot\alpha(\sigma) & \tau(t,\sigma) = 1 \\ & \text{for} \\ -a\cdot(1-\beta(\sigma,\mu)) - c\cdot\beta(\sigma,\mu) & \tau(t,\sigma) = (\bar{1},\mu), \end{cases} \qquad (5\text{-}3)$$

and the conditional expected payoff for the inspectee, given that t is selected, is

$$O_2(\sigma,\tau|t) = \begin{cases} -f\cdot\alpha(\sigma) & \tau(t,\sigma) = 1 \\ & \text{for} \\ -b\cdot(1-\beta(\sigma,\mu)) + t\cdot\beta(\sigma,\mu) & \tau(t,\sigma) = (\bar{1},\mu), \end{cases} \qquad (5\text{-}4)$$

respectively. The expected payoff to the inspector is defined by

$$I_2(\sigma,\tau) = \int_0^\infty I_2(\sigma,\tau|t)\; dF(t),$$

whereas that to the inspectee is (5-5)

$$O_2(\sigma,\tau) = \int_0^\infty O_2(\sigma,\tau|t)\; dF(t).$$

An equilibrium point of the game Γ_2 is defined as follows:

Definition 2

A strategy combination (σ^*,τ^*) of Γ_2 is a (subgame perfect) equilibrium point of Γ_2 if and only if

(1) $EI_2(\sigma^*,\tau^*) \geq EI_2(\sigma,\tau^*)$ for all $\sigma \in \Sigma$,

(2) $O_2(\sigma^*,\tau^*(t,\sigma^*)|t) \geq O_2(\sigma^*,\tau|t)$ for all $\tau \in \mathcal{T}_0$ and for all $t \in [0,\infty)$. ■

In order to simplify the subsequent Theorem and its proof, in line with previous assumptions we use

Assumption 3

1) The distribution F(.) has the density function f(.), i.e., it can be written as

$$F(t) = \int_0^t f(t')dt' \quad \text{for all} \quad t \in [0,\infty).$$

2) For every $\alpha \in [0,1]$, the minmax problem

$$\min_{\sigma \in \Sigma_\alpha} \max_{\mu \in \Delta_M} \beta(\sigma,\mu)$$

has a solution. ■

We denote by $\beta(\sigma)$ the maximum and by $\beta(\alpha)$ the minmax value of $\beta(\sigma,\mu)$,

$$\beta(\sigma) = \max_{\mu \in \Delta_M} \beta(\sigma,\mu) \tag{5-6}$$

$$\beta(\alpha) = \min_{\sigma \in \Sigma_\alpha} \max_{\mu \in \Delta_M} \beta(\sigma,\mu) = \min_{\sigma \in \Sigma_\alpha} \beta(\sigma). \tag{5-7}$$

The next theorem characterizes a (subgame perfect) equilibrium point of the game Γ_2.

Theorem 3

Under Assumption 4.1, a pure strategy combination (σ^*, τ^*) of the inspector leadership game Γ_2 with incomplete information is an equilibrium point of Γ_2 if and only if it satisfies the following conditions:

(1) The false alarm probability $\alpha^* = \alpha(\sigma^*)$ is given by

$$\alpha^* = \arg \max_{\alpha \in [0,1]} \Big[(-e\cdot\alpha)\cdot F(K(\alpha)) + + (-a\cdot(1-\beta(\alpha))-c\cdot\beta(\alpha))\cdot(1-F(K(\alpha))\Big] \qquad (5\text{-}8)$$

where $K(\alpha)$ is given by

$$K(\alpha) = \frac{b-\alpha\cdot f}{\beta(\alpha)} - b \qquad (5\text{-}9)$$

and where $\beta(\alpha)$ is given by (5-6).

(2) σ^* is an optimal strategy of the inspector in the auxiliary game G_{α^*}, i.e., it is given by

$$\sigma^* = \arg \max_{\sigma \in \Sigma_{\alpha^*}} \min_{\mu \in \Delta_M} (1-\beta(\sigma,\mu)).$$

