Final Report

August 1975

A NEW METHODOLOGY TO INTEGRATE PLANETARY QUARANTINE REQUIREMENTS INTO MISSION PLANNING, WITH APPLICATION TO A JUPITER ORBITER

Prepared for:

JET PROPULSION LABORATORY
PASADENA, CALIFORNIA 91103

CONTRACT 954025 UNDER
NASA Contract NAS7-100
Task Order No. RD-4

(NASA-CR-146072) A NEW METHODOLOGY TO INTEGRATE PLANETARY QUARANTINE REQUIREMENTS INTO MISSION PLANNING, WITH APPLICATION TO A JUPITER ORBITER Final Report (Stanford Research Inst.) 124 p HC $5.50 CSCL 06M G3/91

N76-15965

Unclas
08808

STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

ACKNOWLEDGMENTS

The authors wish to express their gratitude to Dr. Charles W. Craven, Project Manager of Planetary Quarantine, Jet Propulsion Laboratory, who sponsored this study and with associated personnel provided continuous assistance. We are especially grateful to Dr. Allan R. Hoffman and Dr. Wayne F. Brady of the Planetary Quarantine Group for providing us with essential scientific information, making constructive criticisms, and finally, reviewing thoroughly the draft report.

We would like also to acknowledge the generous contribution of other JPL staff members: Dr. Ralph F. Miles Jr., who provided advice throughout the study and carefully reviewed the entire report; Dr. John C. Beckman, who discussed with us the planning problems posed by a Jupiter Orbiter Mission; and Dr. Robert G. Chamberlain, who assisted us in comparing the merits of alternative planning methodologies.

ABSTRACT

A new methodology is proposed for integrating planetary quarantine objectives into space exploration planning. This methodology is designed to remedy the major weaknesses inherent in the current formulation of planetary quarantine requirements. Application of the new methodology is illustrated by a tutorial analysis of a proposed Jupiter Orbiter mission.

Present NASA methods express planetary quarantine requirements in the form of maximum probability of contamination constraints; nominal mission plans are then designed to meet these requirements. With the advent of complex missions of long duration, this method is increasingly felt to be inadequate and often too restrictive. It offers no assurance of reaching an adequate balance among numerous mission objectives, and it does not provide rational decision criteria should a mission flight depart from its nominal plan.

The proposed methodology reformulates planetary quarantine planning as a sequential decision problem. Rather than concentrating on a nominal plan, all decision alternatives and possible consequences are laid out in a decision tree. Probabilities and values are associated with the outcomes, including the outcome of contamination. The process of allocating probabilities, which could not be made perfectly unambiguous and systematic, is replaced by decomposition and optimization techniques based on widely known principles of dynamic programming. Thus, the new methodology provides logical integration of all available information and allows selection of the best strategy consistent with quarantine and other space exploration goals.

Ideally, the values (or penalties) associated with planetary contamination should be assessed by an international body such as COSPAR. In the meantime, contamination penalties can be inferred through a simple iterative process to obtain optimal strategies compatible with probability constraints imposed on projects. These inferred penalties are useful signals for coordinating the assignment of current probability constraints.

A Jupiter Orbiter mission has been selected to illustrate the new methodology. A tutorial analysis of the insertion maneuver into Jupiter orbit demonstrates how realistic elements can be taken into account. A preliminary examination of the remainder of the mission has revealed two

major unresolved issues. One is finding a safe method for disposing of the spacecraft at the end of the multiple-encounter exploration phase, assuming that the spacecraft remains under control. The other is finding an emergency disposal system if control of the spacecraft trajectory is partially lost during the exploration phase. The resolution of both issues will require a careful assessment of the probabilities of contamination given impact on Jupiter and its satellites.

CONTENTS

4 A JUPITER ORBITER MISSION (continued)

ILLUSTRATIONS

TABLES

1 THE REPORT IN BRIEF

1.1 Objectives of the Report, Methodological Findings and Recommendations

We have examined in this report the concern over planetary contamination, critically reviewed its expression in NASA's planetary quarantine policy, and identified the weaknesses of current implementation practices, particularly as they apply to the optimization of mission guidance decisions. A new methodology and analytical procedures that avoid these weaknesses are proposed, and are illustrated using a Jupiter Orbiter (JO) mission.

Although present procedures have served a useful purpose for planning early planetary exploration missions, there is no assurance that they can determine optimal strategies for the more complex missions contemplated in the 1980s. These procedures, which assign maximum permissible probabilities of contamination--called probability suballocations--to various mission events, have two major shortcomings.

(1) Trade-offs between quarantine and other space exploration objectives remain implicit. It is therefore difficult to take into account new information, including that obtained during the mission, in developing mission strategies.

(2) There is no systematic means of allocating probabilities of contamination to events of a complex mission, and therefore no assurance that a set of suballocations will lead to an optimal strategy.

The proposed methodology is designed to alleviate these weaknesses; it represents a next logical step in the evolution of planetary quarantine planning.

Our major recommendation is that planetary quarantine planning be reformulated as a sequential decision problem. In such a form, it may be addressed by systematic decomposition and optimization techniques based on widely known dynamic programming principles.

Such reformulation requires that (1) a tree of decision alternatives and possible consequences be laid out, (2) probabilities be associated with consequences, and (3) values be assigned to ultimate outcomes. Where explicit values are not available, as, for example, with contamination penalties, it may be possible to infer them from probability constraints. These values can be applied systematically to the determination of optimal strategies. In the future, however, all values, and in particular contamination penalties, should be established directly by NASA and perhaps ultimately by an international body such as COSPAR.

In a direct assessment of contamination penalties, it is possible to take into account characteristics of the planet, the contaminating agent, and the extent of contamination. Having these values would bring a greater consistency to the determination of optimal mission strategies and space exploration programs.

1.2 Current Issues in Planetary Quarantine Planning

1.2.1 The Evolution of Planetary Exploration

Planetary exploration has evolved from the first Mars and Venus fly-bys of the early 1960s to ever more complex missions such as orbiters, landers, and now multiple-encounter missions. The recent successes of Mariner 10, which flew by Venus and Mercury, and Pioneer 11, which is now en route from Jupiter to Saturn, have demonstrated the feasibility and the potential of this new method of space exploration. The navigational concept is to use the gravitational assistance of one celestial body during a close fly-by to send a spacecraft toward another body. Thus the massive planet Jupiter may be used as a key relay for the exploration of more distant outer planets that would otherwise require enormous quantities of fuel and extremely long flight durations.

These advanced missions pose new planning problems. Many guidance and control decisions must be made during flights that may last for years. During this time, information is being obtained about the state of the spacecraft, its trajectory, and the region of space being explored. To take full advantage of this information, contingency plans must be developed encompassing the variety of events that may take place.

1.2.2 Planetary Quarantine in Mission Planning

Two fundamental goals of NASA's planetary program are to search for the origin of life and to understand the formation of the solar system. The discovery of extraterrestrial life, whether it be like that on Earth or different, would indeed be of the greatest interest to biologists and others. However, this search might be compromised by carelessness in early space ventures.

Viable terrestrial organisms deposited on another life-bearing planet might create rapid and profound changes in that planet's biology before sophisticated life detection experiments could be carried out. Furthermore, results from life detection experiments might be hopelessly confused if terrestrial life forms were discovered without the assurance that they had not been brought from Earth by accident. In a more distant future, attempts at controlling the development of an exobiology might be jeopardized if contamination by undesirable terrestrial organisms had already taken place.

Caution must therefore be exercised in man's exploration of the solar system lest one important object of the inquiry be destroyed in the search process. The crux of the problem is to define a criterion to distinguish between cautious and careless exploration.

1.2.3 Current Planetary Quarantine Planning Practice

For most missions, the risk of contamination cannot be eliminated entirely (absolute sterilization is as unfeasible as absolute reliability), so a compromise is made. A maximum probability of contamination, known as the mission allocation, is set, and the mission planner must demonstrate that this upper limit will not be exceeded.

Current procedures consist of analyzing the potential sources of contamination and making probability suballocations, based on the mission allocation, for each of them. For instance, in a typical fly-by mission, contamination may be caused by accidental impact of the spacecraft or by ejecta flux. An impact of the spacecraft can be attributed to the failure of one of the guidance maneuvers, which can in turn be attributed to another failure, and so forth. Then, based implicitly on the relative difficulty and importance of each maneuver for the success of the mission, a maximum probability of contamination is assigned to each potential source of contamination. This allocation process is often limited to one sequence of guidance maneuvers and likely consequences, i.e., a nominal mission plan. Moreover, in many simple mission plans there is a one-to-one correspondence between a guidance decision and a resulting probability of contamination; hence, the widespread but sometimes confusing habit of considering probabilities of contamination as decision variables.

1.2.4 The Need for Revision of Planetary Quarantine Plans

Some mission planners have indicated that the suballocation procedure is inadequate and may be too restrictive, especially where it is limited to a nominal plan. The actual development of most complex missions will not follow the nominal plan. Hypotheses or circumstances under which the mission contamination analysis originally was conducted will become obsolete. Contamination-related outcomes will become known as the mission proceeds, and since the probability of contamination associated with each potential source was very small to start with, the impression will be most of the time that expected (planned) risks of contamination are being avoided and that the nominal case probability suballocations for the rest of the mission could be relaxed.

The question often raised by mission planners is how to revise the contamination analysis and, if necessary, modify the probability suballocations during a flight. The nominal plan is supposed to lead to the most valuable mission that still satisfies the planetary quarantine requirements under nominal conditions. But to provide optimal results under all foreseeable circumstances, new procedures are needed that can assess probabilities of contamination conditional upon those circumstances.

The same question could be raised at higher levels in the decision hierarchy for space exploration. At the space program level, should mission allocations be revised to take into account the results of earlier missions, and if so, how? At the international level, should scientists revise the planetary quarantine policy as a result of new technological developments and new advances in scientific knowledge, and if so, how? Although the revision procedures would vary in practice from one decision level to another, the same conceptual framework should apply. The conceptual framework must be based on a clear interpretation of the concept of probability and the issue of planetary contamination.

1.3 Weaknesses of a Probability Constraint Formulation as a Basis for Planetary Quarantine Planning

The scientific community expressed concern over possible interplanetary contamination before the space exploration program began. Thus, planetary quarantine research was initiated at an early stage and an international agreement was achieved to insure that future planetary explorations would not be compromised by carelessness or oversight during the first missions. The remarkable results obtained by missions flown to date, with very small probabilities of contamination, demonstrate the usefulness and the sensibility of this early agreement.

However, the use of probability constraints to implement quarantine policy has fundamental weaknesses. In the past, these weaknesses may have had relatively minor effect. In a mission to a single planet, consisting of a small number of guidance maneuvers in a nominal plan, other mission values such as fly-by altitude and fuel conservation could be taken into account more or less intuitively. For gravity-assisted multiple-encounter missions of long duration, the interaction between quarantine considerations and other space exploration objectives will be considerably more complex. For these missions, the weaknesses in a probability constraint formulation may be a serious liability in mission planning.

The three fundamental weaknesses in the probability constraint formulation are:

(1) Implicit Value Judgments--Trade-offs between quarantine and other space exploration objectives are implicitly rather than explicitly stated. It is difficult to see how information on planetary science, spacecraft capability, mission achievements, or future programs is brought into the planning process when the basis for establishing the quarantine constraint is not evident. Of necessity, the maximum acceptable probability of contaminating a planet during the period of biological exploration now represents a value judgment. Its justification must be as follows:

(a) A smaller probability of contamination would entail an unacceptable increase in the cost of exploration or would curtail projects that remain of interest despite their associated risks of contamination.

(b) A larger probability of contamination is not acceptable even though it may permit mission cost reductions or additions of valuable exploration projects.

Therefore, a set of possible planetary exploration programs and a trade-off between the value of these programs and the risks of contamination must be implicit in the planetary quarantine policy. Explicitly stated trade-offs would be a better guide for selecting specific missions within a space exploration program than the specification of a maximum allowable probability of contamination for a planet or satellite.

(2) Negative Value to Resolving Uncertainty--The use of probability constraints can lead to a conclusion that

further information to resolve uncertainty about contamination has a negative vlue. This conclusion contradicts the basic, common sense assumption that more information will improve rather than degrade the planning process, and its logical extension suggests that it might be worthwhile to suppress information that could show the constraint will be violated.

(3) Ambiguous Method of Allocation--The procedure for allocating the probability constraint among missions in a program, or guidance decisions within a mission, is ambiguous and allows for varying interpretations. Furthermore, the definition of probability suballocations may be ambiguous, allowing for several interpretations.

1.4 The Proposed New Methodology

1.4.1 Principle and Scope

Although current procedures have served a useful purpose, we believe that a new methodology is needed to insure that planning for complex future missions is soundly formulated. The new methodology should be applicable to all contamination-related decisions at all levels of the planetary quarantine program. We believe that, to that end, NASA's procedures should be reformulated in terms of explicit values (e.g., scientific values, and contamination penalties) rather than probability constraints. Ideally, the values should be set explicitly by an international body such as COSPAR. However, to be immediately useful to mission planners, the new methodology must be applicable in the context of current NASA quarantine procedures, which use probability constraints.

We therefore propose the following methodology:

(1) Probability Consistency--To use the consistency conditions inherent in probability theory to assess and revise probabilities, including the probabilities of contamination.

(2) Decision Formulation--To define decision sequences for projects (or missions or programs) and reformulate planetary quarantine planning as a sequential decision problem, which can then be addressed by appropriate optimization procedures.

(3) Iterative Reconciliation--To use a simple iterative process to obtain an optimal strategy compatible with the probability constraints imposed on the project.

This methodology assures logical integration of the available information and consistent decision making within projects. It also assures compatibility of the optimal strategy with the probability constraints imposed upon each project. In addition, while probability constraints are imposed upon projects, the method provides useful signals for coordinating these constraints among projects.

1.4.2 The Relation Between Probability Constraints and Value Assessments

The essential departure from the current procedure is that no probability suballocations are necessary; rather, values must be assigned to the various possible outcomes of a project, including that of contamination.

As we have seen, the assessment of a probability constraint requires an implicit judgment: a trade-off between expected value in meeting scientific and other goals of space exploration and the probability of contaminating planets or satellites of biological interest. It is this trade-off that we believe should be addressed by NASA policy makers and the concerned scientific community. The assessment may be conveniently made in terms of a dollar value, applicable perhaps only to a range of probabilities below some threshold.

