

NUREG/CR-3926  
Vol. 2

---

# Strategic Analysis for Safeguards Systems: A Feasibility Study

Appendix

---

Prepared by A. J. Goldman

The MAXIMA Corporation

Prepared for  
U.S. Nuclear Regulatory  
Commission

8501070511 841231  
PDR NUREG  
CR-3926 R PDR

## NOTICE

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, or any of their employees, makes any warranty, expressed or implied, or assumes any legal liability of responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights.

## NOTICE

### Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 1717 H Street, N.W.  
Washington, DC 20555
2. The NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission,  
Washington, DC 20555
3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC Office of Inspection and Enforcement bulletins, circulars, information notices, inspection and investigation notices; Licensee Event Reports; vendor reports and correspondence; Commission papers, and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the NRC/GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, and NRC booklets and brochures. Also available are Regulatory Guides, NRC regulations in the *Code of Federal Regulations*, and *Nuclear Regulatory Commission Issuances*.

Documents available from the National Technical Information Service include NUREG series reports and technical reports prepared by other federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal and periodical articles, and transactions. *Federal Register* notices, federal and state legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free, to the extent of supply, upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, and are available there for reference use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

NUREG/CR-3926  
Vol. 2  
RS, 1S

---

# Strategic Analysis for Safeguards Systems: A Feasibility Study

## Appendix

---

Manuscript Completed: September 1984  
Date Published: December 1984

Prepared by  
A. J. Goldman\*

The MAXIMA Corporation  
7315 Wisconsin Avenue - Suite 900N  
Bethesda, MD 20814

\*Consultant  
The Johns Hopkins University  
Baltimore, MD 21218

Prepared for  
Division of Risk Analysis and Operations  
Office of Nuclear Regulatory Research  
U.S. Nuclear Regulatory Commission  
Washington, D.C. 20555  
NRC FIN B6763

## ABSTRACT

This appendix provides detailed information regarding game theory (strategic analysis) and its potential role in safeguards to supplement the main body of this report. In particular, it includes an extensive, though not comprehensive review of literature on game theory and on other topics that relate to the formulation of a game-theoretic model (e.g. the payoff functions). The appendix describes the basic form and components of game theory models, and the solvability of various models. It then discusses three basic issues related to the use of strategic analysis in material accounting: (1) its understandability; (2) its viability in regulatory settings; and (3) difficulties in the use of mixed strategies. Each of the components of a game theoretic model are then discussed and related to the present context.

## TABLE OF CONTENTS

|   | <u>Page</u> |
|---|-------------|
| ABSTRACT . . . . .  | iii         |
| ACKNOWLEDGEMENTS . . . . .                                      | vii         |
| EXECUTIVE SUMMARY . . . . .                                     | 1           |
| 1.0 BACKGROUND AND OUTLINE . . . . .                            | 3           |
| 2.0 SOME CONCEPTS OF GAME THEORY . . . . .                      | 9           |
| 2.1 The Players and Their Relations . . . . .                   | 10          |
| 2.2 The Strategy Spaces . . . . .                               | 11          |
| 2.3 The Payoff Functions . . . . .                              | 12          |
| 2.4 Solution Concepts . . . . .                                 | 15          |
| 2.5 The Extensive Form . . . . .                                | 19          |
| 2.6 Stochastic and Repeated Games . . . . .                     | 21          |
| 3.0 SOLVABILITY OF GAME-THEORETIC MODELS . . . . .              | 27          |
| 3.1 Matrix Games . . . . .                                      | 27          |
| 3.2 Other Two-Person Zero Sum Games . . . . .                   | 30          |
| 3.3 Bimatrix Games . . . . .                                    | 33          |
| 3.4 Multi-Player Games . . . . .                                | 36          |
| 4.0 GENERIC ISSUES OF FEASIBILITY . . . . .                     | 40          |
| 4.1 Understandibility of Game Theoretic<br>Techniques . . . . . | 41          |
| 4.2 Viability in Regulatory Settings . . . . .                  | 42          |
| 4.3 Possible Difficulties with Mixed<br>Strategies . . . . .    | 48          |
| 5.0 THE PLAYERS . . . . .                                       | 56          |
| 5.1 More Than Two? . . . . .                                    | 58          |
| 5.2 Less Than Two? . . . . .                                    | 59          |
| 6.0 THE STRATEGY . . . . .                                      | 63          |
| 6.1 The Diverter's Strategies . . . . .                         | 63          |
| 6.2 The Defender's Strategies . . . . .                         | 67          |
| 7.0 THE DIVERTER'S STRATEGIES . . . . .                         | 76          |
| 7.1 Review of Literature . . . . .                              | 77          |

|     |  |    |
|-----|--|----|
| 7.2 | The Zero Sum Assumption: Alternatives. . . . | 79 |
| 7.3 | The Zero Sum Assumption: Pros & Cons . . . . | 85 |
| 8.0 | THE DEFENDER'S PAYOFF FUNCTION. . . . .      | 89 |
| 8.1 | Review of Literature. . . . .                | 89 |
| 8.2 | Further Discussion. . . . .                  | 93 |
|     | REFERENCES . . . . .                         | 99 |

## ACKNOWLEDGEMENTS

This appendix describes the results of the writer's research and reflections during the project's duration on a number of the issues under study. These views have benefitted from the project team's two review meetings with our NRC project officer and interested colleagues and, of course, from many informal exchanges with other members of the project team; but they do not reflect the final intense integrative effort that will meld these results (perhaps modified or enlarged) with those of other team members to achieve a full "team response" to the series of specific questions used by the NRC in structuring its formulation.- At the Johns Hopkins University, numerous discussion with Professor J.A. Filar have been extremely helpful. Errors of fact, interpretation or omission remain the writer's responsibility.

## EXECUTIVE SUMMARY

The procedures for setting alarm thresholds based on inventory differences (IDs) in the accounting of special nuclear material have traditionally been based on concepts of statistical quality control and hypothesis testing. This approach has come under some criticism because it is not specifically sensitive to diversion by intelligent adversaries. The theory of games provides a modeling framework which can explicitly identify the "best" course of action against an intelligent adversary.

The NRC has previously undertaken research to develop preliminary game theory models for the material accounting context. These and other related efforts are used as points of reference in the discussions here regarding the applicability of game theory in this context, and the development of specific components of the model.

The basic elements of a game theoretic model are:

- o the players and their allowed relations,
- o the players' strategy spaces, and
- o the players' payoff functions.

The first element includes the number of players, and their ability to communicate and collaborate. Strategy spaces refer to the courses of action that are available to the players, including probabilistic mixtures of strategies. Payoff functions are the mathematical representation of the value that a player receives based on the outcome of a game. These elements must be supplemented by a solution concept. The most general solution for noncooperative games is that of an equilibrium point--a solution (action by each player) from which no player has any incentive to deviate unilaterally, i.e. without arrangement that another player will also deviate. Other solutions are possible, depending on the form of the game. Basic game theory can be extended in several ways to provide more realistic (but often more difficult to solve) models, for example, stochastic games in which players move from game to game, and repeated games with incomplete information players may be playing any one of several possible games but do not know for sure which one.

Three specific issues have been raised regarding the applicability of game theory in setting ID alarm thresholds:

- o the understandability of game theory,
- o the viability of game theory in regulatory settings, and



o possible difficulties of using mixed strategies in implementatation.

The first issues does not appear to be a problem for three reasons: (1) Game theory has a long and successful history of popularization for "lay" persons. (2) At a more technical level, game theory is not regarded as son complex that its study need be deferred to graduate school. (3) It is the responsibility of those engaged in safeguards activities to provide needed technical expertise; and given (1) and (2), this is not a costly requirement. Regarding the viability of game theory in regulatory settings, mathematical models in general have been found to be acceptable in regulatory settings, provided they are not unreasonable, arbitrary, and capricious, and do not contradict relevant data or well-established theory.

Mixed strategies--those in which various strategies are selected with specified probabilities--have the advantage of denying the adversary information regarding just what the player will do under a specific set of circumstances. This can result in higher payoffs to the player using a mixed strategy. On the other hand, use of mixed strategies entails planning--with its associated costs--for more possible actions. It may be possible to develop "near-optimal" pure strategies that are an acceptable compromise.

In addition to these practical issues, several technical issues are also addressed. With respect to the number of players, it is concluded that two players provide an adequate representation of the context, particularly since games with more than two players are more difficult to solve. The inclusion of multiple sites and multiple accounting periods appears to be practical in the specification of strategy spaces. A serious difficulty arises in the specification of the payoff functions. While, the zero sum assumption, in which the diverter's payoff functin is assumed to be the negative of the defender's, seems to be reasonable--though not necessarily uncontroversial--the development of the payoff function is difficult. It must take into account a wide range of motives for diversion as well as the potential uses of any diverted material. This process involves both predicting outcomes and attributing value to those outcomes. Analytic methods, e.g. multiattribute utility theory, exist for developing such function, but they involve subjective judgment which maybe difficult to obtain or justify.

## 1.0 BACKGROUND AND OUTLINE

The safeguards program whose direction is vested in the Nuclear Regulatory Commission (NRC) is a body of regulatory, operational, and research activities aimed at protecting society from the danger implicit in having sensitive nuclear material fall into "the wrong hands." It is common, and conceptually rather natural, to regard the program as composed of three mutually reinforcing but distinctive subprograms:

- Physical Security, involving (1) controls (checkpoints, physical barriers, etc.) over access to and egress from the material, (2) surveillance and alarm systems, and (3) active responses to intrusions;
- Material Control, involving the governance of and responsibility for current movements, locations, and status of the material; and
- Material Accounting, involving the measurement and assay of material quantities and the recording/analysis/reporting of resultant information as a check against loss or diversion.

These "functional" subprograms and their integration require, of course, a variety of supportive activities: managerial, evaluative, analytical, regulation-promulgating, and the like.

Our focus in this document is on the third of these functional subprograms, material accounting. Its after-the-fact nature, and its preoccupation with data rather than explicitly with people or with nuclear material, make its role less dramatic or palpable than those of the other two subprograms. That this role is nevertheless essential, is established in a careful analysis [1;Section 5]<sup>1</sup> of the contributions of material accounting to articulated objectives of the safeguards program.

Although the definition of material accounting might perhaps be construed to include the analysis of data and records generated in the normal management of a facility's operations (batch yields, quality control figures, etc.), our main concern will be with the evaluation of data from additional material balances, inventories, and records provided specifically for "safeguards" purposes. Thus, the characteristic situation to be considered involves "striking a balance" in the customary accounting sense for a particular material balance area (MBA) at the end of a time period; that is, checking the "balance" equation:

---

<sup>1</sup>Numbers in squared braces refer to the list of references.

$$(\text{current contents}) = (\text{prior contents}) + (\text{inputs}) - (\text{outflows}) \quad (1.1)$$

where the left hand term represents the result of a current physical inventory of the material within the MBA, the first term on the right-hand side of the equation is the current estimate of that material's quantity at the start of the time period, and the second and third terms on the right are sums of recorded measured values referring to movements of material into and out of the MBA during the period.

The fact that actual physical measurement processes have inherently limited precision, together with the possibility of human error in any accounting/inventory effort, make it most unlikely that equation 1.1 will hold exactly. The initial estimation of the right-hand side, corresponding to the striking of a trial balance in a double-entry set of accounts, may identify some anomalies whose reconciliation leads to a revised right-hand side, the book inventory. But it remains highly unlikely that even this "improved" value will check perfectly with the physical inventory figure on the left. The discrepancy is presently termed the inventory difference (ID).

If an ID value is "sufficiently small" (i.e., sufficiently close to zero), it can plausibly be regarded as arising simply from the inevitable imperfection of the measurement processes involved. But a "sufficiently large" ID in one time period or over several periods, suggests that the measurement and recording system may have drifted below an acceptable quality of performance, that some material "sinks" or process-loss modes have gone unrecognized, that significant discrete errors may have occurred during the time period, or even (if the ID has the appropriate sign) that a theft or diversion of material has taken place. Such possibilities in turn lead to vigorous and often expensive reactions: an intensified scrutiny of measurement and bookkeeping procedures and of security and control records, a search for material possibly missed in the physical inventory (this can require slowdown or even shutdown of the MBA's normal operations, which in a "bottleneck" case could paralyze much of the facility), and possibly the notification or actual involvement of security and external law-enforcement authorities.

The last paragraph's weasel words "sufficiently small" and "sufficiently large" point up the underlying issue: how and where to set the "alarm threshold" for ID values that separates the satisfyingly-small values calling for no response (and providing evidence for a "hoax" classification of some claimed diversion of material) from the response-requiring larger values. A low threshold may lead to disruptively frequent, unnecessary interruptions of the plant's operations (a high false alarm rate or "Type 1 error probability"), a high threshold, to excessive risk (miss rate or "Type 2 error probability") of failing to initiate a desirable corrective response. Thus, threshold-setting

in the ID context presents, in its own distinctive way, the risk-benefit trade off problem generic in modern regulatory analysis.

The traditional and still-prevalent conceptual framework that presently governs the setting of these alarm thresholds is that of statistical quality control and hypothesis testing [2]. Estimates are made of the probability distributions for errors in the measurement processes whose results enters the terms of equation 1.1. Appropriate mathematical operations on these distributions yield an estimated probability distribution for the difference of the equation's two sides, i.e., for the ID. Confronting this distribution with the actual numerical value calculated for the ID yields an estimate of the probability that a value so different from zero could arise by sheer chance if the measurement systems were operating as postulated and no other sources of error were active. Should that probability be high enough (the conventional level is 95%), then "sheer chance" is accepted as an adequate explanation of the calculated ID-value; if not, the presumption is that there are contributors to the imbalance other than random measurement error. The alarm threshold, then, is the "95th percentile level" of the estimated probability distribution for ID.<sup>2</sup>

This simple idea admits refinement and strengthening in a variety of ways. Approaches to estimating and tracking the error probability distributions of measurement processes can be improved. Distributions of errors from additional sources (e.g., recordkeeping) might be estimated and incorporated in the analysis. Initially neglected interactions between and among error sources can be identified and then properly reflected in the "appropriate mathematical operations" mentioned above. The multi-time period nature of the situation being analyzed might be better exploited (cf. [3;pp. VI 27-32]). A substantial investment of analytical, managerial, and expository effort and expertise has gone into elaborating this approach to attain greater discriminatory power and realism (cf. [4] for one recent example), and into making its application smooth-running and well-understood. Furthermore, its use in the safeguards program has the comforting advantage of ample precedent, by analogy with its common use (for example) in industrial process control and in comparison of scientific hypotheses with data.

A challenge to this well-established paradigm has arisen in the past few years. On the critical side, the challenge observes that although the safeguards program is fundamentally concerned with the possibilities of diversion and theft--threats posed by an intelligent adversary--the statistical methodology

---

<sup>2</sup> For simplicity, we ignore distinctions between one-sided and two-sided tests.

has no feature specifically sensitizing it to such threats; i.e., no conceptual element distinguishing contributions to ID by "innocent chance" from the more serious possible ones due to a malevolent act. Nor does it give explicit per se consideration in setting the alarm threshold to the nature, effectiveness and costs of the responses actuated by an "alarm"--though some informal consideration of these points must be reflected in the above-mentioned 95% figure, corresponding to a 5% false-alarm rate. On the positive side, the challenge notes the existence of a branch of applied mathematics targeted directly at identifying "best" courses of action versus an intelligent adversary, namely the theory of games ("strategic analysis"). In its most aggressively-advanced form, then, the gist of the challenge is that an appropriate game-theoretic analysis should supplant the previously described approach as the basis for alarm-threshold setting.

Exploration of this alternative methodology required efforts to develop an "appropriate game-theoretic analysis." NRC-supported research with this aim reached a first milestone with the appearance of Siri, et. al. [5] and its subsequent journal-paper version [6] by Dresher and Moglewer. A further extension, in which the alarm threshold is no longer assumed fixed prior to the inventory (and hence is no longer knowable in advance by the adversary), is formulated and analyzed in Siri, Ruderman, and Dresher [7]. The work of Avenhaus and various collaborators which is in the somewhat different (but clearly relevant) context of safeguards issues faced by the International Atomic Energy Agency, has appeared in a number of publications, of which we note here only the monograph [8], the journal papers [9, 10], and the course notes [11]; other European literature includes Beinbauer and Bierlein [12] and Hopfinger [13]. (It is professionally disturbing to see that neither of these two lines of research show awareness of the other.)

The novelty of the game-theoretic approach, relative to prior practice, led the NRC-related Material Control and Material Accounting Task Force to conclude [<sup>3</sup>vol. 1, p. 5-33] that it lacked time for a proper evaluation of this "significant area of current technical assistance effort." The Task Force recommended that the NRC undertake a peer review of this methodology by a group of suitable government, academic and industrial professionals. A Peer Review Group, on which the present writer served, was then formed; it was organized by J. H. Opelka (Argonne National Laboratory) chaired by R. F. Lumb (NUSAC, Inc.), and included major contributors to the statistical methodology described above (C. A. Bennett and J. L. Jaech), noted safeguards experts (W.A. Higginbotham and Lumb) and a distinguished academic game-theorist (W. A. Lucas). The Group's consensus conclusions reported in [14], along with detailed supporting discussions plus additional

---

<sup>3</sup>Bibliographic completeness has not been attempted.

viewpoints of individual Group members, included the following (this writer's paraphrase):

- Because of its direct consideration of the antiadversary objective of the safeguards program, game theory is an especially promising tool for use in that program, specifically (but not exclusively) in material accounting to develop a rationale for action in response to ID values.
- However, the particular game-theoretic formulation proposed in [5-7] was not convincing as to validity and therefore not recommended for application; the NRC was encouraged to undertake research and development activities needed to achieve a formulation satisfactory in both validity and in actual workability.
- Issues of "workability" might include the relative unfamiliarity of the approach, its possible call for probabilistic mixtures of responses, and its need for information (e.g., on response costs) and for judgments (e.g., on quantifying society's concern with identification and estimation of diversions) beyond the requirements of the current methodology.
- A successful game-theory formulation, though achieving a higher probability of "alarm when there should be" on a cost-effective, diversion-sensitized basis, would not replace the statistical methodology in the latter's role of "quality control" assurance relative to NRC-licensees' material accounting measurements and procedures.

In response to these recommendations, the NRC solicited proposals for "Strategic Analysis of Safeguards Systems: A Feasibility Study" to analyze further the potential practicability of the game-theoretic approach in the proposed regulatory setting and to assess the likely cost, value, and success chances associated with further research into particular technical and implementation issues. Embodied in the NRC's Request for Proposal No. RS-RES-82-022 (July 12, 1982), the solicitation led to award of the study to a project team assembled by the MAXIMA Corporation, involving senior staff from that organization and from International Energy Associates Limited (IEAL), as well as this writer.

The body of the document is organized as follows. The following Section 2, in order to establish a common vocabulary for the balance of the text, reviews some basic concepts of game theory (and thus the basic ingredients of a game-theoretic formulation of a decision situation). Section 3 takes up a question of both technical and practical nature: the possibility

and difficulty, mathematically and computationally, of actually "solving" a game-theoretic "model" to determine an "optimal" course of action (e.g., setting an alarm-threshold or choosing a response to a particular ID-value). At that point it is possible, in Section 4, to address some issues of feasibility (for a game-theoretic approach) which are relatively "generic" in that their discussion can precede an analysis of the technical specifics of applying game-theory to the particular class of situation at hand.

The second part of the document addresses these technical specifics one by one, in Section 5 through 8, respectively. The NRC-sponsored models [5-7], those of Avenhaus and Frick [9-10] and others from the literature, are used as points of reference throughout. But because our focus is on discussing issues and ideas pertinent to developing an operational game-theoretic model, and not on presenting a "literature review" as such, each of these models appears in disjointed parts in the text--for example, the "adversary's strategy spaces" for all of them are described together in Section 7.1, in connection with that particular element of a game-theoretic model. Our discussions of modeling "the adversary," though based only on modest library research rather than original thought or established experience, go far beyond the efforts of [5-7].

Topics requiring treatment have ranged over a variety of disciplines and subdisciplines, each with a substantial and growing literature. Thus there was no hope, within the scope of this study, of attaining bibliographic "completeness" for the references. A distinct effort has been made, however, to assist the reader or potential modeler in gaining a rapid foothold in the literature of possible unfamiliar areas. One useful tactic in the references for this purpose is to examine an entire listed paper rather than just the particular passage or pages cited, or an entire listed collection rather than just the particular paper cited.

## 2.0 SOME BASIC CONCEPTS OF GAME THEORY

A number of branches of mathematics, for example, differential calculus and linear programming, provide techniques useful in maximizing or minimizing a mathematical function of one or several variables, perhaps subject to more-or-less complicated constraints. If we regard these variables or quantities as under the control of some decisionmaker (who sets their numerical levels), the constraints as defining the limits on this level-setter's freedom of choice (due, e.g., to the scarcity of some resource), and the function to be extremized as "scoring" the utility or cost to the decisionmaker of each possible combination of numerical values for the values; then the mathematical problem just described can be interpreted as seeking a best course of action for the decisionmaker. The general terms "mathematical optimization" and "mathematical programming" are often applied to such problems and to the methods used to attack them.

The theory of games ("game theory") deals with a significantly broader and more difficult class of situations: those in which two or more decisionmakers ("players") are involved, in which each controls some of the variables concerned, and in which the utility or cost experienced by each player depends at least in part on what choices other players make. Thus, the optimization problems of the preceding paragraph might be viewed as degenerate "merely single-player" special cases of the multi-player situation. If all the players' interests run parallel and they are free to cooperate and coordinate their actions, then of course they can, in principle, operate together as a single "big player," and their decision problem can, in principle, be treated by some mathematical optimization technique. Thus, the really characteristic issues of game theory emerge only in the presence of (at least partially) conflicting interests, often accompanied by impediments to whatever cooperation the conflict might otherwise permit. The desire to calculate "optimal" or "rational" behavior for a participant in such a scenario of conflictual inter-dependence with others must first confront the conceptual problems of defining "optimality" or "rationality" in this setting.

Although preceded by mathematical analyses due to Borel [15], Ville [16], and von Neumann [17], it was the celebrated 1944 treatise of von Neumann and Morgenstern [18] that brought this body of problems and major steps toward their resolution to the attention of both the relevant technical communities and a broad intellectual public. The frivolous connotation of the term "game" is distinctly deceptive, although "games" in the ordinary sense are indeed among the situations to which the theory applies. The title and preface of [18] reveal an intense motivating concern with application to economics and the behavioral sciences, disciplines whose enrichment by game theory is now well recognized (cf. for example Shubik [19],



Lucas, et al [20]). The obvious relevance of the field's subject-matter for military analyses is reflected, for example, in the long-term role of the RAND Corporation (then predominantly an Air Force "think tank") as a major center of game-theoretic research, much of it subsequently embodied in [21] and [22]. In short, the multi-player problems addressed by the theory include profoundly serious and practically important situations. Technical publications on game theory and its applications appear in professional journals of the many disciplines impacted by this versatile area (economics, political science, psychology, operations research, mathematics, and others (cf. [23] and [24])) with many notable papers collected in several volumes (Numbers 24, 28, 39, 40, 52) of Princeton University's Annals of Mathematics Studies series. Since 1972 the field has also enjoyed a dedicated journal, the International Journal of Game Theory.

To establish a common vocabulary for what follows, we shall next sketch some of the basic concepts of game theory, at least of those parts to be employed later. These concepts yield the elements of any formal game-theoretic "model"; for application, appropriate entities in the real-world situation must be associated with each of these formal elements and, the usefulness and validity of the model will of course depend heavily on the skill and care with which this match-up is performed. The concepts to be discussed here are:

- the players and their allowed relations
- the players' "strategy spaces"
- the players' "payoff functions"

In addition, we will need to specify what is meant by a "solution" of the model since this will correspond to the previous notion of "rational" or "optimal" behavior by the players. It will be convenient first to describe "strategies" and "payoffs" in the context of what might be called a "one-move" picture of the situation and only then to sketch a more dynamic picture (the extensive form) which takes into account the sequential aspect over time of the game's play.

## 2.1 The Players and their Relations

Obviously, the number of players (decisionmakers) is one of the data specifying any particular game-theoretic model. In a purely formal setting one might as well identify a set of "p" players with the integers  $\{1, 2, \dots, p\}$ . But for application, one would want to indicate the players' "identities" in a way giving at least a rough initial idea of their respective objectives and "degrees of freedom." If the number of players in an initial formation is so large as to be unwieldy, one might hope to alleviate

this difficulty by aggregating players with sufficiently similar interests (cf. Goldman and Shier [25] and Goldman [26]).

By the "relations" among the players we mean in particular their ability to communicate, collaborate, and coordinate their actions. Impediments to such collaboration might be physical, cultural (e.g., taboos), legal (where the effectiveness of the legal sanction is not at question), etc. The cooperative theory of games is very rich in phenomena to be considered (coalition formulation, bargaining and threats, division of spoils among coalition-partners, side payments, etc.); the solution concept originally advanced for it by von Neumann and Morgenstern [18] was shown nearly 25 years later (Lucas [27]) not to be universally applicable (i.e., not every game has such a "solution") and so the current literature presents a variety of possible "solution" concepts with differing advantages and disadvantages. We will evade the need for an exposition (necessarily somewhat lengthy and technical) of these interesting but complicated matters by quickly limiting the scope of the discussion:

- Only non-cooperative games will be considered

The reasonableness of this restriction, for our particular purposes, will be addressed in Section 5.

## 2.2 The Strategy Spaces

"Strategy space" is technical jargon for the set of courses of action ("strategies") among which a player can choose. In a purely formal setting, the numbers of a strategy space of size "m" can simply be referred to "by number," and so the space can be identified with the set  $\{1, 2, \dots, m\}$  of integers. For applications, of course, strategies must be described in terms that are meaningful for the situation being analyzed, which means in turn that this situation must be described in terms adequate for mathematical modeling; we shall have some criticism on this point to make of [5-7] later. A strategy space can be either finite (e.g., the possible mountain passes over which the troops might be marched) or infinite; in the latter case the "infinity" in question is usually the "continuum" type involved in selecting real-number values for one or more continuous variables from certain intervals (e.g., the adversary's target value for diversion-quantity during a particular time period) rather than the "discrete" variety illustrated by the choice of a positive integer from the (infinite) set of all such integers.

In many cases, the theory also requires consideration of probabilistic mixtures of the strategies described above, e.g., "choose action A with probability 0.4 and action B with probability 0.6." This should not be surprising; for example, strong poker players will not always behave the same way when holding the

same hand, and will attempt to randomize among alternative responses in order to avoid revealing a pattern (e.g., regularly alternating between two behaviors) that could be observed and exploited by an opponent. The term mixed strategy is applied to such a mixture, with the original or "pure" strategies viewable as a kind of degenerate special case. Note that choice of the particular mixed strategy mentioned above is determinate insofar as specific probability weights (0.4 and 0.6) are associated with specific actions (here, A and B), yet exactly which action will actually be undertaken remains indeterminate--hence unknowable to opponents until the final moment when the "wheel of fortune" (or whatever random device has been set to incorporate the 40-60 odds) is spun. If the strategy-space of the player in question also included a third pure strategy C, then the preceding mixed-strategy description would be formally rounded out by the redundant addition "and action C with probability 0."

For an infinite pure-strategy space, the formulation of mixed strategies requires a bit more in the way of mathematical statistics apparatus, specifically, the notion of "cumulative distribution function" and (when applicable) the accompanying notion of "probability density (or frequency) function." (The still more general notion of "probability measure" is not needed for our purposes here.)

### 2.3 The Payoff Functions

Suppose each player has chosen a particular course of action from the appropriate strategy space. With these decisions made, the situation under study will evolve in a definite way (subject to a proviso noted below), leading to a definite "outcome." A player will, in general, not be indifferent as to which of these potential outcomes actually occurs, preferring some over others, perhaps even regarding some as extremely satisfactory and some as disastrous. It is therefore assumed that each player can give a numerical score to any potential outcome, a higher score corresponding to greater desirability. ("Costs," or more generally "disutilities," might be represented by negative-valued scores.) Since the outcome depends on the strategy choices of all players, this score is (in the mathematical sense) a function of all these choices. For each player this function is vividly but crassly called the payoff function of that player; its specification in a game-theoretic model is a representation in numerical ("cardinal") terms of that player's objectives.

Symbolically, if  $p$  is the number of players, if  $S_k$  represents the  $k$ -th player's strategy space and  $s_k$  is a generic member of that space (where  $k=1,2,\dots,p$ ), then the payoff function for the  $k$ -th player can be written as  $f_k(s_1,s_2,\dots,s_p)$ . We emphasize again that the payoff to the  $k$ -th player is not under that player's sole control, but instead depends also on the

choices of the other players; formally,  $f_k$  is not a function of  $s_k$  alone, but in general has all of  $(s_1, s_2, \dots, s_p)$  as arguments. If  $p=2$  and the two players' strategy spaces are both finite, then their payoff functions can be conveniently written as payoff matrices, where the matrix for player  $k$  ( $k=1,2$ ) contains, at the intersection of its  $i$ -th row and  $j$ -th column, the score player  $k$  would assign to the outcome resulting from player 1's choice of "his"  $i$ -th course of action from  $s_1$  and player 2's choice of the  $j$ -th course of action (strategy) from  $s_2$ . Such games ( $p=2$ , finite strategy spaces) are therefore called bimatrix games.

It was assumed above that definite choices of strategies by the players would always lead to one and the same definite outcome of the game situation and therefore to definite payoff levels for the players. But this need not be true if the "playing out" of the chosen strategies involves some random elements; for example, dice rolls or the random measurement errors arising in the process of determining an ID value. Such randomness will lead (with various probabilities) to different outcomes, in general not all equally desirable to the players. Thus, each player's actual payoff viewed in advance becomes a "random variable" with a probability distribution of possible values. To obtain a well-defined payoff function it is necessary to find a single number which summarizes the overall desirability or "utility" to the player, of this probabilistic situation.

One natural choice for this summarizing number is the expected value (or "means" or "average value") of the random payoff, obtained by multiplying each possible payoff value by its probability of occurrence and then summing the results (e.g., equally-likely payoffs of 1 and 3 yield an expected value of 2). Thus, the  $k$ -th player's payoff function  $f_k(s_1, s_2, \dots, s_p)$  is now taken to be the expected value of the desirability score attributed by that player to the outcome of the partly random process that follows the respective players' choices of strategies  $s_1, s_2, \dots, s_p$ . (If the random elements involve a continuum of possibilities rather than a finite set, then these expected values are given by integrals rather than finite sums, but this technical complication is not a significant conceptual distinction.)

This selection of expected values as "summarizing numbers" turns out to simplify greatly the mathematics of the theory and rests largely on an influential analysis given in the treatise of von Neumann and Morgenstern [16, Appendix to Second Edition (1947)], which shows that expected values are the only correct summarizing numbers if players' utilities satisfy certain plausible axioms. Those axioms, however, embody the implicit assumption that the players are or should be (in current parlance) risk-neutral, e.g., indifferent between a sure-thing payoff of 2 and a 50-50 gamble between payoffs of sizes 1 and 3. That assumption is somewhat questionable for many applications, including the

ones of particular concern here. A major theme in modern decision science and "economics of uncertainty" research (cf. [28]) is the treatment of "risk-averse" decisionmakers with a well-known early precedent supplied by Markowitz's monograph on portfolio selection [29] in which a positive multiple of the variance of the (random-variable) payoff is subtracted from the expected value to obtain the "summarizing number."

The issue of "summarizing numbers" as payoff functions arises in much the same way when the use of mixed strategies is envisaged, with (deliberate) randomness now entering the actual strategy choices of the players and thereby being injected into the resultant outcome and thus into the payoffs to the players. Again, the use of expected values is traditional (and will be assumed later whenever nothing to the contrary is said), supplying the same advantages and raising the same questions. At a purely abstract level, this use is unnecessary; one can speak of the overall desirability (to a player) of some probabilistic distribution of possible game outcomes, note that this desirability is a function of the players' mixed-strategy choices giving rise to that distribution, assign a symbol to that (payoff) function and operate conceptually with it, all without specifying any concrete formula or calculation procedure for exactly how this overall desirability is built-up from the desirabilities and probabilities of the individual outcomes involved. But for any application purpose, some specific build-up formula (or algorithm) must in fact be specified--either the especially simple linear formula representing the expected value notion, or something else. It is of course much easier to criticize the expected value formulation than to propose and justify some particular alternative, and for many serious applications it may be appropriate to formulate and compare use of expected values with use of several alternatives (in terms of theoretical rationale, ease of use, and plausibility of results in "test cases") rather than making an initial a priori commitment to a single one.