(3) For every pair $(t,\sigma) \in [0,\infty) \otimes \Sigma$,

$$\tau^*(t,\sigma) = \begin{cases} 1 & h(t,\sigma) < 0 \\ 1 \text{ or } (\bar{I},\mu^*(t,\sigma)) \quad \text{if} & h(t,\sigma) = 0 \\ (\bar{I},\mu^*(t,\sigma)) & h(t,\sigma) > 0 \end{cases} \qquad (5\text{-}10)$$

where

$$h(t,\sigma) = -b + (b+t)\cdot\beta(\sigma) + f\cdot\alpha(\sigma)$$

and

$$\mu^*(t,\sigma) = \arg \max_{\mu \in \Delta_M} \beta(\sigma,\mu).$$

Proof

We will prove this theorem by "backward induction": Let (σ^*, τ^*) be an equilibrium point of Γ_2. First, for every pair $(t,\sigma) \in [0,\infty) \otimes \Sigma$ and every $\tau = (\bar{I},\mu)$ we have from (5-4)

$$O_2(\sigma,\tau|t) = (b+t)\cdot\beta(\sigma,\mu) - b.$$

Since $b + t > 0$, the left hand side is maximized with respect to μ by maximizing $\beta(\sigma,\mu)$, and this maximum is attained by Assumption 4.2.

Second, given a pair $(t,\sigma) \in [0,\infty) \otimes \Sigma$, the inspectee's expected payoff is

$$\begin{matrix} -f\cdot\alpha(\sigma) & & 1 \\ & \text{if he selects} & \\ -b\cdot(1-\beta(\sigma)) + t\cdot\beta(\sigma) & & \bar{1}. \end{matrix}$$

Choice of the alternative which gives a higher payoff yields (5-9).

To verify condition (1) in Definition 2 we have to show that

$$\sigma^* = \arg \max_{\sigma \in \Sigma} \int_0^\infty I_2(\sigma,\tau^* | t)\ dF(t) \qquad (5\text{-}11)$$

where by (5-3) and (5-9) $I_2(\sigma,\tau^*|t)$ is

$$I_2(\sigma,\tau^*|t) = \begin{cases} -e\cdot\alpha & \text{if} \quad h(t,\sigma) < 0 \\ -a\cdot(1-\beta(\sigma)) - c\cdot\beta(\sigma) & \text{if} \quad h(t,\sigma) > 0. \end{cases}$$

To do this, we first prove the following

Proposition

If (σ,τ) is Nash Equilibrium and if we let

$A = \{t | h(t,\sigma) = 0,\ 0 < q(t,\sigma) < 1\}$,

then $\text{prob}(A) = 0$.

Proof

Let D be given by

$$D = \inf_{\alpha,\beta}(-e\cdot\alpha-[-a\cdot(1-\beta)-c\cdot\beta]) \gtrsim a - e > 0;$$

in words: D is a lower bound for the inspector's payoff difference between 1 and $\bar{1}$ (keeping his own strategy unchanged). For $\varepsilon > 0$, denote by σ_ε the modification of σ such that $\alpha(\sigma_\varepsilon) = \alpha(\sigma) + \varepsilon$. The consequences of changing σ to σ_ε are the following:

- At any value of t in which the inspectee played 1 against σ, he will certainly play 1 against σ_ε, and the inspector's payoff difference will change by the order of ε.

- At any value of t in the event A the inspectee strictly prefers $\vec{1}$ over $\vec{1}$ against σ_ε, while he was indifferent against σ where he played $\vec{1}$ with probability $q(t,\sigma)$. The inspector's payoff thus increases by at least $q(t,\sigma)\cdot D$.

Summing up, the change in the inspector's payoff resulting from changing σ to σ_ε is at least of the order $\text{prob}(A)\cdot q(t,\sigma)\cdot D - k\cdot\varepsilon$ which is positive for ε small enough contradicting the equilibrium property of (σ,τ). ■

Returning to (5-11) and noting

$$\Sigma = \bigcup_{\alpha\,\varepsilon\,[0,1]} \Sigma_\alpha$$

we can rewrite the maximization problem posed by (5-11) as

$$\max_{\alpha\,\varepsilon\,[0,1]} \;\max_{\sigma\,\varepsilon\,\Sigma_\alpha} \int_0^\infty I_2(\sigma,\tau^*\,|\,t)\; dF(t). \qquad (5\text{-}12)$$

We first consider the first maximization in (5-12):

$$\max_{\sigma\,\varepsilon\,\Sigma_\alpha} \int_0^\infty I_2(\sigma,\tau^*\,|\,t)\; dF(t). \qquad (5\text{-}13)$$