The situation is in many respects parallel to that of assigning a value to human life in the context of decisions on automobile or aircraft safety. This value is useful for making trade-offs in a consistent way between small probabilities of loss of life and the increased costs imposed by automobile seat belts, improved aircraft navigation systems, and similar safety measures. It would be improper to assert that the value of life used in this context represents an amount an individual would pay to avoid certain death. Similarly, it would be improper to assert that the contamination penalty used in the context of a contamination analysis represents an amount society would pay to avoid certain planetary contamination.

We find using an explicit penalty preferable because it permits the weaknesses in an allocation approach to be avoided. The process of allocating probabilities, which cannot be made unambiguous and systematic, is replaced by decomposition and optimization techniques using widely known principles of dynamic programming. This permits selection of the best mission or program strategy consistent with quarantine and other space exploration goals. New information may be easily incorporated in this formulation, and complex sequential strategies or contingency plans may be assessed and evaluated.

1.5 Application to a Jupiter Orbiter Mission

1.5.1 Objective and Scope

A Jupiter Orbiter (JO) mission offers a particularly appropriate opportunity to apply the proposed methodology. From a planetary quarantine point of view, the mission seems unprecedented. The spacecraft will repeatedly fly by several celestial bodies that, although not well-known, are considered of potential biological interest.*

For analysis, the mission can be divided into three phases: a maneuver to insert the spacecraft into Jupiter orbit, a multiple-encounter exploratory phase, and a disposal phase.

The first phase has been used for a tutorial application of the proposed methodology. It is treated as a complete mission consisting of a single fly-by of Ganymede. The focus is on determining an optimal fly-by altitude that will be low enough for fuel conservation but not so low as to cause excessive danger of contamination. The analysis does not pretend to solve the problem; rather it demonstrates how realistic elements can be taken into account with the new methodology.

1.5.2 Major Findings with Respect to a JO Mission

There are two major unresolved issues with the mission designs and strategies that have been proposed to date. One is how to dispose of the spacecraft at the end of the multiple-encounter exploration phase, assuming that the spacecraft remains under control. The second is how to dispose of the spacecraft if control of its trajectory is lost during the exploration phase. In that case, the spacecraft is nearly certain to impact one of the major bodies in the Jovian system sometime during the period of planetary quarantine. Given the current estimates of

---

*The biological interest of the Galilean satellites has not yet been defined. Recent guidelines proposed by the NASA Planetary Quarantine Office indicate strict quarantine policy is under consideration [1].

spacecraft reliability and present assessments of the probability of contamination if the spacecraft impacts one of the Galilean satellites, the probability of contamination of the JO mission exceeds by far the contemplated mission allocation.

The contamination analysis for the multiple-encounter exploration phase can be summarized in a single approximate expression for the probability of contamination:

$$P(C) = \begin{pmatrix}\text{Probability of failure} \\ \text{and consequently of impact}\end{pmatrix} \times p_c \tag{1-1}$$

where

$P(C)$ = probability of contamination

p_c = probability of contamination given impact.

The probability of failure being small compared to 1 is approximately equal to the sum of the probabilities of failure at each encounter. Failure at each encounter can, in turn, be decomposed into failure of the main propulsion system (probability f) and failure of an emergency maneuver to correct the trajectory (probability q) if the fly-by maneuver sets the spacecraft on an impact course (probability p_i); that is,

$$\begin{pmatrix}\text{Probability of failure} \\ \text{at each encounter}\end{pmatrix} \simeq f + p_i q \quad . \tag{1-2}$$

Assuming for simplicity that the exploration phase consists of n similar encounters, the probability of contamination during this phase can therefore be written as

$$P(C) \simeq n(f + p_i q) \quad . \tag{1-3}$$

It would be desirable to perform up to 50 encounters without taking a risk of contamination much greater than 5×10^{-5} (risks of contamination will also be incurred during the insertion maneuver and the standard disposal maneuver). Therefore, the product $(f + p_i q)p_c$ must be limited to approximately 10^{-6}.

The probability of contamination following the execution of a retargeting maneuver, $p_i q$, can be made much smaller than f. Therefore, the two critical parameters are the main propulsion system reliability and the probability of contamination given impact. Their product must be less than 1×10^{-6}.

1.5.3 Assessment of the Probability of Contamination Given Impact

The first major task will be to improve the assessment of the probability of contamination given impact and, if possible, to reduce this probability.

The probability of contamination given impact depends essentially on two factors: the total number of viable terrestrial organisms (VTOs) that may be released on a planet and the probability of growth, p_g, of these organisms on the planet. Assuming that each VTO has an independent probability p_g of survival, the probability of contamination given impact can be calculated as

$$p_c = Np_g \text{ (as long as } Np_g < 1\text{)} \quad .^*$$

However, at the other extreme, a complete dependence assumption could be made: All VTOs will either survive with probability p_g or die. In that case, the probability of contamination given impact is

$$p_c = p_g \quad ,$$

which is usually a much smaller probability (typical values of N may range from 10^3 to 10^6).

Quarantine planning for a JO mission may require careful modeling of factors affecting the probability of contamination given impact.

1.5.4 The Need for an Emergency Disposal System

Current estimates of the main propulsion system reliability range from 10^{-2} to 10^{-4}. It does not seem possible to improve this reliability

---

*This equation is the usual Sagan-Coleman formula. For a discussion of this formula and its limitations, see Judd, North, and Pezier [2].

to the point where the planetary quarantine requirement would be met. Given that current estimates of p_c are close to 1 for the Galilean satellites, an emergency disposal system therefore seems necessary.

It is beyond the scope of this study to try to characterize such a system; however it should meet some strict design specifications. In particular, the emergency disposal system should have (1) a probability of failure of 10^{-4} or better, and (2) given that it functions, a resulting probability of contamination for subsequent events not exceeding 10^{-4}.*

*See derivation on page 94, and Figure 4.11, page 95.

2 A CRITICAL REVIEW OF CURRENT PROCEDURES FOR PLANETARY QUARANTINE PLANNING

2.1 Plan of the Review

The scientific community expressed a concern over possible planetary contamination several years before actual planetary explorations were undertaken. In these days, little was known about the decontamination and sterilization techniques that could be used on spacecraft hardware. Likewise, the goals and possibilities of planetary exploration were only vaguely defined. Under these circumstances, scientists first wanted to require absolute sterilization. However, it was quickly recognized that such a requirement would be inoperative and that a small probability of contamination would have to be accepted if planetary exploration were to be allowed. Planetary quarantine planning procedures were then developed on the concept of a maximum permissible probability of contamination.

The purpose of Section 2.2 is to review the evolution of planetary quarantine planning procedures in order to develop an understanding of the present system. We will show, in particular, how the use of maximum permissible probabilities of contamination has affected the planning process.

Unfortunately, this historical development has led to a danger of confusion about the concept of probability; in particular, it has raised ambiguities about the correct interpretation of probability allocations. Section 2.3 will review the concept of probability and the consistency rules of probability theory as they apply to the assessment and revision of contamination probabilities.

For simple missions, an experienced planner could probably avoid the difficulties inherent in the current planning procedures and determine mission strategies that would be optimal for all practical purposes. With the complex missions contemplated for the 1980s, such an outcome becomes very unlikely. Intuition and judgment will usually be inadequate to trace the consequences of multiple guidance maneuvers and to apply consistent trade-offs between multiple mission objectives. It becomes crucial to found planetary quarantine procedures on logically sound conceptual bases and systematic optimization techniques. Section 2.4 outlines the elements required to develop such planning procedures.

2.2 The Evolution of NASA's Planetary Quarantine Policy

2.2.1 From Sterilization to Planetary Quarantine

The first formal NASA policy for preventing lunar and planetary contamination can be found in letters written by Dr. Abe Silverstein, then director of space flight programs, in October 1959. These letters, addressed to major NASA subcontractors, all stated that payloads that might impact a celestial body must be sterilized before launching.

It was not long before the feasibility of this statement was questioned and the need to make some allowance for a small probability of contamination was realized. Sterility is an absolute state, extremely difficult to achieve when considering large and complex equipment, and almost impossible to guarantee. Thus, the idea of planetary quarantine replaced that of complete sterilization. In a joint paper published in 1959, R. W. Davies and M. G. Comuntzis [3] recommended, among other goals, that the probability of landing a viable terrestrial organism on Mars and Venus should be kept below one chance in a million per mission during the early phases of exploration of these planets. This number was rather arbitrary, and was later revised.

2.2.2 Planetary Quarantine Research

These early directives and recommendations fostered a major effort in spacecraft sterilization and biological research. At that time, there was little experience pertinent to the new problems that had to be solved. For example, it was soon recognized that viable microorganisms could be trapped within closed cavities, between mated surfaces, or even encapsulated into solid materials such as plastics. These organisms could be released upon impact of a spacecraft on a planet and during the material's subsequent fracture and erosion. Classical sterilization techniques such as washing with sporicidal agents could only sterilize exposed surfaces; they were inefficient against trapped organisms. New techniques of decontamination had to be developed using penetrating agents such as dry heat and radiation.

Efforts were also made to reduce the probability of accidental impact by improving the accuracy and reliability of guidance and propulsion systems. New strategies biasing spacecraft trajectories away from the target planet were introduced to prevent an impact following a rocket or spacecraft malfunction.

At the same time, new knowledge in biology and planetology permitted more accurate assessments of the probability that terrestrial microorganisms, if released on a given planet in a viable state, would grow and multiply.

2.2.3 Present International Quarantine Policy

Through all these developments, the same basic formulation of planetary quarantine policy has been retained. The international agreement reached at the meeting of the Committee of Space Research of the International Council of Scientific Unions (COSPAR) in 1966 has not since been fundamentally modified. The formulation follows the inspiration of the earlier recommendation by Davies and Comuntzis [3] and is illustrated by a statement from NASA Policy Directive 8020-10, dated September 6, 1967:

> Biological Contamination: the basic probability of one in one thousand (1×10^{-3}) that a planet of biological interest will be contaminated shall be used as the guiding criterion during the period of biological exploration of Mars, Venus, Mercury, Jupiter, and other planets and their satellites that are deemed important for the exploration of life, life precursors and remnants thereof.

The two major features of this policy are that:

(1) A limit to contamination is expressed in terms of a maximum probability that certain celestial bodies will be contaminated during a given period.

(2) The maximum probability of contamination to be shared by all space-faring nations is expressed by a number internationally agreed upon.

At the international level, this may be the most useful and practical formulation of planetary quarantine to date. However, a careful examination of the needs for a planetary quarantine policy reveals three sources of difficulties with the current formulation:

(1) Although the maximum probability of contamination indicated in the policy can be revised by a new international agreement, there is no guideline to relate this probability to the basic benefits, costs, and contamination consequences that may result from space exploration.

(2) Using probability constraints may violate the basic axioms for rational decision making under uncertainty as they are stated in modern economic and statistical theory (see discussion in the appendix).

(3) It is impossible to formulate the determination of a planetary exploration strategy as a unique sequential decision problem (see explanation in Section 2.4.2).

In addition, because the present policy fails to define "biological interest," "period of biological exploration," "important for the exploration of life," and "contaminated," ambiguities hamper the already complicated task of implementing it.

2.2.4 Implementation of International Policy in U.S. Program Planning by the Planetary Quarantine Office of NASA

A Planetary Quarantine Office, attached to the Planetary Programs Office, Office of Space Science and Applications, NASA, was created in 1963 to carry out the U.S. planetary quarantine program under the supervision of the Space Science Board. NASA directive NHB 8020-12, "Planetary Quarantine Provisions for Unmanned Planetary Missions," April 1969, specifies that mission plans must be submitted to the Planetary Quarantine Officer (PQO) for approval. The details of the procedures for submission and approval of plans are being revised to take into account the degree of biological interest of each contemplated mission. The basic functions of the PQO, however, are clear, and can be grouped into three tasks:

(1) Promote and coordinate planetary quarantine research.

(2) Provide standards, methodologies, and other necessary supports for assessing probabilities of contamination.

(3) Set criterion for approval of mission plans.

The purposes of tasks 1 and 2 are easily understood. Task 3 is more interesting to describe here because it parallels, at a higher level in the space exploration decision hierarchy, the task of the mission planner.

The problem of selecting the optimal space exploration program for biological exploration of the solar planets and satellites is too formidable to be tackled as a whole. Decomposition methods must be used to reduce the problem to a series of simpler subproblems. At the same time, a system of coordination signals must be devised to insure that the optimization of each subproblem will lead to the optimum for the entire program. Thus the analysis of a space exploration program can be decomposed into analyses of mission projects, which in turn can be decomposed into analyses of guidance and control decisions.

2.2.5 NASA's Probability Allocation Process: The Pie Analogy

The decomposition and coordination mechanism used by NASA is to allocate maximum probabilities of contamination to each component of a space exploration program. The PQO sets a maximum probability of contamination to each mission project, the mission allocation; project planners further divide this allocation among the various guidance and control decisions as probability suballocations. The mechanism is complicated by the dynamic aspect of the problem: outcomes become known as decisions are made.

It is tempting to describe this allocation process using a pie to represent the total probability of contamination available to the PQO for all U.S. missions. A mission manager submits his plans to the PQO and asks for a piece of the pie for his mission. Based on the mission description and the estimated total number and types of missions that will be flown during the period of biological interest, the PQO decides how large a piece of pie (mission allocation) the mission manager should receive. Following the flight, the mission manager must report how much of his piece he has actually consumed. Any amount left over from the mission allocation will be returned to the PQO to be shared by other missions.

The same analogy could be used for the mission planner trying to design his mission by first subdividing his piece of pie among numerous decisions.

There are many issues involved in partitioning the probability pie; for example:

- What is a mission's fair share of the pie?
- How can we know what has been left over?
- What happens if the mission eats the whole pie, perhaps a thousand times over?
- What provisions should be made for unexpected guests?

The analogy of the pie, although definitely an oversimplification, may at first seem appropriate, when, in reality it is dangerous and confusing. The concept of probability is simple, but subtle--more so than pie.

2.3 Probability in the Context of Quarantine Planning

2.3.1 The Concept of Probability

A probability is a number expressing a state of knowledge about the occurrence of an uncertain event; it does not represent a physical entity, like a weight or volume. Thus, two individuals who do not share the same information will generally assess two different probabilities to the occurrence of the same event.