There is a special class of games, defined by a particular property of their payoff functions, which will be important in the sequel and can naturally be introduced at this point. Suppose that for every set  $s=(s_1, s_2, \dots, s_p)$  of strategies by the players-- $s_1$  chosen by player 1,  $s_2$  by player 2, etc., the resultant payoffs to the players total 0, i.e.:

$$f_1(s) + f_2(s) + \dots + f_p(s) = 0. \quad (2.1)$$

The interpretation is obvious: whatever some players win (positive-valued payoffs) is at the expense of an equal total loss (negative-valued payoffs) by other players. A game with this property

(2.1) is called zero sum. In the case of just two players ( $p=2$ ), (2.1) implies that:

$$f_2(s) = -f_1(s) \quad (2.2)$$

so that the payoff function  $f_2$  of the player in a two-player zero-sum game need not be specified separately; it is fully determined by that of Player 1 ( $f_1$ ). Furthermore, if both players have strategy spaces that are finite so that we have a "bimatrix game" as defined a little earlier, then it follows that only player 1's payoff matrix need be given (its negative gives player 2's payoffs). It is for this restrictive but common class of games, with its total opposition of players' interests as expressed in equation 2.2, that both theory and the availability of computational solution methods are in an especially satisfactory state.

We close this subsection by noting a non-obvious assumption hidden in equation 2.1 and its special case (2.2); that the "scores" or "utilities" of the different players have somehow been expressed on a common scale which makes their addition meaningful. If only the payoffs to player 1, who hopes for a set  $s$  of strategy choices that will make  $f_1(s)$  "large", are considered, then the function  $f_1(s)$  could just as well have been  $3f_1(s)$  or  $f_1(s) = 10$ . But such changes could change (2.1) from true to false or vice versa, so that the zero sum condition requires some sort of prior consistent normalizing of player's payoff functions.

#### 2.4 Solution Concepts

With the notions of "strategy" and "payoff" now at hand, we return to the effort to define "rational" or "optimal" play. For simplicity, we begin with the two-player case ( $p=2$ ). Consider some pair  $(s_1^*, s_2^*)$  of strategies, the first for player 1 and second for player 2. Imagine that the game is to be played repeatedly (e.g., inventories are struck and "alarm-or-not?" decisions made over a series of time periods). When would it be reasonable for players 1 and 2 to retain strategies  $s_1^*$  and  $s_2^*$  as their respective choices throughout such a series of plays?

It is easy to describe a scenario in which this would not be reasonable. Namely, suppose player 1 has some strategy  $A_1$ , (necessarily) different from  $A_1^*$  for which

$$f_1(A_1, A_2^*) > f_1(A_1^*, A_2^*). \quad (2.3)$$

Furthermore, assume that in the course of the repeated plays, player 1 is able to infer that player 2 is using the particular strategy  $s_2^*$  or at any rate is using some strategy for which

(2.3) or its analog holds. Since under our "noncooperative" hypothesis (see the end of subsection 2.1) player 1 is unable to communicate with player 2 over possible changes in the latter's strategy, the natural working hypothesis for player 1 is that player 2's behavior will persist. Hence, in view of (2.3) it appears to player 1's advantage to change strategy from  $s_1^*$  to  $s_1$ ; and it would appear unreasonable for player 1 to persist in using  $s_1^*$  when such an advantage-promising change is available. Similarly, if player 2 had some strategy  $s_2$  necessarily different from  $s_2^*$  for which

$$f_2(s_1^*, s_2) > f_2(s_1^*, s_2^*)$$

then it would appear unreasonable for player 2 to persist in using  $s_2^*$ .

Reversing the negatives in the last paragraph, we can say that stable choice of the strategies  $(s_1^*, s_2^*)$  by the respective players 1 and 2, is reasonable only if

$$f_1(s_1, s_2^*) \leq f_1(s_1^*, s_2^*) \text{ for all } s_1 \text{ in } S_1 \quad (2.4a)$$

and

$$f_2(s_1^*, s_2) \leq f_2(s_1^*, s_2^*) \text{ for all } s_2 \text{ in } S_2. \quad (2.4b)$$

These conditions assert that neither player has any incentive to deviate unilaterally from his present strategy ( $s_1^*$  or  $s_2^*$ ), "unilaterally" meaning "without arrangement that the other player would also deviate." (Under our "noncooperative" hypothesis, the only deviations possible are such unilateral ones.) Conditions (2.4) define the pair  $(s_1^*, s_2^*)$  to be what is called a (Nash) Equilibrium Point for the game (Nash [30]). The extension to more than two players involves the same underlying ideas; thus a set  $(s_1^*, s_2^*, \dots, s_p^*)$  of strategies for the respective players of a p-player game is defined to be an equilibrium point<sup>4</sup> if, for each player  $i$ ,

$$f_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_p^*) \leq f_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_p^*)$$

for all strategies  $s_i$  in  $S_i$ .

This "equilibrium point" concept is essentially the only solution notion available for general noncooperative games. On the positive side, it is clear (I hope) that the definition of this concept embodies a "stability against deviations" requirement which really does appear to be an essential criterion for a "solution." Also on the positive side, it is known that all

---

<sup>4</sup>Some readers may wish to note the conceptual relation and mathematical similarity of the following conditions to those for a Pareto Optimum in mathematical economics.

games likely to be of applied interest do in fact have such a "solution." More precisely (Nash [30]), any game in which all player's spaces of (pure) strategies are finite does possess at least one equilibrium point, though in general it may involve mixed strategies. This result remains true for some but not all games with infinite strategy spaces (e.g., Owen [31]). Under a further hypothesis of "concavity" of the payoff function, it can be strengthened to assure an equilibrium point using only pure strategies (Rosen [32] Ponstein [33]). Infinite strategy spaces arising in applications can typically be approximated by finite strategy spaces (e.g., the real-number interval  $[0,1]$  replaced by the finite set  $0, 1/100, 2/100, \dots, 99/100, 1$ ), yielding an "approximating game" whose equilibria are "approximate equilibrium points" for the original game.

Use of "equilibrium point" as a solution concept involves some serious risks. First, a game may have more than one equilibrium point, indeed many of them, and they may involve different payoffs to the players. Without introducing a considerable body of problematical assumptions additional to the original game description (e.g., concerning the dynamic course of repeated plays, the "personalities" of the players, etc.), there is in general no natural way to single out one among these alternative "solutions."<sup>5</sup> Furthermore, suppose for example that  $p=2$  and the  $(s_1^*, s_2^*)$  and  $(s_1^{**}, s_2^{**})$  are two distinct equilibrium points. If we pair of  $s_1^*$  with  $s_2^{**}$  rather than  $s_2^*$ , in general we will not obtain an equilibrium point, and so we cannot speak of  $s_1^*$  simply as "a rational strategy for player 1" without restricting it to a combination with the specific choice  $s_2^*$ . These inabilities tend to assure (1) a meaningful specification for rational play by individual players (rather than by the collective of all players) and (2) a definite set of player payoffs arising from play "according to the solution," are substantial drawbacks for application of the theory, despite its helpfulness in focusing attention on the set of equilibrium points.

Note, however, that in any particular application these drawbacks might not arise. Suppose, for example, that our analysis is intended to advise player 1 on a suitable course of action. If that analysis shows the game-theoretic model to have only a single equilibrium point, then the appropriate advice (insofar as it is purely model-based) is relatively clear-cut. If there are multiple equilibrium points, but they all happen to involve

---

<sup>5</sup>One might prefer, however some equilibrium points to others by virtue of possessing additional "stability" properties formulated, for example, by Williams [34], Wu [35], Selten [36], Myerson [37], and Okada [38]. The most concentrated effort to resolve this ambiguity, though a "bargaining" context tangential to our motivating applications, is the work of Harsanyi integrated in [39].



one and the same strategy  $s_1^*$  for player 1, then at least some of the situation ambiguities are resolved. Similarly, if there are multiple equilibrium points but they all happen to involve the same payoff to player 1. Considerations like these can provide criteria for selecting among otherwise plausible game-theoretic models for an application; we will see later that some of these felicitous coincidences have actually arisen in the context of the present study.

No lucky chances are needed, however, if we restrict attention to the zero sum 2-player case, which we recall requires the relation  $f_2 = (-f_1)$  between the player's payoff functions. Here (2.4a,b) can be written

$$f_1(s_1, s_2^*) \leq f_1(s_1^*, s_2^*) \leq f_1(s_1^*, s_2) \text{ for all } s_1 \text{ in } S_1, \\ s_2 \text{ in } S_2 \quad (2.5)$$

From this relation, a number of desirable properties can be deduced. Consider any two equilibrium points  $(A_1^*, A_2^*)$  and  $(A_1^{**}, A_2^{**})$ . First,  $(A_1^*, A_2^{**})$  and  $(A_1^{**}, A_2^*)$  are also equilibrium points; this interchangeability permits us to speak of  $s_1^*$ ,  $(s_2^*)$  as being an "equilibrium strategy" for player 1 (player 2) without specifying a pairing with a particular strategy of the other player. Second, the two equilibrium points yield player 1 the same payoff (since  $f_2 = (-f_1)$ , the same is true for player 2); this common payoff to player 1 from all equilibrium points is called the value of the game. Third, this game-value is the largest payoff that player 1 can assure himself of (through his choice of strategy) despite player 2's efforts; from this conservative viewpoint, player 1's equilibrium strategies, which do in fact assure this largest payoff, merit (and are given) the term optimal. The analog of this last statement for player 2 also holds. (Aumann [40] has extended these results to a class of games wider than zero sum, but apparently difficult to recognize usefully early in their analysis.)

Thus, the theory of zero sum two-player games is very satisfactory. There is a convincing notion of "optimal strategy" for each player, and a "solution" of such a game is given by such a pair of optimal strategies, the corresponding payoff to player 1 (with its negative the payoff of player 2) then giving the (unique) value of the game. (For a more complete analysis, it might be desirable to determine all the optimal strategies for one or both players.) The existence of such a solution when strategy spaces are finite follows as a special case of the more general theorem cited earlier (which did not involve the zero sum assumption). A number of more advanced "minimax" theorems, assuring existence of a solution for many classes of situations involving infinite strategy spaces have appeared in the technical literature (cf. Chapter 5 of Parthasarathy and Raghavan [41]).

As will be seen later, the practicality of actually calculating a solution for such games is also rather high. Thus, there are powerful incentives for modeling a given situation as a zero sum two-person game rather than a more general game, if such a model gives an acceptable representation of reality.

## 2.5 The Extensive Form

The descriptions and definitions given so far are not designed to reflect the fact that a game-theoretic situation may well unfold over an extended period of time, proceeding in stages, and calling on players to make not just one but perhaps a sequence of decisions based on varying levels of information about the prior decisions of the other players. These aspects of the situation are highlighted in the so-called extensive form, a stylized representation of how the "rules" of a game structure the evolution of its activities.

In the extensive-form model, the game is represented as a tree-like network branching out progressively from a "root" node. Any one "play" of the game in effect traces a unique path in this network from the root to an outmost or "leaf" node; the payoffs to the players depend on which leaf-node is the terminus of the path that actually occurs.

Each non-leaf node "belongs" to a particular player whose turn to "move" it will be when and if the growing path reaches that node. The player's move consists of specifying one of the alternative actions available at the node, i.e., which of the edges branching from the node will be "pursued" by the player and added to the growing path. However, the choosing player may not know at exactly which of "his" nodes the play now stands, since this might require unavailable knowledge of what branchings ("moves") had been previously chosen by other players. Instead, the rules partition the nodes "belonging" to each player into certain "information sets." A player about to move will know in which information set the play stands (i.e., in which the tip of the growing path lies), but knows nothing further about which one of the nodes in that information set actually marks the current state of the play. The extensive form has certain axioms enforcing this role in the information sets; for example every node in the same information set must have the same number of edges branching from it, since otherwise the player about to move--and necessarily aware of the number of choices available--could use that awareness to narrow his possible node-locations to some proper subset of the information set involved. The influence of chance events (e.g., random measurement errors or the results of a dice-roll or card-shuffle) is represented by letting certain nodes belong to the "chance player" an automaton who must select from among the available alternatives (edges

branching from the node) in accordance with a prescribed probability distribution associated with that node.

The rich formalism of the extensive form can be collapsed to the "strategic form" described earlier by the following device: define a strategy for a player to be a "complete contingency plan" which specifies in advance for each of the player's information sets (say, I) which numbered branch the player will choose if the course of play should call for a move from some node in I. Note that by the axiom described in the last paragraph, all nodes in I have the same number (say 4) of edges branching from them, hence all can have their branching edges labeled by (say) the first four positive integers so a strategy clause "if in I, choose branch 2" makes sense.

A choice by each player of a "strategy" as just defined will--apart from the effects of moves by the "chance player" if present--collectively determine a definite course of play and a definite play through the tree, hence definite payoff values to the players. The chance player's moves may lead to a probability distribution of payoff values rather than a definite value and this distribution needs to be encapsulated by a single "summarizing number" as explained earlier. With these understandings, any extensive form game is converted to a game of the ("strategy") form discussed in the preceding sections.

A "mixed strategy" is, as before, a probabilistic mixture of the ("pure") strategies defined in the next-to-last paragraphs. The first of the two main results in the theory of extensive form games says that mixed strategies are unnecessary whenever the game is of "perfect information"--i.e., if the player whose turn it is to move knows all prior moves by the other players (including the chance player). The formal version of this condition, by the way, is that each of the players' information sets consists of just a single node. For such a game, assuming its tree network is finite, the result assures the existence of at least one equilibrium point involving pure strategies only (Chapter 15 of [18]; also Kuhn [42].)

For games of even moderate complexity, a strategy as a complete contingency plan is likely to involve an unpleasantly long list of instructions. For any one play of the game, most of these instructions are likely to refer to contingencies not actually encountered during that play (i.e., to information sets not actually entered by the growing path describing the progress or play). Mixed strategies, as probabilistic mixtures of these cumbersome pure strategies, will be even more unwieldy for any use beyond conceptual analysis. It would be much more convenient if probabilistic behavior by a player could be described by what are called behavior strategies. These specify for each of the player's information sets the probability distribution of his choices among the alternatives (branching edges) available

when moving from a node in that set. Unfortunately, it is not true that the effect of any mixed strategy can be duplicated by some behavior strategy and so restricting the analysis of a game to the more tractable behavior strategies might improperly reduce the set of options open to the players. The second of the two main results in extensive form game theory (Kuhn [44]) identifies an important class of games for which this difficulty does not arise and which can therefore safely be analyzed using behavior strategies. These are the games of perfect recall in which a player about to move knows all of his prior moves in that play of the game. This assumption, a natural one for many applications, rules out (for example) situations in which a "player" consists of a team with imperfect communication, as in bridge.

## 2.6 Stochastic and Repeated Games

Although the extensive-form structure is in principle capable of accomodating most of the multi-time-period (briefly, "multi-stage") models we will encounter later, that accomodation can be rather awkward for models with certain features. For example, in some situations a player will receive a "spot" payoff at each stage (e.g., an increment of diverted material), whereas the extensive form associates payoffs only with the final stage of play--i.e., the last node of the path traced through the game-tree by the execution of the chosen strategies. If there are more than a very few stages, the translation of intermediate payoffs to the ends of the (numerous) paths can be laborious, and perhaps destructive of insight about the game.

Another "awkward" class of multi-stage situations involves repeated play of a game in which a player may not initially be fully informed about the capabilities and values (i.e., payoff functions) of the other players. Thus, at each stage, a player's motivation for immediate payoff is confounded with a motivation to act in ways that elicit more information about the opponents, providing a better basis for decisions in later stages. (Reports of boxing matches often note a process of "feeling out the opponent" in the early rounds.) Although the "information sets" of the extensive-form concept can presumably be set up in a way that captures this feature of incomplete information,<sup>6</sup> doing so might prove unnatural and unrewarding.

This explains why certain classes of multi-stage games have been subjected to concentrated study in their own right. Two of these--stochastic games and repeated games with incomplete

---

<sup>6</sup>"Incomplete" should not be confused with "imperfect", the denial of "perfect" as defined in the next-to-last paragraph of Section 2.5.

information--appear in the literature relevant to this report, and in particular are noted in the Peer Review Group report [14; pp. 30-33 and 42] as possibilities for multi-period ID analysis. They are labelled "Advanced Techniques" [14; p.30], a term which accurately reflects both the somewhat greater complexity of the mathematics involved in their analysis, and their present status as active research topics in contrast with many of the "by now classified" subjects described earlier in Section 2.5. Their greater recency also implies less experience with their use as modeling tools, i.e. with learning what features of potential applications lend themselves nicely (or badly) to modelling by such games.

To flesh out the comments in [14], we shall provide brief sketches of the main concepts of these two classes of games, along with a sample of pertinent references. For simplicity, we will confine attention to the two-player case with only finitely many strategies at every stage of play.

Stochastic games, though mostly neglected until the 1970's, are generally attributed to a 1953 paper by Shapley [45]. Also called Markov games, they are multi-stage processes which, at each stage, are in one of a finite set of states. With each state is associated both a particular bimatrix game (in the zero-sum case, a matrix game) and a set of transition probabilities. At each stage the two players, knowing the current state and its associated bimatrix game and transition probabilities, choose their strategies in the current bimatrix game. This choice determines both their immediate payoffs (from the payoff matrices) and, from the transition probabilities, the probability distribution governing the identity of the next state.

The sequence of stage-by-stage payoffs for a player can be accumulated in either of two ways: as the limit over long times of the average payoff per stage (the undiscounted case), or as the sum of the discounted stream of payoffs (the discounted case, where the solution in general depends on the discount factor). The desired solution concept has been limited to the rather natural class of stationary strategies, those in which in a player's behavior at any stage depends only on the current state and not on prior history.

For the discounted case, the existence of a solution was proved by Shapley [45]. However, no finite exact solution algorithm can exist, since (an analogous argument with more detail appears in Section 3.4) it is possible for a problem with rational-number data to have a solution involving irrational numbers (Parthasarathy and Raghavan [46]). If, however, the transition probabilities depend only on the strategy choices of one of the two players (the "controller"), then such an algorithm becomes possible [46]. This remains true (Filar [47] and an analog of the material in Vrieze et al [48]) if the "single controller" hypothesis

is weakened to that of "switching control," in which the identity of the controller may be Player 1 in some states, Player 2 in others.

For the undiscounted case, progress is clouded by the knowledge (Blackwell and Ferguson [49]) that exactly-optimal strategies may fail to exist, although approximate ones will (Mertens and Neyman [50]). When they do exist, they can be in principle be calculated by solving an optimization problem formulated by Vrieze ([51], building on the study of the discounted case by Rothblum [52]; this problem has a linear objective function, but quadratic constraints as well as some linear ones. An alternative "successive approximations" approach is given by Federgruen [53]. Finite solution algorithms for the "single controller" and "switching control" cases are given in [46] and [48] respectively; other cases in which solutions are known to exist are given by Gillette [55] and Hoffman and Karp [56].

Applications of stochastic-game theory have been proposed in the areas of military tactics and weapons development (Charnes and Schroeder [51], Winston [58]), advertising (Albright and Winston [59]), natural-resource management (Sobel [61]) and oligopoly analysis (Kirman and Sobel [62]), as well as inspection (Filar [63]). We may note the remark of Sobel [64; p. 995] that "the tendency to model phenomena as stochastic games has been curbed by inadequate computational procedures"; however specific applications or subcategories, like many of those cited above, prove exceptions to this general observation; cf. also Filar and Schultz [65].

If a "critical diverted quantity" concept is applicable for the multi-stage ID-analysis problem, then modelling that problem as a stochastic game encounters certain difficulties. An adversary's situation at a certain stage should presumably depend, in part, on how much has already been diverted. But since in a stochastic game (with a stationary-strategy solution) behavior is supposed to depend only on the current state, it follows that the "diverted so far" quantity should be part of the specification of a state. And since in such a game both players are aware of the current state, this would require the adversary's opponent to know the cumulative diversion, which is unrealistic. Perhaps some modelling artifice can evade this difficulty, but it offers at least an initial obstacle.

A currently active line of research in a cognate field is suggestive here. A Markov Decision Process (MDP) is essentially the one-player analog of a stochastic game; in the well-developed "classical" theory of such problems (given major impetus by Howard [66]; for modernity see, e.g. Ross [67]), the single decision-maker is assumed at every stage to know the current state. Recently, however, there has been increasing interest in situations in which the current state is only "partially

observable" by the decision-maker: the available information permits him only to place (say, via Bayes Theorem from probability theory) a probability distribution over the alternative possibilities for that state. (An application, and a useful review of prior literature, are given in the 1980 paper of Monahan [68]; basic theory appears in Majumdar [69].) This development might well--and probably will--be extended to the case of two (or more) players, i.e. to "stochastic games with partially observable states."

We turn now to repeated games of incomplete information. They are generally attributed to Harsanyi [70] and to Aumann and Maschler [71-73]. The former is "notationally heavy"; an expository sketch is given by Shubik [19; Section 9.3]. As noted by Kohlberg [74], papers [72, 73] "were never published and are hard to obtain," but their contents are described in the secondary sources consulted.

The basic scenario involves a set of bimatrix games (matrix games, in the zero sum case), only one which will actually be played. This set is split into subsets in two different ways, one relevant for player 1 and the other for player 2; these split-ups are both known to both players, as is an initial probability distribution over the set of games. Initially the "chance player" uses this distribution to determine which of the games in the set will actually be played; player 1 is not told which game this is, but only in which of the subsets of "his" split-up the chosen game lies,<sup>7</sup> and similarly for player 2. (The sizes of the various games' payoff matrices must be such as not to "give away" any information about the identity of the "real game.") The players then repeatedly play the game (not knowing exactly which one it is), their payoffs accumulating but remaining unknown to them until the whole process is done.

There is one other important but rather complicated apparatus in the scenario. Associated with each of the possible games, in addition to the players' payoff matrices, is a second pair of matrices--the information matrices for each player, assumed known by both of them. The  $(i, j)$  entry of player's information matrix contains an information signal which is revealed to player 1 if and when the players choose respective pure strategies at some stage; this signal might for example communicate to player 1 some partial or precise knowledge of player 2's previous strategy choice  $(j)$ , or part of player 2's partial knowledge about which game is being played, or both--this formalism is very general and flexible. Player 2's information matrix is similar. Thus in selecting a strategy at each stage of play, a player must be concerned not only with the resultant payoff,

---

<sup>7</sup>If player 1's split-up is fine enough, this could in fact determine the chosen game.

but also with the long-term effects of the resultant release of information to the other player. (In the non-zero sum case, these effects need not all be negative.)

Most of the literature on these models has been confined to the zero sum case; one exception is Sorin [75]. It has been mainly concerned with establishing the existence of solutions, and the rate of convergence of the solutions finite-stage truncations to a solution (in particular the value) of the infinite-stage process; cf., for example, Kohlberg [74, 76], Mertens [77], Zamir [78-80], Mertens and Zamir [81-84], Ponssard and Zamir [85], Ponssard and Sorin [86, 87], Sorin [75, 88-9], Waternaux [90]. Anticipating the type of concern to be emphasized in Section 3, we note that none of these papers explicitly addresses computational issues, and the examples considered are extremely small (2x2 payoff matrices, at most two possible games). It appears to the writer that the computational aspects of these models are rather unexplored, perhaps because they have appeared intimidating; the only note which is encouraging (for reasons detailed in Section 3.1) is the references in [86, 87] to linear programming formulations.

From the viewpoint of model appropriateness for our intended area of application, three aspects of these structures may create difficulties. One is the notion that stage-by-stage payoffs are kept concealed until the end of the multi-stage process; however, total or partial information about such payoffs could perhaps be transmitted via the information matrices, and the results of Megiddo [91] suggest that this restriction may not be critical. Second, much of the analysis is for the infinite-stage case<sup>8</sup>, so that some care would be needed in establishing proper relations with the finite truncations arising in application. Third--and apparently most serious--is the requirement that the same game be played at every stage (though possibly with changing levels of information available to the players). This can make it hard to differentiate a stage at which an adversary is just starting his pilferage, from one in which just one more "good haul" could bring his total to a critical level. Again some ingenious modelling trick might evade this problem, but it appears that a naturally applicable theoretical framework would involve a mating of stochastic games (in which the "state" or "current game" can vary from stage to stage), with repeated games of incomplete information (in which the players have only partial information at each stage about the current situation). The reference at the end of Sorin's 1984 paper [89] to an emergent confirms the writer's impression that the time is ripe for such a research development, but that it has not yet occurred (cf.,

---

<sup>8</sup>In principle this is also true of some stochastic games, but the latter can have "absorbing states" which in effect would assure finite termination.



however, [92])). Its genesis might well be stimulated by the specific application context of safeguards problems, just as the development of repeated games of incomplete information was spurred [71-73] by the anticipated methodological needs of nuclear test-ban treaty inspection. This suggestion goes beyond the recommendation of [14; p. 42].

### 3.0 SOLVABILITY OF GAME-THEORETIC MODELS

In Section 2, we described the class of situations to which the theory of games is directed, and then proceeded to discuss some basic aspects of game-theoretic models: the elements of such a model and what might be meant by a "solution." Some theoretical results assuring the existence of a solution were cited. In keeping with our motivating application in the safeguards program, attention was confined to non-cooperative games.

Here we turn to a more practical question: how (and with what difficulty) can a game-theoretic model be solved? Thus, our main concern is with computational processes for actually determining a solution to a model given in numerical form.

Two other auxiliary subjects will also concern us, but only briefly. (There is not too much to be said in general terms about either of them.) One stems from the fact that the data of a game-theoretic model (or any other serious decision-aiding mathematical model) are unlikely to be perfectly accurate or reliable. We will therefore want to consider the possibilities for sensitivity or parametric analysis, i.e., for examining how the solution varies with changes in the problem data ("parameters") and to which of these parameteres the results are especially sensitive. The hope, of course, is that such questions can be explored by means more economical and insight-giving than simply perturbing the data in various ways and solving each of the revised problems "from scratch." The second "auxiliary topic" is the non-numerical version of our main concern: logical-symbolic processes for determining or facilitating a closed-form solution of a game-theoretic model in terms of its parameters (the latter appearing as "literals" rather than with prescribed numerical values).

#### 3.1 Matrix Games

These are the two-player zero sum games with only finitely many (pure) strategies for each player, whose theory was described in Section 2.4 as being in particularly satisfactory form. We proceed to show why the same is true for the computational treatment of such games.

Suppose the payoff matrix has entries  $a(i, j)$ ; that is, if players 1 and 2 choose their  $i$ -th and  $j$ -th pure strategies, respectively, then player 2 pays player 1 the amount  $a(i, j)$ . Suppose next that player 1 selected the mixed strategy "x" which chooses his first course of action with probability  $x_1$ , his

second course of action with probability  $x_2, \dots$ , and his last (m-th) course of action with probability  $x_m$ . It follows that

$$\begin{aligned} x_1 + x_2 + \dots + x_m &= 1; \\ x_1 \geq 0, x_2 \geq 0, \dots, x_m &\geq 0. \end{aligned} \quad (3.1)$$

Against the j-th pure strategy of player 2, the mixed strategy  $x$  will yield player 1 an expected payoff

$$a(1,j)x_1 + a(2,j)x_2 + \dots + a(m,j)x_m.$$

Thus, the greatest expected payoff player 1 can be assured of while using  $x$ , no matter which pure strategy (j) player 2 chooses, is the largest number  $v$  for which

$$a(1,j)x_1 + a(2,j)x_2 + \dots + a(m,j)x_m \geq v \quad (3.2)$$

It is not hard to show that the very same number  $v$  remains achievable by player 1 (using  $x$ ) even if player 2 is permitted to use mixed strategies and is therefore the best player 1 (using  $x$ ) can do despite any effort by player 2. According to the theoretical development sketched in Section 2.4 (and assuming expected values of payoffs are appropriate objectives), player 1 should choose  $x$  so as to maximize  $v$ , i.e., should choose the decision variables  $x_1, \dots, x_m$  so as to maximize

$$v = 1v + 0x_1 + 0x_2 + \dots + 0x_m \quad (3.3)$$

subject to (3.1) and (3.2).

This last optimization is a special case of the following more general problems: choose values for a finite set of decision variables, so as to maximize some linear function of those variables subject to a finite set of constraints each of which--like (3.1) and (3.2)--is linear (equation or inequality). Such a problem is called a linear program; the situation from player 2's viewpoint also yields a linear program. Fortunately, linear programming is an extremely well-developed field of mathematical optimization; massive computer programs based on the "simplex method" of G.B. Dantzig stand ready to solve rather enormous linear programs with great rapidity. (This rapidity, known for decades as an empirical fact--e.g., McCall [93]--has more recently received intensive theoretical investigation and verification; as in the prize-winning research of Borgwardt [94, 95].)

Thus, finding a numerical solution of a matrix game of any reasonable size (i.e., the value, and an optimal value for each of the two players) can be regarded as a "well-solved problem." Finding all optimal strategies for one or both players is distinctly more laborious, but a systematic finite procedure for doing so is known (Shapley and Snow [92]). To be precise, if a player

has more than one optimal strategy then these strategies can be represented geometrically as a convex polyhedron, e.g., in two dimensions, a polygon without "holes" or "dents," and the cited method finds the "extreme points" (or "corners" or "vertices") of this polyhedron one by one, thereby implicitly determining all the points in it.

We turn now to sensitivity and parametric analysis. Here there are two distinct bodies of work to be cited. The first of these, due to Mills [99], deals specifically with matrix games. Suppose we have such a game, in which player 1 has  $m$  pure strategies and player 2 has  $n$ , so that the game is described by an  $m \times n$  payoff matrix  $A$ . Let  $D$  be another  $m \times n$  matrix; then the matrix  $A + tD$ , for small real numbers  $t > 0$ , can be regarded as the payoff matrix of a game obtained by "perturbing" the original game  $A$  in the "direction"  $D$ . What [99] gives is a recipe for finding the "directional derivative" (i.e., rate of change) of the value of  $A$  with respect to such perturbations. Specifically, this recipe is the value of the game which has payoff matrix  $D$ , but has the players' mixed strategies restricted to their optimal strategies in the original game  $A$ . The solution of such a "restricted matrix game," like that of an ordinary matrix game, can be carried out by translation to appropriate linear programs; the main computational labor lies in determining the coefficients of the associated "constraints," a task which involves the finite calculation procedure mentioned at the end of the last paragraph.

The second line of work on sensitivity (and parametric) analysis for matrix games uses the fact that such games be translated into linear programs as in (3.1) - (3.3) above. Like so many other aspects of linear programming, its sensitivity analysis techniques have received substantial and successful attention, and such techniques typically accompany the "massive computer programs" mentioned earlier. (Skipping the intervening years we cite only the initiating work of Gass and Saaty [100-102] and a fairly recent comprehensive monograph by Gal [103].) Such techniques can indeed be used to analyze the sensitivity of a matrix game's solution (especially its value) to systematic change in a single entry of the payoff matrix, or even to broader patterns of changes, mostly conveniently introduced "one row at a time" or "one column at a time." Though practicable, they are not quite as computationally efficient as might be desired, essentially because their forte is dealing with changes either in the coefficients of the "maximand" (like (3.3)) of a linear program--cf. the titles of [100-102]--or like (3.1) or (3.2). In matrix game analysis, however, they are called upon for the harder job of treating changes in the coefficients of (3.2), i.e., the payoff entries.

Most of the preceding remarks have stressed the use of linear programming's "simplex method" to solve matrix games.

There is a second solution approach that merits attention, that of fictitious play. It is an iterative procedure, beginning with an arbitrary choice of pure strategies by each of the two players. At any later stage, player 1 reviews the relative frequencies with which player 2's various pure strategies have been chosen in the past, regards these frequencies as the probabilities of a mixed strategy for player 2, and selects a pure strategy which best responds to (i.e., achieves largest expected payoff against) that mixed strategy; player 2 behaves analogously. This process, formulated by Brown [104] as a discrete analog to a differential equations-based solution method formulation by Brown and von Neumann [105], generates for each player an infinite sequence of mixed strategies and expected payoffs. Those sequences have been shown (Robinson [106]) to converge to the game-value and (in an appropriate sense) to the players' sets of optimal strategies. Convergence is reputed to be generally slow (although at least one practitioner solving large military-game models reported good results [107]), parametric analysis is not possible, and since the iterative process must be terminated at some finite stage with only an approximate solution, this method is not usually competitive. But its ease of computer-coding and its strong intuitive basis may at times be compensating advantages.