For every $\alpha\,\varepsilon\,[0,1]$ and every $\sigma\,\varepsilon\,\Sigma_\alpha$ we have, using the Proposition,

$$\int_0^\infty I_2(\sigma,\tau^*\,|\,t)\; dF(t) = \int_{t:h(t,\sigma)\,<\,0} (-e\cdot\alpha)\cdot dF(t) + \int_{t:h(t,\sigma)\,>\,0} (-a\cdot(1-\beta)-c\cdot\beta)\; dF(t) = \qquad (5\text{-}14)$$

$$= \int_0^{\frac{b-\alpha\cdot f}{\beta}-b} (-e\cdot\alpha)\cdot dF(t) + \int_{\frac{b-\alpha\cdot f}{\beta}-b}^{\infty} (-a\cdot(1-\beta)-c\cdot\beta)\cdot dF(t)$$

where $\beta = \beta(\sigma)$. Since the sum of the two integrals in (5-14) is motonically decreasing as a function of β, $0 \le \beta \le 1$, see AVENHAUS and OKADA (1988), the subproblem (5-13) has the same solution as

$$\min_{\sigma \in \Sigma_\alpha} \beta(\sigma). \tag{5-15}$$

According to (5-7) we have denoted this minimum by $\beta(\alpha)$. Then the maximization problem (5-10) is equivalent to the maximization problem

$$\max_{\alpha \in [0,1]} \int_0^{K(\alpha)} (-e\cdot\alpha)\cdot dF(t) + \int_{K(\alpha)}^{\infty} (-a\cdot(1-\beta(\alpha))-c\cdot\beta(\alpha))\cdot dF(t) =$$

$$= \max_{\alpha \in [0,1]} (-e\cdot\alpha)\cdot F(K(\alpha)) + (-a\cdot(1-\beta(\alpha))-c\cdot\beta(\alpha))\cdot(1-F(K(\alpha)))$$

where $K(\alpha)$ is given by (5-9). Therefore, we can prove (1) and (2) in the theorem. Conversely, we can prove without much difficulty that a pure strategy combination (σ^*, τ^*) satisfying (1) to (3) in the Theorem is an equilibrium point of Γ_2. ■

Theorem 3 gives us the following decisions for the inspector and the inspectee at an equilibrium point of the inspector leadership game Γ_2 with incomplete information.

For any $\sigma \in \Sigma$, the inspectee

(1) selects the illegal strategy μ which maximizes the non-detection probability $\beta(\sigma,\mu)$, and

(2) behaves legally if his payoff t in case of no detection is smaller than the critical level

$$K(\sigma) = \frac{b-f\cdot\alpha(\sigma)}{b(\sigma)} - b$$

and behaves illegally if t is greater than $K(\sigma)$.

The inspector

(3) selects the false alarm probability α^* which maximizes his expected payoff given in (5-8), and

(4) selects the test procedure $\sigma \in \Sigma_{\alpha^*}$ which minimizes the non-detection probability $\beta(\sigma,\mu)$ under the assumption that the inspectee maximizes $\beta(\sigma,\mu)$.

When the inspector employs a test procedure $\sigma(\alpha)$ which is a solution of

$$\min_{\sigma \in \Sigma_\alpha} \max_{\mu \in \Delta_M} \beta(\sigma,\mu), \quad 0 \le \alpha \le 1,$$

the inspectee behaves legally if $t < K(\alpha)$ and illegally if $t > K(\alpha)$ where $K(\alpha)$ is given by (5-9). This means that, from the inspector's point of view, the inspectee behaves legally with probability $F(K(\alpha))$.

For any false alarm probability α, $K(\alpha)$ represents the critical value of the inspectee's payoff in case of no detection which determines his behavior. We can show that $K(\alpha)$ is monotonically increasing in α (see AVENHAUS and OKADA (1988)). This implies that the probability $F(F(\alpha))$ for legal behavior, when the inspector selects a false alarm probability α, is also monotonically increasing in α.

We now assume that the distribution $F(.)$ has a bounded support $[d_0, d_1]$ where $0 < d_0 \le d_1$. Figure 11 illustrates the critical value $K(\alpha)$ in (5-9) between legal behavior 1 and illegal behavior 1, the probability $F(K(\alpha))$ for legal behavior 1, and the inspector's expected payoff function in both cases of complete information ($d=d_0=d_1$) and incomplete information ($d_0 < d_1$). With the help of Figure 11, we can see that the discontinuity of the inspector's payoff with respect to the false alarm probability in case of complete information is removed by the introduction of incomplete information.