For example, assume that experts reach a consensus on a definition of life. We shall say L represents that planet Mars bears this kind of life today. Assessments of the probability of L would certainly vary widely among experts because it is practically impossible for two experts to share exactly the same information about Martian life. Note, however, that the subjective aspect of the probability concept does not prevent decisions being based on probability judgments. Indeed, the possible existence of Martian life is being investigated, and the resources spent to that end can only be justified by subjective probabilities of obtaining certain answers.

2.3.2 Probability Updating: Bayes' Rule

This example leads us to a second consequence of the subjective character of the probability concept: A probability assessment must be revised when new information becomes available. The qualifications "prior" and "posterior" are used to distinguish the probabilities of the same event before and after receiving new information. By their very nature, probabilities are subject to the rules of logic, which guarantee a consistency between probabilities reflecting different states of information. The revision of a prior probability to a posterior probability must obey a precise logical rule known as Bayes' Rule.

To illustrate this rule, suppose a life detection experiment is carried out on Mars and does not reveal the presence of life (as defined by L). We shall call N this negative result. The negative result does not necessarily mean that life does not exist: the experimental result is highly localized and may be subject to distortions. What then, is the posterior probability of life (L) given the negative result (N)?

If we denote by p(L) the prior probability of an L and by $p(L|N)$ the posterior probability of L given the occurrence of N, Bayes' Rule indicates that

$$p(L|N) = p(L) \times \frac{p(N|L)}{p(N)} \quad .^{*} \tag{2-1}$$

In words, the posterior probability of life given the negative experimental result must be equal to the prior probability of life multiplied by the resolution of the life detector as expressed by the likelihood ratio $p(N|L)/p(N)$. In general the denominator of this ratio cannot be assessed directly but can be further decomposed into the probability of a negative result given either that life exists (L) or that life does not exist ($\bar{L}$). Mathematically:

$$p(N) = p(N|L)\ p(L) + p(N|\bar{L})\ p(\bar{L}) \quad . \tag{2-2}$$

As an illustration, if

$$p(L) = 0.01 \qquad \text{and therefore } p(\bar{L}) = 1 - p(L) = 0.99$$

$$p(N|L) = 0.2$$

$$p(N|\bar{L}) = 0.999$$

we obtain

$$p(L|N) = p(L) \times \frac{p(N|L)}{p(N|L)\ p(L) + p(N|\bar{L})\ p(\bar{L})}$$

$$= 0.01 \times \frac{0.2}{0.01 \times 0.01 + 0.999 \times 0.99}$$

$$p(L|N) = 0.01 \times 0.2 = 0.002.$$

---

*Bayes' Rule is a straightforward extension of the definition of conditional probability. Consider the joint event N,L:

$$P(N,L) = P(L|N)P(N)$$
$$= P(N|L)P(L) \quad .$$

If $P(N) \neq 0$, we can divide by it, giving Eq. (2-1). Bayes' Rule is nothing more than a consistency requirement on probability assessments reflecting different states of information.

The posterior probability of life given the negative result is about five times smaller than the prior probability of life.

2.3.3 The Prior Interpretation of a Probability Allocation

A probability allocation (or probability constraint), as defined in current procedures, should not be confused with the probability of an event: A probability allocation is a constraint arbitrarily imposed on a decision or a sequence of decisions. As such, it is not subject to any specific revision rule, and is only required to have a clear meaning. In particular, we must be able to verify whether or not a decision or a sequence of decisions exceeds a probability allocation.

The meaning of a probability constraint is simple when applied to a single decision: Based on the state of knowledge prevailing at the time the decision is made, the probability that the decision will lead to contamination should be less than the constraint value.

It is more difficult to verify that a sequence of decisions does not violate a probability constraint. To describe a sequence of decisions, we need to use the concept of a strategy. A strategy is a complete specification as to the decision to be made at each stage in the sequence; for later stages the decision may be contingent on outcomes of (or information obtained after) earlier decisions. Hence, the first decision in the sequence may be consistent with many alternative strategies that differ in the choices to be made at subsequent stages.

The probability that a sequence of decisions will cause contamination is assessed initially with the state of information prevailing at the time the first decision must be made. This decision must take into account the choice among strategies available for subsequent decisions. It requires a model of the sequence of decisions and probabilistic consequences and a rule specifying which alternative will be taken at each subsequent decision stage.

If the probability of contamination associated with the initial decision does not exceed the probability constraint, and if the intention is to follow the strategy as set forth in the prior analysis, it can be said that the probability constraint has been satisfied. We shall refer to this interpretation that the constraint has been satisfied as the prior interpretation principle. We believe it is the only consistent rule under which a probability constraint is meaningful in a sequence of decisions made under uncertainty.

The initial strategy may, however, be abandoned for a new strategy during the execution of the decision sequence. A new strategy may become available, or new information not included in the prior analysis may modify the state of knowledge about future events in an unforeseen way. The planning process cannot verify that the probability constraint will not be exceeded in future states of information that may arise. Good planning cannot guarantee good outcomes in all cases.

2.3.4 Illustration: Arbitrariness of Probability Suballocations

A simple example, shown in Figure 2.1, will illustrate the prior interpretation principle. A mission calls for the exploration of a planet by probes. The mission planner proposes the following strategy: Send a first probe with a probability of contamination of 5×10^{-5} and a probability of satisfactory results of 0.5. Stop if the first probe provides the desired results; otherwise, send a second probe with a larger probability of obtaining results and a probability of contamination of 10^{-4}. The prior probability of contamination for the entire mission is

$$5 \times 10^{-5} + 0.5 \times 10^{-4} = 10^{-4} \quad .$$

When a sequence of decisions is subject to a probability constraint, there is generally an infinity of possible probability suballocations for each decision. Figure 2.1 illustrates this point. In it, Q denotes the mission allocation for the two probes; q_1, the probability suballocation for the first probe; and q_2 and q_3, the probability suballocations for the second probe knowing that the first has been successful (probability p) or unsuccessful (probability 1 - p), respectively. The mission allocation imposes a single constraint on the three suballocations, namely,

$$q_1 + pq_2 + (1 - p)q_3 \leq Q \quad . \qquad (2\text{-}3)$$

There is therefore a wide choice of probability suballocations q_1, q_2, and q_3. To each set of suballocations will correspond a choice of probes and a resulting expected scientific value for the mission. Unfortunately, the planetary quarantine requirement (mission allocation) offers no general rule indicating which assignment of probability suballocations will lead to the most valuable mission. Likewise, at a higher level, there is no indication about which choice of mission allocations will lead to the most valuable space exploration program.

FIGURE 2.1

PRIOR INTERPRETATION: A TWO-PROBE MISSION

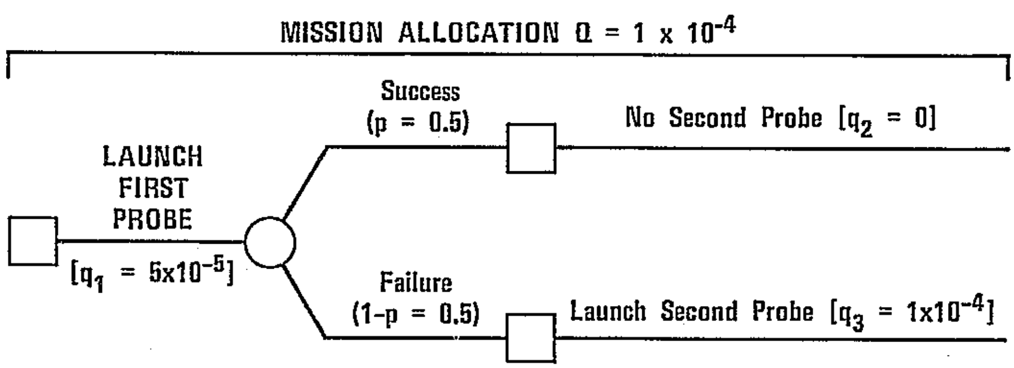

□ = Decision node; ○ = Chance node; ( ) = Probability;

[ ] = Probability of contamination

NOTE: The prior probability of contamination is: $5 \times 10^{-5} + 0.5 \times 1 \times 10^{-4} = 1 \times 10^{-4}$

2.3.5 Revision of Probabilities: Reallocation is Wrong

How can we allocate, revise, and reallocate probabilities of contamination? The answers to these questions depend on the interpretation given to the maximum probability of contamination inscribed in NASA's planetary quarantine policy.

That the probability of contamination can be revised after and even during a mission flight is evident from the definition of a probability as representing a state of knowledge. If a mission flight is monitored and information is gained on its results, the posterior probability that the mission has caused contamination will generally differ from the prior probability. To be sure, mission controllers will often know with certainty whether the spacecraft has or has not impacted the target planet, thus reducing considerably the uncertainty of the contamination event. For example, the probability that contamination of Mars resulted from Mariners 3, 4, 6, 7, and 8 has been revised toward zero because indications are that none of these spacecraft touched Mars or its atmosphere.

Can the original mission allocations for these Mariner flights be declared "unused" and reallocated to new missions? The proponents of such a reallocation should first recognize that it will not necessarily result in less restrictive planetary quarantine constraints. For if reallocation is valid when a probability allocation has not been "used," it should also be valid when a probability allocation has been exceeded.

Imagine that one of the early Mariner fly-bys of Mars had crashed on the red planet and that the revised probability of contamination were larger than 1×10^{-3}. Would the proponents of reallocation admit to having "used" more than their allocation and therefore violated the international planetary quarantine agreement? Or will they hold that they had made their decision in good faith and can therefore launch further missions to Mars? The first interpretation seems absurd: How would the PQO be held responsible for excessive posterior probabilities? Doubtless, it would not be the first instance of somebody being judged for bad outcomes rather than bad decisions but this practice should certainly not be encouraged.

If retained, this interpretation would only leave two options to the PQO:

(1) Reject all missions having the slightest chance of contamination.

(2) Ignore the previous missions results.

The first option is a return to the inoperational absolute sterilization standard; the second option requires some explanations.

A reallocation process could conceivably place negative incentives on gathering information on mission results every time a chance existed that the probability of contamination had exceeded the mission allocation. Knowing such information, the mission manager would have to make a costly corrective maneuver, if possible, or the PQO would have to impose more severe constraints on subsequent missions. (See the appendix for more details).

Reallocation based on posterior probabilities is also clearly inconsistent with the prior interpretation of planetary quarantine. To illustrate this point, imagine someone willing to play forever at Russian roulette with a six-chamber revolver and a single bullet. Although this individual faces one chance in six of being killed every time he plays, his posterior probability of being killed if the gun doesn't fire is always null. Most onlookers will agree, however, that in the long run the probability that he will shoot himself is 1 and not 1/6 as the proponents of reallocation based on posterior probabilities would have us believe. With or without the prior interpretation of planetary quarantine, the idea of a reallocation process based on posterior probabilities must therefore be rejected.

2.3.6 Illustration With a Single-Decision Fly-By Mission

The example illustrated in Figure 2.2 will help our understanding of reallocations based on posterior probabilities. The mission is a single-decision fly-by mission, which upon launching has a 6×10^{-5} probability of ending on an impact trajectory. The probability of contamination on impact is assumed to be 1. The prior probability of contamination associated with the mission is therefore simply 6×10^{-5}. If the mission allocation is larger than 6×10^{-5}, say 10^{-4}, the mission can be launched.

Proponents of reallocation based on posterior probabilities claim that if the mission has not resulted in an impact, the revised probability of contamination is zero and the unused mission allocation can be reallocated to other missions. Thus, it is likely that, say, 20 such fly-bys could be sent without causing contamination. In fact, the probability that such a program will not cause contamination is exactly the product of the probabilities that each mission will not cause contamination, that is,

$$(1 - 6 \times 10^{-5})^{20} \simeq 1 - 1.2 \times 10^{-3} \quad .$$

FIGURE 2.2

A SINGLE-DECISION FLY-BY MISSION

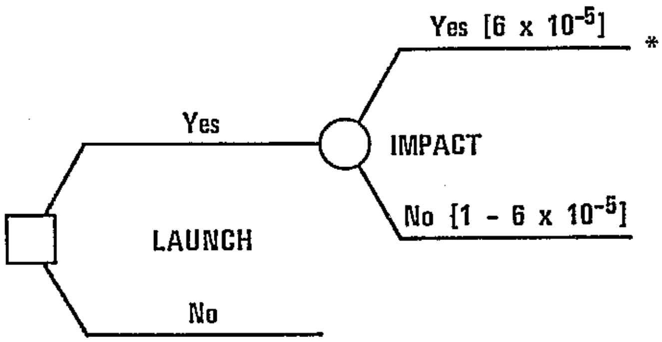

□ = Decision node; ○ = Chance node; * = Contamination;
[ ] = Probability of contamination

In other words, the strategy consisting of sending a series of up to 20 similar fly-bys as long as none has caused contamination <u>will</u> cause contamination with a probability of nearly 1.2×10^{-3}. Proponents of reallocations cannot share the prior interpretation of the planetary quarantine policy; the two views are incompatible.

2.3.7 Posterior Probabilities and Planetary Quarantine Policy

It is felt intuitively that knowing whether or not a planet has been contaminated (posterior probability) should influence the planetary quarantine requirements imposed on future missions. Because posterior probabilities cannot be used for making reallocations, what then should be the effect of posterior probabilities on planetary quarantine requirements?

The answer lies in an extension of the international agreements to cover multiple contaminations of a planet. As long as the probability of a contamination is very small, the probabilities of multiple contaminations are negligible. However, should a posterior probability of contamination be close to 1, chances are that the next contamination event will be the second one. A second contamination may have different consequences from a first. So far, concerns over multiple contaminations have not been expressed in planetary quarantine policies.

2.3.8 Suggestions by Previous Authors

J. O. Light [4] was among the first to recognize the shortcomings of the "nominal plan - preset suballocations" analysis. The prior interpretation of planetary quarantine requirements does not impose such an analytical procedure. To quote from Light:

> The entire formulation of the non-contamination constraint is nonsense unless we devise the strategy for making in-flight decisions before the flight, correctly incorporate this strategy into our probability model, and then adhere to this strategy. A beneficial by-product of this formulation of the non-contamination policy will be the necessity to carefully consider all the standard and non-standard conditions which we might encounter....

Light makes two additional points: (1) planetary quarantine requirements are insufficient to select a unique mission strategy, and (2) some feasible strategies are preferable to others.