As might be expected, a particularly elementary calculation method is available for those special cases in which it is known or suspected that the game has a solution using only pure strategies. One determines the positions of the smallest entries on each row of the payoff matrix and then tries to find among these entries one which is also a largest entry in its column. If an entry with this "saddlepoint" property exists, then its value gives the value of the game, and the associated row and column corresponds to optimal pure strategies for the two players. If no such entry exists, then mixed strategies must be involved to obtain a solution.

### 3.2 Other Two-Person Zero Sum Games

These will be games in which one or both players have infinite spaces of pure strategies. The class most extensively studied is that in which the strategy space of each of the two players is a continuous, real-number interval which, by rescaling, can be taken as  $[0,1]$ . Then the possible strategy-pairs  $(s_1, s_2)$  fill out the "unit square" in the  $(s_1, s_2)$ -plane, whence the term game on the square.

Applications exist in which the payoff function  $f_1(s_1, s_2)$  to player 1--whose negative gives the payoff to player 2--is not continuous throughout the square. This occurs, for example, in "games of timing," e.g., models of duels in which payoffs depend on "who fires first" in a way introducing discontinuities

at those strategy-pairs representing simultaneous firing by the opponents. But for applications of the type underlying this report, it seems reasonable to expect continuous payoff functions. Under that assumption, a game on the square is guaranteed to have a solution, though mixed strategies will usually be involved. Also, an approximate solution of any specified quality can be obtained by approximating the square with a sufficiently fine, finite two-dimensional grid of its points, regarding the values of  $F_1$  at the grid-points as entries of a payoff matrix, and then solving the resultant matrix game. While this is generally the numerical method of choice, the method of fictitious play can also validly be applied (indeed, it can be applied for a much wider class--"compact metric squares"--of infinite strategy spaces; Danskin [108]).

It is for special classes of continuous games on the square that there was notable progress in obtaining "better-than-numerical" solution methods, i.e., methods which approach the ideal of yielding closed-form solutions in terms of the parameters in the payoff functions. (I use "was" rather than "has been" because the field has not been fashionable for some time; perhaps indicating that further advances would be distinctly more difficult.) This progress is embodied in theoretical analyses which both provide detailed information on the mathematical form of a solution, and provide and justify procedures for determining the specifics of that form in any particular instance. The latter procedures are, in general, numerical (though involving far less computation than a purely numerical approach, unaided by information about the solution's form). In especially "nice" cases, however, they can be carried out in symbolic or "closed" fashion to yield an entirely closed-form solution.

An extensive account of most of this work is collected in Volume 2 of Karlin [21]. As a sampler, two of its subdomains are sketched in the next three paragraphs.

One of these subdomains involves continuous games on the square for which the payoff function,  $f_1(a_1, a_2)$ , is "separable," i.e., built up from one-variable functions of the individual strategies  $s_1$  and  $s_2$  in the manner given by the formula

$$f_1(s_1, s_2) = \sum_{i=0}^M \sum_{j=0}^N a_{ij} g_i(s_1) h_j(s_2), \quad (3.4)$$

with the functions  $g_i$  and  $h_j$  assumed continuous. Then the available theory assures us, for example, that each player has at least one optimal strategy which "mixes" only a finite number of pure strategies (thus cumulative distribution functions are not needed), in fact, a number not exceeding the smaller of  $M + 1$  and  $N + 1$ .

1 where M and N are as in (3.4) and in further fact (a sharper limit) not exceeding the "rank" of the matrix  $(a_{ij})$  of coefficients appearing in (3.4). Considerably more specific results can be given for the polynomial case of (3.4), i.e.,

$$f_1(s_1, s_2) = \sum_{i=0}^M \sum_{j=0}^N a_{ij} t_i^i(s_1) t_j^j(s_2). \quad (3.5)$$

To introduce the second subdomain, we define a function to be convex if it is never underestimated by linear interpolation between two of its values, concave if it is never overestimated by linear interpolation. Consider a continuous game on the square in which  $f_1(s_1, s_2)$  for each  $s_1$ , is a convex function of  $s_2$ . Then the theory tells us that player 2 has a pure optimal strategy, while player 1 has at least one optimal strategy that mixes at most two pure strategies. If also  $f_1(s_1, s_2)$  is for each  $s_2$  a concave function of  $s_1$ , then player 1 also has a pure optimal strategy. In addition, the theory provides information that aids in actually determining the various pure strategies just mentioned, as well as their "mixing weights" where appropriate.

The reader may recall from calculus that convexity and concavity are characterized by the signs of the second derivative of the function involved (" $\geq 0$ " for convexity, " $\leq 0$ " for concavity). This suggests that the preceding result might be generalized by stipulating the sign of some derivative of higher order than the second. That turns out to be the case; the theory assures us that if

$$\partial^N f_1 / \partial^N s_2 \geq 0 \quad \text{for all } (s_1, s_2) \quad (3.6)$$

where the order N of the partial derivative obeys  $N > 2$ , then player 2 has an optimal strategy which mixes at most  $N/2$  pure strategies (here the endpoints 0 and 1 of the pure-strategy space count only 1/2 if they occur, while player 1 has an optimal strategy which mixes at most N pure strategies.

We conclude this section by noting, by way of illustration, a relevant class of games other than the continuous games on the square. These are the "S-games" introduced by Blackwell and Girschick [109], in which player 1 chooses an integer  $i$  from the finite set  $\{1, 2, \dots, n\}$  while player 2 chooses a point P in a given subset S of n-dimensional space; the  $i$ -th coordinate of P is then the payoff to player 1 from player 2. Perhaps surprisingly, this artificial-looking situation can be used to model some interesting applications (cf. [109]). Recently, Filar and Raghaven [110] have given an iterative solution method for such games, under natural hypotheses ("closed and bounded") on the set S. If the method terminates, it does so

with a solution to the game; if not, then continuing it long enough yields approximate estimates of the game-value and approximately-optimal strategies to any desired degree of approximation. The method alternates between solving matrix games (which grow in size at each iteration), and solving certain special nonlinear optimization problems (whose difficulty depends largely on how complicated the set  $S$  is).

### 3.3 Bimatrix Games

Recall that these are the games in which each of the two players has only finitely many pure strategies, but the zero sum assumption is dropped so that a separate payoff matrix for each player is required.

An initial comment on solution methods for such games is that fictitious play does not in general work; this was shown by Shapley [112; Section 5]. Nor does it appear that the powerful computational methods of linear programming can be brought directly to bear, as they can for matrix games.

In view of the last remark, it was regarded as somewhat of a triumph when finite schemes<sup>9</sup> for finding a solution (i.e., an equilibrium point, generally involving mixed strategies) were found. One such method, described by its authors as lying "within the usual format of linear programming computations," was given by Lembe and Howson [113]; its discovery had a major role in originating an area of mathematical operations research ("linear complementarity theory") which is now prospering on its own, without particular reference to game theory. The results of [113] also show that (apart from "degenerate cases"--which might, however, be more likely to arise in the structured non-random data of an application) the number of solutions is finite (in fact, odd). Improvements to the method and an adaptation to games in extensive form (Section 2.5) were proposed by Wilson [114].

An alternative approach rests on the observation (Mills [115]) that finding a solution of a bimatrix game can be translated into a quadratic programming problem: maximizing a quadratic function subject to linear constraints on the variables. Mangasarian and Stone [116], for example, have exploited this observation by applying one of the (several) available numerical methods for quadratic programming; some gaps in the associated theoretical justification--still unfilled today, so far as I know--were not reflected in any observed difficulties with the computational experiments reported [116; p. 352, para. 3]. One favorable

---

<sup>9</sup>By "finite" we mean that the method obtains an exact solution in finitely many steps.



feature of this approach is that it might permit application of known methods for sensitivity and parametric analysis of quadratic and more general nonlinear optimization problems. That body of techniques, though by no means as powerful and complete as for linear programming, has seen considerable development (cf. Mine et al [117], Fiacco and Hutzler [118], and the extensive recent monograph by Bank et al [119] with coverage including linear complementarity theory).

As noted earlier (Section 2.4), fixation on a single one of a bimatrix game's (possibly) many solutions is, in general, unacceptably arbitrary, so that it may well be necessary to determine all such equilibrium points. There was some skepticism and confusion about whether and how the method of [113] could be extended to a finite technique for "finding all solutions" in an appropriate explicit or implicit sense; cf. Aggarwal [120], Todd [121], and especially p. 183 of Shapley [122]. Whatever the resolution of this question, a somewhat different line of development which does yield a finite solution method for the "all equilibrium points problem" has been given in successive papers by Vorobiev [123], Kuhn [124], and Mangasarian [125] with recent improvements by Winkels [127]. All known approaches involve or are akin to finding all vertices of a polyhedron (described by linear inequalities and equations) in a high-dimensional--a finite but potentially formidable computational task (Dyer [128]). Thus the practicality of these methods for games of application-interesting complexity must be regarded a priori with some suspicion, although ingenious exploitation of the special mathematical features of a particular model is always a possibility. It appears to the writer that these methods might be extended further to yield (at high computational cost) some minimal degree of parametric analysis capability for the "all equilibrium points" problem; however, such extensions have apparently not been pursued,<sup>10</sup> perhaps because their cost-benefit prognoses are so discouraging.

It is perhaps worth noting explicitly a main difficulty in parametric analysis of a model which admits multiple "solutions." Suppose one is for the moment concentrating on a particular solution of some initial "base-line" version of the model, and now asks how that solution varies as some parameter appearing in the model changes from its base-line value. As this change progresses, the solutions might split into two or more distinct solutions ("bifurcation" or "multifurcation"); alternately, some distinct solutions might coalesce into a single one. This possibility obviously presents both conceptual and computational hazards to any simplistic notion of "solution-tracking." If

---

<sup>10</sup>We remind the reader again (cf. ftn 3) that completeness of literature review could not be attempted within the scale of this effort.

one can somehow begin with a solution which is stable in the sense of Jansen [129] or some of the authors cited in Section 2.4, then such phenomena would be ruled out for "sufficiently small" excursions of the parameter in question. For the application content motivating this report, however, the need to analyze a substantial range of model versions would probably make a "small perturbation" limitation unacceptable.

The considerations of the last paragraph emphasize once again the strong desirability, for applications, of a model which yields a unique solution, or at least whose solutions are "equivalent" in the sense of yielding equal payoffs to the player in whose interests the analysis is being conducted. Bimatrix games with unique solutions have been studied, for example, by Millham [130] and Heurer [131], but from the viewpoint--absolutely perverse for our purposes--of taking as given the strategy-pair which is to be the "unique equilibrium point of the game," and then constructing a game for which this is true. A more useful result, most naturally posed for two-player games with infinite strategy spaces but also adaptable to their bimatrix-game discretizations, is that of Rosen [32]. It uses the notion of "concavity" of a function defined in Section 3.2 ("never overestimated by linear interpolation") sharpened to "strict concavity" by ruling out regions of "flatness," so that the condition reads "always underestimated by linear interpolation." If, for each possible choice of player 2's strategy  $s_2$ , player 1's payoff  $f_1(s_1, s_2)$  is a strictly concave function<sup>11</sup> of player 1's strategy  $s_1$ , and similarly with the two players reversed, then (according to [32]) the game has a unique equilibrium point.

This is perhaps the best place to mention a promising further concept described by Filar [132], that of a semi-antagonistic equilibrium point (SAEP). We begin with a bimatrix game B, with payoff functions  $f_1$  and  $f_2$  for the respective players; and recall that B is in general not zero sum (i.e.,  $f_1 + f_2 \neq 0$ ). With B can be associated two hypothetical matrix (hence, zero sum) games  $M_1$  and  $M_2$  defined as follows: in  $M_1$  the payoff function for player 2 is  $f_2$  so that (by the zero sum condition) the payoff function  $f_2$  rather than to maximize "his" own original payoff function  $f_1$ . Similarly, in  $M_2$  the payoff function of player 1 is  $f_1$  so that the payoff function of player 2 is  $-f_1$ ; i.e., player 2's objective has been shifted from maximizing  $f_2$  (as in B) to "hurting" player 1 by minimizing  $f_1$ . An equilibrium point  $(s_1^*, s_2^*)$  of B is called an SAEP of B if  $s_1^*$  is an optimal strategy for player 1 in the matrix game  $M_1$ , and  $s_2^*$  is an optimal strategy for player 2 in the matrix game  $M_2$ . Not every bimatrix game B possesses a SAEP, but when one exists it "represents

---

<sup>11</sup>It is technically easy to extend the concept of "linear interpolation" from functions of one variable to functions of several variables (e.g., the entries of a mixed strategy).

an equilibrium situation which is reasonable when the players suspect each other of vindictiveness" [132]. Filar also (1) gives a finite algorithm for determining a SAEP when one exists (it could be extended to determine all of them) or else verifying that none exists; (2) shows that any two SAEP's give the same payoffs to the players; and (3) shows that SAEP's have the interchangeability property--if  $(s_1^*, s_2^*)$  and  $(s_1^{**}, s_2^{**})$  are SAEP's then so are  $(s_1^*, s_2^{**})$  and  $(s_1^{**}, s_2^*)$ --so that one can speak of a strategy for one player alone as being "SAEP." Issues of parametric and sensitivity analysis are not addressed in [132], but their possibilities appear more favorable for SAEP's than for equilibrium points in general. In short, when SAEP's do exist and when the situation being modeled has the features of "mutual suspicion of vindictiveness" that make SAEP's especially attractive as a "solution" concept, their use regains many of the advantages typically lost in passing from matrix to bimatrix games.

### 3.4 Multi-Player Games

One purpose here is to discuss the solvability of p-player game models with  $p > 2$ . Before doing so, we pause to note that the section headings of this chapter omit one class of games; 2-player non-zero sum games with infinite strategy spaces. There is no general theory known that applies to broad classes of these games, in the spirit of the results for zero sum games described in Section 3.2. For purposes of numerical solution, they can be approximated ("discretized") by suitable bimatrix games, so that the content of Section 3.3 becomes relevant.

The preceding digression completed, we turn to the multi-player case. For simplicity, we assume initially that the (pure) strategy space of each of the  $p$  players ( $p > 2$ ) is finite.

The first fact to be noted about such games is a depressing one: no finite general algorithm for their (exact) solution is known, and in fact none can exist. To see how so sweeping an assertion can be justified, first recall the classification of real numbers into those that are rational (i.e., expressible as the ratio of two integers) and those that are not (the irrational numbers, such as  $\sqrt{2}$ ). Notice that if we begin with two rational numbers, then their sum, difference, product, and quotient are also rational numbers. Thus, no finite algorithm using the standard arithmetic operations (addition, subtraction, multiplication, division), when applied to an initial set of rational-number data, can possibly lead to an irrational number.

Now suppose  $X$  is some irrational number and  $G$  is a 3-player game whose data (the payoff values  $f_1(s_1, s_2, s_3)$ ,  $f_2(s_1, s_2, s_3)$ , and  $f_3(s_1, s_2, s_3)$  for all possible pure-strategy choices  $(s_1, s_2, s_3)$  by the three players) are all rational numbers.

Furthermore suppose that  $G$  has a unique equilibrium point and that the strategies of that equilibrium point (in general, mixed) yield  $X$  as the payoff to one of the players. Since no finite algorithm applied to the data of  $G$  can yield the payoff  $X$  corresponding to the unique solution of  $G$ , it follows that no finite algorithm can solve  $G$ , hence that no general finite algorithm for 3-person games can exist. In 1979, Bubelis [133] showed that for any irrational number  $X$  from a certain infinite class of irrational numbers (the algebraic irrationals--those which, like  $\sqrt{2}$  but unlike  $\pi$  or the base  $e$  of natural logarithms--arise as a solution of a polynomial equation with integer data), a 3-person game  $G$  with the properties stated above can indeed be constructed, and so the "no finite algorithm" conclusion follows.

The extension of the above argument, from three to more than three players, turns out to be easy. A more serious question arises if one permits, as single steps in a "finite algorithm" not only the four standard arithmetic operations listed above, but also the extraction of roots (square roots, cube roots, etc.). Then taking  $X$  as  $\sqrt{2}$  in the above argument would no longer yield a contradiction to the existence of a finite solution algorithm. But it is a classical mathematical result, due (1826) to the tragically short-lived Abel (for background, cf. for example, Chapter 10 of Tietze [134]) and generally referred to as "unsolvability by radicals," that there are algebraic numbers  $X$  which cannot be reached from rational data by any finite algorithm even if root-taking is permitted. With  $X$  so chosen, the argument remains in force.

A somewhat more cheerful note is the fact that the 1964 solution algorithm of Lemke and Howson [113] for the bimatrix case (i.e.,  $p=2$ ), can be generalized to  $p$ -player games with  $p > 2$ . Such generalizations were published back-to-back in 1971 by Rosenmuller [135] and Wilson [136]. Since the algorithm of [65] is finite, this might appear to contradict the "no finite algorithm" result cited above. Resolving this apparent paradox requires a little more detail on the nature of the algorithm in [113]. That algorithm can be viewed as having finitely many "stages." At some of these stages a set of simultaneous linear equations needs to be solved, but since that is a well-known finite computational task, the overall labor is finite. When the algorithm is generalized to games with more than two players, it continues to have only finitely many stages; but now at some of these stages, a set of simultaneous nonlinear equations of a particular type ("multilinear") needs to be solved, a difficult task which does not admit a finite exact solution method and can only be done approximately. Wilson [136] initially remarks that "presumably there are or will be numerical methods adequate to this task" but later observes "this is by no means a trivial presumption."

Further progress along these lines is reported by Garcia, Lemke, and Lueth [138] and more recently by Vander Laan and Talman [139]. The former includes results of computational experiments with games of 3 and 4 players (7-140 sec. on an IBM 360/50), the latter with games of 3 players; but since only 3 or 4 pure strategies per player were permitted, no extrapolation can be made to games of more application-interesting size. Colleagues have indicated that additional refinement of the techniques in [139] is thought (no doubt not unanimously!) to be the most promising direction for improved numerical solution methods for general multi-player games. But the problem appears intrinsically quite difficult, and the cited methods still find only one out of a possible multiplicity of equilibrium-point solutions, involving in general different payoffs and strategies.

A variety of particular multi-player game models have been solved, sometimes in closed form, but typically by ad hoc ingenuity so that no useful generalizations are apparent. One broader class of nicely-solvable games, identified by Howson [140], are the polymatrix games. These games are describable by a set of  $p(p-1)/2$  matrices  $A_{ij}$ , one for each ordered pair of distinct players  $i$  and  $j$ . The entry  $A_{ij}(s_1, s_2)$  in the  $s_i$ -th row and the  $s_j$ -th column of  $A_{ij}$ , represents the contribution to the total payoff of player  $i$  due to that player's choice of strategy  $s_i$  and player  $j$ 's choice of strategy  $s_j$ . When  $p=3$ , for example, the three players' payoff functions would be

$$f_1(s_1, s_2, s_3) = A_{12}(s_1, s_2) + A_{13}(s_1, s_3),$$

$$f_2(s_1, s_2, s_3) = A_{21}(s_2, s_1) + A_{23}(s_2, s_3), \text{ and}$$

$$f_3(s_1, s_2, s_3) = A_{31}(s_3, s_1) + A_{32}(s_3, s_2).$$

Such games have a natural modeling interpretation (payoffs arising as sums of returns from 2-player interactions) and are shown in [140] to admit a finite solution algorithm based on that of [113] for bimatrix games. I suspect that the methods of [123, 124, 127] can be extended from bimatrix games to find all equilibrium points of polymatrix games but have not encountered this in the literature.

In winding up this section, we now drop the restriction of finiteness on the players' strategy-spaces. The resulting class of models includes many of interest in the study of oligopolistic competition (cf. [19; pp. 370-374]). If each player's payoff is a strictly concave function of that player's strategy for each possible set of strategy choices by the other  $p-1$  players, then the previously cited result of Rosen [32] assures that there is only one solution. So far as numerical techniques are concerned, if the games (i.e., its payoff function and strategy spaces) have no special mathematical structures then there seems

nothing more to suggest than to pass to a discrete approximation and employ the methods described earlier in this section.

A rather different approach, which does however require some special properties of the payoff functions  $f_1, f_2, \dots, f_p$  has been studied by mathematical economists, e.g., Arrow and Hurwicz [141]. Suppose, hypothetically, that the game is to be played repeatedly, and let  $s_i(t)$  denote the strategy chosen by player  $i$  at the  $t$ -th play. As a "short-memory" analog of the fictitious-play concept described in Section 3.1, we might imagine that player  $i$ 's choice  $s_i(t+1)$  at trial  $t+1$  would be such as to maximize

$$f_i(s_1(t), s_2(t), \dots, s_{i-1}(t), s_i(t+1), s_{i+1}(t) \dots, s_p(t)) \quad (3.7)$$

since in the absence of communication or cooperation, "he" has no way of predicting how or whether the other players' choices would change from the previous trial. A more cautious approach would be to shift from  $s_i(t)$  in the direction of the maximum of (3.7), thus increasing player  $i$ 's payoff if all other players stand pat, but not moving all the way to the maximum. If the time-parameter  $t$  is made continuous, one way of doing this is to choose  $s_i(t)$  as a function of  $t$ , to satisfy the gradient-following condition

$$ds_i(t)/dt = c_i(\partial f_i / \partial s_i)(s_i(t), \dots, s_p(t)) \quad (3.8)$$

with each  $c_i$  a positive constant, and with additional rules to cover the cases where (3.8) would lead  $s_i$  outside the strategy space for player  $i$ . Rescaling each  $s_i$  permits changing every  $c_i$  to 1. Numerical solution of the  $p$  simultaneous differential equations (3.8) one equation per player, then yields a computational procedure which might be hoped as time progresses (i.e.,  $t \rightarrow \infty$ ), to converge to an equilibrium point. Conditions under which this hope is justified are presented and verified in [141]. Without repeating the specifics, we note that these conditions involve strict concavity, or convexity, of various auxiliary functions assembled from the payoff function  $f_1, \dots, f_p$  by various partitionings of the players into two groups; regardable as opposing aggregated players in a related two-player, zero sum game.

#### 4.0 SOME GENERIC ISSUES OF FEASIBILITY

In the preceding two sections, we have given an overview of the main ingredients of a game-theoretic model of a noncooperative decision problem, including the principal "solution" concepts involved. We have also reviewed what is known in general about such "solutions"--their nature (i.e., typically involving mixed strategies), existence, uniqueness or lack thereof, and the ease or difficulty with which they can be calculated. With this background developed, our aim in the present section is to address several issues bearing on explicit questions posed in the Nuclear Regulatory Commission's Request for Proposal [142] that led to the present study.

These issues all pertain to the feasibility of utilizing a game-theoretical model to aid the material accounting function described in Section 1; that of setting the "alarm thresholds," for values of inventory discrepancies in quantities of special nuclear material, in a way that properly balances the costs of over-frequent "false alarms" against the obvious risks of an alarm policy that is too "relaxed." Furthermore, the issues to be discussed below will be of a relatively generic nature: they do not refer to some specific game-theoretic model or technically delimited class of such models, nor to the particulars of the alarm threshold-setting problem within a spectrum of risk-benefit analysis questions arising from the Commission's responsibilities. Thus, the level of discussion here will be appropriate to the level of generality ("feasible to apply (game theory) in a regulatory framework?) suggested by A-1 and B-2 in [142; p.21].

Concretely, the topics to be treated are:

- Understandability of game-theoretic techniques,
- Viability of game-theoretic models in regulatory settings, and
- Mixed strategies as a source of possible difficulties in implementation.

The first of these topics is directly responsive to items A-2 and B-2 of the RFP [142; p. 21], the third to Task 1a identified there [142; p. 21]; while the second, as noted above, pertains to items A-1 and B-1. Our reactions to the first topic are unequivocally reassuring. Those to the second topic are generally positive, though necessarily more diffident (the writer is not an expert in administrative law or protocols) and paying greater attention to necessary provisos. As will be seen below, the third topic raises some cost-benefit tradeoff questions of its own that are susceptible to analytical treatment but, more importantly, it can play an important, useful role in providing guidelines

and criteria to assist the development of an operational game-theoretic model for the "alarm threshold" problem.

#### 4.1 Understandability of Game-Theoretic Techniques

This issue is raised in [142] with particular reference to the feasibility of achieving "understandability by licensees." It appears to express a concern that game theory, a subject not typically part of the educational background of today's senior engineers or managers,<sup>12</sup> might be so very esoteric or abstruse as to be beyond adequate comprehension (and, therefore, informed acceptance) by suitable licensee personnel. Fortunately, this is not the case. That conclusion will be supported by three lines of reasoning.

First, the subject has enjoyed a long and successful history of popularization for "lay" persons. J. McDonald integrated his prior articles on the topic in Fortune into an acclaimed "low-tech" introduction to the field [143] which, "by popular demand," has passed through a number of incarnations (e.g., [144]). A second well-received and reprinted early work of the same genre, described in its preface as "a primer-for home study," intended for "the intelligent layman who happens not to have acquired a mathematical vocabulary," was that of Williams [145] at the RAND Corporation. A more up-to-date but still untechnical treatment is given by Davis [146], while a yet more recent entry is that of Jones [147].

Second, even at a more technical level the topic is not currently regarded as so complex that its study need be deferred to graduate school or to the senior undergraduate years. Many colleges and universities (my own included) offer introductory courses or course-modules in the field, which do not impose specialized mathematical prerequisites beyond the initial courses common for most technical students. The Mathematical Association of America and the Educational Development Center, typically with funding from the National Science Foundation, had long-term efforts during the 1970s to develop and disseminate modular instructional material addressed to undergraduate students of science, technology, and engineering; many of these modules dealt with aspects of game theory. An extremely influential mid-1950s freshman text (Kemeny et al [148]) devoted a chapter to linear programming and the theory of games. By the usual processes, parts of these materials have gradually seeped down into the "enrichment matter" offered by some secondary schools to their more advanced students. Learning has generally been

---

<sup>12</sup>Though in fact it is now part of the curriculum in many leading business schools (not necessarily as a separate course).



facilitated by the extra motivation imparted by the intriguing nature of the questions addressed in game theory.

It follows from the last two paragraphs that acquiring an understanding of game-theoretic concepts and methods should not prove beyond the competence of suitably educated and motivated licensee personnel. Indeed, some junior staff members may already have received education in the field. Statistics-trained employees would already have a good grasp of many of the ideas involved. I regard as most remote any necessity of developing as a training aid, an application-oriented monograph like that for statistical methodology [2], but the above material supports the practicality of such a step as an unlikely fall-back.

This leads to my third reason for classifying "licensee understanding of game-theoretic techniques" as a dismissable issue. The licensees are engaged in an operation for which, as a matter of settled policy, the possibility of deliberate diversion is of serious concern (as evidenced by expensive and intensive precautions). That is, they are engaged in an activity whose considerations include intelligent reaction to the possible presence of an "adversary." Since decision analysis in such situations is precisely the subject of game theory, game theory is one of the relevant technical disciplines for the operation. If the NRC concurs, then--to be politely hard nosed about it--it becomes the licensee's responsibility to hire or train personnel to a suitable level of proficiency, just as it would be for some relevant branch of nuclear engineering. Note that no considerable expense is involved--there is obviously no need for 24-hour on-site coverage by "the game theorist." And as has been indicated above, the necessary expertise is not so rare or arcane as to make this requirement a really burdensome one.

#### 4.2 Viability in Regulatory Settings

Here the term "viability" has been used as short-hand for robustness against accusations of violating the constitutional guarantee of procedural due process of law. Thus, the question is whether a game-theoretic model can be of such a nature, and its use in aiding alarm threshold setting can be of such a mode, as not to render this element of the regulatory process "unreasonable, arbitrary, or capricious." Before offering some impressions on this point, the writer must acknowledge his lack of expertise in the areas of administrative law, and regulatory practice, plus awareness that even for experts the topic of "due process" is by no means straightforward (cf. [149]).

It seems useful to begin, more generally, with the commissioning and use of mathematical models by public agencies, especially at the Federal level. This subject was extensively surveyed in the early and mid-1970s; cf. Fromm et al [150] and Gass and

Sisson [151]. The volume of such activities was found to be very substantial (and spot-checks at the time, by the present writer, indicated it was probably considerably underestimated). Actual utilization of the models was found to be relatively high when model development had been carried out organizationally close to (or in close contact with) the potential-user branch of the sponsoring agency, but typically much less frequent and successful in other cases. Further critical observations, focused on the quality of documentation and verification (and corroborated for military-oriented models by Brewer and Shubik [152], are less relevant for present purposes. The titles and sources (e.g., EPA, SEC) of many of these models clearly indicate their intended use as analytical aids in regulatory procedures.

Unfortunately, these broad surveys of a decade ago have not been repeated or updated, although excellent reviews in a few narrower areas have been carried out (e.g., Friedman/OTA [153]). Nevertheless, it seems safe to conclude--given the continuing general trends toward "mathematicalization" and "computerization"--that mathematical-model efforts related to regulatory policies and practices remain substantial. The pages of the Bell Journal of Economics, for example, offer ample testimony to this conclusion, though they normally do not specify the degree and nature of actual adoption of the proposed models. Another illustration is given by the large-scale modeling systems--more or less descendants of the Project Independence Energy System (PIES)--developed and operated by the Department of Energy; cf. for example Gass [154] and Gass et al [155]. Johnson [157] describes operational application in fisheries regulation. More generally, it seems clear that many quantitative issues involved in a regulatory judgment--the effect on competition of a proposed merger, the effect on regional pollution levels of a proposed change in processing fuels or technology, the adequacy of continuing current charges to provide a reasonable rate of return for a public utility--involve complexities and data-volumes that must be receiving formal mathematical treatment. As noted in [153; pp. 7-8], "models--are often the method of choice to meet the requirements of legislation," and further "in translating legislative requirements into management practices, agencies often recommend procedures that depend on the use of models." (Page 186 of [153] sketches a number of water-resource-related model uses by the NRC in support of its regulatory activities.)

This is not to say, of course, that "any old mathematical model" should or could prove viable in a regulatory setting. Common sense, good professional modeling practice, and the obvious spirit of "due process" all suggest, for example, that a suitable model should do the following:

- a. not contradict reliable, relevant data and to the extent possible should, in an appropriate way, be based on such data, and consistent with Steele [158,p.23]
- b. not contradict well-established, high consensus theory (if such exists) in the clearly relevant discipline (e.g., nuclear physics or economics) and to the extent possible should, in an appropriate way, be based on such a theory.

At (a), in one case where a model dealing in part with complex hydraulic flows in plastic manufacturing plants produced flow rates departing from those observed by a factor of 10, the court (not surprisingly!) rejected the challenged regulation [153; p. 62]. It is impossible to resist quoting [157; p. 87] the list of criticisms leveled by natural-gas industry experts, and evidently supported to the satisfaction of the hearing examiner, at the developers of an early econometric model offered by staff of the Federal Power Commission: "false sophistication, non-professional performance, faulty use of data, incorrect identification of variables, statistical ineptitude, and conceptual inconsistency." (The examples of Finkelstein [159] are of related interest.)

But these horror stories are quite atypical today; as subsequent events showed in the latter instance, the rejection of a particular model in no way signaled rejection of modeling in general. To the contrary,<sup>13</sup> a relatively narrow standard has evolved for judicial review of models used for agency rulemaking: that the court, though conducting a searching and careful inquiry, should not substitute its judgment for that of the agency; that the agency's use of the model should be accorded a "presumption of regularity"; that the model's documentation must provide an adequate explanation of it's being a "rational choice" as basis for the regulation without the court's then undertaking to determine whether it was the "best possible approach." [153; p. 62] also notes that the Federal courts have proved relatively flexible in applying the "reasonable basis" test to disputed regulations and have displayed a reluctance to involve themselves in evaluating models per se. This exhibits sensible recognition of the point emphasized by Steele [158; p. 23]<sup>14</sup> that "a reasonable approach does not mean 'the correct' approach because there is no way to define in detail 'the correct' (modeling) approach toward the solution to a specific regulatory problem."