We have two possibilities about the equilibrium false alarm probability α^* in case of incomplete information, i.e., (1) an interior solution in (α_0, α_1) where $d_0 = K(\alpha_0)$ and $d_1 = K(\alpha_1)$, and (2) the boundary solution $\alpha^* = \alpha_1$. From the viewpoint of the inspector, the boundary solution $\alpha^* = \alpha_1$ is important because the probability for illegal behavior $\bar{1}$ is zero if he selects α^*. We call an equilibrium point with the boundary solution $\alpha^* = \alpha_1$ the legal behavior equilibrium point. In the following, we provide a sufficient condition for the legal behavior equilibrium point.

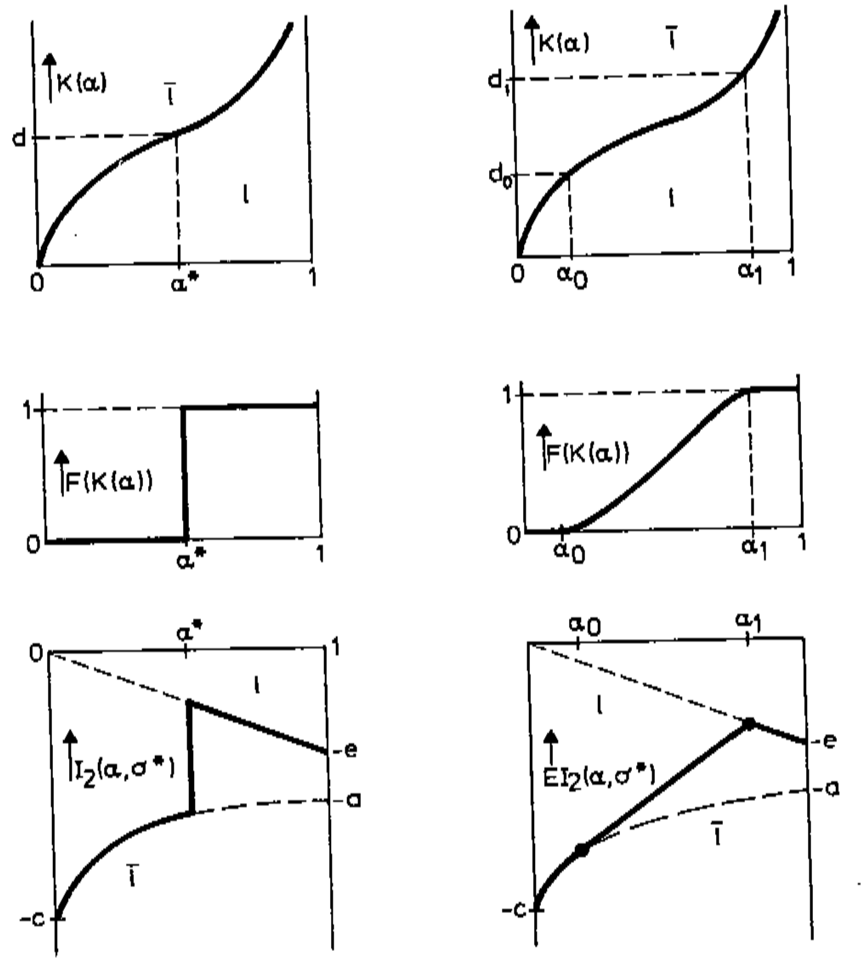

Figure 11: Critical value K(α) and probability F(K(α)) for legal behavior and inspectee's payoff as functions of α for complete (left hand side) and incomplete (right hand side) information: $EI_2(\alpha,\tau^*) = \max_{\sigma \in \Sigma_\alpha} EI_2(\sigma,\tau^*)$

Theorem 4

Assume that $F(x)$ is the uniform distribution on $[d_0, d_1]$ and that $K(\alpha)$ is differentiable in (d_0, d_1) where $K(\alpha_0) = d_0$ and $K(\alpha_1) = d_1$. If the condition

$$\frac{1}{d_1 - d_0} \cdot \left(\min_{\alpha_0 \le \alpha \le \alpha_1} \frac{dK(\alpha)}{d\alpha} \right) \ge \frac{e}{a-e} \qquad (5\text{-}16)$$

holds, then the equilibrium false alarm probability α^* is $\alpha^* = \alpha_1$, and the probability for legal behavior is

$$F(K(\alpha^*)) = \int_{d_0}^{K(\alpha^*)} dF(t) = 1.$$

Proof

We can write $F(t)$ as

$$F(t) = \frac{t - d_0}{d_1 - d_0}, \quad d_0 \le t \le d_1 .$$

Then the inspector's expected payoff is given by

$$G(\alpha) = -e \cdot \alpha \cdot \frac{K(\alpha) - d_0}{d_1 - d_0} + (-a \cdot (1 - \beta(\alpha)) - c \cdot \beta(\alpha)) \cdot \frac{d_1 - K(\alpha)}{d_1 - d_0} .$$