Considerations other than strictly planetary quarantine constraints should be introduced into mission planning to determine the optimal strategy among all feasible strategies. Thus, the scientific value of a mission should play an important role. Everything else being equal, a strategy with a large expected scientific value and a low prior probability of contamination is desirable. However, a smaller scientific value must often be traded for a lower probability of contamination; hence, the desirability of assessing a trade-off value.

Of course, one might also go as far as to propose not only supplemental objectives but also additional constraints to remedy the laxity of prior interpretation, which would eliminate heretofore feasible strategies.

Thus, under prior interpretation, there may be acceptable strategies involving conditional probabilities of contamination far exceeding the mission allocation. If the mission allocation is interpreted as representing a trade-off between expected scientific return and risk of contamination, how can these higher conditional probabilities be justified?

An example may illustrate the importance of the probability that a contingent event will occur. Suppose that a guidance decision, if taken, gives a spacecraft a probability of 10^{-6} that both its propulsion system will fail and it will be on an impact trajectory. Since this probability is below the mission allocation of 10^{-4}, the decision is taken; subsequently it is learned that the spacecraft is on an impact trajectory and the propulsion system has failed. Clearly this would be a bad outcome, but would it be the result of a bad decision? If the probability assessment of 10^{-6} is a reasonable reflection of the information available at the time the guidance decision was made, we would say the decision was good. The conditional probability of contamination given an impact trajectory and rocket failure was approximately equal to 1, but the probability that the condition would occur was so small that the constraint was met.

Some authors, such as Light and Chamberlain, have been more concerned about the rigidity of the probability suballocation process. R. G. Chamberlain [5] has proposed a method intended to free mission planners from problems associated with probability reallocations. Chamberlain's method is based on an identification of states of nature and a decision criterion that does not take into account the probabilities of these states. Our understanding of his method is as follows:

- Suppose that a mission consists of m stages at which we will need or have the opportunity to execute a guidance decision. Since each decision may be either executed or not, there are at most 2^m subsets of possible executions.* At the planning stage, we are uncertain about which of these subsets we will obtain; each subset will have an associated probability.

- Following Chamberlain's proposed procedure, we make the guidance decisions so that the mission allocation will not be exceeded for any execution subset, regardless of its probability.

Chamberlain's proposed procedure was intended to be used only within certain bounds, for in some cases it would be much more restrictive than the prior interpretation principle. An example will illustrate this problem. Consider the last two maneuvers of a fly-by mission. Following the first maneuver the propulsion system will fail with probability p. If the propulsion system does not fail, we assume the second maneuver is executable. We denote by q_1 and q_2 the probabilities of putting the spacecraft on an impact trajectory as a result of each of the two maneuvers. The probability of contamination given impact is approximately 1. The probability allocation for this sequence of 2 maneuvers is denoted by Q.

According to Chamberlain's proposed procedure the probability allocation must not be exceeded whether or not the propulsion system fails. Hence the two constraints

$$q_1 \leq Q$$

$$q_2 \leq Q \quad .$$

These two constraints are always more restrictive than the single constraint that the prior probability of contamination for the two maneuvers be less than the probability allocation, that is,

$$pq_1 + (1 - p)\, q_2 \leq Q \quad .$$

Typical numerical values might be $Q = 2 \times 10^{-5}$ and $p = 10^{-4}$ and therefore the constraint imposed by Chamberlain's procedure on the first maneuver

---

*It may be known in advance that some combination of guidance decisions are unfeasible.

($q_1 < 2 \times 10^{-5}$) might force to bias the spacecraft trajectory away from the target planet in a manner that is not imposed by the prior interpretation principle (e.g., $q_1 = 0.1$ and $q_2 = 10^{-5}$ is feasible).

Chamberlain's method should therefore not be used without care. Chamberlain believes that for the majority of the planning problems encountered in the past, his procedure, if carefully applied, was capable of bringing some simplifications without unduly overconstraining the missions. We believe that the main advantage of Chamberlain's method is to force explicit consideration of contingent states of nature. However, any planning simplification brought about by this method is obtained at the expense of adding new constraints and ignoring potentially important information about the probability of occurrence of each state of nature. This may become inappropriate for the planning of more complex missions such as multiple-encounter missions.

2.3.9 Probabilities in the Context of Quarantine Planning: A Summary

Concern over planetary contamination has resulted in the assignment of probability constraints for mission planning. The intentions behind the policy are admirable, but the procedures for implementing the policy need revision.

Probabilities are the reflection of states of knowledge that may evolve in time; it is a subtle task to keep them within limits. To give an operational meaning to a probability constraint, one must specify the probabilistic event to which it applies and the state of knowledge used as a basis for assessing the probability of that event.

When a single decision or a sequential decision is subject to a maximum probability of contamination constraint, it is clear that the constraint can be applied only to a prior probability of contamination, as it may be assessed when the mission strategy is planned. To state the contrary, i.e., to judge a decision on its outcome, would be inoperational, and policies of this type lead to contradiction from given objectives.

We have stated a prior interpretation principle that defines precisely how the probability of contamination of a decision sequence can be assessed. This prior interpretation is widely accepted, although not always defined precisely nor applied systematically.

2.4 The Elements of Rational Quarantine Planning

2.4.1 Explicit Goals and Planning Horizon

Exploration of the planets has become an ongoing business rather than a dramatic new scientific capability. Ideally, the planning of exploration efforts should be carried out through a formal process that makes goals explicit and allows a particular mission to be seen in the context of a comprehensive program. The desirability of such planning was clearly evident to those who first proposed a quarantine policy [6]. Yet the annual nature of Congressional appropriations and the large role of scientific advisory groups in NASA has led away from a formal process for long-range planning. This reality complicates the job of the mission planner because he does not receive clear direction on goals and objectives. The problem is compounded by the effect of his present mission on future exploration efforts, which may be important yet difficult to define.

A major difficulty is that there is no established timetable for when certain types of missions may be attempted, even where the feasibility of such missions may be easily assessed. Attractive five-year programs may be potentially undesirable if considered as the first five years of a 20-year program. Conversely, the first five years of an optimal 20-year program may seem to provide little scientific information at a high cost because they prepare the way for the next 15 years. Selecting a distant planning horizon is usually preferable to selecting a close one; it provides a less restrictive framework for planning space exploration programs.

A second difficulty is to identify the possible outcomes of space exploration programs and to assign values to them. Unfortunately, the difficulty increases with the span of the planning horizon. Assigning values to well-defined, short-term research efforts is already an arduous task. The economic value of scientific information is not directly measurable; its assessment requires consideration of all potential decisions that might benefit from the new knowledge. An additional complexity of long term research efforts is that early findings will raise new questions whose answers might be deemed important enough to redirect the research.

2.4.2 Formal Criteria for Selecting Among Space Exploration Programs

To summarize; translating concerns over possible contamination into a specific criterion for accepting or rejecting planetary exploration programs requires:

(1) The specification of a planning horizon.

(2) The identification of outcomes over the planning horizon and the assessment of values, at least on a relative scale, to these outcomes (e.g., contamination penalty relative to scientific value).

A space exploration program may then be represented as a sequence of decisions and probabilistic outcomes.* Outcomes comprise the successive states of the system (achievements in the exploration program) and information affecting the state of knowledge about the spacecraft and its dynamics and about celestial objects of interest. The selection of an optimal program then becomes a well-defined dynamic optimization problem, amenable to solution by standard methods such as dynamic programming.

Interestingly, solutions of dynamic problems share a fundamental property: The optimal course of action at a given time depends only on the state of knowledge about the state of the system and its dynamics at that time; it does not depend on how this state of knowledge was reached. Thus the present probabilities that a planet has been contaminated or will be contaminated are factors pertinent to the planning of new missions. Prior probabilities of contamination are only important as they affect the current state of knowledge, i.e., the current probability that a planet has been contaminated or will be contaminated. Probability allocation procedures currently in use do not share this property.

Dynamic programming provides powerful concepts and a logic for solving certain types of sequential decision problems, i.e., problems in which a sequence of decisions must be made with each decision affecting future decisions. The basic principle of dynamic programming is that in an optimal strategy for n sequential decisions, the last (n - 1) decisions also form an optimal strategy. An optimal strategy might therefore be constructed in a recursive manner if we start by determining the optimal strategy for the last decision.

To that end, state variables, i.e., variables whose values completely specify the instantaneous situations of the sequential decision problem at any stage, must be defined. The maximum permissible probability of contamination is part of the specifications of space exploration decisions and should therefore be included among the state variables.

---

*For an example of this formulation, see Matheson and Roths' paper [7] and a report on the Voyager project by SRI [8].

However, at any decision stage except perhaps the first one, there is no unique assignment of probability constraints on the remaining decision. Instead, there is an infinite number of assignments satisfying the probability constraint imposed upon the complete sequence of decisions (recall Section 2.3.4 and Figure 2.1). Therefore, not one but an infinite number of dynamic programming problems exist and there is no general rule to indicate which choice of probability constraints leads to the most valuable solution.

In Section 3 we propose a new methodology that assesses values to all the outcomes of a mission or a space exploration program. An optimal strategy can then be determined by dynamic programming. This new methodology can be applied with limits imposed upon strategies by probability constraints, but it does not require such limits.

3 A NEW METHODOLOGY

3.1 Design Principles for a New Methodology

3.1.1 The Rationale for a New Methodology

Current and contemplated planetary exploration missions are raising intricate new planning problems. Early planetary missions consisted of fly-bys with little flexibility in the choice of trajectories, limited scientific payloads, and often relatively short and well-defined useful lives. Nominal plans were devised for these missions and closely followed in the absence of unexpected catastrophic failures.

Orbiters, landers, and now multiple-encounter missions require more than a nominal plan. The number of guidance maneuvers has increased. Scientific payloads have become more sophisticated, and many experiments are competing for the choice of trajectories and the use of common resources such as energy. Planning horizons must be extended to several years for the exploration of outer planets and sometimes become difficult to define. At the same time, a greater number of failure modes and execution errors of varying severity can be expected. Under these circumstances, strict adherence to a nominal mission plan may be less than optimal, if not impossible. It is preferable to develop contingency plans, that is, to analyze foreseeable circumstances and determine the best course of action in each specific case.

The need to consider contingent events and decisions is particularly apparent in the assessment of contamination probabilities. A simple contamination analysis of a nominal mission plan is sufficient, provided that the nominal plan be followed and that no new information pertinent to the risk of contamination be gained during the mission flight. These conditions led to the development of standard expressions of planetary quarantine requirements and standard contamination analysis procedures. But advanced missions do not meet these conditions; we must now analyze contingency plans that take into account not only failure modes and execution errors but also information about the region of space being explored. For example, the discovery of high radiation belts surrounding a planet can affect both the bioload and the reliability of a spacecraft traversing them. The probability of the spacecraft impacting the planet and the probability of contamination given impact must be revised accordingly.

For these advanced missions, the planetary quarantine policy in its current formulation is difficult to implement both conceptually and practically. Conceptually, the meaning of a probability constraint imposed on a sequence of decisions is subtle. A detailed analysis of the various states of knowledge and strategies intended at the time each decision is executed must be carried out. Practically, the assignment of a probability constraint may be inconsistent with the logic of sequential decision making.

We believe that, with the experience of current procedures, it is now time to develop a new approach to planetary quarantine on sound bases. Although many interpretations of the planetary quarantine policy can be made, none of these can eradicate the fundamental difficulties inherent in the chance-constrained formulation. These difficulties will become increasingly apparent in planning the complex missions contemplated for the 1980s. A fresh approach is therefore required.

The current policy, however, cannot be modified completely and instantaneously. It has permitted, so far, remarkable scientific achievements and very low probabilities of contamination. A revision of the U.S. planetary quarantine program and ultimately of the COSPAR agreement will require extensive debate among all interested parties and preferably some trial cases of the proposed revisions.

A new approach must therefore be general enough to supersede ultimately the current chance-constrained formulation and, in the meantime, be compatible with it. That is, a new methodology should permit determining a feasible and rational strategy for a project now subject to a maximum permissible probability of contamination.

3.1.2 A Rational Allocation of Resources Through Explicit Trade-Offs

There are many ways to reduce the probability that a mission will lead to contamination. For example, we can reduce the bioload of the spacecraft, reduce the probability of an accidental impact by biasing the trajectory away from the capture circle of the target planet, improve the reliability of the propulsion and guidance systems, and so forth. All these actions tend to increase costs and time delays or reduce flexibility and expected scientific returns, i.e., they correspond to the use of resources measured on a mission value scale. On the other hand, the reduction in contamination probability may have a direct scientific value (e.g., minimizing the probability of a false result in a life detection experiment) or, more often, an indirect value for subsequent missions.

The first problem faced in mission planning is therefore to develop a strategy that will achieve a given reduction in contamination probability at a minimum cost, i.e., for a minimum reduction of the expected value of a mission exclusive of contamination considerations.

A solution to this problem requires that the costs and scientific values associated with alternative strategies be measured on a common scale, a difficult but inescapable task. A solution technique must rely on trade-offs well-defined in simple situations if it is to help balance the use of resources in complex situations. At the guidance decision level, a relative scale may be sufficient; the use of each resource may be expressed as a fraction of the expected scientific value of a nominal flight. At the mission design level, a dollar scale may be preferable because direct costs play a paramount role in the choice among mission design alternatives.

As the probability of contamination of a mission is decreased, each incremental reduction becomes more difficult (costly) to achieve. The second problem faced in mission planning is therefore to determine at what point a marginal reduction in the probability of contamination would entail an unacceptable increase in mission cost (or reduction in mission value). Determining this break-even point requires that the dangers of contamination be measured on the same scale as the resources used to avoid contamination.

3.1.3 Illustration of the Trade-Off Between Expected Scientific Value and Probability of Contamination

We can illustrate the relationship between the use of a contamination value and a probability constraint by the representation of mission strategies given in Figure 3.1. Each strategy is characterized by two numbers, a probability of contamination (on the horizontal axis) and the expected value of scientific achievement and other exploration objectives, excluding quarantine (on the vertical axis). We consider all possible mission strategies, and we identify the set of strategies so that for each of them, no other strategy lies above or to the left; that is, for the same probability of contamination, no other strategy has a higher expected value of scientific and other objectives, and for the same expected value of scientific and other objectives, no other strategy has a lower probability of contamination. This set of strategies is widely known in the economic theory that deals with trade-offs among different objectives, and following the practice in economics, we shall refer to it as the Pareto-optimal border.