---

<sup>13</sup>The following passage is based on [153; pp. 61-62] which gives the specific citations including some apparently relating to [156].

<sup>14</sup>To better match the present context, "econometric" was broadened to "(modeling)" in the following extract.

Chapters 6 through 8 of [154] deal specifically with model roles and presentation strategies in the context of "procedural due process." One piece of advice is to avoid the "oversell," of attempting to envelop the model in "an aura of verity." Instead, model documentation and presentation should stress the rationality and reasonableness both of the entire approach and of the major choices made (necessarily on partly subjective "best judgment" grounds) at key points in the model-development process. A second implicit recommendation ([154], pp. 87-88) is that the model should be capable of appropriate particularization to the specific cases at hand (e.g., by setting parameters to the relevant values) instead, for example, of being irrevocably wedded to aggregate average data that might be demonstrably significantly inapplicable to the particular instance. (This reinforces our stress, in Section 3 above, on the possibilities for sensitivity and parametric analysis of game-theoretic models.) A third observation ([154], pp. 3, 74-75) is that the severity of the "due process" criterion will naturally depend on the salience of model-use in arriving at the regulatory decision already substantially determined, to a maximum role in which model outputs become the determinate guidelines for decisions. Here it is useful and encouraging to note the variety of regulatory instruments and modes available to the NRC (see Section 2.0 of the main report) in determining how best it might utilize the game-theoretic approach.

Some distinctions among different possible functions of a mathematical model should be raised here. (They are discussed in a game-theoretic context by Wiberg [160].) A model may simply provide a compact representation or encoding of some body of observations (descriptive function). It might be viewed as a theory of the phenomena in question, giving intellectual insight into why the observations turned out as they did (explanatory function). If the description or theory is thought to have validity extending beyond the observations or data already at hand, it may be employed for the predictive function of forecasting the nature of future occurrences of the phenomena, perhaps in hypothetical alternative futures reflecting different possible "states of nature," or different choices of policy or design by one or more decisionmakers. (This predictive capability is the most common motive for applied model-building.) Finally, it may be explicitly intended to help recommend a "best" policy or design (the optimization or normative function, arising for example in linear programming and game theory). These distinctions are useful despite the inevitable overlaps among their neat-sounding categories; e.g., the predictive mode would be used to generate tentative forecasts whose accuracy can be checked to evaluate a model's success as an explanatory theory or a predictive model (plus an explicit "scoring criterion") might be used in lieu of formal optimization to decide which is the best among a stipulated set of alternative decisions--which need not, of course, happen to include "the best" one.

Although most of the models alluded to in the preceding references are of the predictive variety, the comments do encompass normative models as well. For example, [153] with its positive findings on "viability" explicitly includes this category ([106], p. 155-157). Examples of an agricultural type (e.g., pest management and pesticide-use control) are given by Rovinsky and Shoemaker [161]. Energy-related models like those cited above [154, 155] often are (or have major submodels) of the optimization type. Some members of the judicial branch may have acquired a sympathetic familiarity with optimization concepts from past expositions (e.g., Nagel with Neef [162], Nagel and Neef [163]) of those concepts in the context of applicability to the legal process. We note that some checks for "reasonableness" available for predictive models are no longer so clearly applicable to normative models, since observations at variance with model outputs may simply reflect non-optimal behavior, improvement on which may have been the very motive for developing the model.

There seems no special reason to modify the preceding general comments about "viability," when the focus is further narrowed to game-theoretic models. But specific citable references have proven (for the writer) hard to come by; lack of a bibliographic survey of real (i.e., operational and accepted) applications of game theory is a serious gap in the literature. At the conference documented in [164], it was stated in connection with Shapley [165] that the courts had showed "intelligent sympathetic interest" in the use of game-theoretic constructs in judging the fairness of voting and representation schemes, a topic whose mathematical analysis was initiated largely within the legal profession itself (e.g., Banzhaf [166-168]). Another regulatory-pertinent area of application is to equitable allocation of costs (Lucas [169]); here there is at least one documented operational use (Billera et al [170]) as well as a substantial number of potential ones whose status is less clear--to water management, urban-transportation subsidies, airport landing fees, etc.--and to a case "which has been argued in the U.S. courts," concerning how to allocate taxes for accounting purposes, and involving the U.S. Government and the McDonnell Douglas Corporation. These developments, however, unlike the models developed by Goldman and Pearl [171-172] in the context of "weights and measures" inspections and income-tax-return auditing, stem from the "cooperative game" theory rather than the "noncooperative" branch most relevant for the present study.

There are two additional points, plausibly assignable either to this section or the next, which we choose to address here. Both reflect the fact that a "solution" to a game-theoretic model may require a player to adopt a "mixed strategy," i.e., a probabilistic mixture of courses of action.

The first point is that achieving such a mixture obviously requires a deliberate act of randomization. Might not the outcome

of such a "dice-toss" be held to be intrinsically "arbitrary or capricious," and therefore a failure of procedural due process? In light of the previous comments concerning standards of judicial review, this does not seem to the writer to be a serious threat, so long as the randomization was carried out (as well as the very use of randomization) systematically derived from a respected relevant theory (here, game theory). A telling precedent appears to be offered by the "random lottery" version of the Selective Service Draft, which did not succumb to any legal challenge of "arbitrary because random."

[Apropos the phrase "respected relevant theory" in the last paragraph, it should be noted [14, Section IV] that the (diverse) Peer Review Group had no qualms about the relevance of game theory to the suggested application and to possible related ones in the "safeguards" context. Its relevance for other types of application, especially some which are predictive rather than normative, might in particular cases prove arguable; the possibilities for controversy are illustrated by two recent lively--but, the writer feels, minority view--papers by Kadane and Larkey [173, 174].]

The second point, more hypothetical in a way, stems from the nature of the optimal mixed strategy calculated in an illustrative numerical exercise of the model proposed to the NRC in [7]. One "component" of that mixture [7; p. 40] to be employed with probability 0.058--i.e., on the average, in about 1 out of every 17 inventory periods--calls for setting the alarm threshold at a slightly negative level so the "maybe a diversion!" signal could sound even when the ID calculation indicated a (sufficiently small) excess of nuclear material at hand. The Peer Review Group [14, pp. 20-22] was suspicious of this result, which might reflect an inappropriateness in the model or an error in obtaining its solution, and which admits a quick (though intellectually unsatisfactory) "fix" by denying negative levels to the strategy space of the threshold-setting player. But if, after careful consideration and analysis, such an apparently counter-intuitive solution-component remained, would that in itself run afoul of the "unreasonable, arbitrary, or capricious" criterion?

With some trepidation, the writer suggests that this would probably not prove true. An action that would be unreasonable if taken invariably (e.g., every inventory period), may quite rationally be taken with some low frequency for purposes of deterrence or deception. For example, scarce police resources may be assigned to occasional extra random patrol, through low-crime neighborhoods, as the best means of keeping them "low-crime" and optimizing overall protection. Illustrations from poker (the need for infrequent but persistent "bluffing" on some weak hands, and "folding" on some stronger ones) are part of common folk-wisdom as well as verified consequences of game-theoretic analyses. Such arguments, it is suggested, are likely

to reverse an initial impression of absurdity for negative alarm thresholds (as limited-probability components of an optimal "mix") assuming, of course, that the underlying model otherwise appeared reasonably-based. The equitability of how the costs associated with such a policy are allocated might perhaps come under judicial scrutiny. But that is a different question.

#### 4.3 Possible Difficulties with Mixed Strategies

At the end of the preceding section, we addressed two possible threats to the acceptability of a mixed-strategy solution produced by a game-theoretic model: that the randomization aspect of a mixed strategy might be held "arbitrary and capricious," and that objections might arise if the optimal mix contained--with some small but positive probability that could conceivably lead to its activation--a rather counter-intuitive pure strategy. Reasons were given for expecting such threats to be surmountable without undue difficulty.

There is, however, a further class of possible difficulties associated with mixed strategies. Recall the source of such a strategy's advantage: that the opponent, when if in touch with a well-entrenched "mole," cannot exploit advance knowledge of your course of action (pure strategy) because that course is not actually determined until the random "device" implicit in the mixed strategy is exercised to choose (with the appropriate relative odds) among the menu of pure strategies involved. The other side of the coin, of course, is that you cannot benefit, either, from such advance knowledge.

Most of us presumably value, though to different degrees, order and predictability in the important parts of our lives. It is plausible that facility operators and managers, because both of occupational traits and of the personal attributes leading to their senior status, are especially likely to set high store on controllability and predictability of the operations for which they are professionally responsible. For such persons, a "wait until the dice are rolled to determine the response" dictum might be a particularly galling aggravation of the normal uncertainty inherent at the start of an ID determination.

These possible psychic costs have a more tangible counterpart. Each possible response (except the "all clear-do nothing" one) presumably involves a nontrivial sequence of activities, where timely and efficient execution upon demand may well require prior development of plans, instruction sets, stocking of particular equipment at particular points, practice drills, and the like. Attaining and maintaining a "ready state" of preparedness for each of a number of responses (those pure strategies which enter the optimal mixed strategy with positive or "sufficiently positive"

probability-weights), all perhaps associated with the very same range of ID-values, will in general be distinctly more expensive and strenuous than standing semper paratus for only a single course of action. (The optimal mixed strategy given for an illustrative case in [7; p. 40] contains 11 pure strategies, 8 of them with probability-weights exceeding 0.05.)

The severity of this effect cannot be estimated without a more concrete understanding of the various responses, and of their associated preparation steps and costs. But if it were so severe as to require alleviation, how might this be done?

The most simple (and simplistic) approach would be to eschew the use of mixed strategies, confining the "allowable" choices to the class of pure strategies. How much of a sacrifice in "protection" might this limitation involve, in particular for a player representing NRC-like interests in the kind of anti-diversion context motivating this study? In the absence of a specific validated model, only crude and tentative answers to such a question can be attempted. Such attempts, described in the next few paragraphs offer conflicting evidence but on balance indicate that the sacrifice may well prove acceptable. In case it does not, we will then go on to sketch some more sophisticated approaches.

Our initial rough-and-ready analysis of the "pure strategies only" approach involves a zero sum two-player game, i.e. a matrix game, in which the "NRC player" is the row-choosing player 1. Such a game is described by some m-by-n payoff matrix A, whose entry  $a(i,j)$  represents the payoff to player 1 if that player chooses the course of action represented by row i, while the opponent chooses the pure strategy symbolized by column j. Adding a constant to every entry of A does not change the strategic analysis of the game, and so without loss of generality we will limit attention to the case in which all  $a(i,j)$  are positive. If mixed strategies are permitted, then from player 1's viewpoint the game can be solved by solving the linear program (3.1-3.3) given in Section 3.1; the resulting optimal value, the "game-value" in the sense defined in Section 2.4, will be denoted  $v_{opt}$  and is the greatest (expected) payoff of which player 1 can assure himself. On the other hand, if player 1 can only use pure strategies, then his choice of any particular row i could result in his receiving that row's smallest entry as payoff, symbolically the quantity  $\min_j \{a(i,j)\}$ , and so the best payoff of which he can assure himself is obtained by choosing i to make that quantity as large as possible, i.e. to achieve at least the payoff:

$$v_{pure} = \max_i \{ \min_j \{ a(i,j) \} \}. \quad (4.1)$$



The ratio:

$$R = (v_{opt} - v_{pure}) / v_{opt}, \quad (4.2)$$

which is bounded above by 1 and is positive unless the original game happened to have a pure-strategy solution (in which case  $R = 0$ ), is then a reasonable normalized index of player 1's sacrifice in protection by confining himself to pure strategies.

A computer program was written<sup>15</sup> to carry out a Monte Carlo study of the magnitude of  $R$ . The entries of payoff matrix  $A$  were chosen independently and at random from the interval  $[0,1]$ . Specifically, a small scale study generated 400  $5 \times 5$  payoff matrices, determined  $v_{opt}$  for each by solving the associated linear program, then determined  $v_{pure}$  and  $R$  from equations (4.1) and (4.2) respectively. The resulting mean value of  $R$  was approximately  $1/3$ , i.e. on the average the restriction to pure strategies "cost" player 1 roughly one-third of the "protection-value" level available when mixed strategies were permitted. In roughly  $1/5$  of the cases, half or more of the mixed-strategy value was lost when only pure strategies were allowed.

These results appear quite discouraging for the advisability of the "pure strategies only" approach. Moreover, there is some reason to conjecture that the results would be significantly worse if the random payoff matrices of the above Monte Carlo experiment were replaced by matrices more specifically representative of our anti-diversion situation. The critical point is that  $R$  is small for games in which player 1 could make his choice known in advance ("and do your worst, you villain!") at relatively little cost, i.e. in which player 2 would not be materially assisted by advance information about player 1's chosen course of action. But our situation appears to lie at just the opposite extreme: expert opinion indicates<sup>16</sup> that accurate intelligence is regarded as "pivotal" (deLeon et al, [176; p. xii]) and uncertainty about the security systems is "abhorrent" (Jenkins [177; p. 7]) for adversaries, so that [176] "the deliberate creation of uncertainty" would appear to present the greatest obstacle to potential adversaries in planning and executing their acts."

A more careful conceptual analysis, however, reveals a flaw in the preceding reasoning and leads to a more encouraging prognosis. To explain the flaw, it is useful for concreteness

---

<sup>15</sup>I am grateful to my student E.S. Won for performing this task.

<sup>16</sup>Though the language of the following citations and their contexts could be narrowly construed as referring only to physical security, I believe a broader interpretation accurately reflects the writers' view; cf. the cited need (Bass et al, [175, p. 15] for more than a "castle and moats" concept of "security."

to turn to the simple but prototype scenario sketched in NUREG-0290 [5; p. 18] and in [6]:

Move 1: Diverter removes  $x$  grams of SNM.

Move 2: Inventory-taking leads to a figure of  $u$  grams of SNM as unaccounted for.

Move 3: Defender, knowing  $u$ , chooses one of the available courses of action.

The preceding paragraph treated Move 3's "course of action" as a (pure) "strategy" for the Defender (player 1), and noted accurately that advance knowledge of this choice would be advantageous to the Diverter (player 2). But in fact the above scenario describes a game in extensive form as defined in Section 2.5 (Move 1 "belongs" to the Diverter, Move 2 to the "chance player" who generates the random measurement/recording error in the reported ID-figure, and Move 3 belongs to the Defender). Thus a "strategy" for the Defender should be defined, as in Section 2.5, as a "complete contingency plan"--not as a specific course of action or response, but rather as a response rule (or decision rule) which specifies, as a function of the observed ID-value  $u$ , what action will be taken. Note that  $u$  might plausibly be (and is, in the existing models) taken to be the sum of the measurement error and the diversion amount:  $u = e + x$ . If the measurement process is rather imprecise (e.g., if  $e$  has large variance, then even if the Diverter knew the Defender's response rule (i.e. strategy), he could not confidently predict the Defender's actual response to any specific diversion-level  $x$ , because that response would depend on  $u = e + x$  which would be only poorly predictable from  $x$ . Thus the imprecision of the measurement process, a drawback to the quality-control aspects of safeguards activity, ironically offers some comfort to the Defender intent on using a pure strategy: it creates uncertainty for the Diverter even if the latter learns in advance of the former's strategy. For clarity, we emphasize that what would be "learned" is the Defender's response rule, not the specific response to be made, which depends on the random error  $e$  and so is not determined in advance.<sup>17</sup>

To what extent can the "pure strategies for the Defender might not be so bad" argument of the last paragraph overcome the "restriction to pure strategies looks bad" arguments that preceded it? To the writer's surprise, the answer appears rather promising for a restriction to pure strategies. In the context of the simple scenario given above, the argument runs as follows:

---

<sup>17</sup>In a more elaborate model, the Diverter might find it advantageous to learn or infer the Defender's response as soon as that response is determinable, i.e. after the ID-determination. But that goes beyond any of the cited models.

Let  $F$  denote the probability distribution of the random-error component,  $e$ , of the ID-quantity  $u$ . To avoid technical complications assume for the moment that the possible diversion amounts  $x$  are limited to a finite set, whose  $i$ -th member is denoted  $x_i$ . It is reasonable to assume that  $F$  is "atomless" or "absolutely continuous", i.e. that it attributes zero probability to exactly attaining any particular numerical value for  $e$ . (In the cited models, for example,  $F$  is generally taken to be the normal (Gaussian) distribution typically used to represent measurement errors.) Thus, if the Diverter chooses value  $x_i$  for  $x$ , the probability distribution of the ID  $u = e + x_i$  will be a simple "translation by  $x_i$ " of  $F$ , say  $F_i$ , which will again be atomless. Finally, still for technical simplicity, assume that the set of possible response rules available to the Defender is finite. These circumstances satisfy the conditions of an old (1951) theorem by Dvoretzky et al [178; Section 9], which assures that the Defender--but not, in general, the Diverter--will have an optimal "nonrandomized" (i.e., pure) strategy. Intuitively, the point of the "atomless" assumption is that knowing  $x_i$  does not permit the Diverter to single out any single value or finite set of values of the ID quantity  $u = e + x_i$  as particularly likely, and therefore limits the Diverter's ability to predict the Defender's response (to  $u$ ) even given knowledge of the latter's response rule. (Our "u" is the "x" of [178].)

As just indicated, the "atomless" assumption is critical to the above argument. The simplifying assumptions of finite sets of diversion levels and of response rules for the Defender turn out to be less critical; with a little care (see [178; Section 4], and the recent papers of Radner and Rosenthal [179] and of Aumann et al [180] extending this line of research), it appears that they can be relaxed at the cost of weakening an "optimal pure strategy" conclusion to one of "approximately-optimal pure strategy" to any desired degree of approximation. Although the cited results do not seem to apply explicitly to more complicated scenarios for the "ID-alarm" problem, the writer finds it likely that they can be extended so as to apply in many cases. The "bottom line" consequence is that in the context of a future effort to develop an operational game-theoretic model for this problem area, there are good grounds for believing that the design goal and criterion

o the model should admit optimal or near-optimal pure strategies for the Defender

is satisfiable (rather than merely desirable), and therefore reasonably adoptable as an initial guideline. (The weasel word "initial" corresponds to the "in many cases" a few lines earlier.) As relevant evidence, we note that [5, 6, 9] report pure-strategy solutions for the Defender. An apparent exception, mentioned earlier, is the "highly mixed" optimal Defender strategy of [7, p. 40]. But as noted in [14; pp. 21-22, p. 38], the Peer

Review Group was suspicious of the technical analysis leading to that purported solution, and in the present context it is especially pertinent to note that the suspected flaw lay in treating the Defender's strategy (in part) as a response rather than a response rule.

If the desired scenario for an operational model does turn out to differ from the simple one above in ways that frustrate the hope expressed in the last paragraph, one might seek to limit the number of distinct pure strategies present in the recommended mixed strategy. This can be done, for example, by modifying the linear program (3.1-3.3) given in Section 3.1. Note that " $x_i$ ", unlike its usage a few paragraphs ago, will now stand for the probability-weight assigned by a player to his  $i$ -th pure strategy. These continuous variables  $x_1, \dots, x_m$ , plus the variable  $v$  must now be supplemented by discrete variables, say  $q_1, \dots, q_m$ , with the desired interpretation that

$$q_i = 1 \text{ if } x_i > 0, \quad q_i = 0 \text{ if } x_i = 0 \quad (4.3)$$

so that  $q_1 + q_2 + \dots + q_m$  counts the number of pure strategies present in the mixed strategy represented by  $x_1, \dots, x_m$ . If this number is to be at most (say)  $L$ , then we adjoin to the linear program the additional linear constraint

$$q_1 + q_2 + \dots + q_m \leq L. \quad (4.4)$$

To enforce (4.3), we also impose the conditions

$$x_i \leq q_i \quad (\text{all } i), \quad (4.5)$$

$$0 \leq q_i \leq 1 \quad (\text{all } i), \quad \text{and} \quad (4.6)$$

$$q_i \text{ is an integer} \quad (\text{all } i). \quad (4.7)$$

The new optimization problem involves the maximization (3.3) subject to constraints (3.1), (3.2), (4.4), (4.5), (4.6)--so far, still a linear program--and finally (4.7), which puts the problem into the class of (mixed) integer linear programs ("mixed" because both continuous and integer variables are present). Finite solution methods for such problems (and implementing computer codes) exist, but in general are distinctly more laborious than for ordinary linear programs. The special way in which the discrete variables  $q_i$  figure in the constraints can probably be exploited to yield a solution algorithm (perhaps of "lagrangian" type) more efficient than those for the general run of such problems. (One staple reference on integer programs and their solution is the text by Garfinkel and Nemhauser [181] with its rather unorthodox dedication ("To the knicks"); an update by Nemhauser and L. Wolsey is forthcoming.)

Several variations on the last modeling theme are readily possible. For example, one might not make substantial preparations for the  $i$ -th course of action unless its probability-weight

in the mixed strategy exceeded some threshold-level  $t_i$ . This can be handled simply by replacing (4.6) with

$$x_i - t_i \leq q_i \quad (\text{all } i). \quad (4.6)$$

Or, instead of placing an a priori limit on the allowed number of pure strategies, one might want to let the optimization process balance the costs of their use (say,  $c_i$  for the  $i$ -th pure strategy) against the benefits in achieving the game's objective versus the opponent. This can be represented by dropping (4.4), but replacing the simple objective (3.3) with

$$\text{maximize } v - (c_1q_1 + c_2q_2 + \dots + c_mq_m). \quad (4.8)$$

Unfortunately, these linear integer-program models fail to capture an important aspect of the situation. Consider, for example, the term  $c_1q_1 + c_2q_2$  which is subtracted from  $v$  in equation (4.8). It indicates a cost of  $c_1$  if the chosen mixed strategy "uses" pure strategy 1 but not pure strategy 2 (i.e.,  $q_1 = 1$  and  $q_2 = 0$ ), a cost of  $c_2$  if the mixed strategy uses pure strategy 2 but not pure strategy 1, and a cost of  $c_1 + c_2$  if pure strategies 1 and 2 are both used. But if pure strategies 1 and 2 are both present in the chosen mixed strategy, the correct associated cost might be either distinctly less than  $c_1 + c_2$  (e.g., if the advance planning and preparation for the two pure strategies have significant overlap or economies of joint performance), or distinctly more (e.g., if the two sets of preparation are such as to interfere or to compete for scarce resources). Thus the indicated terms in (4.8) should really be replaced by  $c_1q_1 + c_2q_2 - c_{12}q_1q_2$ , where the "interaction coefficient"  $c_{12}$  has the appropriate sign. Note that the cross-product term  $q_1q_2$  makes the optimization problem nonlinear; still higher-order nonlinearities will arise, analogously, from considering combinations of three or more pure strategies.

Nonlinear integer programs can be treated by increasingly ingenious and efficient "linearization" techniques (e.g. Glover and Wolsey [182], Glover [183]) or by direct algorithmic approaches (cf. the survey by Cooper [184]; its restriction to "pure integer" rather than "mixed" problems is not too important for our purposes). These treatments do, of course, involve greater computational effort than for the corresponding linear cases. What is more chilling is the prospect of having to determine suitable interaction coefficients, like  $c_{12}$  above, for each of the numerous relevant combinations of mixed strategies. These considerations suggest retaining both the expanded maxim and (4.8) augmented by suitable nonlinear terms, and the limitation (4.4) with  $L$  chosen to keep the number of nonlinearities within acceptable limits for purposes of cost-estimation and computation. The trade-off issue between the two terms of the augmented (4.8) is likely to prove acute, because sets of pure strategies that admit substantial joint preparation are by that very token prone to lack the "diversity

of response" that provides good protection against the varied options of the opponent.

Although the issues introduced in this Section's fourth paragraph and treated in the Section's body appear significant for applied game-theoretic modeling in general the writer cannot recall seeing discussions of them in the previous technical literature.

## 5.0 THE PLAYERS

In this section we begin a more focussed discussion of the development of a game-theoretic model for the particular class of situations motivating this study. The reader is reminded that our study's objective is not the ambitious one of actually developing such a model, but rather an analysis of the issues bearing on the feasibility and preferable directions of such a research effort.

The initial decision in creating a suitable model is, as indicated in Section 2.1, an identification of the number, identities, and relations of the players in the "game." Possible participants, with language sometimes chosen to "personify" groups as if they were individuals, include

- one or more potential "diverters" or their agents
- the NRC
- the facility operator
- one or more public-interest group

Several comments about this list are in order. First, the reviews in Sections 2 and 3 show that games with three or more players offer considerable difficulty, both theoretically (as regards assurance of a conceptually compelling unique "solution") and computationally (bearing in mind that sensitivity analysis will require solution of the model for a number of sets of parameter-values, not just one). Thus there are strong practical and intellectual reasons for paring the above list down to just two players--if, of course, that can be accomplished without distorting reality in a way vitiating the usefulness of the model. This incentive for parsimony will color all that follows. It is relevant to observe that the proposed models encountered so far by the writer (e.g. [5-13], and also Avenhaus<sup>18</sup> [186-189], Bierlein [190-191]) are all in fact limited to two players; the sole exception is Bierlein [192], which has a somewhat different viewpoint to be reviewed later. The same restriction (to 2 players) prevails in the substantial body of work (e.g. Dresner [193], Anscombe et al [194], Aumann et al [195]) performed during the 1960's on game-theoretical analyses of inspection problems arising from possible arms-control agreements, a natural consequence of the essentially bipolar nature of the international strategic-weapons power balance then existing.

Second, even though public-interest groups (1) might well play a role in discussing the appropriateness of current or proposed regulatory decision-aids such as a game-theory model in the setting of alarm thresholds, and (b) might become involved

<sup>18</sup>I appreciate Prof. Avenhaus making available a pre-publication of his transparencies for the 1984 paper [189].

post hoc in commenting on particular outcomes of such steps, it is nevertheless hard (for the writer) to envisage a role for them as separate "players" in the context of the model itself. Assuming a suitable payoff function for the NRC player--an assumption which in a sense begs the question--it seems reasonable to regard that player's role as incorporating the relevant public-interest concern.

Third, we consider the need for a separate "facility operator" player. In the models proposed in the context of IAEA operations, e.g. those in the cited papers of Avenhaus [8-11, 186-189] and Bierlein [12, 190-192], the operator is explicitly or implicitly identified with the Diverter. Apart from innuendos [196] concerning Israel's acquisition of certain SNM a number of years ago, the writer knows of no suggestions that such an identification would be appropriate for the installations under the NRC's responsibility. It has been observed [175, p. 18] that some post-diversion scenarios might provide "cover-up" incentives for the operator to have common interest with the Diverter, and also (Willrich and Taylor, [197; p. 116-7]) that management might like "to have some clandestine material on hand simply as a convenient way to remove material accountancy anomalies as they arise--an easy way to balance the books." But given the purposes and priorities of developing a game-theoretic model the NRC's ID-analysis problem, it seems on balance (despite the considerations just noted) that such a development effort can properly treat the facility operator as belonging in the "anti-diversion" camp.

This does not mean, however, that the facility operator should be regarded as the essential persona of the Defender. The injurious potentialities of a successful diversion or diversion-hoax can extend far beyond the facility concerned, or even the relevant industry. Furthermore, these extensions are by not means of "second-order" importance relative to those effects local to the facility and its management. Thus the societal interests being protected go substantially beyond what would naturally figure most prominently in an operator's payoff function. On the other hand, because the operator's expenses for alarm-induced activities and process interruptions will presumably, for the most part, be passed along to the public in one form or another (prices, taxes), it is plausible to incorporate those expenses into the payoff function of the "NRC cum public" player already described. This is not entirely satisfying: for example, actual or potential unreliability in meeting supply schedules involves facility disutilities that may be difficult to quantify in terms commensurate with other contributions to the composite player's payoff function (cf. the remarks of Edlow in [198; p. 92].) But such modeling tradeoffs between detailed realism and tractability are rarely entirely satisfactory; under the present institutional arrangement, the indicated trade-off does seem advisable at least as the initial modeling strategy.



## 5.1 More Than Two?

Our discussion so far has led to recommending use of a single Defender player, identifiable mainly with the NRC, but with a payoff function that also does justice to the legitimate concerns of the facility operator. The next part of that discussion will concern the analogous issue of aggregation for diverter groups.

For readers who can entertain only with impatience the notion of any diverter's presence as a serious possibility, considering the presence of several such groups must seem downright farcical. A partial rejoinder is that whatever particular features (nature of material, vulnerability to penetration) might make a facility or MBA an especially attractive target to one diverter group, could prove enticing to other groups as well.

Suppose for the moment that two such groups were "co-present" at a facility. If their efforts are collaborative or supportive, then for a game-theoretic model that can probably be aggregated into a single player. If their efforts in effect interfere (e.g., their diversions trigger an alarm threshold that either one alone would have dribbled under), then it is conservative from the Defender's viewpoint to proceed with the analysis as if only one of the two were at hand. While these rather simplistic arguments should be replaced by a more careful treatment (to which [25-26] might contribute useful theory), they indicate in a rough way that for modeling purposes it is likely superfluous to postulate more than one "diverter at a single site."

The situation for a multi-site model seems more problematical. If such a situation could simply be treated as a collection of independent single-site games, no particular difficulty would arise. But there are two obstacles to such a decomposition. One is that the Defender's reactions may be system-wide rather than local--e.g., the response to an actual or perceived diversion attempt at one site might not be confined to that site. The other reason is that strictly speaking, diversion-seeking groups at different sites should be regarded as belonging to the same player only if they would pool their booty. It would be pointless for the writer to speculate on the reliability of possible intelligence on "who would pool with whom," and on whether that intelligence would suggest that rigorous application of the last sentence's criterion would yield just one Diverter, or more; the range of potential adversaries listed in [175] seems too diverse to permit an easy "just one" conclusion. It would of course be conservative for the Defender to regard his adversaries as all working together against him--and the resulting two-player format provides a natural start for model development--but the issue is one that should be flagged for further analysis (e.g., to quantify the "conservatism" in some rough way, especially in the light of a "threshold quantity" concept) when an operational multi-site model is attempted.

## 5.2 Less Than Two?

Our initial list of possible players has now dwindled to a single Defender and (at least for a single-site model) a single Diverter. There is no question about the Defender's inclusion in the prospective game-theoretic model, especially since the model is aimed at assisting "him." And since the threat of the Diverter is a major raison d'etre for the safeguards system, one might think that the inclusion of that player in the model would also be accepted without question. But in fact, the appropriate nature of this inclusion proved a distinctly controversial element in the deliberations of the Peer Review Group [14]. Because the points at question may prove important for the development of an operational model, we will review them here in the light of reflective hindsight.<sup>19</sup>

The NRC-supported models [5-7] presented to the group for review, as well as the European literature (Avenhaus, Bierlein) cited earlier in this report, make the classical game-theoretic assumption that both players are in fact "present" for the play of the game--this is part of the normal meaning of "player." Most members of the Peer Review Group regarded the assumption of an always-present Diverter as "overly conservative," especially given the "lack of hard evidence" that there had ever been a diversion [14; p.20]. The evidentiary basis for this evaluation is marred by the "Wilmington incident" (hurriedly appended as a footnote to [14; p.20]), and quite possibly by the larger-scale rumored "Apollo diversion" [196] as well, but the assumption of an ever-ready Diverter might still appear extreme: the term "paranoid" was in fact suggested. A proposed counter-argument (not by the present writer) was that it would be prudent to lock one's door against burglars each night, even if it were paranoid to believe that your door was in fact tried by a burglar every single night. Exceptions to the majority view are stated in [14; pp. B-2 and B-6].

There are two matters of terminology which may obscure the fundamental issue, and can therefore usefully be cleared out of the way in advance. One of them pertains to language like "always present"; those who criticize a model's attributing this behavior to the Diverter are not asserting that the Diverter is intermittently present, but rather that this presence is less than certain, i.e. that the probability of presence is less than unity. An uncritical "frequency" interpretation of the probabilities usually does no harm in applications, but

---

<sup>19</sup>In what follow, material not explicitly keyed to the Peer Review Group's report [14] is largely based on the writer's files noting the informal exchanges of information and views within the Group. Enough time has passed since the Group's activities (in 1978-9) to permit further reflection, but the reflections of other Group members may of course have developed along quite different lines.

here may prove misleading. The second linguistic point is that opponents of the "certainly present" model assumption do not really think that proponents of the assumption regard the presence of a Diverter at every facility as certain. Instead, the issue is whether modeling as if this were the case, given the underlying concerns of the safeguards program, should be viewed as "appropriately prudent" or as "overly conservative."