It can be seen without much difficulty that

$$\frac{dG(\alpha)}{d\alpha} = -e \cdot \frac{K(\alpha) - d_0}{d_1 - d_0} - (c-a) \cdot \frac{d\beta(\alpha)}{d\alpha} \cdot \frac{d_1 - K(\alpha)}{d_1 - d_0} +$$

$$+ (a \cdot (1 - \beta(\alpha)) + c \cdot \beta(\alpha) - e \cdot \alpha) \cdot \frac{dK(\alpha)/d\alpha}{d_1 - d_0} .$$

Therefore, we have

$$\frac{dG(\alpha)/d\alpha}{d_1 - d_0} > -e + (a \cdot (1 - \beta(\alpha)) + c \cdot \beta(\alpha) - e \cdot \alpha) \cdot \frac{dK(\alpha)/d\alpha}{d_1 - d_0}$$

for any α, $\alpha_0 \le \alpha \le \alpha_1$. Since

$$a \cdot (1 - \beta(\alpha)) + c \cdot \beta(\alpha) - e \cdot \alpha$$

is monotonically decreasing in α, we obtain

$$\frac{dG(\alpha)}{d\alpha} > -e + (a-e) \cdot \frac{dK(\alpha)/d\alpha}{d_1 - d_0} \quad \text{for all} \quad \alpha : \alpha_0 \le \alpha \le \alpha_1 .$$

The condition (5-16) implies that

$$dG(\alpha)/d\alpha > 0 \quad \text{for all} \quad \alpha : \alpha_0 \le \alpha \le \alpha_1 .$$

Therefore, the boundary value $\alpha^* = \alpha_1$ is the solution of the maximization problem (5-8) in Theorem 3. ∎

Using some properties of the function $K(\alpha)$ which are given in AVENHAUS and OKADA (1988), the condition (5-16) shows us two situations in which we can obtain the legal behavior equilibrium point:

(1) the inspector's cost e in case of false alarm is small, and

(2) the possible values of the inspectee's benefit t in case of no detection are high under the condition that the difference $d_1 - d_0$ is fixed.

We can interpret our results in both cases as follows: In case (1), the inspector does not worry very much about his loss in case of false alarm. Therefore, he can employ the test procedure with the high false alarm probability α_1. In case (2), the inspectee has a high incentive to behave illegally. Therefore, the inspector employs the test procedure with the high false alarm probability α_1 in order to prevent the inspectee from behaving illegally.

6. Applications

Even though we assumed that the inspector is incompletely informed about the inspectee's gain in case of undetected illegal behavior, in practice, it will be difficult if not impossible for him to get any information about the inspectee's payoff parameters. Therefore, it is very important that Theorem 3 gives us an advise for the construction of best test procedures which do not depend on the players' payoff parameters, once the value of the false alarm probability is given. In fact, this theorem establishes a bridge between the game theoretical approach which is necessary for the appropriate description of the inspection problem and the traditional statistical treatment where the error second kind probability is minimized for given value of the error first kind probability.

According to Theorem 3, the optimal decision scheme of the inspector for a given value of the false alarm probability α is the scheme guaranteeing the *maxmin value* of the auxiliary zero-sum game $G_\alpha = (\Sigma_\alpha, \Delta_M, 1-\beta)$ in which the inspector as player 1 chooses the test $\sigma \in \Sigma_\alpha$, the inspectee as player 2 chooses $\mu \in M$, and the payoff from player 2 to player 1 is the detection probability $1 - \beta(\sigma,\mu)$. So, one has to find

$$\max_{\sigma} \min_{\mu} (1-\beta(\sigma,\mu)) \quad (6\text{-}1)$$

and the coresponding strategies. Operationally, it is much more convenient to compute

$$\min_{\mu} \max_{\sigma} (1-\beta(\sigma,\mu)) \quad (6\text{-}2)$$

which, of course, would be the same if G_α had a value. However, the existence of a value G_α does not follow from a standard minmax theorem, since the payoff function $1 - \beta(\sigma,\mu)$ does not satisfy the usual convexity requirements. Therefore, in applications one establishes the existence of this value by finding it as well as the optimal strategies.