FIGURE 3.1

REPRESENTATION OF MISSION STRATEGIES: EXPECTED VALUE VERSUS PROBABILITY OF CONTAMINATION

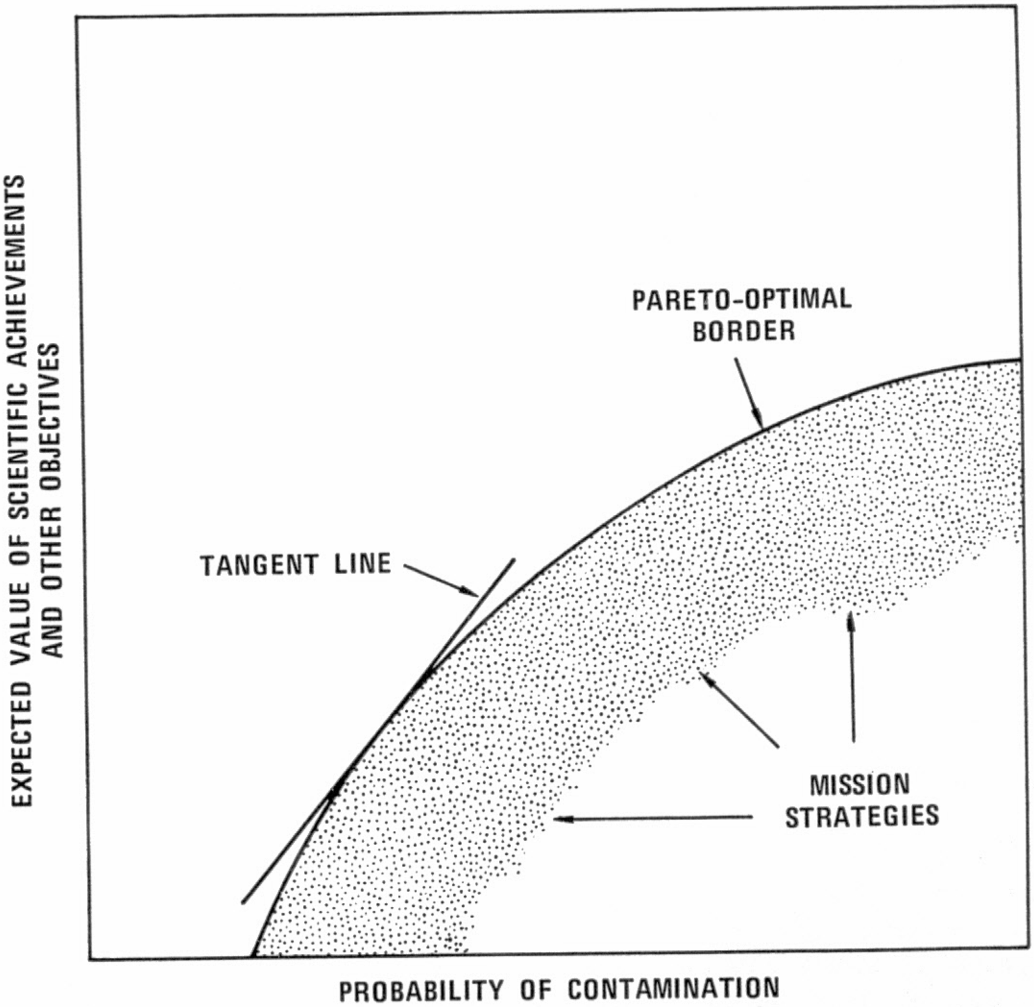

NOTE: Slope of tangent line is incremental change in expected value divided by incremental change in probability of contamination.

In choosing among strategies for the mission, we will want certainly to choose one on the Pareto-optimal border; for any strategy that does not belong to this set, one can find a strategy on the border with either a lower probability of contamination or a higher expected value by moving up and to the left.*

Which one do we choose? Our choice involves a trade-off between increments of expected value on the vertical axis and increments in the probability of contamination on the horizontal axis; the ratio of these increments is the slope of the tangent line to the Pareto-optimal border. In turn, this slope at a point on the Pareto-optimal border defines a contamination penalty if it is assumed that the values of the increments balance each other. By dividing the increment in expected value by the increment in contamination probability, we obtain a loss per unit of contamination probability; if this unit is made equal to one, it is the penalty associated with contamination. The increasing slopes from right to left on the Pareto-optimal border correspond to increasing contamination penalties; a given marginal decrease in contamination probability corresponds to larger and larger marginal decreases in expected value.

3.1.4 The Relation Between Value Trade-Off and Probability Constraint

An assessment of a contamination penalty is implicit in the current procedures. If the equivalent of \$10 million is spent in sterilization costs--including costs ascribed to any resulting time delays and reduction in reliability--to decrease the probability of contamination of a given mission, say from 1×10^{-4} to 5×10^{-5}, then the implicit penalty associated with contamination is

$$K = 10/(5 \times 10^{-5}) = 2 \times 10^{5} \text{ million} = \$200 \text{ billion} \quad .$$

More generally, if a contamination penalty is not available for planning but a probability constraint is available and a set of possible strategies is given, then a penalty can be inferred from the constraint.

---

*Practically, the assessment of an expected scientific value and a probability of contamination for each strategy is always proximal, and strategies close to the optimal border should not be neglected. Rather, a more detailed assessment of strategies on or close to the optimal border may be justified before a final decision is made.

Suppose an allocation Q is given as the maximum permissible probability of contamination. Then the strategy space of Figure 3.1 is divided into "accept" and "reject" regions as shown in Figure 3.2. The mission planner then wishes to select the highest expected value in the "accept" region as the preferred strategy. This strategy has the highest value consistent with the probability constraint. The slope of the Pareto-optimal border at the selected strategy indicates the contamination penalty implied by the mission allocation.

An explicit assessment of the contamination penalty would be preferable. When only one space exploration project is considered and has a well-defined design (fixed set of physically possible strategies), the assessment of a probability allocation implies a contamination penalty. If, however, an uncertain number of projects is contemplated over an ill-defined time horizon, the assignment of coherent probability allocations (coherent in the sense that the use of resources will be well balanced) is practically impossible. On the other hand, the assessment of a contamination penalty should not be influenced by the number and type of missions planned during a given period. It should depend only on knowledge about the target planet and may be revised upwards or downwards as planetary exploration progresses and new information is gained. For example, it may be that at some point contamination may become desirable as we attempt to create or modify a planet's biota.

The assessment of the value of contamination penalty is particularly challenging. To start with, contamination must be defined. Contamination might be a catastrophic, or irrelevant, outcome, or even an interesting development, depending on the planet being contaminated, the extent and nature of the contamination, and its timing in the exploration program. The assessment of the value of contamination is doubly indirect; it depends on the value of scientific information that may or may not be obtained as a result of contamination. As a practical matter, the ratio of the contamination penalty to the expected value of the scientific returns of a space exploration program may be easier to determine than either of these two terms taken independently.

3.1.5 A Powerful Solution Technique Using Dynamic Programming

Another advantage of assessing a contamination penalty rather than a probability allocation is that it simplifies determining an optimal strategy. A chance constraint in itself is insufficient to define the value of a strategy. Furthermore, it offers no guide to decompose a complex problem into simpler subproblems. On the other hand, a contamination penalty together with explicit values assigned to other mission

FIGURE 3.2

ROLE OF CONTAMINATION PENALTY IN THE NEW METHODOLOGY

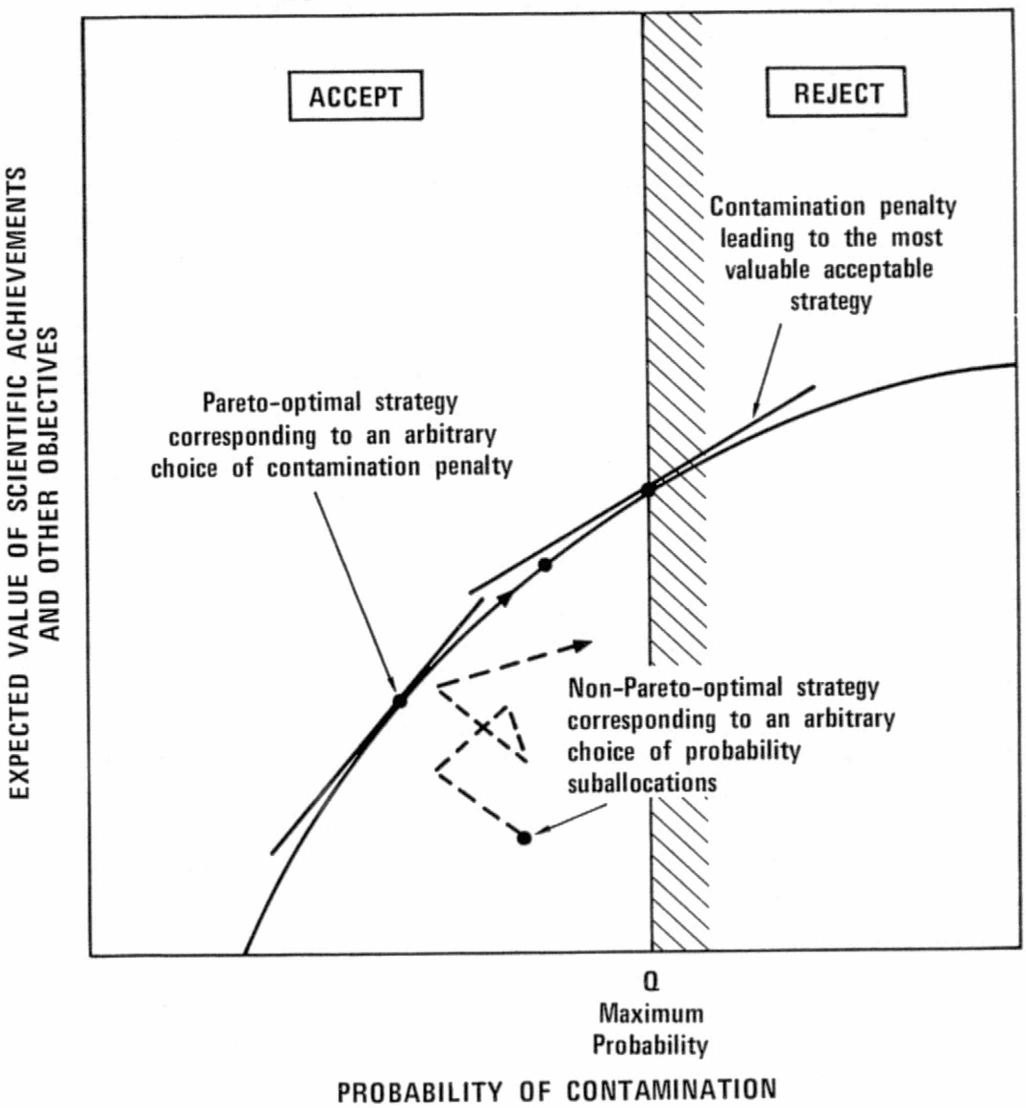

outcomes permits strategies to be evaluated. Mission planning can be characterized as a sequential decision problem, and standard methods can be used to determine the best strategy. Such methods include decision trees, which are a special case of dynamic programming (see R. A. Howard [9] for an introduction to dynamic programming).

In a sequential decision problem, each decision may affect future decisions; hence, the apparent difficulty of determining the first decision. With decision trees, in which alternatives and outcomes are discrete as opposed to continuous, this problem is easily solved. In chronological order: (1) all the decisions and their probabilistic consequences are laid out, (2) values are assigned to the outcomes, and (3) each decision is optimized, starting with the final decisions.* Then, progressing from last to first, each decision is made knowing the optimal strategy for the rest of the sequence. This recursive induction process terminates when the initial decision has been reached. Thus, instead of having to solve simultaneously for all the decisions in a decision sequence, each decision can be analyzed separately.

3.1.6 Iterative Reconciliation with Current Chance-Constraints

If guidance to the mission planner is available not in the form of explicit values but only in the form of a chance-constraint, the optimal strategy he obtains by assuming values for the outcomes and using dynamic programming must be checked against these constraints.

The general principle is simple. If the optimal strategy first obtained exceeds the maximum permissible probability of contamination, the contamination penalty has been underestimated and must be revised upwards. It is clear that, if the contamination penalty is increased, strategies with large probabilities of contamination will become less and less desirable. The optimal strategies will therefore have lower probabilities of contamination until finally, for large enough penalty values, strategies with zero probability of contamination (e.g., "do nothing" alternative) will become best.

---

*Practically, the steps are carried out first at a low level of detail. A first analysis will reveal the variables whose possible variations will most affect the entire results. These variables are then carefully reassessed (usually by building a more detailed model) and the problem is reanalyzed. This iterative process stops when the cost of additional modeling and analysis is no longer justified by potential improvements of the optimal strategy.

Conversely, if the optimal strategy first obtained does not exceed the probabilitity allocation, the contamination penalty may be revised downward in the hope of obtaining a strategy with a larger expected scientific value and a still-acceptable probability of contamination.

The convergence of this process is studied in more detail in Section 3.2.2.

3.2 Description of the Proposed Methodology

Based on the needs and the design principles we have just discussed, the methodology we propose is:

(1) Probability Consistency--To use the rules of probability theory, such as Bayes' Rule, to assess and revise probabilities, including the probabilities of contamination.

(2) Decision Formulation--To value explicitly project outcomes, including the possible event of contamination, and thus permit the use of dynamic programming to select optimal strategies.

(3) Iterative Reconciliation--To use a simple iterative process to obtain an optimal strategy compatible with a constraint on the probability of contamination of target planets or satellites while projects still remain subject to maximum permissible probabilities of contamination.

3.2.1 The Task Sequence

A step-by-step implementation of the proposed methodology includes the following tasks:

(1) Lay out the project decision tree. As noted above, a decision tree is a general description of the sequences of decisions and probabilistic outcomes, including the possibility of obtaining additional information. Care should be taken to incorporate all significant decision alternatives and to describe the information available and pertinent to the decision-making process at each decision stage.

(2) Assess probabilities to outcomes. The probability of each outcome is assessed on the basis of all the information available before the outcome occurs. (Bayes' Rule is used where applicable.)

(3) Assign values to outcomes. Value judgments fall into three classes: scientific values, costs, and contamination penalties for each planet (or satellite) visited during the project and subject to planetary quarantine requirements. If planning is carried out within an overall probability constraint, the contamination penalty will only serve as a coordinating signal, and the results will be insensitive to the choice of an initial contamination penalty. At the guidance decision level, the value assignments can be made relative to a nominal mission value. At the mission design level, dollar values are necessary.