A fairly natural idea, in this setting, would be to seek to develop a model which included as a parameter a quantity  $p(D)$  denoting the probability of the Diverter's presence. (If  $p = 0.5$ , for example, one might figuratively speak of the game as having "1.5 players.") Such a model could then be analyzed for its sensitivity to the value of  $p(D)$ . One would need to develop two payoff matrices for the (row-choosing) Defender--one of them (say  $A_1$ ) to serve in the presence of the Diverter, the other ( $A_2$ , constant across each row) applying in the Diverter's absence. Then one approach would regard  $p(D)A_1 + (1 - p(D))A_2$  as the Defender's effective payoff matrix; a more sophisticated one, suggested in [14, p. 32], would take the probabilities  $p(D)$  and  $1 - p(D)$ , and the payoff matrices  $A_1$  and  $A_2$ , as the ingredients of a repeated game of incomplete information as defined in Section 2.6.

Although such an approach has considerable appeal and merits exploration, the writer is suspicious of it. One reason is that  $p(D)$  might really be a derived strategic variable rather than a (constant) parameter of the model--e.g., the probability of the Diverter's presence at a site might reflect the intensity of efforts to penetrate that site, which might reflect the attractiveness of that site as a target, which might reflect the alarm-threshold in force there. A more important reason, perhaps, is the  $p(D)$  is a "soft" parameter that might too readily be manipulated to coax model outputs into a desired region. Consider, for example, the process by which a value or a "reasonable" value-range for  $p(D)$  would be estimated. If that process rested heavily on the "hard evidence of diversion" criterion mentioned above, then low values of  $p(D)$  would presumably emerge. But suppose the criterion were reversed, to require "hard evidence of non-diversions," or at least "non-diversion beyond a reasonable doubt." It is dubious [199] that most past investigations of "triggering" ID-levels, with their honest motivation for identifying possible non-malignant loss mechanisms as "explanations" so that normal operation might resume, could nearly satisfy the latter criterion. That it is much harder to satisfy, does not automatically imply it should not be used; after all, absence of after-the-fact "hard evidence" of competent clandestine activity might by the same token be too easy to satisfy. So the question of where the "burden of proof" should lie, not really a technical question and certainly not an easy one, could heavily impact the estimation of the ostensibly "objective quantity  $p(D)$ ." (Issues in the nature and decision-theoretic use of subjective probabilities are indicated, e.g., by Kyburg [200].)

Retrospectively, the writer has been led to speculate that the Peer Review Group's majority-objection to the "certainly present Diverter" scenario marked an objection to a still more extreme assumption hidden in the reviewed models [5-7]. In these models, a pure strategy for the Diverter is the selection of an amount to divert. The implicit assumption is that having chosen such an amount, the Diverter can in fact successfully abduct that quantity. We might call this the assumption of the "certainly present and perfectly capable" Diverter. My conjecture is that the presumption of a "perfect capability," disregarding as it does the effectiveness of safeguards elements other than material accounting, is what really stuck in the throats of the reviewers, but that because this source of irritation was less readily identified, their wrath was displaced onto the more visible "certainly present" attribute. An implicit assumption of this type--that a chosen decision can and will be executed accurately--is an unrealistic imperfection in many decision-aiding mathematical models; it might well be called the "Ko-Ko fallacy" after the characters's excuse in Act II of The Mikado:

"It's like this: when Your Majesty says 'Let a thing be done,' it's as good as done--particularly, it is done, because Your Majesty's will is law--so why not say no?"

This suggests that efforts to develop an operational game-theoretical model for our problem might well accept the "certainly present" convention as appropriately prudential, and instead concentrate on properly crediting other safeguards elements by modifying the "perfectly capable Diverter" assumption. Specifically, the effectiveness of those elements might be represented through the parameters of a conditional probability distribution describing how much (possibly zero) material would actually be diverted if the Diverter sets  $x$  as his "target" amount. Although such parameters would inevitably share to some degree the "softness" of which  $p(D)$  was accused above, it is anticipated that this degree could be substantially lessened because of the considerable body of analysis to which the other elements have presumably been subjected. That analysis is hoped to be realistic and knowledge-based in its treatment both of equipment reliability and of human fallibility, e.g. in vigilance, and in the design, fabrication, installation, operation and maintenance of equipment; cf. Marshall [201], Green [202; pp, 312-313].

We have arrived at a recommendation for a two-player model (Defender and Diverter), except for the possibility of more than one (non-cooperating) Diverters in a multi-site analysis. The general nature of the Defender's interests is implicit in the preceding description of that player as corresponding roughly to the NRC acting in the public behalf, with due sensitivity to the needs of facility operators. We need not for the moment elaborate on the Diverter's interests (there will be more about this in Section 7), beyond the extent to which they are explicit

in his role-title. It seems clear that the two players' interest are broadly antithetical, and that communication or coordination between them would not be natural in our materials accounting setting of alarm-setting and responses--it would be natural in a scenario of negotiation over return of diverted material, or over a threat based on claimed possession of such material. This "cool and distant" relationship provides the justification for the earlier decision, at the end of Section 2.1, to confine discussion to non-cooperative games.

## 6.0 THE STRATEGY SPACES

Our purpose here includes a review of the pure strategies regarded as available to the players in the various models cited earlier. This plus additional discussion, will indicate some of the possibilities in developing an operational model for the application at hand.

Before turning to "strategies" proper, however, it is necessary to note the different possible settings in which these strategies would be chosen, utilized, and would interact. By the term "setting" we mean the spatial and temporal extents of the model, i.e.

- single-site or multi-site
- single (accounting) period or multi-period.

A multi-period model defines what in Section 2.5 was called a "game in extensive form"; as indicated in that Section's mention of "information sets," care will be necessary in specifying what information will be accessible to the player at different points in time. This specification, also, can be properly regarded as part of the "setting." Analogous issues include (i) the choice of total time-span for a multi-period model and (ii) the timeliness and accuracy with which player-agents at different sites can share information and coordinate efforts.

Models of the single-site single-period type are the natural starting point for a "less to more complex" progression in model development, and are likely to be building-blocks in later more sophisticated constructs. For brevity, we shall refer to such models as simple.

### 6.1 The Diverter's Strategies

The models developed for the NRC (Siri et al [5,7], Dresher and Moglewer [6]) are all "simple" in the sense just defined. In each of them, a pure strategy for the Diverter is a number  $x$ , the amount to be diverted, which is selected from a specified interval  $[0, K]$ . The upper limit  $K$  was interpreted [7; p. 17] as the smaller of the site's stock-level and a "credible threat amount." One Peer Group member suggested that  $K$  might depend on the physical and chemical properties of the materials at the site in a way representing the time and effort necessary to produce a clandestine fissile explosive. As noted in Section 5.0, the models also assume the Diverter to be "perfectly capable" of achieving his desired diversion-level  $x$ , but since that is not a necessary assumption in connection with this strategy-space, it is irrelevant for us here.

This choice of strategy-space is not unreasonable, but it does raise some questions. If  $K$  is "small," it may be implausible to model the diverter as unable or unmotivated to seek no more material than  $K$ ; if  $K$  is "large," it may be unrealistic to conceive the diverter as believing himself able to make off with so much at a single site during a single time-period. Other elements of the model may compensate for the second criticism; i.e., optimal play might forbid the Diverter from choosing excessive diversion-levels despite their formal availability in his strategy space.

The restriction to a single time-period, implying inability to represent the important possibility of "dribble" (i.e., a little at a time) diversion, is very dubious for an operational model. This was recognized in [7; p. 48], and indeed the Peer Review Group [14, p. 42] explicitly recommended development of a multi-period model. Ignoring for the moment other possible dimensions of diverter-action, let us ask how the preceding strategy-space would generalize to a multi-period model.

Such a generalization would presumably replace the single target diversion-level  $x$  by a sequence  $(x_1, x_2, \dots, x_T)$ , where  $x_t$  denotes the Diverter's target-level in the  $t$ -th time period. This is the approach taken by Avenhaus and Frick [9, 10]. The upper bound  $K$  has an analog in their model, namely a constant appearing in a condition

$$x_1 + x_2 + \dots + x_T = K \quad (6.1)$$

that delimits the Diverter's strategy-space. (Our  $x_t$ ,  $T$ ,  $K$  correspond respectively to the  $M_i$ ,  $n$ ,  $M$  of [9, 10].) Clearly  $K$  here represents a "critical total quantity" which the Diverter is committed to secure at any costs receiving no "extra credit" for gains beyond this level. The "perfect capability" assumption for the Diverter is obviously present, as is the implicit presumption that he can retain and cumulate his stash, period by period, despite the Defender's recovery efforts and other possible loss processes.

The existence of a quantity  $K$ , playing the role indicated by equation (6.1) and assumed known by both players, is somewhat discomfoting. It turns out however that under this assumption the model of [9, 10] has a unique solution in pure strategies, in which the Defender's strategy (the one of main interest for our purposes) is in fact independent of  $K$ , so that apparently the Defender need not know  $K$ . Less comfoting is the explicit remark [9; p. 123] that the model admits negative  $x_t$ -value ("anti-diversions") in equation 6.1; the test includes no assertion that these anomalies are absent from the model's unique solution, which might therefore exhibit bizarre credibility-damaging Diverter behavior. The noting of various limitations, in this paragraph and the last one, is not intended as disparagement of a pioneering

research contribution, but their alleviation does suggest specific technical direction for further efforts to achieve an operational model.

The diverter-strategies described so far are of quantitative nature: how much to steal (or aim at stealing)? In contrast, other related models postulate what we will call dichotomous strategies for the Diverter: at each time-period, a binary decision as to whether or not to undertake a diversion, presumably of some unspecified "canonical" amount. In these multi-period models, a pure strategy for the Diverter is a selection of which time periods to be active in. For the relevant Models II and III of Bierlein [190], the total number of diversions is fixed in advance, yielding an analog of condition 6.1 if the strategy variables  $x_t$  are confined to two values (0,1) corresponding respectively to non-diversion and to diversion in period  $t$ . (The notation of [190] uses "r" rather than "k," and "illegal action" rather than "diversion.") Model II differs from Model III in that the latter requires the periods of active diversion to be consecutive; the text notes explicitly Model II's assumption that during the intervals between diversion, past depredations cannot be detected by the Defender. The Diverter's strategy is chosen once-and-for-all at the start of play, and as before the Diverter is assumed "perfectly capable" (this last assumption will no longer be explicitly mentioned since it is ubiquitous in the cited literature). We note in passing the relatively "simple" models due to Borch [203, 204], in which the Insured decides whether or not actually to invest in loss-reducing measure promised to secure a reduced insurance premium.

The models in Bierlein [191], Beinbauer and Bierlein [12], and Bierlein [192] are rather different. For one thing, they do not all fix the number of time-periods for action (still dichotomous) by the Diverter. They involve a critical detection-time  $t$ ; the Diverter "succeeds" if some diversion goes undetected for  $t$  time-periods. (Footnote 1 in [191] interprets  $t$  as the time required to fashion a weapon from diverted material.) The diverter is assumed to know the "given" limit on the number or mean frequency of the Defender's inspections. In the model of Part I of [12], a pure strategy for the Diverter is the choice of a single time period in which to divert (or, a decision not to divert at all); this choice is made at the start of play. The scenario is more interesting in the other references just cited and in Hopfinger [13]; a pure strategy for the Diverter consists (roughly speaking) of a decision at each time-period, knowing the timing of all prior inspections, as to whether or not to divert during the current period. (The term in [13] is "aggression" rather than diversion.)

Having considered models for the multi-period single-site case, we turn next to the single-period multi-site case, beginning with those models in which the Diverter's strategies are dicho-



tomous. These are the models of Goldman and Pearl [170, 171]. In their context of intended application, it is most natural (though not necessary) to equate the diverters with facility-operators; the results of [25] justify aggregating these into a single Diverter player, for whom a pure strategy is a subset of the sites chosen as the scenes of "cheating" (i.e., diversion). The expressive limitations of dichotomous strategies ("divert or not") are mitigated by differentiating the sites as to their rewards from a diversion. Pages 192-3 of [170] note possible directions for model extension, one of them addressed in the later [171], but generalization to a multi-period model is not mentioned. Rumball [205] has suggested applying these models to the patrolling of New Zealand's territorial waters against illicit fishing operations.

At this point we introduce a second possible mode of behavior by the Diverter, noted in [7; p. 48] but especially relevant in the "operator as diverter" contexts emphasized by Avenhaus in [206, 185-8]. This behavior is the falsification of accounting or measurement data, a play which, if successful, could significantly ease the subsequent removal of material. Such a theme is the game-theoretic (i.e., adversary-conscious) analog of an important principle in control-system designs, the need to make special provision assuring that stressful incidents at a plant will not damage the sensors and displays providing the very status information whose integrity is vital for managing the incident (Young [207; pp. 7-11]). In the models of [186], for example, the Diverter seeks a given total amount (analogous to the  $K$  of (6.1)), must decide how to partition his effort between data falsification and immediate removal of material (doing exclusively the former is asserted to be optimal [186, p. 303]), and then how to apportion the former among the various sites.<sup>20</sup>

There have been only a few treatments of the combined complications of multiple period and multiple sites. Recently Avenhaus has reported [189] a multi-period extension of the model of [186] mentioned just above. For models with dichotomous Diverter-strategies, Model I of Bierlein [190] requires the Diverter to allocate a prescribed number of diversion-acts among the available (site, time) combinations in advance of play, whereas Model III of Bierlein [192] appears to permit the Diverter--knowing the times of prior inspections at all sites--to choose both the time and the set of sites for an act (if any) of simultaneous diversion. A broad set of Diverter strategies is admitted in the prototype "Travelling Inspector Model" of Filar [63]: in each time period, the Diverter, knowing where the Defender

---

<sup>20</sup>We have been somewhat cavalier in classifying this work as multi-site, since its actual wording refers to various classes and batches of material at a single site. But the mathematics appears to admit multi-site interpretation.

(the single "inspector") spent the previous period, must decide for each site among a finite set of possible levels (perhaps null) for diversion. This model is a single-controller stochastic game in the sense defined in Section 2.6, the "state" corresponding to the Inspector's location.

In winding up this review of proposed strategy-spaces for the Diverter, some final remarks are appropriate. One is that the work of Dresner [193] and some of the papers in [194] (e.g. Anscombe [208], Kuhn [209], Davis [210], Maschler [211,212]) are relevant, but do not appear to add significantly to the ideas about modeling Diverter behavior given in the papers already cited. Another concerns the multi-period case: from the literature consulted during this study, it appears that an adversary with sufficient "insider" status to attempt covert diversion would by the same token be able to gain knowledge of some past actions by the Defender. Some of the previously described strategy spaces seem unrealistic in this respect, requiring the Diverter to hew to a preplanned schedule even though adaptations to the course of events would be advantageous. Thus, this point should be kept in mind in designing an operational model. A third observation is that dichotomous strategies (i.e., "steal or no-steal" decisions without regard to quantity), though possible springboards for useful generalizations, do not themselves appear an appropriate modeling construct in our context of material accounting. A possible exception, related to Section 5's criticism of the "perfectly capable Diverter" scenario, would be a submodel in which the strengths and weaknesses of other safeguard elements would be reflected in a probability distribution used to translate a "decision to divert" into a particular quantity diverted. The writer lacks the "feel," for the concrete particulars of a diversion opportunity and activity needed to assess the promise of this possibility.

A concluding idea, probably "wild" but recorded for completeness is that the Diverter could leave some of the diverted material in an easily-found position suggestive of accidental misplacement or overlook, thereby reducing suspicion and promoting premature termination of recovery/search efforts.

## 6.2 The Defender's Strategies

Because some subgroups of the cited models treat Defender strategies in ways unlikely to be useful for our particular purposes, we can dismiss these treatments with relatively brief mention. One such group consists of the previously mentioned IAEA-related models analyzed by Avenhaus et al [206, 186-189] in the context of a Diverter who can indulge in measurement-data falsification as well as in material-removal affecting the data. In these scenarios (e.g. [187; p. 313]), during an inspection shut-down the Defender can make independent measurement

to check some of the data reported by the facility operator (the potential Diverter). The defender's "strategy" decision includes how to allocate his (limited) sampling effort, and what mathematical combination of the discrepancies between the Defender-measured and operator-reported values should serve as a "best statistic" to be compared against a threshold level.

For the NRC setting, with the presumption (cf. Section 5.0) that the facility-operator is Defender-oriented rather than Diverter-oriented<sup>21</sup>, the notion of regular remeasurement by the Defender does not seem to the writer to fit very naturally. (This could change if future incidents or intelligence heightened concern about data-tampering as a Diverter tactic.) Where the cited analyses could more likely prove useful in the present context, is for the submodeling of reinventory aspects of the Defender's response to an "alarm" situation. With that suggestion we drop the "remeasurement" theme, referring the reader to [186, 187] for references (e.g. by Avenhaus, Frick, Jaech, Stewart) additional to those already listed.

A second group of Defender strategy-spaces requiring only brief description are those involving dichotomous decisions "inspect or not," translated in our situation to "alarm or not," which are unrelated to any indicator of possible diversion (e.g., an ID-level) but instead are based on distributing limited inspection effort to achieve optimal "risk coverage." The "distribution" takes place over the possible time-periods (if the model is multi-period) and/or sites (if the model is multi-site). The "limit" might reflect criteria of inspection-resources or intrusiveness-constraints; it might fix the actual number of inspections (Dresher [193], Bierlein [190], Beinhauer and Bierlein [12]), or bound the probabilistic-average number or cost of these inspections (Bierlein [192]) or the average interval between successive inspections (Bierlien [191]). In Hopfinger [13], the number of inspections is randomly chosen at the start of play from a probability distribution known to both players, but the chosen number is revealed only to the Defender. Most of these models assume that inspection is sure to detect a diversion, but non-unit probabilities of detection (assumed known) are considered for example in Beinhauer and Bierlein [12; Section 3.3], Goldman and Pearl [170-171; cf. p. 192 of the former], and Anscombe [208]. Note that whenever the limited stock of "allowed inspections" suffers a draw-down the situation can be regarded as having entered a "new state," a consideration reinforcing Section 2.6's suggestion that "repeated games of

---

<sup>21</sup>This could also be the case in those IAEA settings where the principal diversion threat is attributed not to the "host" government but to foreign or political-faction or subnational separatist groups; cf. Willrich and Taylor [197; pp. 117-8], Lovett [213; p. 210], Taylor [215], Dunn [216].

incomplete information" need to be extended to multiple-site scenarios in order to be naturally applicable.

In the multi-period multi-site Travelling Inspector Model of Filar [63], inspection is non-dichotomous--i.e., it can take place at any of a finite set of levels of intensity. This choice, however, does not depend on any prior indicator of suspiciousness like an ID-determination. The model's rather general payoff function could permit the probability of an inspection's detecting a diversion (and the accuracy of estimating that diversion) to depend both on the intensity of inspection and on the level of diversion, but no such submodel or interpretation is made explicit.

We turn now to models in which the Defender's decisions are based (in whole or part) on ID-levels or something analogous to them. A first example is the formulation by Kuhn [209] of a multi-period situation; the number of inspections is fixed, and each inspection has the same (known) probability of detecting a "violation" if one has occurred. At each stage, the Defender receives a signal indicating either "no violation" or "violation" or "doubtful"; if the third case arises, an "inspect or not" decision must be made. An actual diversion would have produced the "doubtful" signal (rather than the "violation" one) with a known probability, and similarly for an actual "no diversion." (The "signals" here are based on seismic data, with "no diversion" corresponding to a natural earthquake, "diversion" to a nuclear test.) We might regard this as a three-level discretization of the possible ID-values, with the boundaries separating the three regions fixed in advance, rather than subject to optimization by the Defender. A similar structure is studied by Schleicher [217] in the context of income-tax evasion, with the further feature that some "doubtful" cases may lie in a designated "uninspectable" class (corresponding, e.g., to Swiss bank accounts). Other game-theoretic analyses related to income-tax evasion and auditing include Hoffman et al [218] and the recent paper of Greenberg [219].

The highly pertinent multi-period model of Avenhaus and Frick [9,10] raises two points of particular interest here. One, which brings us back to the basic equation 1.1 early in Section 1.0, is how the "prior contents" term in that equation should be estimated when past time-periods have left non-zero (though non-alarmable) discrepancies between "book" and "physical" inventory figures. The particular approach adopted in [9,10] is to employ a particular variance-weighted average of the book and inventory figures; this gives a minimum-variance unbiased estimate of "prior contents," a desirable feature from a statistical viewpoint, and also substantially simplifies the game-theoretic model by leading to ID-values for the different time periods that are uncorrelated. It is disconcerting, however, to find in [9; p. 120] the acknowledgement that this estimate "is not

necessarily the best one from the point of view of detecting missing material." That illustrates the fact, also emphasized for example by Klein et al, [220], that "statistically optimal" estimates of quantities will in general not be the optimal ones for decision-problem uses.

The second especially interesting aspect of the Avenhaus-Frick model is its treatment of the theme of limited resources or intrusiveness for the Defender, a theme which reduces this player's strategy space to a subset of what it would be without such limits. That treatment in [9, 10] is to assume a fixed over-all false-alarm rate (FAR), i.e. a fixed probability that in the absence of diversion, the random measurement errors in ID would trigger the alarm in at least one time period. The Defender-chosen alarm thresholds and their corresponding false-alarm probabilities can (and do) vary from period to period, so long as their multi-period composition yields the stipulated over-all FAR value.

This approach has considerable appeal. Along with the model's other assumptions, it yields substantial benefits of analytical tractability, leading to a provably unique solution and to a fairly simple iterative numerical solution procedure. It encapsulates the "cost" element of the problem's "risk-cost" tradeoff quite neatly in a single parameter, the FAR level. One apparent difficulty is the need to know (or have a good estimate of) the probability distribution of ID in the absence of diversion<sup>22</sup>; this need, however, seems common to all efforts to improve the modeling aspects of material accounting (whether game-theoretic or purely statistical), and so should not be held against a particular model or methodology. (It would be important to study--via sensitivity analysis--how critical these accuracy-needs are for different approaches.)

Some other difficulties may be more serious. For example, it might be thought more natural to fix the overall miss rate (probability of not detecting a diversion effort-- an "acceptable risk" concept. Notice, furthermore, that under a fixed FAR value a less precise measurement system would lead to higher alarm thresholds, i.e. (roughly speaking) to greater inhibitions again sounding the alarm. It is questionable to the writer that a model with this property built-in would properly capture the public-interest concerns we have attributed (in Section 5.0) to the Defender. A natural rejoinder is that the "fixed" FAR-value is intended to be constant relative to the player's decisions, but not with respect to changes in other model parameters such as (say) the variance of a normal measurement-error distribution

---

<sup>22</sup>One reason this is a "difficulty" is lack of certainty that the data used in reaching such an estimate really did come from diversion-free environments.

representing the precision of ID-determination. This illustrates the underlying weakness: in being silent on how the "fixed FAR value" should be chosen or determined (in [221; p. 630] the same authors term this choice "subjective"), the approach "cops out" on addressing the balancing of risk and cost, the very purpose of model development. One might hope to remedy this by coupling with an auxiliary submodel relating cost (of "alarming") to FAR, but the second sentence of this paragraph indicates that such a submodel would need to address risk considerations as well as cost.

On balance, the writer's inclination is for a model that more directly and explicitly tackles the intrinsic difficulty of representing the risk-cost tradeoff in its payoff structure. But the clever research tactic of the "fixed FAR" concept should be kept in mind as a fallback position. Similar comments apply to other model's restriction of "inspection effort" to some externally determined limit.

In working towards an operational game-theoretic model, it is important that players' strategies be conceptualized in sufficiently concrete terms to guide empirical and analytical efforts to quantify how payoffs depend on these strategies. With some possible partial exceptions (e.g. [170, 171, 203, 204, 218]), this level of specificity has not been attempted in the cited literature. In particular, the abstract encoding of Defender options as "inspect or not" or "alarm or not" does not in itself describe the consequent search and recovery procedures in a way aiding the quantitative representation of these measures' costs and effectiveness. The same is true of the different "levels of effort" that can be exerted by Filar's Travelling Inspector [63]. These abstract encodings should not prove incompatible with later efforts to make them "operational" through more detailed application-specific modeling, but they do not provide any initial hints or steps to assist such further efforts.

We will return to this point at the end of the present section, but want first to conclude our review of published "Defender strategy spaces" with the models [5-7] developed for the NRC. They are praiseworthy in achieving a higher (though still not high enough, in the writer's opinion) degree of explicitness in describing the Defender's responses. Recall from Section 4.3 that in these single-period single-site models, the Diverter first chooses and achieves a diversion-level of  $x$  units of SNM. Then the measurement processes involved in ID-determination introduce a random error  $e$ , so that an ID-level of  $u = x + e$  is reported to the Defender. The Defender compares  $u$  with an alarm threshold,  $z$ . If  $u$  is greater than  $z$ , the Defender sets a quantity  $y_2$  as the target level for an intensive "alarm conditions" search and recovery operation; if  $u$  is less than or equal to  $z$ , a target level  $y_1$ , (possibly zero) is set for presumably less intensive "no alarm" search and recovery actions. From

the payoff functions (described later) for these models, it is apparent that  $y_1$  and  $y_2$  are to be construed literally as recovery target-levels and not just as indices of search-effort. For example, the expression  $\min(x,y)$ --where  $y$  is  $y_1$  or  $y_2$ --is used to designate the quantity of material recovered by a successful search. (Thus a multi-period extension of this model would have to modify the Diverter's period-by-period accretion of material to allow for recoveries.)

The greater explicitness of these models lies in their introduction of the auxiliary quantities  $y_1$  and  $y_2$  to provide a somewhat more concrete picture of the Defender's responses. In the model of [5, 6], the alarm threshold  $z$  is fixed and is known to both players, so that a pure strategy for the Defender consists in selecting the pair  $(y_1, y_2)$ . In the further model introduced in [7], the alarm threshold is also regarded as part of the Defender's strategy  $(z, y_1, y_2)$ , and therefore as unknown to the Diverter. This change proves (in a numerical example) to alter the game-value very significantly in the Defender's favor [7; p. 39]. While skeptical of the validity of the particular payoff functions employed, the writer would like to emphasize the qualitative point illustrated here: the importance of information to the adversary of a safeguards system, and the particular suitability of game-theoretic models for representing alternative "information scenarios" and quantifying the effects of their differences.

As described above, the Defender's choice of  $(y_1, y_2)$  or  $(z, y_1, y_2)$  can be made after the ID-value  $u$  is reported to him. Thus a Defender strategy should describe how these choices would be made in response to any particular value of  $u$ , i.e. it should in general be a response rule specifying functions of  $u$ ,  $(y_1(u), y_2(u))$  or  $(z(u), y_1(u), y_2(u))$  rather than specific numerical responses  $(y_1, y_2)$  or  $(z, y_1, y_2)$  which could in principle be chosen before the inventory is taken. The availability of response rules to the Defender is explicitly recognized in both [5; pp. A-2, 3] and [7; p. 17], but then the analyses of the models go on to treat  $(y_1, y_2)$  or  $(z, y_1, y_2)$  as numbers rather than functions of  $u$ . It is suggested in [5; p. A-4] that this restriction is made to match the operational nature of an "alarm" concept, which implies a response sensitive to the distinction between "above-threshold" and "below-threshold" ID values. On the other hand, the restriction is introduced in [7; pp. 17-18] with an "explanation" which seems a non sequitur to the writer, and which assesses that this limitation is not damaging to the Defender--an assertion, repeated verbally to the Peer Review Group, which the writer believes erroneous, an impression shared by the Group [14; pp. 22, 38] as noted in Section 4.3.

Before the extension of this strategy-space concept to multi-period and/or multi-site games is attempted, its improvement for the "simple" case should be considered. We refer in particular

to the Peer Review Group's recommendation [14; p. 41] that "multiple action criteria for varying amounts of ID should be incorporated."

One step in that direction would be to re-analyze the models of [5-7] without the restriction noted in the next-to-last paragraph. It would be interesting to learn whether this alone (i.e., without changes in the payoff function) would correct some of the counter-intuitive model outputs criticized by the Peer Review Group [14]. At any rate, the separation of all possible ID-levels into just two classes (by a single alarm threshold) seems too coarse a classification to afford the Defender a full exercise of his safeguarding capabilities. Models with more than one alarm-level should be considered; it is hoped that their formulation could be linked with a natural hierarchy of increasingly urgent responses (in terms of degree of extent of shut-down, involvement of external agencies, etc.). Alternatively, one can imagine dispensing altogether with the concept of a discrete set of alarm levels: if for example the intensity of a Defender response could be adequately expressed--for purposes of payoff modeling--in terms of a single quantity  $y$  (e.g., the target-level of a search-recovery effort), then a pure strategy for the Defender could be described simply as a response rule function  $y(u)$  expressing how response-intensity ( $y$ ) would vary with ID-level ( $u$ ). But perhaps high-intensity and low-intensity responses would have such qualitatively different features as to preclude unified representation through a single quantity  $y$ . Indeed, the variety of responses appropriate for inclusion in an operational model may not lend itself to a merely uni-dimensional depiction along a scale of "intensity" or whatever.

We have spoken above of "search and recovery" operations, with an implication that their effectiveness would be judged by success in regaining or uncovering missing material. But for extension to a multi-period model, it would also be important to take into account the extent to which different Defender responses in one time period might inhibit the Diverter's opportunities in later periods, perhaps by identifying and apprehending "him," or by imposing stricter practices that could close off some diversion loophole. Some recovery-oriented steps, by destroying the authentic environment in which a possible diversion took place, might hamper identification of a culprit.

This brings us back to a point mentioned earlier: the need to describe Defender responses concretely enough that mathematical expressions for their cost and effectiveness could be developed for payoff-function use. At the outset of the present study, the writer expected to encounter documents that would greatly aid subsequent modeling in this regard, documents that would set out "standard operating procedures" (SOP's) for post-alarm situations much more specifically than such brief generalities as "one would do A and B, and in extreme cases even C." (For example, the "which might include" of [14; p. 22, 24] falls



short of the "specific investigative action" of the same document's p.22.)

That expectation has not been realized. There are several possible explanations for this. One is that inadequate interrogatory zeal and persistence and document-persual efforts were applied. A second is that well-articulated plans and procedures do indeed exist, but the present study's "need to know" was deemed insufficient to warrant access to information of such potential value to an adversary. (Such considerations are noted, for example, by Willrich and Taylor [197; p. 126, also pp. 152, 89]; the analogous inhibition for studies of terrorism is observed by Wardlaw [222; pp. ix-x].) If this is the case, it should be taken into account in setting up the "logistics" (here, security arrangement) for an effort to develop an operational game-theoretic model.

However, the explanation just mentioned was not in fact offered to the writer, who therefore remains unable to eliminate a third possible explanation: that generic SOP's at the NRC level have not been formulated. Perhaps individual licensees have detailed contingency plans which have been presented to the NRC in advance for approval; perhaps procedures are improvised for individual incidents, and presented to the NRC for approval before or after their execution; quite likely, procedures are in part built up incrementally by precedent and experience with past (presumed) false-alarm situations. But the writer is not aware of a systematic analytical basis adopted by the NRC for consistent evaluation of possible operator-proposed response plans, or for suggesting alternatives or improvements to such proposals. (pp. VI-4 through VI-19 of [3] may be relevant.)

Such a situation is consistent with the one described by Avenhaus [187; p. 322] as applying to the international scene--following an alarm, a "second action level" should come into play, but for this "there are not precise procedures, at least for the case of nuclear material safeguards." And it also matches well an observation recurring in the recent literature on "technological accidents" and their management: that after so much in the way of resources and dedication have been devoted to the prevention of such accidents (here, "diversions"), too little may be mustered for carefully planning and preparing the responses in case the unhappy event does occur. (See, e.g., Fischer [224; pp. 11-12], Lathrop [226; pp. 8-9]. Such advanced planning (and practice, and maintenance of readiness) is regarded as particularly important in the "so many unexpected things can go wrong" situations when responses may need to be made in atmospheres of uncertainty and stress, where special issues of communication with the public media may arise, and where coordination is required among organizations with unclear demarcations of authority and responsibility. (Cf. Fischer [227], Lathrop [228], Marrett [230], and many of the other papers in [223, 225] and

Moss and Sills [229]; related concerns and recommendations appear in Jenkins [232; pp. 17, 22], Macnair [233; p. 274], Bass et al [234, pp. 10-11], Sloan [235].) These exacerbating conditions might well apply to a serious ID-alarm incident.