There are many applications of this procedure (see AVENHAUS (1986)). As examples, we consider once more the two decision problems which we formulated in the second section. In the following, we consider only the auxiliary game G_α which is an important element of the solution of both games, with or without commitment power.

In doing so, we make full use of the multinormal distribution properties in conjunction with the Neyman-Pearson test. We will formulate a general theorem which covers both problems given in the second section and, thereafter, apply it to these problems, specifically.

Theorem 5

Given the multivariate normally distributed random vector X and the two alternative hypotheses

$$H_0(1) \; : X \sim N(0,\Sigma) \quad (6\text{-}3a)$$

$$H_1(\bar{1}) \; : X \sim N(\mu,\Sigma), \; \mu'\cdot e = M. \quad (6\text{-}3b)$$

(Do not confuse the covariance matrix Σ with the inspector's strategy set Σ_α now!) Consider the game $G_\alpha = (\Sigma_\alpha, \Delta_M, 1-\beta)$ in which the inspector as player 1 chooses the test $\sigma \in \Sigma_\alpha$ with a given false alarm probability α for a decision beween H_0 and H_1, player 2 chooses a diversion strategy $\mu \in \Delta_M$, and where the payoff from player 2 to player 1 is the detection probability $1 - \beta(\sigma,\mu)$.

ıen the game G_α has a value v given by

$$v = 1 - \beta = \phi\left(\frac{M}{\sqrt{e'\cdot\Sigma\cdot e}} - U(1-\alpha)\right) \quad (6\text{-}4)$$

ıere ϕ is the cumulative standard normal distribution, and U its ıverse. The optimal strategy σ_α^* of the inspector is given by the ʟarm set

$$\{x \mid x'\cdot e > \sqrt{e\cdot\Sigma\cdot e}\cdot U(1-\alpha)\}. \quad (6\text{-}5)$$

ıe optimal strategy of the operator is the (pure) diversion vector μ^* iven by

$$\mu^* = \frac{M}{e'\cdot\Sigma\cdot e}\cdot\Sigma\cdot e. \quad (6\text{-}6)$$

roof

e shall prove this theorem by showing that (σ_α^*, μ^*) is in fact a addle point for $1 - \beta(\sigma_\alpha, \mu)$ which means for all σ_α and μ

$$\beta(\sigma_\alpha^*, \mu) \le \beta(\sigma_\alpha^*, \mu^*) \le \beta(\sigma_\alpha, \mu^*). \quad (6\text{-}7)$$

e start by proving the right hand side inequality.

f the inspectee behaves legally, the density of the random variable X s according to (6-3a)

$$f_0(x) = (2\Pi)^{-n/2}\cdot|\Sigma|^{-1/2}\cdot\exp(-\frac{1}{2}\cdot x'\cdot\Sigma^{-1}\cdot x) \quad (6\text{-}8a)$$

'hile, if he makes a deviation μ, the density is according to (6-3b)

$$f_\mu(x) = (2\Pi)^{-n/2}\cdot|\Sigma|^{-1/2}\cdot\exp(-\frac{1}{2}\cdot(x-\mu)'\cdot\Sigma^{-1}\cdot(x-\mu)). \quad (6\text{-}8b)$$

ıs it is well known that, for any given μ, the most powerful test σ_μ :or testing (the alternative hypothesis) f_μ versus (the null ıypothesis) f_0 with type I error probability α is the Neyman-Pearson :est the alarm set of which is given by:

$$\{x \mid \frac{f_\mu(x)}{f_0(x)} > K_\alpha''\} = \{x \mid x'\cdot\Sigma^{-1}\cdot\mu > K_\alpha'\}$$

vhere K_α'' and K_α' are constants determined by α (only). This means :hat for all μ and the corresponding Neyman-Pearson test σ_μ we have

$$\beta(\sigma_\mu, \mu) \le \beta(\sigma, \mu) \quad \text{for all} \quad \sigma.$$

?or $\mu = \mu^*$ denote the corresponding σ_{μ^*} by σ^* whose alarm set is therefore:

$$\{x \mid \frac{M}{e'\cdot\Sigma\cdot e}\cdot x'\cdot e > K_\alpha'\} \quad \text{or} \quad \{x \mid x'\cdot e > K_\alpha\}$$

for some other constant K_α.