(4) Find the optimal project strategy using dynamic programming. (A criterion of maximizing net expected value will be used, but a more general risk attitude may be incorporated into the framework of dynamic programming by using expected utility theory. See [10].)

(5) Compute the prior probabilities of contamination (for each planet or satellite) associated with the optimal strategy.

    (a) If a probability constraint is exceeded, increase the corresponding contamination penalty and return to step 4.

    (b) If a probability constraint is not reached, decrease the corresponding contamination penalty to allow for a possible larger scientific value and return to step 4.

(6) Terminate the procedure when the contamination penalties have reached the minimum nonnegative values compatible with the probability constraints.

3.2.2 Convergence of the Iterative Process

Mathematically, the selection of an optimal guidance strategy according to the currei chance-constrained formulation can be written as

$$\left.\begin{array}{l}\text{maximize } \overline{v(d)} \\ \text{over } d \text{ in } S \\ \text{subject to } p_j(d) \le Q_j \text{ for } j = 1 \text{ to } m\end{array}\right\} \qquad (3\text{-}1)$$

where

$d = (d_1, d_2, \ldots, d_n)$ = a strategy (vector of guidance decisions) for conducting a mission

S = the set of physically possible strategies

$v(d)$ = the value of the mission with strategy d [$\overline{v(d)}$ represents the expected value]

$p_j(d)$ = the probability of contamination of planet j with strategy d

Q_j = the mission allocation for planet j

m = the number of planets visited during the mission and subject to planetary quarantine.

Figure 3.3 graphically represents this problem when only one planet is subject to planetary quarantine. The set S of physically possible strategies has been plotted in an expected value (exclusive of contamination penalty) versus probability of contamination space. The mission allocation Q divides this set into a subset of feasible strategies (to the right of A). The strategy having the maximum expected value in the feasible subset, say d*(Q), is the optimal strategy. We will call the locus of d*(Q) as a function of Q the Pareto-optimal border (POB) of the set S. To any strategy not on the POB, there is a corresponding and preferable strategy on the POB; either a strategy on the POB will have a larger expected value for the same probability of contamination, or a lower probability of contamination for the same expected value, or both. The POB of set S drawn in Figure 3.3 is the curve OABCDE (we incorporate the "do nothing" alternative O in the set of strategies S).

The proposed value iteration procedure can be formulated mathematically as

FIGURE 3.3
DETERMINATION OF OPTIMAL MISSION STRATEGY

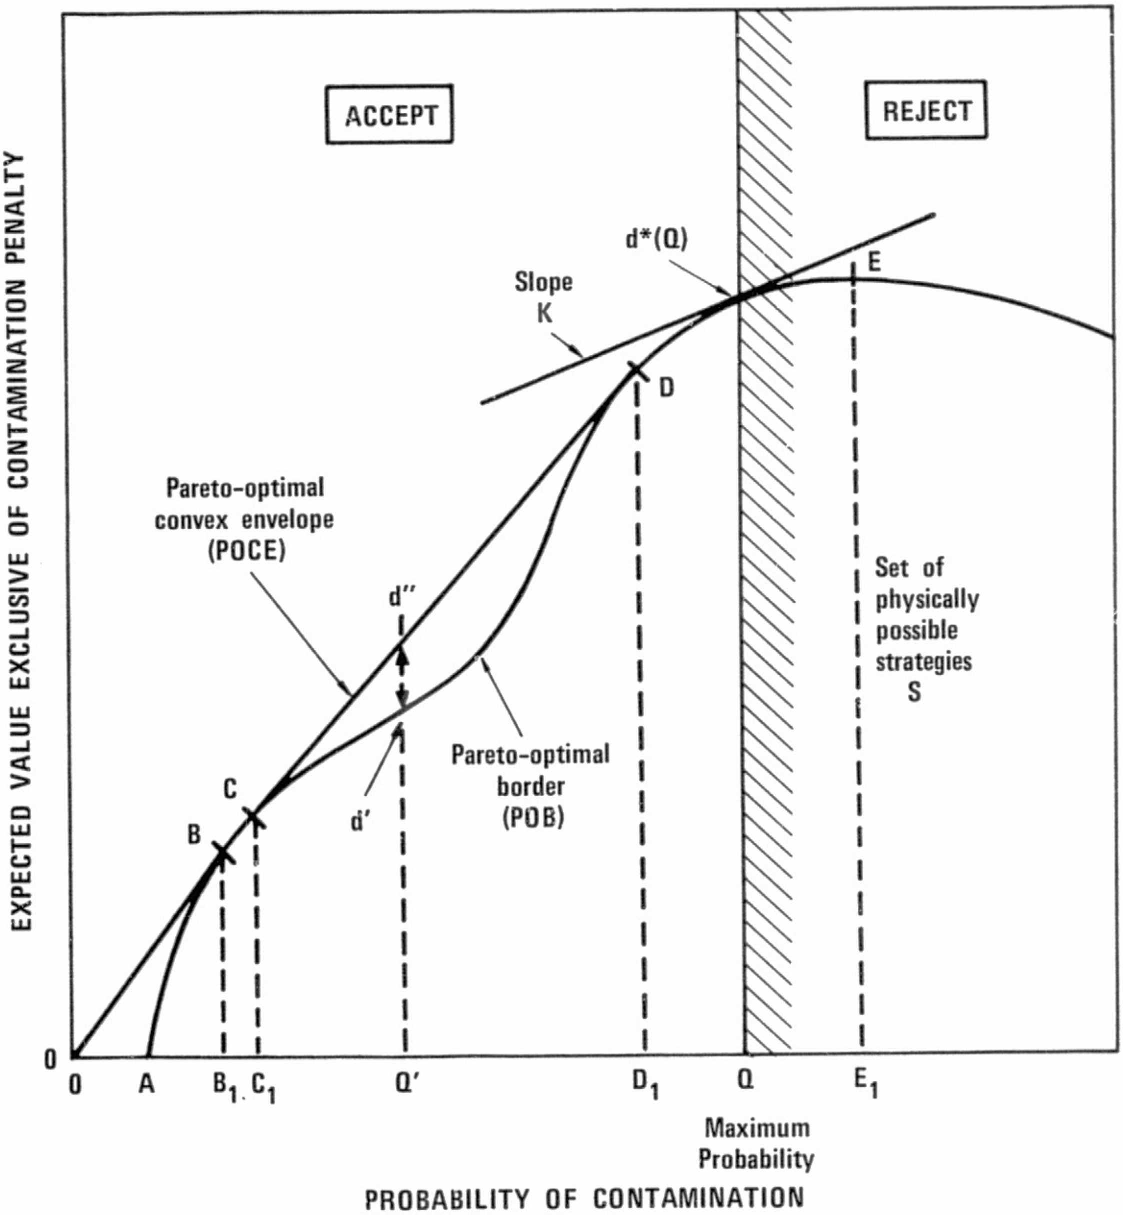

$$\left.\begin{aligned}&\text{maximize } \overline{v(d)} - \sum_j K_j p_j(d)\\&\text{over } d \text{ in } S\\&\text{where the contamination penalties}\\&K_j \geq 0 \text{ are chosen so that at the maximum } d^*\\&\qquad p_j(d^*) \leq Q_j \text{ for } j = 1 \text{ to } m \quad .\end{aligned}\right\} \tag{3-2}$$

This procedure is often referred to as a Lagrangian formulation, where the K_j's are the Lagrangian multipliers.

The proposed value iteration procedure has an easy interpretation. Define expected net value as the expected value of a mission minus the expected contamination penalty, that is, $\overline{v(d)} - Kp(d)$. In Figure 3.3, we see that strategies of equal expected net value will lie on a straight line of slope K. Consider, then, all nonintersecting tangents to set S with nonnegative slopes K. We shall call their envelope the Pareto-optimal convex envelope (POCE) of set S. It can be shown that the optimal strategies according to the proposed value iteration procedure must lie on this envelope. Their locus is therefore O, curve segment BC, and curve segment DE. Consequently, the chance-constrained formulation and the proposed value iteration procedure do not always lead to the same optimal strategy.

A sufficient condition for always obtaining the same optimal strategies with both procedures is that the POB and the POCE coincide, i.e., that the POB be convex. Mathematically, sufficient conditions for always obtaining the same optimal strategies with Eq. (3-1) and (3-2) is that (Kuhn-Tucker conditions)

$$K_j(p_j(d) - Q_j) = 0 \text{ for } j = 1 \text{ to } m \quad . \tag{3-3}$$

(See for example, [11], page 63, Section 1.3, for a discussion of convexity, Kuhn-Tucker conditions and their relations to Lagrange multipliers.)

3.2.3 The Issue of Convexity

If strategy set S is not convex, there are mission allocations for which the POB does not coincide with the POCE. For example, to mission allocation Q′ in Figure 3.3 corresponds Pareto-optimal strategy d′, below the POCE. Current NASA procedures might lead to the selection of strategy d′ given mission allocation Q′.

No contamination penalty, however, will lead to the selection of strategy d′. If the contamination penalty is slightly larger than the slope of the POCE at Q′ (slope of line segment CD), a strategy close to C on curve BC will be selected. If the contamination penalty is slightly smaller than the slope of the POCE at Q′, a strategy close to D on curve DE will be selected. If the contamination penalty is exactly equal to the slope of the POCE at Q′, strategies C and D are equally desirable.

Trade-off considerations permit verifying directly that a strategy on the POB that is not also on the POCE is always a poor choice. Consider again strategy d′ in Figure 3.3 and strategies C and D bordering the gap where the POB does not coincide with the POCE. The ratio of incremental value to incremental probability of contamination is larger between d′ and D than it is between C and d′ (the slope of line segment d′D is larger than the slope of line segment Cd′). Therefore, if strategy d′ is preferred to strategy C, then, a fortiori, strategy D should be preferred to strategy d′. Vice versa, if strategy d′ is preferred to strategy D, then, a fortiori, strategy C should be preferred to strategy d′: the expected value decrement per unit of contamination probability decrement is less between d′ and C than it is between D and d′. Strategy d′ or any strategy on the POB between C and D can never be preferred to both strategies C and D, i.e., the former are always poor choices.

To determine how much better strategy C or D is than any intermediate strategy on the POB between C and D, the expected contamination penalty must be taken into account, i.e., Eq. (3-2) and not (3-1) must be used. Graphically, the difference in expected net value is represented by the vertical distance between parallel lines of slope equal to the contamination penalty and drawn through the points C or D and d′. In particular, when the contamination penalty is exactly equal to the slope of line segment CD, both strategies C and D are preferable to strategy d′ by a quantity represented in Figure 3.3 by the vertical distance d′d″.*

3.2.4 Extension of the Strategy Set and the Value of Information

Introducing new strategies can never reduce the value of a project if the old strategies remain unchanged. Either one of the new strategies is optimal and increases the value of the project or an old strategy remains optimal and the value of the project is not affected.

However, the new optimal strategy and its incremental value will depend on what procedure is used to determine the optimum.

Figure 3.4 shows the addition of a new strategy set S_2 to an old strategy set S_1. According to the chance-constrained formulation, optimal strategy A is replaced by optimal strategy C, whereas, using a contamination penalty, strategy A should be replaced by strategy B (using the same trade-off K′ between incremental value and incremental probability of contamination). In the first case, the value increment is only C_2-A_2, whereas in the second case, the value increment is E_2-A_2. Conversely, if K″ is the correct trade-off, C is truly the new optimal strategy, but D and not A should have been the previously optimal strategy: the choice of A first and then C does not correspond to a consistent evaluation of the danger of contamination.

A way to extend the strategy set is to introduce data gathering devices and make decisions conditional upon the new information. The value of the new information will generally be very sensitive to the method used to determine optimal strategies. The contamination penalty approach permits us to determine the value of information in a consistent manner provided the contamination penalty is not changed during the course of the mission design.

---

*Point d″ can be interpreted as the representation of a random mechanism selecting strategy C with probability p and strategy D with probability (1 - p). The arbitrary nature of the probability constraint Q′ can be shown by the fact that a randomized strategy between C and D can offer a higher expected value than strategy d′ for the same probability of contamination. We are not advocating the use of random strategies for space mission planning; we see this result as an artificial and undesirable consequence of the probability constraint formulation.

FIGURE 3.4

EXTENSION OF STRATEGY SPACE AND VALUE OF INFORMATION

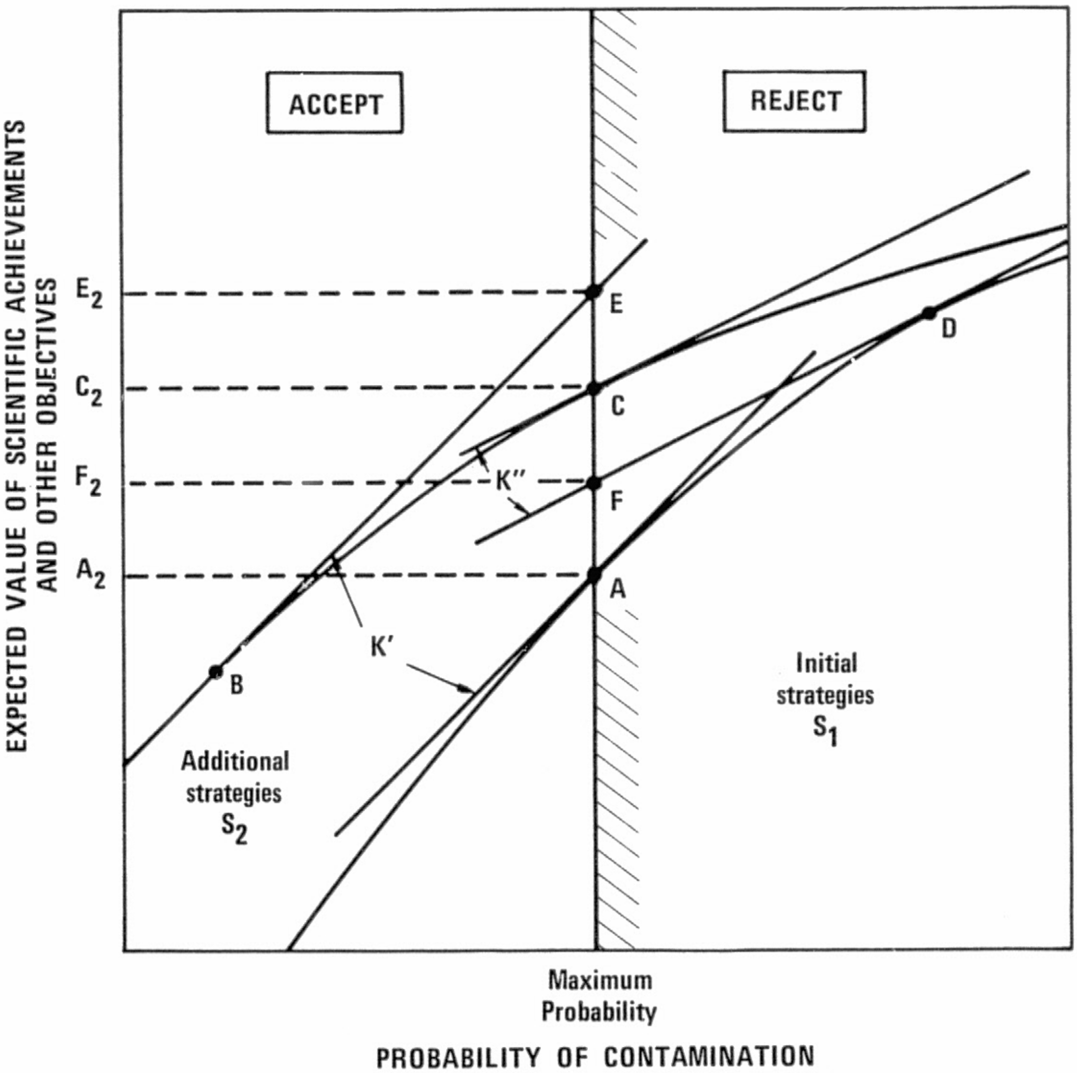

3.3 A Tutorial Application

3.3.1 The Basic Mission

3.3.1.1 Description

The example used in this section is an extremely simplified representation of a two-maneuver fly-by mission. Only those elements that are essential for a demonstration of the new methodology have been retained. In particular, the number of possible strategies has been kept to a minimum so that they may easily be identified and evaluated. A more realistic example is presented in Section 4.