The representation and modeling of Defender strategies might perhaps be able to benefit from the fairly well-advanced field of Optimal Search Theory. The early "classics" of that field have been conveniently collected and revised (Koopman [236]); other references include a prize-winning 1978 monograph by Stone [237] and a recent survey paper by the same author [238]. Complications arising from target movements may not prove relevant for the desired applications: cf. Chapter 2 of Gal [239]. "Search and Surveillance" is now an indexing term for Operations Research/Management Science purposes.

## 7.0 THE DIVERTER'S PAYOFF FUNCTION

A player's payoff function, in a game-theoretic model, should provide a mathematical expression of how the various possible strategy-choices would affect the degree of satisfaction of that player's preferences and objectives. It is therefore useful, in beginning this Section, to attempt in broad qualitative terms to list plausible general objectives for the Diverter. We take these to be the following:

- To divert material--either as much as possible, or some critical quantity.
- To avoid detection or alarm-sounding.
- To leave some basis for a later "hoax" claim--of a successful or "large" diversion when in fact none or a much smaller one occurred.

The threat associated with the third member of the list, "hoaxing" is frequently mentioned in connection with diversion and material accounting. Examples include Messinger [1; pp. ix, 14], NUREG-0450 [3, v.2; P. IV-13], Mengel [240; pp. 218-220], Willrich and Taylor [197; p. 123], Bass et al [175; pp. 2, 8], Jenkins [241; p. 10]. Nevertheless, none of the cited modeling papers addresses this objective explicitly. While it could conceivably be left for attention in a higher-level "threat-negotiation game" (cf. for example Jenkins [242], Selten [244]), this seems a sorry passing of the buck, since aid in "hoax-proofing" is regarded as an important function of material accounting. The Defender could calculate, from the solutions of some of the cited models, a corresponding miss rate (probability of failure to detect or "alarm at" a diversion), or an "expected value" for material diverted without an alarm, and could then point out in reply to a diversion-claim how low these quantities are. Those solutions would come from a game innocent per se of the third objective, but since that is the game the hoaxster is pretending to have played, perhaps this is good enough. Perhaps not, so that expressing this criterion really would require changing the Diverter's strategy space and/or payoff function. The writer has not succeeded in thinking this point through to a conclusion, and therefore "flags" it as possibly needing attention in an effort to form an operational game-theoretic model. From here on, we will generally ignore the third objective.

The second objective, as stated, lumps together three possible Diverter motives which can in principle be teased apart. One is to avoid detection because it may lead to losing back (to the Defender) some or all of the diverted material. Another is to avoid alarm or detection because it reduces chances for diversion in later time periods. The third is to avoid detection

of the diversion because it may lead to "getting caught," with unpleasant consequences.

The first of the listed objectives seems relatively self-explanatory. Its mention of a "critical quantity" illustrates a more general possibility: that the Diverter's value ("utility function") for diverted material might well vary nonlinearly with the quantity diverted. Some scenarios might impose "time pressure" on the Diverter, assigning greater utility to diversion in an early period than to diversion of the same amount in a later time-period--none of the cited models have this particular feature, though those imposing a fixed or minimum diversion quantity implicitly introduce as a deadline the over-all time horizon of the multi-period analysis, raising the question of how this horizon should be chosen. (Such "end effects" issues are common in multi-stage decision models.)

### 7.1 Review of Literature

The roles of the first two listed objectives will be readily recognized as we review the Diverter payoff-functions proposed in the cited models. In the models [190, 191, 12] of Bierlein and of Beinhauer and Bierlein, as well as Hopfinger [13] and Model I of Bierlein [192], the Diverter's payoffs are based on a structure

|   |    |           |       |
|---|----|-----------|-------|
|   | D  | $\bar{D}$ |       |
| A | -c | 0         |       |
| - |    |           |       |
| A | d  | 0         | (7.1) |

where  $c$  is greater than 0 (usually), and  $d$  is greater than 0. Thus "no diversion" ( $\bar{D}$ ) for the Diverter leads to zero payoff. A successful diversion (diversion (D) together with too-late alarm no alarm ( $\bar{A}$ )) yields payoff  $d$ , while a detected attempt leads to the (usually) negative payoff ( $-c$ ). (Some of the cited papers use different symbols, and we have here identified "inspect" with "alarm" and "detect"; modifications to the interpretations of  $c$  and  $d$  can accommodate inspections with imperfect detection probabilities.) It is not hard to see that the strategic analysis of such games, though not their absolute payoffs, depends on the data ( $c, d$ ) only through their ratio  $c/d$ --one can think of measuring payoffs with " $d$ " as revised payoff-unit--thus reducing the number of parameters in the model. (Choosing the normalization  $c + d = 1$  often appears convenient for the particular algebra involved.

The papers just cited all propose zero-sum games. This might raise an objection: since display 7.1's second column does not "credit" the Diverter for causing a false alarm, the zero-sum formulation will not "penalize" the Defender for false alarms. However, the Defender is kept from excessive "alarming" or inspecting by explicit limitations on his strategy space; see Section 6.2.

A technical interjection is convenient at this point. Let  $M$  denote the Defender's miss rate, i.e. the probability of a violation (diversion) going undetected or unalarmed. (We have a success-failure dichotomy in mind; depending on the specific scenario, possibly multi-period and/or multi-site, "a violation" might read "at least one violation" or "every violation.") Thus  $M$  depends on the strategies chosen by the two players, i.e.  $M = M(s_1, s_2)$ . Under the assumption 7.1, the expected payoff to Player 2 (the Diverter), if diverting, is

$$Md + (1 - M)(-c) = (d + c)M - c. \quad (7.2)$$

Since  $c$  and  $d$  are constants, if  $d + c$  is greater than 0 then maximizing the above expression is equivalent to maximizing  $M$ , i.e., replacing the Diverter's payoff function by  $M$  yields a game "strategically equivalent" to the "if diverting" subgame of the original one. And if the model is zero-sum, a similar replacement can be made for the Defender. This could make it unnecessary to arrive at values for  $c$  and  $d$ , much simplifying the modeling task. More generally, the Diverter would have to compare the value of 7.2 using the maximized  $M$ , with the payoff 0 of his "don't divert" strategy; cf. the Lemma of [190; p. 60].

Dresher's zero-sum repeated game [193] also employs the structure 7.1, but (up to its Section 8) with the further assumption  $c = d = 1$ . The generalization by Kuhn [209] passes directly to the use of miss-rate  $M(s_1, s_2)$  as the Diverter's payoff, as do those of Anscombe [208] and Schleicher [217]; these models are also zero-sum. Avenhaus and Frick [9, 10] pass from 7.1 to the use of miss-rate.

In the models of Goldman and Pearl [170, 171] and the multi-site Models II and III of Bierlein [192], a payoff structure like 7.1 is assumed at each site (with site-specific values of  $c$  and  $d$ ), and the total payoff to the Diverter is obtained by summing these "local payoff functions" over the sites. Here the replacement of the payoff function by a single "miss rate" probability is no longer valid. Model I of [170], and those of [171, 192], are zero sum.

A more general payoff structure than 7.1 is given by

|   |    |           |       |
|---|----|-----------|-------|
|   | D  | $\bar{D}$ |       |
| A | -c | -f        |       |
| - |    |           |       |
| A | d  | 0         | (7.3) |

where  $c$  is greater than or equal to  $f$  is greater than 0. Here involvement in a false alarm occasions a negative payoff ( $-f$ ) to the Diverter; this is suggestive of the second broad Diverter's objective listed at the start of Section 7.0. The format (7.3) is employed, for example, by Avenhaus [187] with  $c$  greater than  $f$ ; it also applies to the models of Maschler [211, 212] with the further condition  $c = f$ , after a mathematical operation (subtracting a common constant from each matrix entry) which yields a strategically equivalent game. The models just cited are not zero sum; the same is true of the "insurance-cheating" models of Borch [203, 204], which fall under (7.1) rather than (7.2).

The models [5-7] developed for the NRC are all zero sum. Thus their (common) payoff function for the Diverter is simply the negative of that for the Defender, which will be described in Section 8. It is considerably more concrete than those described above, but that very specificity has opened it to more detailed criticism [14; p. 34].

## 7.2 The Zero Sum Assumption: Alternatives

Imagine for the moment that the field of "stochastic games with incomplete information," wished-for at the end of Section 2.6, had already achieved a satisfactory conceptual, theoretical, and computational status. Besides that pipe-dream about the state of the art, imagine further (now on the empirical side) that we had a reliable inventory of possible "adversary types," could estimate the probability with which each of these types gave a proper classification of the Diverter, and understood each type well enough to be able to delineate a strategy space and a reasonable payoff function for it as potential Diverter. Then we would have at hand both methodology and information sufficient for the relatively easy development of a relatively non-arguable treatment of the Diverter in an operational game-theoretic model.

Obviously, the real situation falls far short of that ideal-- this is to be expected in almost any serious decision problem with a significant behavioral aspect. In particular, the development of a payoff function for the Diverter will involve considerably more in the way of approximation and arguability than was true for the rosy picture painted above. For reasons to be discussed below, the writer believes that the best approach to these difficulties is to

- Adopt a zero sum model; concentrate on developing a "good" payoff function for the Defender (that for the Diverter will be taken as its negative).

The fundamental arguments supporting this model-development tactic are that (a) efforts to do significantly better appear likely to be futile, while (b) in its own right, this approach has much to recommend it. The models developed for the NRC [5-7] adopt this tactic without much ado, and that choice was not criticized by the Peer Review Group [14], but the writer nevertheless thought it important to explore this issue more extensively.

To begin the "futility" arguments, we may note that one element of the "dream scenario" sketched earlier comes close to actuality: existence of a plausible inventory of "adversary types." Such typologies appear implicitly or explicitly for example, in Lovett [213] and in Willrich and Taylor [197; Chapter 6]. Perhaps the most extensive analyses are those performed by the RAND Corporation for the Sandia National Laboratories, and drawing on related RAND research dating back to 1972. The relevant documents include deLeon et al [176], Bass et al [175, 244-246] and Jenkins [177]; they deal with attributes (i.e., capabilities), motivations and possible actions of potential adversaries of U.S. nuclear facilities and programs, of course including diversion among the "possible actions." A principal and persistent conclusion [245; p. v] is that:

"Nuclear defenders must anticipate a surprisingly wide range of threats from an equally wide array of potential adversaries, who may be animated by ideological, economic, or personal motivations, or some combination of these."

The existence of this "surprisingly wide array of potential adversaries" supports the conclusion that it would be futile (or at least inadvisable, as an initial tactic, to seek to develop a correspondingly wide array of Diverter payoff functions in formulating an operational game-theoretic model. A single mathematical form for such a payoff function, with different adversary types representable by different settings of the function's parameters, would of course be desirable but seems only a "long shot" possibility in view of the diversity of motivations involved. Specifically, the categories of adversaries explicitly associated

with diversion in [245; pp. 72-73] and reiterated in the later [246, p. 56], as identified by motivation, are

- hostile (i.e., disgruntled) employees
- psychotics
- individuals acting for idiosyncratic reasons
- mercenaries or foreign agents
- occasional or novice criminals or opportunists.

However, the distinction between "diversion" and "theft by stealth" in these documents (see, e.g. [175, p. 7] and [245; bottom of final fold-out]) limits the former to efforts involving attempts to alter records; this limitation, signalled only in the indicated "fine print" of the documents, is too restrictive for our purposes. Making the necessary correction adds the further categories

- political terrorists
- antinuclear extremists
- philosophical or religious extremists
- professional criminals.

We will not attempt here to summarize the documents' extensive discussions of these adversary types and the probable appeal to them of various actions (emphatically including hoaxes and "faked diversions"). But it is worthwhile to record some explicit implications of those discussions (see all Reinstedt and Westbury [247]):

- Very few possibilities can be confidently ruled out (e.g., [246; p. 46], though some useful judgments of likelihood can be made. Most of them have already occurred [245; p. 76].
- Despite initial and follow-up personnel clearance procedures, "insiders" can fall (or, over time, come to fall) in any of the nine categories above.
- Non-hostile employees may be coerced into acting on behalf of terrorists or professional criminals (the corresponding payoff function should then reflect the interests of the coercer).
- Professional criminals might act "on commission" for foreign agents or terrorists.

These observations further confirm the inadvisability (except as a last resort following failure of other approaches) of setting out to formulate a whole array of Diverter payoff functions for the different adversary types.



An alternative is to find a good reason for singling out some one of the adversary types, and to attempt to develop a well-based payoff function for that type. It seems to the writer that the strongest motivations underlying the safeguards program reflect particular concern that misappropriated SNM (or its pretension) would be used to engender "terror" of nuclear destruction as a basis for some threat or extortion. (Passage of such material through a "foreign agent" to his government is also a heavy concern, but might often reflect fear that the recipient nation would utilize it to support covert "terrorist" activities as surrogates for traditional military confrontations, cf. Dror [248], Jenkins [249-251].) So it seems reasonable to single out, among the various adversary types listed, "the terrorist" as object of payoff-function construction.

There is no lack of literature bearing at least peripherally on such an effort. During the present study, the writer was dismayed to discover the extent to which "terroristics" has become a "growth industry," with its own journal (Terrorism: An International Journal) since 1977, no dearth of monographs and specialized conferences (e.g. Jenkins [231], Crenshaw [252]), and even discussions of data-base duplications and inconsistencies (Fowler [253]), of agendas for quantitative research (e.g. Fowler [254]), and of the use of artificial-intelligence "expert systems" (Waterman and Jenkins [255]). One indicator of this fashionability is the appearance (possible more than once) of Brian M. Jenkins--head of the RAND Corporation's research in this area--on Michael Jackson's nationwide radio-interview program. And the students of my own University have selected Terrorism as the topic for the next of our Milton S. Eisenhower Symposium lecture series, a local "major event."

Exploring the literature shows quickly that specializing from a "general" Diverter to a "terrorist" does not resolve all questions of identification. Problems of definition (what is the range of behaviors that should be labeled "terrorism"?) appear genuinely sticky and can become value-laden--for example (cf. Schelling [256; 49-50]) why does the superpowers' use of the nuclear-deterrence threat fail to qualify? These problems are worried at length by many authors (e.g. Dror [257], Wardlaw [222], Wilkinson [258], Paust [260], Jenkins [261], Devine and Rafalko [263], Taylor and Vanden [264]), with varied results and with (it appeared) different mixes of dutiful exasperation and intellectual pleasure. The purposes of many of these authors and of others (e.g. Shultz [265], Barres [266; pp. 88-92]) leads them further, to taxonomies of terrorist groups. The different "cells" of these classification schemes (six of them in [266]) for example in turn suggest need for different payoff functions, whose "bridgeability" by a common mathematical form (through use of different parameter-values) might prove a difficult,

and conceivably unsolvable, modeling challenge.<sup>23</sup> So narrowing the Diverters to "terrorists" is still not enough to present a clearly-defined adversary type for payoff-function synthesis.

Suppose however that we could arrive at such a type. Recall that in a game-theory model, the players are assumed to be rational payoff-maximizers. Now, Wilkinson [258; p. 127] notes the roles of "hatreds and fanaticisms", sometimes deliberately fostered, in encouraging terrorist violence. May [267] identifies the "ecstatic element" of emotional satisfaction derived by some perpetrators of terrorist acts. Jenkins [268; p.10] mentions a "lunatic fringe" but is skeptical of its effectiveness; however his colleagues Ronfeldt and Sater, in a fascinating study [269] of the "dynamite terrorism" of the late nineteenth century as as a plausible analog for possible nuclear terrorism today, note the theme of millennial redemption through apocalyptic destruction. And Jenkins elsewhere (e.g. [270; p. 4]) describes terrorists as living in a "fantasy world" and waging "fantasy wars," while Fried [271; p. 120] pictures many of them as "functioning on a psychotic level, as attested by delusional thinking and cognitive malfunctioning."

All this may not appear too compatible with game theory's "rational optimizer" picture. But we are reminded that the greatest extremes may not be typical of the more capable terrorist groups which are of primary concern (e.g. [271; pp. 120-21]), and that apparently bizarre statements and behavior may in fact be well-suited to terrorists' need to capture public attention (e.g., Jenkins' often-quoted "terrorism is theater" [272; p.3], Alexander [273]). More important is the general point (Norton and Greenberg [274; pp. 6-7]) that "rationality" is properly defined only relative to a particular set of values and perceptions, so that it is inappropriate (and risky) to regard terrorists as necessarily "mindless" (Jenkins [275; p. 3]), as consisting of "the less intelligent or less able" (Wilkinson [258; p. 132]); or as lacking in dedication, ability for careful planning and operations, possible technical sophistication (including use of computer simulations), and all-around ingenuity ([175; pp. 16-18], Hutchinson [276; p. 158], Mengel [240; p. 192]).<sup>24</sup>

---

<sup>23</sup>In speaking of representation by a "common mathematical form," we mean something more useful than the "cheap trick" that reproduces any two functions  $F(x)$  and  $G(x)$  from the formula  $(1-t)F(x) + tG(x)$  by setting the parameter  $t$  at 0 and 1 respectively.

<sup>24</sup>The opposite extreme, namely the combination in a single adversary of high levels of all the "dangerous" capabilities, was viewed in 1978 by deLeon et al [175; pp. 50-53], as "unlikely," but with cautions to the reader.

Assume then that the terrorist-Diverter can be viewed as "rational" relative to a framework of rationality--a "mind-set quite different from yours or mine or the Defender's. There remains the question of whether that framework can be fathomed well enough to provide the basis for a well-grounded Diverter payoff function. It is interesting to observe the rather rapid transition from Jenkins' 1978 characterization [277] of the terrorist mindset as an "area of ignorance" to the "Satisfied that we can depict the full range of motives and possible actions" of 1980 [174; p. 5], though this contrast reflects an imperfect matching of contexts. It is certain that research on the topic has been intensive, and has significantly increased the factual information available and the level of theoretical analysis possible; see for example Kellen [278], Sundberg [279], Jenkins [231; pp. 12-15, 52-69] as well as [269]. Related studies in the criminological field may prove helpful (cf. Carrol [281]), though the preponderance of crimes do not seem appropriate analogs.

Thus the opinion quoted by Norton and Greenberg [274; p. 13], that "there is no way of studying terrorist ideology in any meaningful way," appears too pessimistic. There is distinct progress towards answering some of the questions listed by Barres [266; pp. 11-13], though the dryly critical literature-review in Section 1 of [277] should be a corrective to premature confidence. The present level and nature of insights in this area might, in the hands of an imaginative modeler, prove useful in suggesting possible general structures for a Diverter payoff function, and could very probably be of value in conceptual testing of a proposed payoff function. More concrete utilization, for constructing such a function, at least using the kind of methodology envisaged in Volume I of this report, would involve estimating a multiattribute utility function; a "classic" mid-1970's vintage account of the relevant theory and procedures, along with selected applications and a major case study, is given by Keeney and Raiffa [282], with some more recent material appearing in the special journal-issue [283]. A critical part of this technique requires ascertaining the preferences of "terrorist decision-makers" between a number of pairs of "pure" alternatives and probabilistic mixtures of alternatives. It seems unlikely (to the writer) that sufficient information of this kind could be inferred from the available writings of these persons and from what is known of their past choices. Indeed, the modeler's frequent procedure in such an analysis involves subjecting the relevant decision-makers to interviews and questionnaires especially designed to elicit the information.

This last notion is not quite so totally ridiculous as it may initially appear. Terrorists and ex-terrorists have granted extensive interviews, and in some cases written their memoirs, so that accessibility might conceivably be achieved. But the issues of veracity and validity do not seem satisfactorily resolvable. It is hard to see why active terrorists would choose

to assist such a study with "honest" answers, and a safeguards model significantly dependent on responses even from "reformed" terrorists--cf. our earlier citations of terms like "fantasy," "psychotic," "delusional"--might not inspire much confidence.<sup>25</sup> The sample of informants would be small; apprehendees and recusants might be distinctly unrepresentative of the wider "population" in question, and in many cases might not have belonged to the decision-making "opinion leaders" elite.

### 7.3 The Zero Sum Assumption: Pros and Cons

The preceding section contained a rather extended discussion of the feasibility of developing "genuine" Diverter's payoff functions. The writer's conclusion from that discussion is that chances for success, relative for example to the "viability" criteria formulated in Section 4.2, are too dubious for this to be the approach of (initial) choice in setting out to construct an operational game-theoretic model. That is especially true in the presence of a much more attractive alternative, namely the zero sum approach mentioned above. We now proceed to offer reasons for regarding this approach as attractive in its own right, while also noting some provisos and limitations.

First and foremost among the affirmative reasons is the notion that, even though the interests of Defender and Diverter will not be "precisely opposite" in the mathematical sense expressed by the zero sum condition (e.g., as in the "I win, you lose" setting of many recreational games), nevertheless the fundamental relation of the two players is one of opposition. Thus the zero sum assumption, in capturing the "essential nature" of the underlying situation, cannot go wrong too badly. This reflects the natural "get the first-order effects right" priority-philosophy of most applied mathematical modeling. Such talk of "fundamental" and "essential" is nonrigorous and subjective; the writer hopes, however, that others would concur with the point just made.

A second reason is the prudential nature of the zero sum approach. By conducting its analysis versus a hypothetical "maximally inimical" adversary, it protects the Defender from the consequences of possibly guessing wrong about the extent and the particular way in which the actual Diverter's interests might not be totally opposed to his own. (The preceding discussion of adversary types does not encourage expectations of "guessing right.") Such conservatism may well be the appropriate stance in a regulatory setting, especially in view of the preceding

---

<sup>25</sup>This sentence may fail to do justice to the role of sensitivity analysis, and to the precautions presumably developed by practitioners of multiattribute utility analysis for dealing with potentially hostile or unreliable respondents.

paragraph's suggestion that not too much bias would be introduced. (We reiterate that the Peer Review Group [14], though explicitly sensitive to symptoms of possible over-conservatism in the models [5-7] developed for the NRC, made no criticism of those models' being zero sum.) To avoid misunderstanding, it should be noted that the characterization of zero sum analysis as "worst case" is correct only insofar as the Diverter's payoff function is concerned; it does not involve any "conservative" expansion of the Diverter's supposed capabilities (i.e., strategy space).

Nor is it true that this conservatism is of a kind automatically leading to a solution in which alarms are more frequent than they would otherwise be. Suppose for example that the Diverter's "true" payoff function were given by the previous matrix 7.1. Suppose further that the Defender's payoffs are given by the matrix

$$\begin{array}{c}
 \phantom{A} \\
 \phantom{-} \\
 A
 \end{array}
 \begin{array}{cc}
 D & \bar{D} \\
 \hline
 \begin{array}{|c|c|}
 \hline
 c'-a & -a \\
 \hline
 -d' & 0 \\
 \hline
 \end{array}
 \end{array}
 \quad (7.4)$$

where  $c'$  is the Defender's analog of  $c$  in (7.1), i.e. the value placed by the Defender on the occurrence of a detected diversion, while  $d'$  is the Defender's analog of  $d$ , i.e. the loss suffered by the Defender from an undetected diversion. The quantity  $a$  represents the cost to the Defender of executing the responses to an alarm; we will assume that  $a$  is less than  $c' + d'$ , a non-restrictive condition which is satisfied if merely a successful diversion is more costly to the Defender than is a false alarm. Note that the differences in structure between (7.1) and the negative of (7.4)--the payoff matrix attributed to the Diverter by the zero sum assumption--simply reflects the idea that the Diverter is "really" indifferent to the costs imposed by "alarming" on the Defender, a fairly plausible idea unless the adversary is an anti-nuclear extremist out to bankrupt the facility.<sup>26</sup> Under these circumstances, the "true" non-zero sum game turns

<sup>26</sup>One could imagine an adversary whose aim was not so much diversion, as the damaging of U.S. weapons programs (or long-term economic health, or energy-independence) through interruption (or cost escalation) of nuclear-material operations. But that is a rather different scenario from those under consideration here.

out to have a unique equilibrium-point solution, in which the mean relative frequency of alarms is  $1/(1 + c/d)$ . For the zero sum version, this frequency becomes  $1/(1 + c'/d')$  in the unique optimal strategy for the Defender. (The value of  $a$  affects the strategies for the Diverter in the solutions of both games, but not those of the Defender.) Thus the zero sum approach could either increase or decrease the frequency of alarms, depending on the relative sizes of  $c/d$  and  $c'/d'$ . These results for a very simple model may not be indicative of those for more realistic cases, but at least warn against accepting apparent "consequences" of zero sum modeling without explicit analysis.

A third reason favoring the zero sum approach is its immense easing of conceptual and computational burdens: as is implied by the material in Sections 2 and 3, adopting this approach leads to a highly convincing "solution" concept, unique solution payoffs, and access to a far superior body of algorithms for numerical solution and sensitivity analysis.

Fourth, the approach has whatever virtues inherent in precedent: it yields the "classical" and most familiar type of game-theoretic model, and has been adopted in a number of the safeguards-related analyses cited earlier.

Fifth, most of the above-mentioned difficulties in fashioning a "genuine" payoff function for the Diverter become much more manageable when attention is shifted to the Defender. Thus the prospects for a well-grounded Defender's payoff function are (relatively) good; since in the zero sum approach the assumed payoff function for the Diverter is directly based on (specifically, is the negative of) that for the Defender, and is obtained from the latter by a process with a substantial rationale involving the four reasons already given, it would "inherit" the latter's "well-groundedness" in a way consistent with the viability criteria noted in Section 4.2.

There are of course arguments against the zero sum approach, additional to those intimated in the course of the previous discussion. One is the existence of countervailing precedent: as already noted, several of the models in the literature are not zero sum. Second, there is the intellectual discomfort (and argumentative disadvantage) in imposing a strong hypothesis--that of a "zero sum adversary"--which is not believed or expected to be literally true. This objection, however, may unduly depreciate the natural roles of approximation and tractability in applied mathematical modeling. Third, the literature contains some explicit rejections of this approach. For example, we find in Kupperman [284; p. 411] "The 'game' between terrorists and government is not zero sum"--but that passage in fact refers to the "threat-and-negotiation" situations mentioned near the start of Section 7.0, rather than to our "ID analysis" context. Avenhaus [221; p. 320] observes "this does not lead necessarily

to a zero sum game," but after noting difficulties, quickly goes on to a zero sum model. He cites the possible nonequality of what we called  $d$  and  $d'$  in matrices (7.1) and (7.4) above, and with Frick in [285; p. 630] points out differences in the two players' evaluations of false alarms as a flaw in zero sum treatment. The objections by (Joint Committee on Atomic Energy member) Hosmer [286; p. 7-8] are not so clear to this writer, but their gist seems to be that the zero sum equating (in effect) of  $d'$  with  $d$  may be "very dangerous in the real world populated by very fallible people, some of whom are very certain to be just no good."

On balance, it seems to the writer that "the Ayes have it" concerning the adoption of the zero sum approach. The premises for that judgment have been laid out in a way which, it is hoped, will facilitate tracing out reasons for agreeing or disagreeing.

## 8.0 THE DEFENDER'S PAYOFF FUNCTION

A central problem in any game-theoretic application is the formulation of payoff functions which achieve an "appropriate" compromise between detailed realism and analytical tractability. See for example<sup>27</sup> Shepard [288; pp. 378, 383-4], or the remark of Fain and Phillips [289; p. 370]:

"The trick then is to get the most in realism compatible with obtainable, understandable answers. It is necessary that one restrict the factors to the most essential. This trade-off is easy to understand but hard to make."

Such generalities of course pertain to any kind of applied mathematical modeling. What makes them especially acute for our situation is that the referent "reality" involves preferences and value-schemes--judgmental and behavioral elements--in addition to more tangible elements (e.g., yields, direct costs) of kinds more readily based on "hard" technical data. The quantitative treatment of the former elements by the behavioral and decision sciences strikes many observers as less advanced and reliable than available treatments of the latter elements by the physical sciences and their related engineering disciplines.<sup>28</sup>

This is not to denigrate the acceptance and usefulness of analyses with such orientations in many fields; e.g. economic policy, market research, and the regulatory contexts cited in Section 4.2. But it does re-emphasize, in the present setting, the particular challenge of developing suitable payoff functions. If the zero sum approach recommended in Section 7.2 and 7.3 is adopted, then this challenge becomes focussed on the Defender's payoff function. In terms of both priority and intrinsic difficulty, addressing this challenge should precede detailed work on computational solution methods, in a "staged" approach to developing an operational game-theoretic model.

### 8.1 Review of Literature

A number of the cited models are zero sum, with their Diverter's payoff function already specified in Section 7.1. Thus their

---

<sup>27</sup>These citations from the older literature, though still germane, do not reflect the last two decades of progress in Game Theory.

<sup>28</sup>This wording is intended to bypass the interdisciplinary and intradisciplinary disputes about whether the behavioral sciences can in principle attain, and should properly aim at, the same kinds of "success" achieved by some areas of the physical sciences.



Defender's payoff function need not be spelled out here. These models include those of Beinhauer and Bierlein [12], Bierlein [189-191], Hopfinger [13], Drescher [193], Kuhn [209], Anscombe [208], Schleicher [217], Avenhaus and Frick [9, 10], and Goldman and Pearl ([171] and Model 1 of [170]). Typically these payoffs are either of an "abstract" (i.e., relatively uninterpreted) nature, or represent the negative of the Defender's miss rate; those of [170, 171] are the negatives of more concretely-described Diverter payoffs.

As noted earlier, many of the models limit the Defender's strategy-space in a way representing a bound on his resources or permitted intrusiveness. This bound then appears as a parameter in the game-theoretic model. Bierlein [190, etc.] and Hopfinger [13] are principally concerned with optimal sizing of an inspection system, in the following specific sense: finding the lowest value of the bound on Defender (i.e., Inspector) resources under which the only optimal strategy for the Diverter in the resulting game is that of "no diversion." Note that this does not involve a balancing by the Defender of risk and cost: the "acceptable risk" is set at zero, and the lowest cost for attaining it is sought. It is conceivable that the above problem could somehow be solved without solving the game-theoretic model explicitly; this is in fact essentially what is done by Hopfinger, who noted [13; p. 9] that his game-theoretic model proved too difficult for explicit solution.

In the non-zero sum model of Avenhaus [187], the Defender's payoffs are assumed given by a matrix of the form

$$\begin{array}{cc}
 & \begin{array}{cc} D & \bar{D} \end{array} \\
 \begin{array}{c} A \\ - \\ A \end{array} & \begin{array}{|c|c|} \hline -b' & -a \\ \hline -d' & 0 \\ \hline \end{array}
 \end{array} \tag{8.1}$$

where notation has been changed to better match the previous matrix 7.4. It is assumed that  $b'$  is less than  $d'$ , which coincides with the weak restriction  $a$  is less than  $c' + d'$  made above in (7.4). After subtraction of a common constant from all entries, the same is true of the models of Borch [203, 204], where the entries have fairly concrete interpretations in terms of insurance premium reductions, cost of loss-reduction measures by the insuree to justify the reduction, etc. The model of Maschler [211, 211] fall into the same class, but with  $a = 0$ . The non-zero sum Model 2 of Goldman [170] adopts as Defender's payoff function the negative of the Diverter's gains from undetected "violations";

this corresponds to setting  $a = b' = 0$  and  $d' = d$  in (7.4) at each site. Unique equilibrium-point solutions exist for [187, 203, 213], and apparently for Maschler's model [211; p. 25]. Model 2 of [170] has a multiplicity of equilibrium points, but the Diverter has the same payoff (though different strategies) in all of them; the payoffs to the Defender also remain the same except under a special "degeneracy" of the model's data.

Model 3 of Goldman [170] deviates from the usual game-theoretic format to study the following scenario: the Defender chooses a strategy (mixed, in general) that will maximize his payoff, under the assumption that the Diverter learns of that strategy and reacts so as to maximize his payoff. Maschler [211, 212] also employs this concept, which apparently arose in economics in studying price competition by a duopoly; hence the "Price Leadership" in the title of [211]. Such "leader-follower" or "Stackelberg" games (von Stackelberg [290]) have received considerable study in recent years; Basar and Oldser [291] is a current treatise containing a substantial treatment.