Now, as a linear combination of multinormal random variables, the random variable $x'\cdot e$ is normally distributed with expectation $E(x'\cdot e) = E(x')\cdot e = \mu'\cdot e = M$ under f_μ and $E(x'\cdot e) = 0$ under f_0. For both alternatives (and indpendently of μ) the variance of this random variable is $var(x'\cdot e) = e'\cdot\Sigma\cdot e$. Therefore, the false alarm probability α is given by

$$1 - \alpha = prob_0(x'\cdot e \le K_\alpha) = \phi\left(\frac{K_\alpha}{\sqrt{e'\cdot\Sigma\cdot e}}\right),$$

and the probability of not detecting the deviation μ is given by

$$\beta^* = prob_{\mu^*}(x'\cdot e \le K_\alpha) = \phi\left(\frac{K_\alpha - M}{\sqrt{e'\cdot\Sigma\cdot e}}\right).$$

Eliminating K_α from these two equations implies (6-4) and (6-5).

In order to complete the proof, it remains to establish the left hand side inequality of (6-7). This, in fact, follows readily, since for all μ with $\mu'\cdot e = M$

$$E(e'\cdot X) = e'\cdot\mu = M,$$

and the variance of $e'\cdot X$ is independent of μ, implying

$$\beta(\sigma_\alpha^*,\mu) = \beta(\sigma_\alpha^*,\mu^*) \quad \text{for all} \quad \mu. \qquad \blacksquare$$

This result shows that in both examples, i.e., independently of the structure of the covariance matrix, the optimal strategy of the inspector is to just add all observed data and to perform a threshold test on the sum. Notice that this procedure is independent of the value of M, the order of magnitude of which the inspector might know, but not its precise value.

In case of the material accountancy example this result is especially interesting: With (5-4) we get

$$\sum_{i=1}^{n} X_i = \sum_{i=1}^{n} (I_{i-1}+R_i-S_i-I_i) = I_0 + \sum_{i=1}^{n} R_i - \sum_{i=1}^{n} S_i - I_n$$

which is the material balance for the whole sequence of inventory periods. Hence, this result means that all intermediate inventories $I_1,\ldots,I_{n-1}$ must not be taken into account in the optimal procedure.

Let us repeat that this surprising result has been derived under the assumption that "time is not valuable" as formulated in the second tion: Obviously, one will have to use the intermediate inventories closing intermediate balances if it is important for the inspector detect *timely* any diversion of material.

case of the data verification example the result shows that the so-led D-statistic

$$D := \sum_{i=1}^{n} X_i$$

ch has been introduced many years ago also in the area of auditing KE, NETER, LEITCH (1982)) is the best test statistic of the pector.

us mention that also in the other extreme of minimal sample size 1) the trivial D-statistic D = x is the best test statistic, ever, the optimal falsification strategy is more complicated 'ENHAUS, BATTENBERG, FALKOWSKI (1983)): If the total falsification smaller than some critical value, then again the equal distribution optimal, otherwise, only one datum has to be falsified. This is uitive: For small total falsification, the falsification can be dden" in the measurement uncertainty, whereas for large total sification this is no longer possible; therefore, here the pectee has to play "vabanque": He has to put everything on one card hope that this card is not drawn. For sample sizes between one and the situation gets very complicated: numerical calculations licate that beyond some total falsification the D-statistic is no ger optimal, however, one can give upper limits of that total sification below which the D-statistic is optimal.

se results are important for practitioners: For many years already IAEA has used the D-statistic for the comparison of operators' and spectors' data; originally, it has been justified by heuristic guments (STEWART (1970)). The theory presented here shows its range applicability and, furthermore, how the best test procedure would ve to be determined in general.

7. References

R. AVENHAUS, Safeguards Systems Analysis - With Applications to Nuclear Safeguards and other Inspection Problems. Plenum Press, New York and London 1986

R. AVENHAUS, H.P. BATTENBERG, B.J. FALKOWSKI, Optimale Testverfahren bei der Datenverifikation (Optimal Data Verification Test Procedures). Methods of Operations Research 50, pp. 154-164, 1985

R. AVENHAUS and H. FRICK, Analyse von Fehlalarmen in Überwachungssystemen mit Hilfe von Zweipersonen-Nichtnullensummenspielen (Analysis of False Alarms in Safeguards Systems by Means of Two-Person Non-Zero-Sum Games). Operations Research Verfahren XXVI, pp. 629-639, 1977