The basic mission depicted in Figure 3.5 shows two corrective maneuvers. Each maneuver offers a choice between a "close" and a "far" fly-by. The first maneuver, d_1, is always performed, but the second maneuver, d_2, may not be performed because of a failure of the guidance system or of the propulsion system. The probability of failure of at least one of these two systems is $f = 1 \times 10^{-2}$ and is assumed to be independent of the first choice of fly-by altitude.

The probabilities of impact after each maneuver are given below:

Maneuver	Decision	Probability of Impact
First maneuver followed by failure of guidance or propulsion system	Close	4×10^{-2}
	Far	4×10^{-4}
Second maneuver	Close	1×10^{-3}
	Far	1×10^{-5}

In all cases, the probability of contamination given impact is 5.5×10^{-2}.

A relative value scale is defined by the two following assignments: 100 is assigned to a nominal flight consisting of two successful "close" fly-by decisions; a value of zero is assigned to the impact outcome. The other values are indicated on the value tree in Figure 3.6. Note that biasing the trajectory away from the target planet (making a first decision to fly by "far") simply penalizes the rest of the mission by 10.

FIGURE 3.5

STRUCTURE OF THE BASIC MISSION

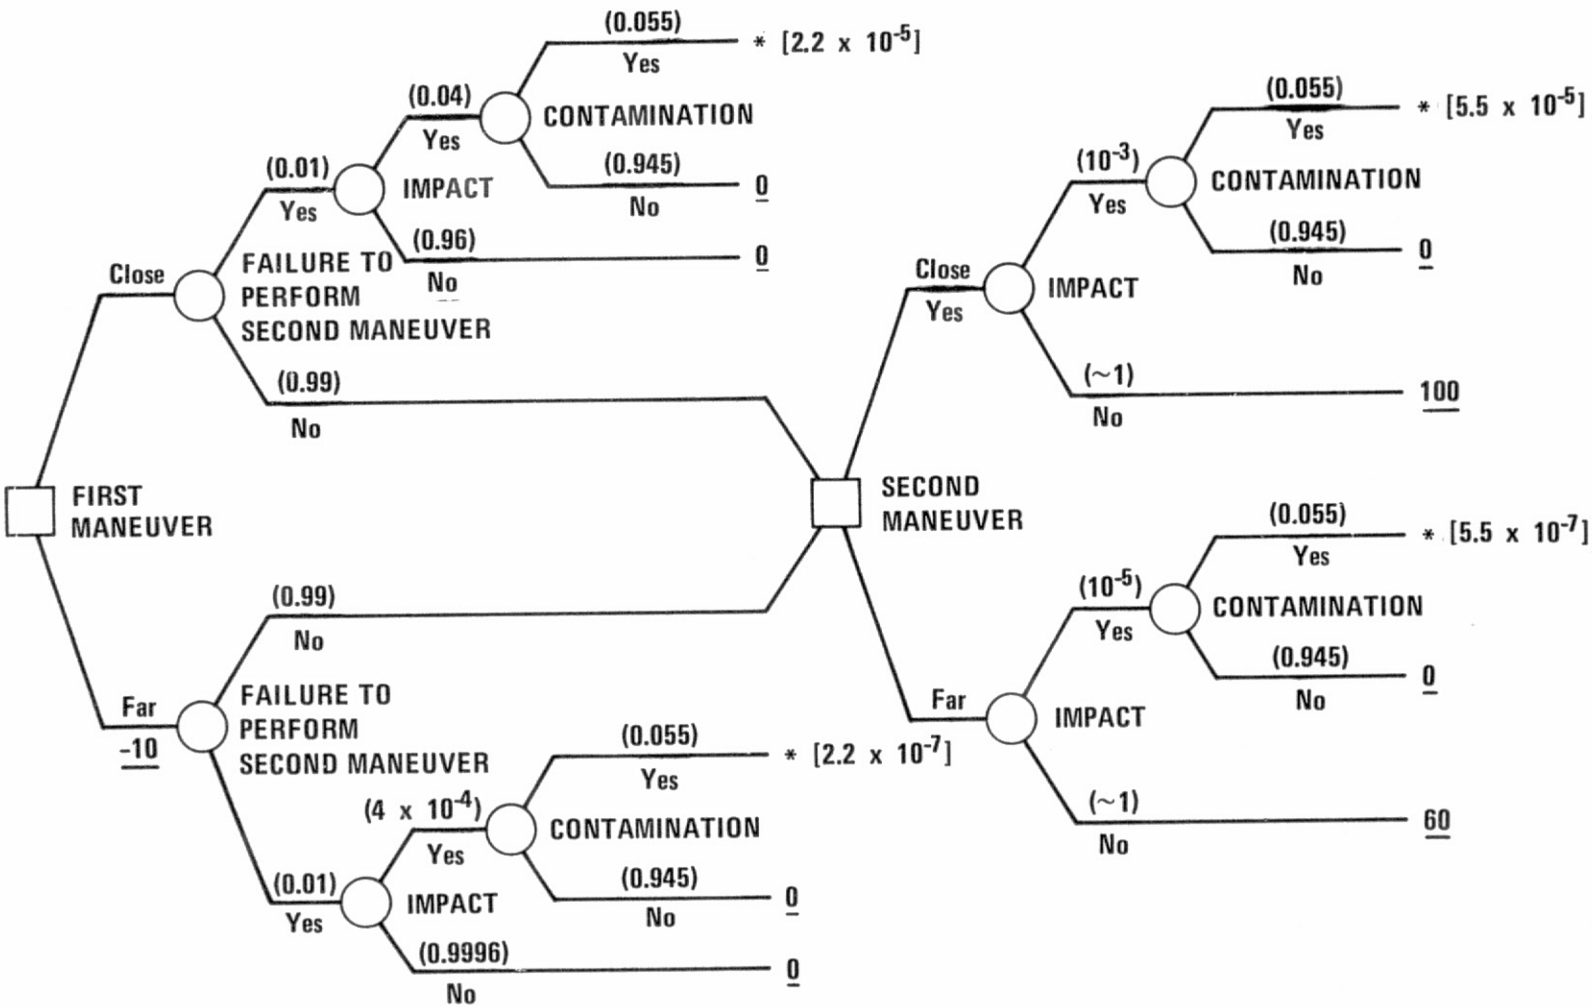

= Decision node; = Chance node; * = Contamination; ( ) = Probability; [ ] = Probability of contamination;

___ = Value

NOTE: Except for the -10 penalty associated with the first maneuver "far", the consequences of the second maneuver are independent of the first. Therefore, the consequences of the first maneuver when a second maneuver is possible have been coalesced, and a single decision tree has been drawn for the second maneuver.

FIGURE 3.6

RELATIVE VALUE TREE FOR THE BASIC MISSION

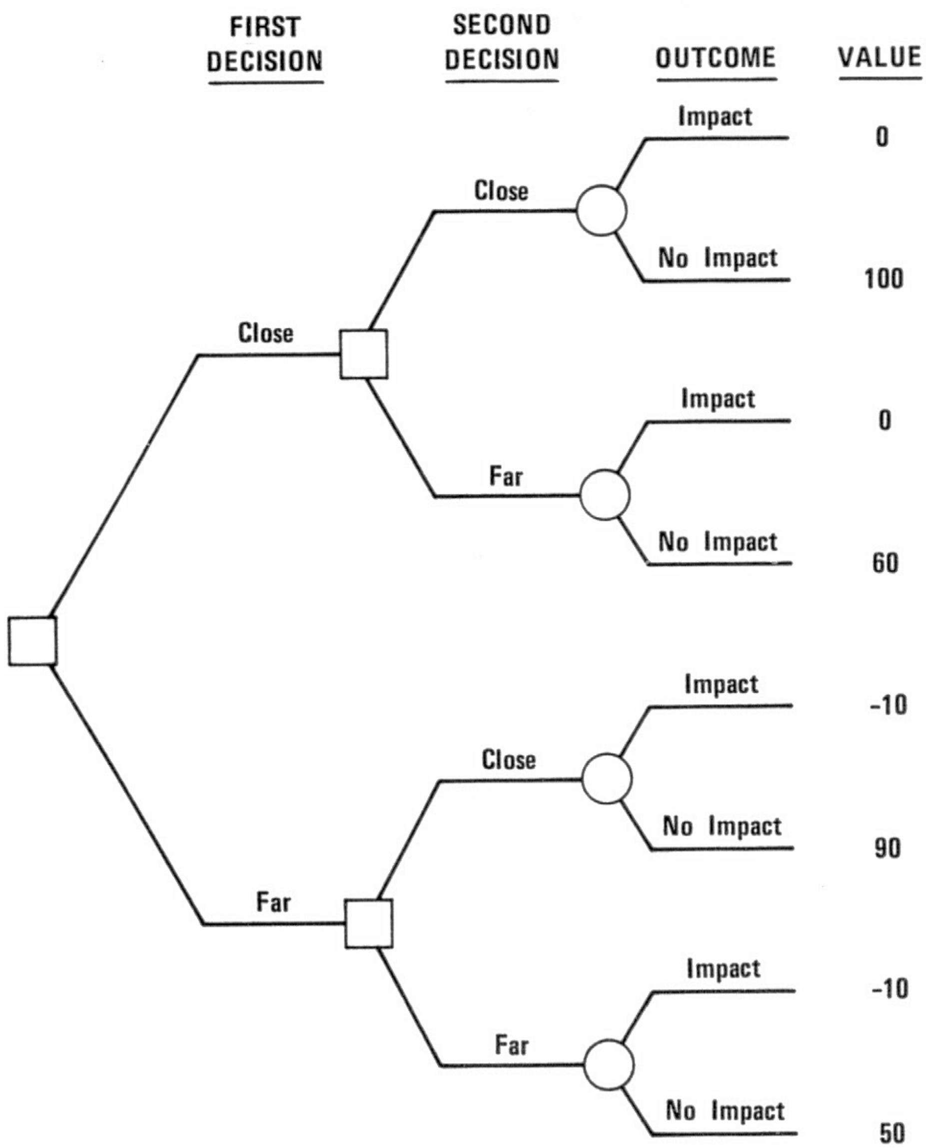

More realistic values could have been assigned, but these are sufficient for the purpose of this illustration and permit simplifying the graphical representation of the mission in Figure 3.5. In that figure, only one tree has been drawn to represent the second maneuver regardless of the decision for the first maneuver; the differential of 10 between the two alternatives for the first maneuver is directly represented by a -10 on the "far" branch of the first decision. If the second maneuver cannot be executed, the expected value of the mission is assumed equal to zero. The results are insensitive to this assumption in view of the small probability of failure.

Finally, we assume the basic mission has been allocated a probability of contamination of 6×10^{-5}.

3.3.1.2 Direct Inspection of Strategies

The four possible strategies for conducting the basic mission are described in Table 3.1 and Figure 3.7. Table 3.1 shows that strategies having the high expected values also have the high probabilities of contamination. For example, Strategy 1, which corresponds to two choices of "close" trajectories, has an expected value of

$$
\begin{aligned}
& 100 \times (\text{Pr. no failure}) \times (\text{Pr. no impact}) \\
&= 100 \times (0.99) \times (0.999) \\
&\simeq 99
\end{aligned}
$$

and a probability of contamination of

$$
\begin{aligned}
& \left[\begin{pmatrix}\text{Pr. impact course} \\ \text{after first maneuver}\end{pmatrix} \times (\text{Pr. failure}) \right. \\
&\left. + (\text{Pr. no failure}) \times \begin{pmatrix}\text{Pr. impact after} \\ \text{second maneuver}\end{pmatrix}\right] \\
&\times (\text{Pr. contamination given impact}) \\
&= [(0.04)(0.01) + (0.99)(0.001)]\ (0.055) \\
&\simeq 7.7 \times 10^{-5} \quad .
\end{aligned}
$$

Strategy 1 is the most valuable of all, but because the probability of contamination exceeds the mission allocation of 6×10^{-5}, it must be rejected. A rapid inspection of Table 3.1 shows that the three other strategies are acceptable, and among them, Strategy 3 has the highest expected value. Strategy 3 is therefore the optimal solution.

Table 3.1

BASIC MISSION STRATEGIES

Strategy Number	Decisions: First Maneuver	Decisions: Second Maneuver	Expected Value	Probability of Contamination
1	close	close	99	7.7×10^{-5}
2	close	far	59	2.2×10^{-5}
3	far	close	89	5.5×10^{-5}
4	far	far	49	7.7×10^{-7}

Note: All the numbers have only two significant digits.