We turn finally to the most explicit and ambitious effort to develop a Defender's payoff function, that in the zero sum models [5-7] formulated for the NRC. It uses the following notation, previously introduced in Section 6.2:

- $x$  = quantity taken by Diverter,
- $e$  = random error in estimating ID
- $u = e + x$  = ID-value reported to Defender,
- $z$  = alarm level
- $y$  = target level of search-recovery efforts by Defender  
( $y = y_1$  if  $u$  is less than or equal to  $z$ ,  $y = y_2$  if  $u$  is greater than  $z$ ).

Thus the Diverter's strategy is given by  $x$ . The Defender's strategy consists in choosing the values of  $y_1$  and  $y_2$ , and in the second model of [7], choosing  $z$  as well. It is convenient to designate the two possible scenarios--"no alarm" corresponding to  $u$  is less than or equal to  $z$  and to  $y = y_1$ , and "alarm" corresponding to  $u$  is greater than  $z$  and to  $y = y_2$ --by an index  $k$  taking the values 1 and 2 for these respective scenarios. The Defender's payoff, under scenario  $k$ , is taken as the negative of the cost-function

$$M_k(x, y_k) = B(k-1) + c_k y_k + x - b_k \min(x, y_k) + e_k |y_k - x_k|. \quad (8.2)$$

Here  $B$  is the fixed cost of a "clean-out inventory," incurred regardless of the target level of the search; note from (8.2) that it is incurred only in case of an alarm ( $k=2$ ). The term  $c_k y_k$  represents the variable part of the cost of the search, with  $c_k$  determinable [7; p. A3] "by engineering estimates of

labor and materials involved." Thus the first two terms in (8.2) are to represent the cost of the search effort. Without a more concrete submodel of the search process (the need for which was noted in Section 6.2), it is unclear that the linearity of the second term--and its independence of  $x$ --are appropriate. In general, these terms seem a very narrow construal of the disutilities (interruptions, relations with outside authorities, effects on confidence, etc.) associated with some "alarm" responses.

Under the given scenario, a fully successful search would recover the amount  $\min(x, y_k)$ --this explanation partly explicates the intended role of the "target level" in delimiting a search effort--leaving the amount  $x - \min(x, y_k)$  unrecovered. This expression resembles the third and fourth terms in (8.2). The coefficient  $b_k$  may then represent a "search quality measure" indicative of how far from "ideal" the search capabilities are. The authors describe  $b_k$  [7; p. A3] as involving "the value to the Defender--of recovering the material diverted and the probability of recovering it," and note that it depends on "societal values"; it is unclear why the third term,  $x$ , should not also be modified in light of those values. One would probably want to replace these two terms by a function of  $x$  and  $y_k$ , say  $D(U(x, y_k))$ , where  $U(x, y_k)$  is the expected quantity of diverted SNM left unrecovered by the search process, and  $D(u)$  is a function--probably nonlinear--expressing the disutility to society of the loss of the quantity  $U$ . If "recovery" has some value in itself other than reducing the loss, that too should be articulated in the model in a clearly-explained way. Note that if the first and second terms appear naturally in "cost" units, then some means of unit-conversion with what replaces the third and fourth terms is required.

The fifth term on the right-hand side of equation 8.2 is described as the "error penalty" for a wrong estimate by the Defender. This is the first indication that  $y$  is meant to serve as an estimate of  $x$ , not merely as a target level for the search. It is not obvious that the two should coincide; for example a "target level" might well include a "safety allowance" over and above the Defender's "estimate" of  $x$ . At any rate, the Peer Review Group [14; p. 25] found the conceptual basis for this term especially unsatisfying.

In view of the "no-alarm, alarm" interpretation of the two scenarios, the Defender's expected (i.e., mean) cost is taken in [5-7] to be

$$M_1(x, y_1) \text{ Prob } (u \text{ is less than or equal to } z) + M_2(x, y_2) \text{ Prob } (u \text{ is greater than } z). \quad (8.3)$$

The probabilities in question are determined using the relation  $u = e + x$  and the assumption that  $e$  has an unbiased (i.e., zero-mean) normal distribution. Thus the frequency function (i.e., probability density) of  $u$  depends in an explicit way on  $x$ ; denote

it by  $p(x,u)$ . For reasons detailed in Section 4.3 and 6.2, relating to the interpretation of a Defender's strategy as a response-rule to ID-values rather than specific responses--and hence as involving a pair of functions  $y_1(U)$  and  $y_2(U)$  rather than numbers  $y_1$  and  $y_2$ --it appears to the writer that (8.3) might be better replaced by

$$\int_{\infty}^Z M_1(x, y_1(u)) p(x, u) du + \int_Z^{\infty} M_2(x, y_2(u)) p(x, u) du. \quad (8.4)$$

The preceding discussion by no means exhausts the criticisms which can be, and have been [14], levelled at the payoff function based on equation 8.2. Such specific criticisms were possible only because the authors of [5-7] went further than others in basing their formulation as a "semi-explicit" picture of the Defender's responses, and are extremely valuable in providing insights for the development of improved payoff functions even in the "simple" case (one time period, one site). Having examined the "competing" models in the literature to an extent not possible during his participation in the work of the Peer Review Group, the writer now wishes that besides concurring (as he did) with the stated criticisms, he had also associated himself with the gracious observation of Higinbotham [14; p. B-5] that the authors of the NRC-supported models [5-7] "deserve considerable credit for their initiative."

## 8.2 Further Discussion

On grounds of professional experience and expertise, the writer would certainly defer to other members of the project team as regards procedures and pitfalls in developing a Defender's payoff function. This understood, it may nevertheless be worthwhile to offer a few general and elementary remarks on the topic, with a particular view to disentangling some of the issues involved. Such "structuring" also has implications for the organization of a model-development effort.

Given any particular strategy-choices by Diverter and Defender, the assignment of an associated payoff-value involves two conceptually distinct operations:

- (a) estimating the outcome or result of the interaction of those strategies, and
- (b) attributing a value (utility or disutility), on behalf of the Defender, to that outcome.

Operation (a) deals with a "what would happen?" question, and so essentially seeks a prediction, at an appropriate level of detail, accuracy and reliability. Operation (b), in contrast,

is essentially evaluative, focussing on the Defender's preferences and value-scheme. The boundaries between the two may not be as easy to fix as the above language suggests; for example, "how much SNM would the Diverter finally get away with?" is a question<sup>29</sup> that clearly belongs to (a), but though "what would the Diverter do with a particular amount of SNM?" is also a "what would happen?" question, the writer suspects that in the present context it might be better treated in conjunction with (b).

This over-simplified but useful separation into (a) and (b)--prediction and evaluation--is illustrated by the expression  $D(U(x, y_k))$  suggested in Section 8.1. Here the formulation of the "unrecovered portion" function  $U(x, y)$  is an instance of (a), while constructing the function  $D(U)$  is a case of (b). The way in which the former appears "nested within" the latter is also typical for the sequential logic of the situation. (An example of such a separation, in the context of selecting bullet types for police handguns, is given in Hammond and Adelman [293].)

The type (a) work will probably require more explicit submodels, of the mechanisms of diversion and response, than the writer has so far encountered; the need for these has already been noted in previous Sections. Quite likely such work will also need informed judgments about some aspects of "what would happen"; for systematizing the gathering and synthesis of such expert opinions, use of the "Delphi Method" (cf. for example Linstone and Turoff [294], Sackman [295]) may merit consideration.

Some of the type (b) work may yield to relatively straightforward "costing out" applied to the submodels mentioned in the last paragraph. But the greater and most difficult part of it, involving the "weighing" of various consequences relative to each other and to tangible costs, is expected to require the estimation of a multiattribute utility function using the specialized techniques of interview and questionnaire alluded to in Section 7.2. (Again we note [282, 283] as samples of a wider literature; Keeney [296] is very readable.) The combination of these steps and approaches--drawing as available on "hard" data, on techniques employed in other applications and having considerable theoretical basis, and on consultation with and solicitation of judgment from a broad spectrum of informed and concerned individuals--provides the best chance (cf. [14; pp. 35, 41]) for satisfying the viability conditions of Section 4.2.

The term "multiattribute" employed above suggests a different useful "cut" in discussing and developing a Defender payoff function. It goes back to the fundamental issue of the underlying

---

<sup>29</sup>perhaps to be answered only in a probabilistic sense.

regulatory problem, namely balancing the "macro" attributes of risk and cost. If players 1 and 2 (Defender and Diverter) choose strategies  $s_1$  and  $s_2$  respectively, then it is plausible that the associated payoff to the Defender is the negative of a "total disutility" roughly decomposable as:

$$C(s_1, s_2) + R(s_1, s_2) \quad (8.5)$$

where the first term represents the cost of the Defender's responses (alarm, interruption, search, etc.) while the second term measures the disutility to him of the risks associated with the successful abduction of material. For example, the sum of the first two terms in equation 8.2 would correspond to the first summand above, while the sum of the third and fourth terms (or its suggested replacement  $D(U(x, y_k))$ ) would illustrate the second summand in the last display. Some means of expressing the two summands on a common scale, so that they can sensibly be added, is also implicit in equation 8.5

The task of deriving a Defender's payoff function can now be crudely regarded as splittable into deriving expressions for each of the two summands. Operations of prediction and evaluation--types "(a)" and "(b)" above--will enter into both of these subtasks, but the relative role of (b) seems likely to be much heavier for R than for C.

The topics of risk assessment and of risk-cost and risk-benefit analysis have accumulated a voluminous literature of their own, involving both prescription and research. (A journal Risk Analysis was initiated in 1981.) One particularly valuable feature of these and related writings is their identification of pitfalls that non-experts might easily overlook and succumb to in the brisk pace of an applied study. For example, given the Defender's responsibility (Section 5.0) to represent the interests and concerns of various elements of the public, there is useful "sensitization" for the modeler in being reminded of how such elements may differ in their rankings of various attributes (Rokeach [297]), so that one's own values should not unthinkingly be inputted to--or adopted on behalf of--the general public. A related comment by Wynne [298; p. 28] concerns overhasty dismissal by "technical experts" of what they regard as extraneous arguments reflecting "selfish, irrational, ignorant or malevolent" behavior. Einhorn [300] points out how inaccurate judgments can bias experience in ways inhibiting collection of the flaws. In the context of flood-disaster insurance, Kunreuther [301] offers startling empirical findings on the under-regard of low-probability high-impact risks. Chapter 2 of Fischhoff et al [302], synthesizing numerous papers by the same authors, discusses in some detail "five generic complexities" in risk-related decision analysis; see also, for example, Salem et al [303]. Issues concerning the elicitation and use of numerical probabilities in such analyses are noted in [304; pp. 41-3]

and explored in Solomon et al [305]. Divergences between actual and perceived risks (Covello et al [306]) became the topic of the first annual meeting of the Society for Risk Analysis.

In view of the notion (Section 7.2) that terrorists may warrant special concern as an "adversary type," the likelihood and possible consequences of terrorists' "going nuclear" command special interest. The matter of likelihood appears controversial. The jolly-titled [307; p. 89] suggested in 1978 that necessary skills, to date, were "beyond the capabilities of contemporary terrorist organizations." Others have been far less sanguine; earlier, Jenkins [275; pp. 9-10] had opined that "the requisite technical knowledge--will spread," noted the "extreme difficulty" [308; p.3] of assessing this threat, and declared himself [269; p. 8] a "prudent agnostic," stating:

"I don't know whether terrorists will go nuclear, but the consequences if they were to do so may be so serious that society cannot afford to take a chance."

The writer's sense is that the more recent literature, though morbid, grows increasingly pessimistic on this topic<sup>30</sup>, noting greater technological sophistication on the part of terrorist groups, and suggesting that the tactical and ideological constraints inhibiting their pursuit of mass destruction may be eroding under conditions of frustration and generational change (e.g., [231; pp. 63-8], [278; pp. 169-70], [270; pp. 6-8], [250; pp. 1-2], [284; p. 50], [310; pp. 227-8] and the implications of [269]). This leads us to the matter of "consequences," where unclassified references to such "nightmare possibilities" (Wilkinson [258; p. 135]) include the deliberately dramatic introduction to Rosenbaum [311] and the paper Kupperman [312]. A particularly systematic discussion is given by Jenkins [314], who lists (among other possibilities) increased security at all facilities, crackdown on political dissidents, intensified disarmament and anti-nuclear energy debates, and of course the destabilizing symbolism of a first post-Nagasaki nuclear detonation. It also seems worthwhile to note the appraisal by Willrich and Taylor [197; p. 107] that "the damage which might result from a nuclear theft is potentially much greater than the damage that could result from the maximum credible accident in the operation of a nuclear power reactor. Yet another observation, not seen in the consulted literature (perhaps because of its indelicacy), is that the disutility to the Defender of a terrorists' diversion could in principle depend on the identity of the terrorists' likely target--but of course we cannot count on having a Diverter who is "an enemy of our enemy."

---

<sup>30</sup>Heavy reliance on numerous and easily-accessible RAND documents may have biased this impression.

The general tenor of the preceding paragraph is perhaps somewhat alarmist, consistent with the "prudential" attitude suggested in Section 7.3. As partial antidote, we note Jenkins' concern [308; p. 2] over exaggerated "threatmongering." It is also true that exaggerated concern and response can be construed as a yielding towards what several authors describe as a "generalized objective" of many terrorist groups: to incite actions by the authorities which indirectly undermine the latter by causing loss of public support and confidence. (Cf. for example Crenshaw [315], Wilkinson [258; pp. 137-8], Wardlaw [222; pp. 66-9]; in the present context the "undermining" might be more military-economic than political.)

This concludes our discussion of preliminary ideas concerning the development of a Defender's payoff function. It seems important to point out that such a development effort would not be valuable only to (and thus should not be regarded as "chargeable" only to) the construction of an operational game-theoretical model. It would also contribute directly to providing an improved basis for practically any broad-scope analytical attempt at evaluating, balancing and enhancing the material accounting function, and quite likely (by analog and extension) other safeguards functions as well. Such a contribution would represent further progress in the direction exemplified by Bennett et al [316].

We began Section 8.0 with the remark that determining suitable payoff functions was a central problem in applied game-theoretic modeling. We close it, by describing a methodological novelty which has been suggested for evading this problem (at a price!) when it appears insurmountable. This device is mentioned only for "just in case" purposes; recourse to it would be made as a "last resort" and is not expected to be necessary.

The fundamental idea is that if it proves too difficult to assign numerical values to the outcomes of various pure-strategy choices by the two players, it should at least be possible to assign from the Defender's viewpoint a preference ranking--perhaps with ties--to those outcomes. Assume that both players have just finitely many pure strategies (a matter of discretization, if not true at the outset); then analogous to the payoff matrix of Section 2.3, the Defender is confronted with an "outcomes matrix" which his preference ranking can convert into a "preference-ranks" matrix. For any particular probability-level  $p$ , say 0.90, the Defender might then ask "What is the highest mean outcome-rank I can assure myself of obtaining no matter what the Diverter does, and how can I do this?"

This approach's reliance on ranks gives it a purely "ordinal" rather than "cardinal" nature, easing its application but sacrificing much of the information present in even a "fuzzy" payoff matrix. While ordinal considerations had entered game theory in various other contexts (Goldman [317], Shapley [112], Goldberg, Goldman



and Newman [318], Drescher [319]) including solvability in pure strategies and eliminability of some pure strategies as "dominated" by others, the determined effort to create a fully ordinal theory based on the question posed above was made by J. Walsh (with G. Kelleher also contributing), in a series of papers [320-327] concentrated in the period 1969-1972 following Walsh's presidency of the Operations Research Society of America.

Surprisingly, these ideas have gone neglected during the intervening years with the single exception of deVries [328]. Possible reasons include the tremendous prestige and elegance of the "classical" cardinal-value theory, and the location of [320-327] in journals not followed as "mainstream" by much of the U.S. and European research committees. Another likely reason is the algorithmic unattractiveness of the (finite) solution method as presented in these papers. The writer's doctoral student Won [329], in a dissertation currently being completed (deo volens), has applied recent algorithmic developments to obtain more efficient means of calculating the Defender's optimal strategy--in general, mixed--in such a "percentile" game ("median" games are the special case  $p = 0.50$ ), and has modified the underlying model so as to control the probability of a very "bad" outcome. Thus ordinal game theory can be kept in mind as a fallback position if the cardinal approach flounders on the determination of a Defender's payoff function.

## REFERENCES

1. M. Messinger. A Conceptual and Analytic Framework for Safeguards Material Accountability. Report to the NRC Material Control and Accounting Task Force, Office of Nuclear Material Safety and Safeguards, May 1977.
2. J.L. Jaech. Statistical Methods in Nuclear Material Control. U.A. Atomic Energy Commission, TID-26298, 1978.
3. Report to the Material Control and Material Accounting Task Force: Blueprint for the Future. NUREG-0450, April 1978.
4. Mathematical Derivation of the Automated Material Accounting Statistics System (AMASS). NUSAC, Inc., Report 591, prepared for the U.S. Nuclear Regulatory Commission under Contract NRC-02-080-024, 1981.
5. W.E. Siri, H. Ruderman, J. Winsen, and M. Dresner. A Study of Nuclear Material Accounting, Final Report (July 1, 1976-April 1, 1977). NUREG-0290, Vol. 1, 1977.
6. M. Dresner and S. Moglewer. Statistical Acceptance Testing in a Competitive Environment. Oper. Res. 28 (1980), 503-511.
7. W.E. Siri, H. Ruderman, and M. Dresner. The application of Game Theory to Nuclear Material Accounting. NUREG/CR-0490, 1978.
8. R. Avenhaus. Material Accountability: Theory, Verification, Applications. IIASA Series on Applied Systems Analysis, Vol. 2, J. Wiley, New York, 1978.
9. R. Avenhaus and H. Frick. Game Theoretical Treatment of Material Accountability Problems, Part I. Int. J. of Game Theory 5, (1977), 117-135.
10. H. Frick. Game Theoretical Treatment of Material Accountability Problems, Part II. Int. J. of Game Theory 6 (1978), 41-49.
11. R. Avenhaus. Game Theoretical Analysis of Inventory Verification. Contribution to Course "Mathematical and Statistical Methods in Nuclear Safeguards," Joint Research Center at Ispra (Italy), November-December 1981.
12. R. Beinhauer and D. Bierlein. Games Theoretical Models for Inspection Procedure. Pp. 425-466 in Symposium on Safeguards Techniques (Vol. 2), International Atomic Energy Agency, Vienna, 1970.

13. E. Hopfinger. Reliable Inspection Strategies: A Game Theoretic Analysis. Math systems in Economics, Vol. 17, Verlag Anton Hain, Meisenheim am Glan, 1975.
14. C.A. Bennett, et al. A Review of the Application of Strategic Analysis to Material Accounting. NUREG/CR-0950, 1979.
15. E. Borel. Applications aux Jeux de Hasard. Vol. 4, No. 2 of Traite du Calcul des Probabilites et de ses Applications, E. Borel, et al. (eds.), Gauthier-Villars, Paris, 1938.
16. J. Ville. Su la Theorie Generale des Jeux ou Intervient l' Habilité des Joueurs. Pp. 105-113 in Vol 2, No. 5 of Traite du Calcul des Probabilites et de ses Applications, E. Borel et al. (eds.), Gauthier-Villars, Paris, 1938.
17. K.J. von Neumann. Zur theorie des Gesellschaftsspiele. Math. Annalen 100 (1928), 295-320.
18. J. von Neumann and O. Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, 1944.
19. M. Shubik. Game Theory and its Applications. Proceedings of Symposia in Applied Mathematics, Vol. 24, American Mathematical Society, 1981.
20. W.F. Lucas, et al. Game Theory and its Applications. Proceedings of Symposia in Applied Mathematics, Vol. 24, American Mathematical Society, 1981.
21. S. Karlin. Mathematical Methods and Theory in Games, Programming, Economics, two volumes, Addison-Wesley, 1959.
22. M. Dresher. Games of Strategy: Theory and Applications. Prentice-Hall, 1961.
23. R.B. Braithwaite. The Theory of Games as a Tool for the Moral Philosophy. Cambridge University Press, 1963.
24. J.M. Smith. Evolution and the Theory of Games. Cambridge University Press, 1982
25. A.J. Goldman and D.R. Shier. Player Aggregation in Noncooperative Games. J. Res. National Bureau of Standards 85 (1980), 391-396.
26. A.J. Goldman. Player Aggregation in Noncooperative Games, II. J. Res. National Bureau of Standards 86 (1981), 383-391.
27. W.F. Lucas. A Game With No Solution. Bull. Amer. Math. Soc. 74 (1968), 237-239.

28. K.J. Arrow. Essays in the Theory of Risk-Bearing. Markham, Chicago, 1971.
29. H. Markowitz. Portfolio Selection. J. Wiley and Sons, 1959.
30. J. Nash. Non-Cooperative Games. *Annals of Math* 54 (1951), 286-295.
31. G. Owen. Existence of Equilibrium and Continuous Games. *Int. J. Game Theory* 5 (1976), 97-105
32. J.B. Rosen. Existence and Uniqueness of Equilibrium Points for concave n-Person Games. *Econometrica* 33 (1965), 520-533.
33. J. Ponstein. Existence of Equilibrium Points in Non-Product Spaces. *J. Siam Appl. Math.* 14 (1966), 181-190.
34. J.B. Williams. Stability in Noncooperative Games. *Oper. Res.* 19 (1971), 774-783.
35. Wu Wen-tsun and Jiang Jia-he. Essential Equilibrium Points of m-Person Non-Cooperative Games. *Scientia Sinica* 11 (1962), 1307-1322.
36. R. Selter. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *Int. J. Game Theory* 9 (1975), 25-55.
37. R.B. Meyerson. Refinements of the Nash Equilibrium Concept. *Int. J. Game Theory* 7 (1978), 73-80.
38. A. Okada. On Stability of Perfect Equilibrium Points. *Int. J. Game Theory* 10 (1981), 67-83.
39. J. Harsanyi. Rational Behavior and Bargaining Equilibrium in Games and Social Situations. Cambridge University Press, London, 1977.
40. R.J. Aumann. Almost Strictly Competitive Games. *J. Siam Appl. Math.* 9 (1961), 544-550.
41. T. Parthasarathy and T. E. S. Raghavan. Some Topics in Two-Person Games. Elsevier, New York, 1972.
42. H.W. Kuhn. Extensive Games. *Proc. Nat. Acad. Sci.* 36 (1950), 570-576.
43. H.W. Kuhn and A.W. Tucker Contributions to the Theory of Games, II *Annals of Math. Study* 28. Princeton University Press, 1953.

44. H.W. Kuhn. Extensive Games and the Problem of Information. Pp. 193-216 in [43].
45. L.S. Shapley. Stochastic Games. Proc. Nat. Acad. Sci. (U.S.A.) 39 (1953), 1095-1100.
46. T. Parthasarathy and T.E.S. Raghavan. An Orderfield Property when One Player Controls Transition Probabilities. J. Optimiz. Theory and Appl. 33 (1981), 375-392.
47. J.A. Filar. Ordered Field Property for Stochastic Games When the Player Who Controls Transitions Changes from State to State. J. Optimiz. Theory and Appl. 34 (1981), 503-515.
48. O.J. Vrieze, S.H. Tijs, T.E.S. Raghavan and J.A. Filar. A Finite Algorithm for the Switching Control Stochastic Game. Oper. Res. Spectrum 5 (1983), 15-24.
49. D. Blackwell and T.S. Ferguson. The Big Match. Ann. Math. Stat. 39 (1968), 159-163.
51. O.J. Vrieze. Stochastic Games with Finite State and Action Spaces. Doctoral dissertation, University of Amsterdam, 1983.
52. U.G. Rothblum. Solving Stopping Stochastic Games by Maximizing a Linear Function Subject to Quadratic Constraints. Pp. 103-105 in [125].
53. A. Federgruen. Successive Approximation Methods in Undiscounted Stochastic Games. Oper. Res. 28 (1980), 794-810.
54. M. Dresher, A.W. Tucker and P. Wolfe (eds.). Contributions to the Theory of Games, III, Annals of Math. Study 39. Princeton University Press, 1957.
55. D. Gillette. Stochastic Games with Zero Stop Probabilities- Pp. 179-187 in [54].
56. A.J. Hoffman and R.M. Karp. On Non-terminating Stochastic Games. Mgmt. Sci. 12 (1966), 359-370.
57. A. Charnes and R.G. Schroeder. On Some Stochastic Tactical Antisubmarine Games. Naval Res. Log. Qt. 14 (1967), 291-311.
58. W. Winston. A Stochastic Game Model of a Weapons Development Competition. SIAM J. Control and Optimiz. 16 (1978), 411-419.
59. S.C. Albright and W. Winston. Optimizing Advertising Expenditures in a Dynamic Duopoly. Oper. Res. 27 (1979), 682-693.

60. I.J. Mirman and D. Spulber (eds.). The Economics of Reversible Resources. Elsevier North-Holland, New York, 1981.
61. M.J. Sobel. Stochastic Fishery Games with Myopic Equilibria. Paper 14 (pp. 259-268) in [60].
62. A.P. Kirkman and M.J. Sobel. Dynamic Oligopoly with Inventories. 52, (1974), 279-287.
63. J.A. Filar. The Travelling Inspector Model. Technical Report 374, Mathematical Sciences Department, Johns Hopkins University, 1983.
64. M.J. Sobel. Myopic Solutions of Markov Decision Processes and Stochastic Games. Oper. Res. 29, (1981), 995-1009.
65. J.A. Filar and T.A. Schultz. Interactive Solutions for the Travelling Inspector Model and Related Problems. The Johns Hopkins University Oper. Research Group Report Series, 83-06 (1983).
66. R.A. Howard. Dynamic Programming and Markov Processes. MIT Press, Cambridge, Massachusetts, 1960.
67. S. Ross. Introduction to Stochastic Dynamic Programming. Academic Press, New York, 1983.
68. G.E. Monahan. Optimal Stopping in a Partially Observable Markov Process with Costly Information. Oper. Res. 28 (1980), 1319-1334.
69. M. Majumdar. A Note on Learning and Optimal Decisions with a Partially Observable State Space. Paper 7 (pp. 141-153) of [60].
70. J. Harsanyi. Games with Incomplete Information Played by "Bayesian" Players. Mgmt. Sci. 14 (1967), 159-182, 320-334, 486-502.
71. R.J. Aumann and M. Maschler. Game Theoretic Aspects of Gradual Disarmament. Paper 5 in [194].
72. R.J. Aumann and M. Maschler. Repeated Games with Incomplete Information: A Survey of Recent Results. Paper 3 (pp. 287-403) in Final Report by Mathematica, Inc. to U.S. Arms Control and Disarmament Agency under Contract ACDA/ST-116, September 1967.
73. R.J. Aumann and M. Maschler. Repeated Games with Incomplete Information. Paper 2 (pp. 25-108) in Final Report by Mathematica, Inc. to Arms Control and Disarmament Agency under Contract ACDA/ST-143, November 1968.

74. E. Kohlberg. Optimal Strategies in Repeated Games with Incomplete Information. *Int. J. Game Theory.* 4 (1975), 7-24.
75. S. Sorin. Some Results on the Existence of Nash Equilibria for Non-Zero Sum Games with Incomplete Information. *Int. J. of Game Theory.* 12 (1983), 193-205.
76. E. Kohlberg. The Information Revealed in Infinitely-Repeated Games of Incomplete Information. *Int. J. Game Theory.* 4 (1975), 57-59.
77. J.F. Mertens. The Value of Two-Person Zero-Sum Repeated Games: The Extensive Case. *Int. J. Game Theory.* 1 (1971-2), 217-227.
78. S. Zamir. On the Notion of Value for Games with Infinitely Many Stages. *Ann. Stat.* 1 (1973), 791-796.
79. S. Zamir. On the Relation between Finitely and Infinitely Repeated Games of Perfect Information. *Int. J. Game Theory* 1 (1971-2), 179-198.
80. S. Zamir. On Repeated Games with General Information Function. *Int. J. Game Theory* 2 (1973), 215-229.
81. J.F. Mertens and S. Zamir. The Value of Two-Person Zero-Sum Repeated Games with Lack of Information on Both Sides. *Int. J. Game Theory* 1 (1971), 39-64.
82. J.F. Mertens and S. Zamir. On a Repeated Games without a Recursive Structure. *Int. J. Game Theory* 5 (1976), 173-182.
83. J.F. Mertens and S. Zamir. The Normal Distribution and Repeated Games. *Int. J. Game Theory* 5 (1976), 187-197.
84. J.F. Mertens and S. Zamir. Minmax and Maxmin of Repeated Games with Incomplete Information. *Int. J. Game Theory* 9 (1980), 201-215.
85. J.P. Ponssard and S. Zamir. Zero-Sum Sequential Games with Incomplete Information. *Int. J. Game Theory* 2 (1973), 99-107.
86. J.P. Ponssard and S. Sorin. The LP Formulation of Finite Zero-Sum Games with Incomplete Information. *Int. J. Game Theory* 9 (1980), 99-105.
87. J.P. Ponssard and S. Sorin. Some Results on Zero-Sum Games with Incomplete Information: The Dependent Case. *Int. J. Game Theory* 9 (1980), 233-245.

88. S. Sorin. A Note on the Value of Zero-Sum Sequential Repeated Games with Incomplete Information. *Int. J. Game Theory*, 8 (1979), 217-223.
89. S. Sorin. On a Repeated Game with State Dependent Signalling Matrices. Paper presented at March 1984 meeting on Mathematical Game Theory at the Oberwolfach Institute.
90. C. Waternaux. Solution for a Class of Repeated Games without a Recursive Structure. *Int. J. Game Theory*, 12 (1983), 129-160.
91. N. Megiddo. On Repeated Games with Incomplete Information Played by Non-Bayesian Players. *Int. J. Game Theory*, 9 (1980), 157-167.
92. S. Lakshmivarahan and K.S. Narendar. Learning Algorithms for  $n$ -Person Zero-Sum Stochastic Games with Incomplete Information. *Math. of Oper. Res.* 6 (1981), 379-386.
93. E.H. McCall. Performance Results of the Simplex Algorithms for a Set of Real World Linear Programming Models. *Comm. Ass. Comput. Mach.* 25 (1982), 207-212.
94. K.H. Borgwardt. Some Distribution-Independent Results About the Asymptotic Order of the Average Number of Pivot Steps of the Simplex Method. *Math. of Oper. Res.* 7 (1982), 441-462.
95. K.H. Borgwardt. The Average Number of Pivot Steps Required by the Simplex Method is Polynomial. *Zeit. f. Oper. Res (Series A)* 26 (1982), 157-177.
96. H.W. Kuhn and A.W. Tucker (eds). Contributions to the Theory of Games, I, *Annals of Math. Study* 24. Princeton University Press, 1950.
97. L.S. Shapley and R.N. Snow. Basic Solutions to Discrete Games. Pp. 27-37 in [96].
98. H.W. Kuhn and A.W. Tucker (eds.). Linear Inequalities and Related Systems, *Annals of Math. Study* 38. Princeton University Press, 1956.
99. H.D. Mills. Marginal Values of Matrix Games and Linear Programs. Pp. 183-193 in [98].
100. T.L. Saaty and S.I. Gass. The Parametric Objective Function, I. *Oper. Res.* 2 (1954), 316-319.
101. S.I. Gass and T.L. Saaty. The Parametric Objective Function, II. *Oper. Res.* 3 (1955), 395-401.



102. S.I. Gass and T.L. Saaty. The Computational Algorithm for the Parametric Objective Function. *Nav. Res. Log. Qt.* 2 (1955), 39-45.
103. T. Gal. Postoptimal Analyses, Parametric Programming and Related Topics. McGraw-Hill, 1979.
104. G.W. Brown. Iterative Solution of Games by Fictitious Play. Pp. 374-376 in Activity Analysis of Production and Allocation, T.C. Koopmans (ed.), Cowles Commission Monograph, 1951.
105. G.W. Brown and J. von Neumann. Solutions of Games by Differential Equations. Pp. 73-79 in [96].
106. J. Robinson. An Iterative Method of Solving a Game. *Annals of Math.* 54 (1951), 296-301.
107. H. Everett. Personal communication.
108. J.M. Danskin. Fictitious Play for Continuous Games. *Naval Res. Log. Qt.* 1 (1954), 313-320.
109. D. Blackwell and M.A. Girshick. Theory of Games and Statistical Decisions. John Wiley, New York, 1954.
110. J.A. Filar and T.E.S. Raghavan. An Algorithm for Solving S-Games and Differential S-Games. *SIAM J. Control and Optimization* 20 (1982), 763-769.
111. M. Dresher, L.S. Shapley and A.W. Tucker (eds.) Advances in Game Theory, *Annals of Math.* Study 52. Princeton University Press, 1964.
112. L.S. Shapley. Some Topics in Two-Person Games. Pp. 10-28 in [111].
113. C.E. Lemke and J.T. Howson, Jr. Equilibrium Points of Bimatrix Games. *J. Soc. Ind. Appl. Math.* 12 (1964), 413-423.
114. R. Wilson. Computing Equilibria of Two-Person Games from the Extensive Form. *Mgmt. Sci.* 18 (1972), 448-460.
115. H. Mills. Equilibrium Points in Finite Games. *J. Soc. Ind. Appl. Math.* 8 (1960), 397-402.
116. O.L. Mangasarian and H. Stone. Two-Person Nonzero-Sum Games and Quadratic Programming. *J. Math. Anal. and Applic.* 9 (1964), 348-355.