R. AVENHAUS and A. OKADA, Inspector Leadership Games with Incomplete Information. Discussion Paper No. 17 of the Zentrum für Interdisziplinäre Forschung, Universität Bielefeld, July 1988

R. AVENHAUS and S. ZAMIR, Safeguards Games with Applications to Material Control. Discussion Paper No. 12 of the Zentrum für Interdisziplinäre Forschung, Universität Bielefeld, May 1988

H.P. BATTENBERG, Optimale Gegenstrategien bei Datenverifikationstests (Optimal Counterstrategies for Data Verification Tests). Ph.D. Dissertation of the Hochschule der Bundeswehr München, 1983

D. BIERLEIN, Direkte Überwachungssysteme (Direct Inspection Systems). Opeations Research Verfahren VI, pp. 57-68, 1969

D. BIERLEIN, Auf Bilanzen und Inventuren basierende Safeguards-Systeme (Safeguards Systems Based on Balances and Inventories). Operations Research Verfahren III, pp. 36-43, 1970

D. BIERLEIN, Game Theoretial Models of Safeguarding Different Types of Illegal Activities. Proceedings of the 4th Formator-Symposium on Mathematical Methods for the Analysis of Large-Scale Systems, Czechosl. Acad. Sc., Prague 1983

S.J. BRAMS, Superpower Games. Yale University Press, New Haven, London 1985

M. DRESHER, A Sampling Inspection Problem in Arms Control Agreements: A Game Theoretic Analysis. Memorandum RM-2972-ARPA, The Rand Corporation, Santa Barbara 1962

G.L. DUKE, J. NETER, P.A. LEITCH, Power Characteristics of Test Statistics in the Auditing Environment: An Empirical Study. Journal of Accounting Research, Vol. 20, No. 1, pp. 42-47, Spring 1982

R.A. DYE, Optimal Monitoring Policies in Agencies. Rand Journal of Economics, Vol. 17, No. 3, pp. 339-350, Autumn 1986

J. FICHTNER, Statistische Tests zur Abschreckung von Fehlverhalten (Statistical Tests for Deterring Illegal Behavior). Ph.D. Dissertation, Universität der Bundeswehr München 1985

H. FRICK, Spieltheoretische Behandlung mehrfacher Inventurprobleme (Game Theoretical Treatment of Multiple Inventory Problems). Ph.D. Dissertation, Universität Fridericiana Karlsruhe 1976

J. HARSANYI, Games with Incomplete Information Played by "Bayesian" Players. Parts I-III, Management Science 14, pp. 159-162, 320-334, and 486-502, 1967/68

E. HOPFINGER, Zuverlässige Inspektionsstrategien (Reliable Inspection Strategies), Z. Wahrscheinlichkeitstheorie verw. Gebiete, pp. 35-46, 1974

International Atomic Energy Agency, The Structure and Content of Agreements between the Agency and States Required in Connection with the Treaty on the Non-Proliferation on Nuclear Weapons, IAEA Document, INF/CIRC/153, Vienna 1971

C.S. KANODIA, Stochastic and Moral Hazard. Journal of Accounting Research, Vol. 23, No. 1, pp. 175-193, Spring 1985

H.W. KUHN, Recursive Inspection Games, Applications of Statistical Metholology to Arms Control and Disarmament. A US ACDA Report under Contract ACDA/ST-3, pp. 169-182, 1963

E.L. LEHMANN, Testing Statistical Hypotheses. Wiley, New York 1959

M. MASCHLER, A Price Leadership Method for Solving the Inspector's Non-Constant-Sum Game. Nav. Res. Logistics Q. 13, pp. 11-33, 1966

M. MASCHLER, The Inspector's Non-Constant-Sume Game: Its Dependence on a System of Detectors. Nav. Res. Logistics Q. 14, pp. 275-290, 1967

R. SELTEN, A Simple Model of Imperfect Competition, Where 4 are Few and 6 are Many. Int. J. Game Theory 3, pp. 141-201, 1973

H. v.STACKELBERG, Marktform und Gleichgewicht (Market Form and Equilibrium). Julius Springer, Berlin 1934

K.B. STEWART, A Cost-Effectiveness Approach to Inventory Verification. Proceedings of the IAEA Symposium Safeguards Techniques in Karlsruhe, Vol. II, pp. 387-409, IAEA, Vienna 1971

E. Van DAMME, Stability and Perfection of Nash Equilibria, Springer-Verlag 1986