117. H. Mine, M. Fukushima and Y.-J. Ryang. Methods of Parametric Non-Linear Programming. *Int. J. Systems Sci.* 12 (1981), 95-110.
118. A.V. Fiacco and W.P. Hutzler. Basic Results in the Development of Sensitivity and Stability Analysis in Nonlinear Programming. *Comput. and Oper. Res.* 9 (1982), 9-28.
119. B. Bank, J. Guddat, D-Klatte, B. Kummer and K. Tammer. Nonlinear Parametric Programming. Birkhauser Verlag, Basel, 1983.
120. V. Aggarwal. On the Generation of all Equilibrium Points for Bimatrix Games through the Lemke-Howson Algorithm. *Math. Prog.* 4 (1973), 233-234.
121. M.J. Todd. Comments on a Note by Aggarwal. *Math. Prog.* 10 (1976), 130-133; also *Math. Prog.* 14 (1978), 112-15.
122. L.S. Shapley. A Note on the Lemke-Howson Algorithm. *Math. Prog. Study* 1 (1974), 175-189.
123. N.N. Vorobiev. Equilibrium Points in Bimatrix Games. *Theory of Prob. and Its Appl.* 3 (1958), 297-309 (transl).
124. H.W. Kuhn. An Algorithm for Equilibrium Points in Bimatrix Games. *Proc. Nat. Acad. Sci.* 47 (1961), 1956-1662.
125. O.L. Mangasarian. Equilibrium Points of Bimatrix Games. *J. Soc. Ind. Appl. Math.* 12 (1964), 778-780.
126. O. Moeschlin and D. Pallaschks (eds.). Game Theory and Related Topics. North-Holland, 1979.
127. H.M. Winkels. An Algorithm to Determine All Equilibrium Points of a Bimatrix Game. Pp. 137-148 in [126].
128. M.E. Dyer. The Complexity of Vertex Enumeration Methods. *Math. of Oper. Res.* 8 (1983), 381-402.
129. M.J.M. Jansen. Regularity and Stability of Equilibrium Points of Bimatrix Games. *Math. of Oper. Res.* 6 (1981), 530-550.
130. C.B. Millham. Constructing Bimatrix Games with Special Properties. *Nav. Res. Log. Qt.* 19 (1972), 709-714.
131. G.A. Heuer. Uniqueness of Equilibrium Points in Bimatrix Games. *Int. J. Game Theory.* 8 (1979), 13-25.

132. J.A. Filar. Semi-Antagonistic Equilibrium Points and Action Costs. Johns Hopkins University Oper. Res. Group Paper 82-01 (1982).
133. V. Bubelis. On Equilibria in Finite Games. Int. J. Game Theory 8 (1979), 65-79.
134. H. Tietze. Famous Problems of Mathematics. Graylock Press, Baltimore, 1965 (trans.).
135. J. Rosenmuller. On a Generalization of the Lemke-Howson Algorithm to Noncooperative N-Person Games. SIAM J. Appl. Math. 2 (1971), 73-79.
136. R. Wilson. Computing Equilibria of N-Person Games. SIAM J. Appl. Math. 21 (1971), 80-87.
137. T.C. Hu and S.M. Robinson (eds.). Mathematical Programming. Academic Press, New York, 1973.
138. C.B. Garcia, C.E. Lemke and H.J. Luthi. Simplicial Approximation of an Equilibrium Point for Noncooperative N-Person Games. Pp. 227-260 in [137].
139. G. van der Laan and A.J.J. Talman. On the Computation of Fixed Points in the Product Space of Unit Simplexes and an Application to Noncooperative N Person Games. Math. of Oper. Res. 7 (1982), 1-13.
140. J.T. Howson, Jr. Equilibria of Polymatrix Games. Mgmt. Sci. 18 (1972), 312-318.
141. K.J. Arrow and L. Hurwicz. Stability of the Gradient Process in n-Person Games. J. Soc. Ind. Appl. Math. 8 (1960), 280-294.
142. Nuclear Regulatory Commission. Request for Proposal R5-RES-82-022, July 12, 1982.
143. J. McDonald. Strategy in Poker, Business and War. W.W. Norton & Co., New York, 1950.
144. J. McDonald. The Game of Business. Doubleday, New York, 1975.
145. J.D. Williams. The Complete Strategist. McGraw-Hill (RAND Series 1, New York, 1954).
146. M.D. Davis. Game Theory: A Non-Technical Introduction. Basic Books, New York, 1976.

147. A.J. Jones. Game Theory: Mathematical Models of Conflict. Halsted Press, New York, 1980.
148. J.G. Kemeny, J.L. Snell and G.L. Thompson. Introduction to Finite Mathematics. Prentice-Hall, Englewood Cliffs (N.J.), 1957.
149. J.R. Pennock and J.W. Chapman (eds.). Due Process: Nomos XVIII. Atherton, New York, 1977.
150. G. Fromm, W.L. Hamilton and D.E. Hamilton. Federally Supported Mathematical Models: Survey and Analysis (2 vol.). National Science Foundation (Contract NSF-C-804), GPO Stock Number 038-000-00221-0, 1975.
151. S.I. Gass and R.L. Sisson. A Guide to Models in Government Planning and Operations. Environmental Protection Agency, 1974.
152. G. Brewer and M. Shubik. The War Game. Harvard University Press, Cambridge MA, 1979.
153. R.M. Friedman (Project Director). Uses of Models for Water Resources Management, Planning and Policy. U.S. Office of Technology Assessment, OTA-0-159, August 1982.
154. S.I. Gass (ed.). Oil and Gas Supply Modeling. National Bureau of Standards Special Publication 631, 1982.
155. S.I. Gass, F.H. Murphy and S.H. Shaw (eds.). Intermediate Future Forecasting System. National Bureau of Standards Special Publication 670, 1983.
156. S.I. Gass, (ed.). Operations Research: Models and Methods. Proceedings of Symposia in Applied Mathematics, v. 25, American Mathematical Society, Providence, R.I., 1981.
157. F.C. Johnson. Practical Aspects of Fishery Management Modeling. Pp. 85-104 in [156].
158. J.L. Steele. The Use of Econometric Models by Federal Regulatory Agencies. D.C. Heath & Co., Lexington MA, 1971.
159. M.O. Finkelstein. The Judicial Reception of Multiple Regression Studies in Race and Sex Discrimination Cases. Columbia Law Review 80 (1980), 737-754.
160. H. Wiberg, Some Comments on the Functions and Meaning of Game Theory. Cooperation and Conflict 3 (1968), 247-253.
161. R.B. Rovinsky and C. Shoemaker. Operations Research: Applications in Agriculture. Pp. 151-174 in [156].

162. S.S. Nagel with M. Neef. Operations Research Methods as Applied to Political Science and the Legal Process. Sage University Paper Series on Quantitative Methods in the Social Sciences 07-002, Sage Publications, Beverley Hills, CA. 1976.
163. S.S. Nagel and M. Neef. Legal Policy Analysis. D.C. Heath and Co., Lexington, MA, 1977.
164. W.F. Lucas (ed.). Game Theory and Its Applications. Proceedings of Symposia in Applied Mathematics, v. 24, American Mathematical Society, Providence, R.I., 1981.
165. L.S. Shapley. Measurement of Power in Political Systems. Pp. 69-81 in [164].
166. J.F. Banzhaf. Weighted Voting Doesn't Work: A Mathematical Analysis. Rutgers Law Rev. 19 (1965), 317-343.
167. J.F. Banzhaf. Multi-member Electoral Districts: Do They Violate the "One Man, One Vote" Principle? Yale Law J. 75 (1966), 1309.
168. J.F. Banzhaf. One Man, 3312 Votes: A Mathematical Analysis of the Electoral College. Villanova Law Review 13 (1968), 304-332.
169. W.F. Lucas. Applications of Cooperative Games to Equitable Allocation. Pp. 19-36 in [164].
170. L.J. Billera, D.C. Heath and J. Raanen. Internal Telephone Billing Rates = A Novel Application of Non-Atomic Game Theory. Oper. Res. 26 (1978), 956-965.
171. A.J. Goldman. A Game-Theoretic Model of Inspection-Resource Allocation. J. Res. Nat. Bur. Stnds. 86 (1981), 193-215.
172. M.H. Pearl and A.J. Goldman. A Game-Theoretic Model of Inspection-Resource Allocation. J. Res. Nat. Bur. Stnds. 86 (1981), 193-215.
173. J.B. Kadane and P.D. Larkey. Subjective Probability and the Theory of Games. Mgmt. Sci. 28 (1982), 113-120. Comments, rejoinders, replies and notes (by several authors): 120-125 and 29 (1982), 1334-1345.
174. J.B. Kadane and P.D. Larkey. The Confusion of "is" and "Ought" in Game Theoretic Contexts. Mgmt. Sci. 29 (1983), 1365-1379. Comment (by M. Shubik), 1380-1383.
175. G. Bass, B. Jenkins, et al. Motivations and Possible Actions of Potential Criminal Adversaries of U.S. Nuclear Programs-- Executive Summary. RAND R-2554/1-SL, February 1980.

176. P. de Leon, B. Jenkins, et al. Attributes of Potential Criminal Adversaries of U.S. Nuclear Programs. RAND R-2225-5L, February 1978.
177. B.M. Jenkins. The Potential Criminal Adversaries of Nuclear Programs: A Portrait. RAND P-6513, July 1980. (Paper at 21st Annual Meeting of the Institute of Nuclear Materials Management.)
178. A. Dvoretzky, A. Wald and J. Wolfowitz. Elimination of Randomization in Certain Decision Procedures and Zero-Sum Two-Person Games. *Ann. Math. Stat.* 22 (1951), 1-21.
179. R. Radner and R. W. Rosenthal. Private Information and Pure-Strategy Equilibria. *Math. of Oper. Res.* 7 (1982), 401-409.
180. R.J. Aumann et al. Approximate Purification of Mixed Strategies. *Math. of Oper. Res.* 8 (1983), 327-341; *Corr.* 9 (1984), 317.
181. R.S. Garfinkel and G.L. Nemhauser. Integer Programming. Wiley and Sons, 1972.
182. F. Glover and E. Woolsey. Converting the 0-1 Polynomial Programming Problem to a Linear 0-1 Program. *Oper. Res.* 22 (1974), 180-182. See also A.J. Goldman, "Linearization in 0-1 Variables: A Clarification", *Oper. Res.* 31 (1983), 946-947.
183. F. Glover. Improved Linear Integer Programming Formulations of Nonlinear Integer Problems. *Mgmt. Sci.* 22 (1975), 455-460.
184. N.W. Cooper. A Survey of Methods for Pure Nonlinear Integer Programming. *Mgmt. Sci.* 27 (1981), 353-361.
185. F. Arientesi, R. Avenhaus, M. Franklin, J.P. Shipley (eds.). Mathematical and Statistical Methods in Nuclear Safeguards (Lectures from a course held at the Joint Research Center, Ispra (Italy), Nov 30-Dec. 4, 1981). Harwood Academic Publishers.
186. R. Avenhaus. Nuclear Material Accountancy and Data Verification (the MUF-D Problem). Pp. 293-315 of [185]; also Proc. 22 Annual Meeting on Nuclear Materials Management, *J. Inst. Nuc. Mat. Mgmt.*, 132-142.

187. R. Avenhaus. Game Theoretic Analysis of Inventory Verification. Pp. 317-350 in [185].
188. R. Avenhaus and E. Leitner. Verification of the Near Real Time Accountancy Data. Proc. July 1983 Annual Meeting of Inst. Nuc. Mat. Mgmt., 255-262.
189. R. Avenhaus. Verification of Material Balance Sequences Data. Paper presented at March 1984 meeting on Mathematical Game Theory at the Oberwolfach Institute.
190. D. Bierlein. Direkte Überwachungssysteme. Oper. Res. Verfahren 6 (1969), 57-69.
191. D. Bierlein. Auf Bilanzen und Inventuren Basierende Safeguards - Systeme. Oper. Res. Verfahren 8 (1970), 36-43. Abbreviated as Part II of [12].
192. D. Bierlein. Game Theoretic Models of Safeguarding Different Types of Illegal Activities. Pp. 69-83 in Proc. Fourth Formata Symposium (May 1982) on Mathematical Methods for the Analysis of Large-Scale Systems, Academia (Publishing House of the Czechoslovak Academy of Sciences), Prague 1983.
193. M. Dresner. A Sampling Inspection Problem in Arms Control Agreements: A Game-Theoretic Analysis. RAND Memo RM-2972-ARPA, February 1962.
194. F.J. Anscombe, M.D. Davis, H.W. Kuhn, M. Maschler, H.D. Mills and F.F. Stephan. Applications of Statistical Methodology to Arms Control and Disarmament. Final Report (2v.) by Mathematica, Inc. to U.S. Arms Control and Disarmament Agency under Contract ACDA/ST-3, June 1963.
195. R.J. Aumann, J.C. Harsanyi, M. Maschler, G. O'Brien, R. Radner and M. Shubik. Development of Utility Theory for Arms Control and Disarmament. Final Report by Mathematica, Inc. to U.S. Arms Control and Disarmament Agency under Contract ACDA/ST-80, June 1966.
196. Facts on File, May 8, 1981, p. 305-A2.
197. M. Willrich and T.B. Taylor. Nuclear Theft: Risks and Safeguards. Ballinger, 1974.
198. R.B. Leachman and P. Althoff. Preventing Nuclear Theft: Guidelines for Industry and Government. Praeger Publishers, New York, 1972.

199. C. Holden. NRC Shuts Down Submarine Fuel Plant. *Science* 206 (5 October 1979), 30-31.
200. H.E. Kyburg, Jr. Subjective Probability: Criticisms, Reflections and Problems. Chapter 6 in his Epistemology and Inference, U. Minnesota Press, 1983.
201. E. Marshall. The Salem Case: A Failure of Nuclear Logic. *Science*, 220 (1983), 280-282.
202. H.P. Green. Lessons of Three Mile Island for the Institutional Management and Regulation of Nuclear Power. Pp. 311-315 in [229].
203. K. Borch. Insurance and Loss Prevention: Producer's Liability. *Greek Economic Review*, 1 (1979), 56-69.
204. K. Borch. The Price of Moral Hazard. *Scand. Actuarial Journal* (1980), 173-176.
205. E.I. Rumball. Modelling the Exclusive Economic Zone Enforcement Problem. *New Zealand J. of Sci.* 21 (1978), 343-349.
206. R. Avenhaus and H. Frick. Statistical Analysis of Alternative Data Evaluation Schemes (The MUF-D-Evaluation Problem). Proc. First ESARDA Symposium, Brussels, 1979, pp. 442-446.
207. R.E. Young. Control in Hazardous Environments. Peter Peregrinus Ltd., Stevenage and New York, 1982.
208. F.J. Anscombe. Inspection against Clandestine Rearmament. Paper 3 (pp. 69-115) in [194].
209. H.W. Kuhn. Recursive Inspection Games. Paper 5 (pp. 169-181) in [194].
210. M.D. Davis. Two Test Ban Models. Paper 6 (pp. 183-219) in [194].
211. M. Maschler. A Price Leadership Solution to the Method for Solving the Inspector's Non-Constant-Sum Game. *Naval Res. Log. Qt.* 13 (1966), 11-33. Also Paper 5 (pp. 289-377) in [194].
212. M. Maschler. The Inspector's Non-Constant-Sum Game: Its Dependence on a system of Detectors. *Naval Res. Log. Qt* 14 (1967), 275-290.
213. J.E. Lovett. Who are the Enemy? Paper 15 (pp. 207-218) in [197].



214. M. Willrich (ed.). International Safeguards and Nuclear Industry. The Johns Hopkins University Press, 1973.
215. T.B. Taylor. Diversion by Non-Governmental Organizations. Chapter 7 (pp. 177-198) in [214].
216. L.A. Dunn. Military Politics, Nuclear Proliferation and the "Nuclear Coup d'Etat." *J. of Strat. Studies*, 1 (1978), 31-48.
217. H. Schleicher. A Recursive Game for Detecting Tax Law Violations. *Economics et Societies*, 5 (1971), 1421-1440.
218. K.L. Hoffman, L.S. Joel and M.H. Pearl. Resource Requirements and Allocations in IRS' Audit Division. National Bureau of Standards Report 79-1711 (1979).
219. J. Greenberg. Avoiding Tax Avoidance: A (Repeated) Game-Theoretic Approach. *J. of Econ. Theory*, 32 (1984), 1-13.
220. R.W. Klein, L.C. Rafsky, D.S. Sibley and R.D. Willig. Decisions with Estimation Uncertainty. *Econometrica*, 46 (1978), 1363-1386.
221. R. Avenhaus. Application of the Theory of Games to Safeguards Problems. Chapter 22 (pp. 319-339) in [198].
222. G. Wardlaw. Political Terrorism: Theory, Tactics and Counter-Measures. Cambridge University Press, 1982.
223. D.W. Fischer (ed.). Managing Technological Accidents: Two Blowouts in the North Sea. Pergamon Press, 1982.
224. D.W. Fischer. An Overview of the North Sea Blowout Problem. Chapter 1 (pp. 9-30) in [223].
225. J.W. Lathrop (ed.). Planning for Rare Events: Nuclear Accident Preparedness and Manangement. Pergamon Press, 1981.
226. J.W. Lathrop. Overview. Chapter 1 (pp. 3-12) in [225].
227. D.W. Fischer. Organizational Issues at Bravo and Three Mile Island: A Comparison. Chapter 5, (pp. 75-91) in [223].
228. J.W. Lathrop. An Open Discussion of Problems in Nuclear Accident Preparedness. Chapter 2 (pp. 13-44) of [225].
229. T.M. Moss and D.L. Sills. The Three Mile Island Nuclear Accident: Lessons and Implications. *Annals N.Y. Academic Science*, v. 365 (1981).

230. C.B. Marrett. The Accident at Three Mile Island and the Problem of Uncertainty. Pp. 280-291 in [229].
231. B.M. Jenkins (ed.). Terrorism and Beyond: An International Conference on Terrorism and Low-Level Conflict. RAND R-2714-DOE/DOJ/DOS/RC, December 1982.
232. B.M. Jenkins. Introduction and Overview. Pp. 3-37 in [231].
233. D.G. Macnair. Terrorism in the Marine Environments. Pp. 273-275 in [231].
234. G. Bass, B.M. Jenkins, K. Kellen, D. Ronfeldt. Options for U.S. Policy on Terrorism. RAND R-2764-RC, July 1981.
235. S. Sloan. Simulating Terrorism. University of Oklahoma Press, Norman, 1981.
236. B.O. Koopman. Search and Screening: General Principles with Historical Applications. Pergamon Press, 1979.
237. L.D. Stone. Theory of Optimal Search. Academic Press, New York, 1975.
238. L.D. Stone. The Process of Search Planning: Current Approaches and Continuing Planning. *Oper. Res.* 31 (1983), 207-233.
239. S. Gal. Search Games. Academic Press, 1980.
240. R.W. Mengel. Terrorism and New Technologies of Destruction: An Overview of the Potential Risk. Pp. 189-245 in [239]. Also Appendix 2 (pp. 443-473) in Report of the Task Force on Disorders and Terrorism, National Advisory Committee on Criminal Justice Standards and Goals, Government Printing Office, Washington, 1976.
241. B.M. Jenkins. Combatting Terrorism: Some Policy Implications. RAND P-6666, August 1981. (Prepared for a State Department conference, "Terrorism in the 1980's", May 21-22, 1981.)
242. B.M. Jenkins. Talking to Terrorists. RAND P-6750, March 1982.
243. R. Henn and O. Moeschlin (eds). Mathematical Economics and Game Theory: Essays in Honor of Oscar Morgenstern. Lecture notes in Economics and Mathematic Systems, v. 141 (1977), Springer-Verlag.

244. R. Selten. A Simple Game Model of Kidnapping. pp. 139-155 in [243].
245. G. Bass, B. Jenkins, K. Kellen et al. Motivations and Possible Actions of Potential Criminal Adversaries of U.S. Nuclear Programs. RAND R-2554-SL, February 1980.
246. G. Bass, B. Jenkins and K. Kellen. The Appeal of Nuclear Crimes to the Spectrum of Potential Adversaries. RAND 2803-SL, February 1982.
247. R.N. Reinstedt and J. Westbury. Major Crimes as Analogs to Potential Threats to Nuclear Facilities and Programs. RAND N-1498-SL, April 1980.
248. Y. Dror. Crazy States: A Counterconventional Strategy Problem. Lexington Heath, 1971 (Second Edition: Millwood, N.Y., Kraus Reprints, 1980).
249. B.M. Jenkins. New Modes of Conflict. RAND R-3009-DNA, June 1983.
250. B.M. Jenkins. Testimony before the Committee on the Judiciary. RAND P-6962, February 1984.
251. B.M. Jenkins. Combatting Terrorism Becomes a War. RAND P-6988, May 1984; also Newsday, 13 May 1984.
252. M. Crenshaw (ed.). Terrorism, Legitimacy and Power: The Consequences of Political Violence. Wesleyan University Press, 1983.
253. W.W. Fowler. Terrorism Data Bases: A Comparison of Missions, Methods and Systems. RAND N-1503-RC, March 1981.
254. W.W. Fowler. An Agenda for Quantitative Research on Terrorism. RAND P-6591, July 1980.
255. D.A. Waterman and B.M. Jenkins. Heuristic Modeling using Rule-Based Computer Systems. RAND P-5811, March 1977. Also pp. 285-330.
256. T.C. Schelling. Who Will Have the Bomb? Pp. 42-56 in [238].
257. Y. Dror. Terrorism as a Challenge to the Democratic Capacity to Govern. Paper 4 (pp. 65-90) in [252].
258. P. Wilkinson. Political Terrorism. John Wiley and Sons, 1974.

259. Y. Alexander and S.M. Finger (eds.). Terrorism: Inter-Disciplinary Perspectives. John Jay Press, New York, 1977.
260. J.J. Paust. A Definitional Focus. Pp. 18-29 of [259].
261. B.M. Jenkins. The Study of Terrorism: Definitional Problems. RAND P-6563, December 1980.
262. M.E. Wolfgang (ed.). International Terrorism. Special issue of Ann. Amer. Acad. Political Social Science, 463 (1982).
263. P.E. Devine and R.J. Rafalko. On Terror. Pp. 39-53 in [262].
264. R.W. Taylor and H.E. Vanden. Defining Terrorism in El Salvador: "La Matanza". Pp. 106-118 in [262].
265. R. Schultz. Conceptualizing Political Terrorism: A Typology. J. of Int. Affairs, 32 (1978), 7-15.
266. L.R. Barres. Terrorism and Global Security: The Nuclear Threat. Westview Press, Boulder, 1979.
267. W.F. May. Terrorism as Strategy and Ecstasy. Soc. Res., 41 (1974), 277-298.
268. B.M. Jenkins. The Potential for Nuclear Terrorism. RAND P-5876, May 1977.
269. D. Ronfeldt and W. Sater. The Mind-Sets of High-Technology Terrorists: Future Implications from an Historical Analog. RAND N-1610-SL, March 1981.
270. B.M. Jenkins. Some Reflections on Recent Trends in Terrorism. RAND P-6897, July 1983.
271. R. Fried. The Psychology of the Terrorist. Pp. 119-124 of [231].
272. B.M. Jenkins. RAND's Research on Terrorism. RAND P-5969, August 1977. Also in Terrorism 1 (1977).
273. Y. Alexander. Terrorism, the Media and the Police. J. of Int. Affairs, 32 (1978), 101-113.
274. A.R. Norton and M.H. Greenberg. Studies in Nuclear Terrorism. G.K. Hull, Boston, 1979.
275. B.M. Jenkins. Will Terrorists Go Nuclear? RAND P-5541, November 1975.

276. M. Hutchinson. Defining Future Threats: Terrorists and Nuclear Proliferation. Pp. 147-188 in [239]; also pp. 298-316 in [259].
277. B.M. Jenkins. The Terrorist Mindset and Terrorist Decision-making: Two Areas of Ignorance. RAND P-6340, June 1979.
278. K. Kellen. Terrorists-What are They Like? How Some Terrorists Describe Their World and Actions. Pp. 125-173 in [231]; also RAND N-1300-SL, November 1979.
279. J. Sundberg. Operation LEO: Description and Analysis of a European Terrorist Operation. Pp. 174-202 of [231].
280. V.J. Konecni and E.B. Ebbesen (eds.). The Criminal Justice System: A Social-Psychological Analysis. W.H. Freeman, San Francisco, 1982.
281. J.S. Carroll. Committing a Crime: The Offender's Decision. Paper 3 (pp. 49-67) in [280].
282. R.L. Keeney and H. Raiffa. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley, 1976.
283. C.W. Kirkwood (ed.). Special Issue on Decision analysis. Oper. Res., 28:1 (January-February 1980).
284. R.H. Kupperman and D.M. Trent. Terrorism: Threat, Reality, Response. Hoover Institution Press, Stanford, 1979.
285. R. Avenhaus and H. Frick. Analyze von Fehlalarmen in Überwachungs-systemen mit Hilfe von Zweipersonen-Nichtnullsummenspielen. Oper. Res. Verfahren, 26 (1971).
286. C. Hosmer. Keynote Re. ks. Chapter 1 (pp. 3-15) in [198].
287. A. Mensch (ed.). Theory of Games: Techniques and Applications (Proceedings of a NATO Conference). American Elsevier, 1966.
288. R.W. Shepard. Some Reasons for Present Limitations in the Application of the Theory of Games to Army Problems. Pp. 373-388 in [287].
289. W.W. Fain and J.B. Phillips. Applications of Game Theory to Real Military Decisions. Pp. 363-372 in [287].
290. H. Von Stackelberg. The Theory of the Market Economy. Oxford University Press, 1952. (Translation of Marktform und Gleichgewicht, Springer, Vienna, 1934.)

291. T. Basar and G.J. Olsder. Dynamic Cooperative Game Theory. Academic Press, 1982.
292. K.R. Hammond (ed.). Judgement and Decision in Public Policy Formation. AAAS Selected Symposium 1, Westview Press, 1978.
293. K.R. Hammond and L. Adelman. Science, Values and Human Judgement. Paper 7 (pp. 120-141) in [292].
294. H.A. Linstone and M. Turoff (eds.). The Delphi Method: Techniques and Applications. Addison-Wesley, 1975.
295. H. Sackman. Delphi Critique: Expert Opinion, Forecasting and Group Process. Lexington-Heath, 1975.
296. R.L. Keeney. The Art of Assessing Multiattribute Utility Functions. *Org. Beh. and Human Perf.*, 19, (1977), 267-310.
297. M. Rokeach. The Nature of Human Values. Free Press, New York, 1973.
298. B. Wynne. Redefining the Issues of Risk and Public Acceptance. *Future*, 16 (1980), 13-32.
299. R.A. Shweder (ed.). Fallible Judgement in Behavioral Research. New Directions for Methodology of Social and Behavioral Science (Jossey-Bass Inc., San Francisco, Number 4, 1980.
300. H.J. Einhorn. Overconfidence in Judgement. Pp. 1-16 in [299].
301. H. Kunreuther. Disaster Insurance Protection: Public Policy Lessons. Wiley-Interscience, 1978.
302. B. Fischhoff, S. Lichtenstein, P. Slovic, S.I. Derby and R.L. Keeney. Acceptable Risk. Cambridge University Press, 1981.
303. S.L. Salem, K.A. Solomon and M.S. Yesley. Issues and Problems in Inferring a Level of Acceptable Risk. RAND R-2561-DOE, August 1980.
304. National Research Council (Committee on Risk and Decision Making, Assembly of Behavioral and Social Sciences). Risk and Decision Making: Perspectives and Research. National Academy Press, 1982.
305. K.A. Solomon, W.E. Kastenberg and P.F. Nelson. Dealing with Uncertainty Arising Out of Probabilistic Risk Assessment. RAND R-3045-ORNL, September 1983.

306. V.T. Covello, W.G. Flamm, J.V. Rodricks and R.G. Tardiff. The Analysis of Actual Versus Perceived Risks. Plenum Press, 1983.
307. R.K. Mullen. Mass Destruction and Terrorism. *J. of Int. Affairs*, 32 (1978), 63-89.
308. B.M. Jenkins. Terrorism and the Nuclear Safeguards Issue. RAND P-5611, March 1976.
309. L.Z. Freedman and Y. Alexander. Perspectives on Terrorism. Scholarly Resources, Wilmington, 1983.
310. Y. Alexander. Terrorism and High-Technology Weapons. Pp. 225-240 in [309].
311. D.M. Rosenbaum. Nuclear Terror. *International Security*, 1 (1977), 140-161.
312. R.H. Kupperman. Facing Tomorrow's Terrorist Incident Today. Pp. 379-441 in [238].
313. J.K. King (ed.). International Political Effects of the Spread of Nuclear Weapons. U.S. Government Printing Office, Stock Number 041-015-00105, April 1979.
314. B.M. Jenkins. The Consequences of Nuclear Terrorism. Pp. 79-102 in [313]; also RAND P-6373, August 1979.
315. M. Crenshaw. Introduction: Reflections on the Effects of Terrorism. Paper 1 (pp. 1-37) in [252].
316. C.A. Bennett, W.M. Murphey and T.S. Shur. Societal Risk Approach to Safeguards Design and Evaluation. ERDA-7, U.S. Energy Research and Development Administration, June 1975.
317. A.J. Goldman. Recognition of Completely Mixed Games. *J. Res. Nat. Bur. of Standards*, 67B, 23-29.
318. K. Goldberg, A.J. Goldman and M. Newman. The Probability of an Equilibrium Point. *J. Res. Nat. Bur. of Standards*, 72B (1968), 93-101.
319. M. Dresner. Probability of a Pure Equilibrium Point in N-Person Games. *J. of Comb. Theory*, 8 (1970), 134-145.
320. J.E. Walsh. Discrete Two-Person Game Theory with Median Payoff Criterion. *Opsearch* 6 (1969), 83-97.
321. J.E. Walsh. Median Two-Person Game Theory for Median Competitive Games. *J. Oper. Res. Soc. Japan*, 12 (1969), 11-20.

322. J.E. Walsh. Generally Applicable Solutions for Two-Person Median Game Theory. *J. Oper. Res. Soc. Japan*, 13 (1970), 1-5.
323. J.E. Walsh. Identification of Situations Where Cooperation is Preferable to Use of Median Game Theory. *Opsearch*, 7 (1970), 89-95.
324. J.E. Walsh and G.J. Kelleher. Generally Applicable Two-Person Percentile Game Theory. *Opsearch*, 8 (1971), 143-152.
325. J.E. Walsh and G.J. Kelleher. Greatly Increased Practical Usefulness of Two-Person Game Theory by Adoption of Median Criterion. *Rev. Franc. d'Informatique et Recherche Operationnelle (RIRO)* V-3 (1970), 81-88.
326. J.E. Walsh and G.J. Kelleher. Identification of Two-Player Situations where Cooperation is Preferable to Use of Percentile Game Theory. *J. Oper. Res. Soc. Japan*, 13, 159-168.
327. G.J. Kelleher. Simplified Solutions for Two-Person Percentile Games. *Opsearch*, 2 (1972), 30-38.
328. H. deVries. Quantile Criteria for the Selection of Strategies in Game Theory. *Int. J. Game Theory*, 3 (1974), 105-114.
329. E.S. Won. Solutions in Non-Cardinal Two-Person Games. Dissertation Research Proposed to the Johns Hopkins University (accepted April 1983).