3.3.1.3 Determination of Optimal Strategy Using the Value Iteration Process

Let us assign a penalty K to the contamination event. The first step is to determine a range of values of K such that the contamination penalty will rule out some of the strategies. If value differences among strategies were counted in tens of millions of dollars, and probabilities of contamination are in the range of 10^{-5} to 10^{-4}, then pertinent values of the contamination penalty would therefore lie in the range $10^7/10^{-4} = 10^{11}$ to $10^7/10^{-5} = 10^{12}$, one hundred billion to one trillion dollars.

The next step is to pick a value of K in the range of interest, say 5×10^5, and determine the best strategy. A quick analysis (see Figure 3.5) shows that following a "close" decision for the first maneuver, the expected value for a "close" decision on the second maneuver is

FIGURE 3.7
BASIC MISSION STRATEGIES

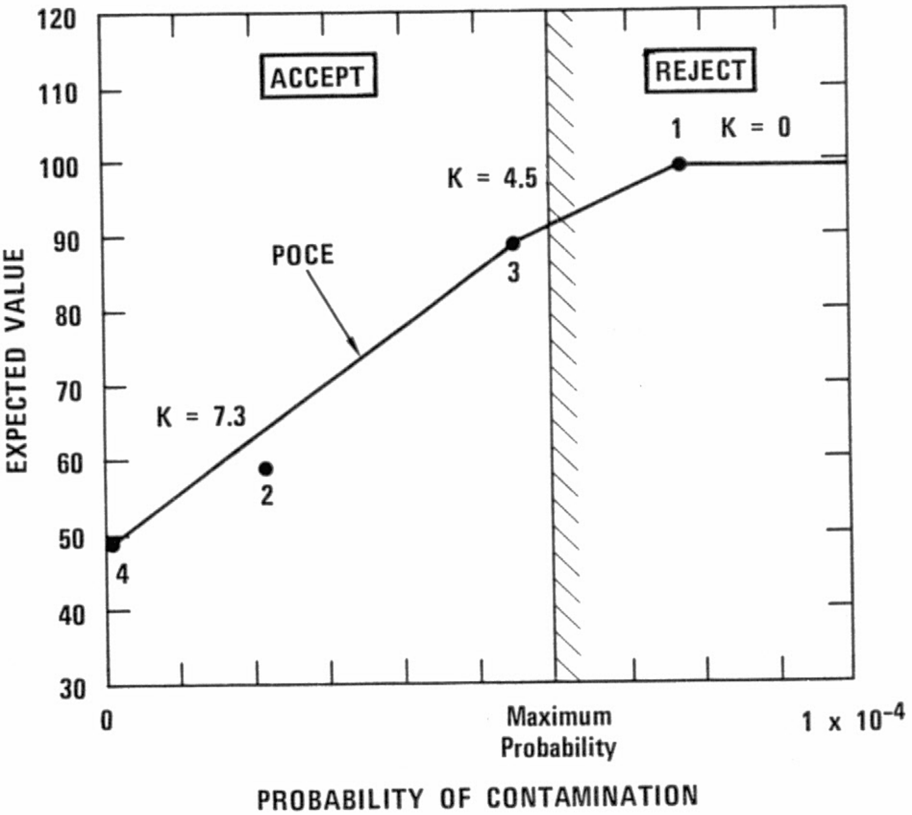

$$(\text{Pr. no impact}) \times (\text{value})$$

$$- (\text{Pr. impact}) \times \begin{pmatrix}\text{Pr. contamination} \\ \text{given impact}\end{pmatrix} \times K$$

$$= (0.999) \times (100) - (0.001) \times (0.055) \times 5 \times 10^5$$

$$\simeq 72.4 \quad .$$

Likewise, the expected value for a "far" decision on the second maneuver is

$$(1 - 10^{-5}) \times (60) - (10^{-5}) \times (0.055) \times 5 \times 10^5$$

$$\simeq 60.01 \quad .$$

Decision d_2 = "close" is therefore preferable when d_1 = "close." The same conclusion holds when d_1 = "far," the only difference being a uniform reduction of 10 in expected value.

The best alternative for the first decision is selected in the same manner, given that the best strategy for the second mission is always to fly by "close." The expected value of the mission with d_1 = "close" is

$$(\text{Pr. no failure}) \times (\text{value})$$

$$- (\text{Pr. failure}) \times (\text{Pr. impact}) \times \begin{pmatrix}\text{Pr. contamination} \\ \text{given impact}\end{pmatrix} \times K$$

$$= (0.99) \times (72.4) - (0.01) \times (0.04) \times (0.055) \times 5 \times 10^5$$

$$\simeq 60.7 \quad .$$

Likewise, the expected value of the mission with d_1 = "far" is

$$(0.99) \times (72.4 - 10) - (0.01) \times (4 \times 10^{-4}) \times (0.055) \times 5 \times 10^5$$

$$\simeq 61.7 \quad .$$

Decision d_1 = "far" is therefore slightly better than d_1 = "close."

The best strategy with a contamination penalty of 5×10^5 is therefore the third strategy listed in Table 3.1. The expected net value is 61.7

If all the values were expressed in millions of dollars and the planetary quarantine policy stipulated that the contamination penalty for the target planet be taken equal to 5×10^5 million dollars, we would have found the optimal strategy and the iteration would be over.

Suppose, however, that the mission has been given a probability allocation of 6×10^{-5}, as we assumed at the beginning of this section. The strategy obtained with $K = 5 \times 10^5$ has a probability of contamination of 5.5×10^{-5}. We could therefore diminish the contamination penalty in the hope of finding a strategy with a larger scientific value and a probability of contamination larger than 5.5×10^{-5}, but still less than 6×10^{-5}.

The choices of strategies and the corresponding expected net values have been represented in Table 3.2 and Figure 3.8 for values of K ranging from 0 to 10^6.

Table 3.2

CHOICES OF STRATEGY FOR THE BASIC MISSION
AS A FUNCTION OF THE CONTAMINATION PENALTY K

Contamination Penalty (10^5)		Optimal
From	To	Strategy
0	4.5	1
4.5	7.3	3
7.3	infinity	4

For K = 0, the strategies have the expected values shown in Table 3.1. When K increases, the expected net values decrease linearly; the slopes correspond to the probabilities of contamination for each strategy. Thus, Strategy 1 starts with the highest expected value for K = 0, but is very sensitive to the contamination penalty. When $K = 4.5 \times 10^5$, Strategy 3 becomes more valuable than Strategy 1; Strategy 3 is finally supplanted by Strategy 4 when K reaches 7.3×10^5. Regardless of the cost of contamination, Strategy 2 is always a poor candidate.

FIGURE 3.8

EXPECTED NET VALUES FOR THE FOUR STRATEGIES OF THE BASIC MISSION

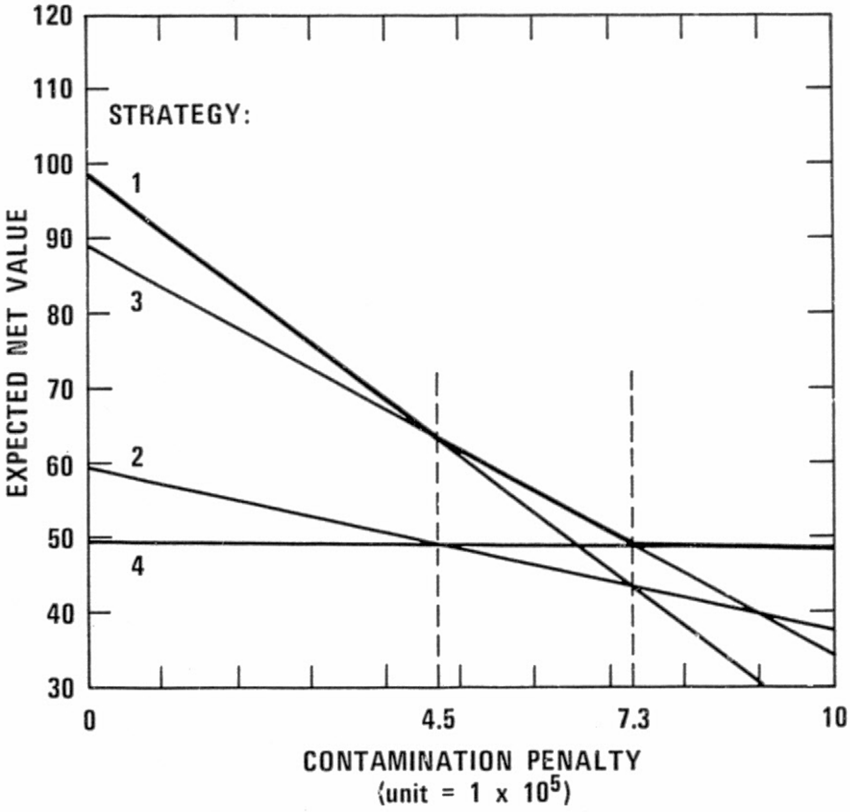

3.3.1.4 Interpretation of the Value Iteration Process in the Strategy Space

The same results could also be read in Figure 3.7. First, let us draw a Pareto-optimal convex envelope (POCE) to the set of available strategies by joining each pair of strategies by line segments, and keeping the upper left edge of the polygon that has been created. In Figure 3.7 we find that the POCE consists of line segments joining Strategy 1 to 3 and 3 to 4. Strategy 2 lies below the line segment joining Strategy 3 to 4.

The various slopes of the POCE can be interpreted as trade-offs between expected value and probability of contamination. Thus the slope of the line segment joining Strategy 1 to 3 is equal to

$$\frac{(99 - 89)}{(7.7 \times 10^{-5} - 5.5 \times 10^{-5}} \simeq 4.5 \times 10^{5} \quad .$$

Strategy 1 should therefore be preferred to Strategy 3 when the contamination penalty is less than 4.5×10^5, and vice versa. The same argument applies to a comparison of Strategies 3 and 4 with, this time, a slope or trade-off value of 7.3×10^5. Strategy 2 lying below the line segment joining Strategies 3 and 4 is not selected with any trade-off value.

3.3.1.5 The Issue of Convexity

If the mission allocation had been less than 5.5×10^{-5} but more than 2.2×10^{-5}, Strategy 2 would have been selected on an expected value basis exclusive of contamination penalty. The iterative reconciliation process would indicate, however, that the contamination penalty derived from the mission allocation is at least equal to 7.3×10^5. Under these circumstances, Strategy 4 is always prefereable to Strategy 2 on an expected net value basis including the contamination penalty. Mathematically,

$$49 - K \times 7.7 \times 10^{-7} > 59 - K \times 2.2 \times 10^{-5}$$

as long as

$$K > (59 - 49)/(2.2 \times 10^{-5} - 7.7 \times 10^{-7}) \simeq 4.5 \times 10^{5} \quad .$$

In fact, given that $K = 7.3 \times 10^5$, the smallest contamination penalty compatible with the mission allocation, the expected net value of Strategy 2 is

$$59 - 7.3 \times 10^5 \times 2.2 \times 10^{-5} \simeq 43 \quad ,$$

whereas the expected net value of Strategy 4 is

$$49 - 7.3 \times 10^5 \times 7.7 \times 10^{-7} \simeq 49 \quad .$$

In other words, given that the contamination penalty is at least equal to 7.3×10^5, it is better to select Strategy 4 instead of 2; a reduction of the probability of contamination from 2.2×10^{-5} to 7.7×10^{-7} is worth at least 16, whereas the corresponding decrease in expected scientific value is only 10. Note also that, if the contamination penalty is exactly equal to 7.3×10^5, Strategy 3 has the same expected net value as Strategy 4 ($89 - 7.3 \times 10^5 \times 5.5 \times 10^{-5} \simeq 49$).

We therefore believe that strategies that lie below the POCE of the strategy set, such as Strategy 2 in Figure 3.7, should never be selected. There are always better strategies on the POCE, and at least one of them has an acceptable probability of contamination.

3.3.2 The New Mission (Basic Mission with Temperature Sensor)

3.3.2.1 Description

The purpose of the refinement of a temperature sensor is to demonstrate how information gained during the flight can be used to improve the guidance decisions. We assume that surface temperature readings of the target planet are available a few days before the second maneuver. The probability of contamination given impact can therefore be reassessed, and the second maneuver can be made on the basis of better information. We will show that, consequently, the expected value of the mission can only be increased.

For the sake of simplicity, we shall assume that there are only two possible temperature readings: high and low. Knowledge of the accuracy of the temperature sensor and prior knowledge of the target planet's surface temperature indicate (using Bayesian inference) the following:

Temperature Reading	Probability	Conditional Probability of Contamination Given Impact and Temperature Reading
High	0.5	0.1
Low	0.5	0.01

Figure 3.9 shows the structure of the basic mission expanded to include the new temperature information. We now distinguish two decisions for the second maneuver, each decision being conditional on a temperature reading. The new probabilities of contamination given impact replace the old average of 5.5×10^{-2}. All the other elements of the new mission are the same as in the basic mission.

3.3.2.2 Direct Inspection of Strategies

There are eight strategies for conducting the new mission. Four of them are identical to the four strategies of the basic mission (where the second maneuver is independent of the temperature reading) and are denoted by the same numbers. In the four new strategies, the second maneuver is conditional upon the temperature reading. These strategies have been numbered from 1′ to 4′, according to their similarities with each of the four unconditional strategies.

Table 3.3 and Figure 3.10 represent the expected values and the probabilities of contamination of each of these eight strategies. A quick inspection reveals that Strategies 2′ and 4′ have much larger expected values and only slightly larger probabilities of contamination than their nonprimed counterparts. In these strategies, advantage is taken of a low temperature reading to fly by "close" instead of always "far." On the other hand, Strategies 1′ and 3′ seem to be worse than their nonprimed counterparts. Indeed, it is counter-intuitive to fly by "far" when the temperature reading is low, instead of always flying by "close."

Strategy 3 still has the highest expected value among the feasible strategies (with a probability of contamination less than or equal to 6×10^{-5}). Therefore, according to the chance-constrained formulation, the same guidance strategy should be selected with or without the temperature information. The scientific value of the surface temperature data might be very high, but the value of the data for improving guidance decisions is null.

FIGURE 3.9

STRUCTURE OF THE NEW MISSION (BASIC MISSION WITH TEMPERATURE SENSOR)

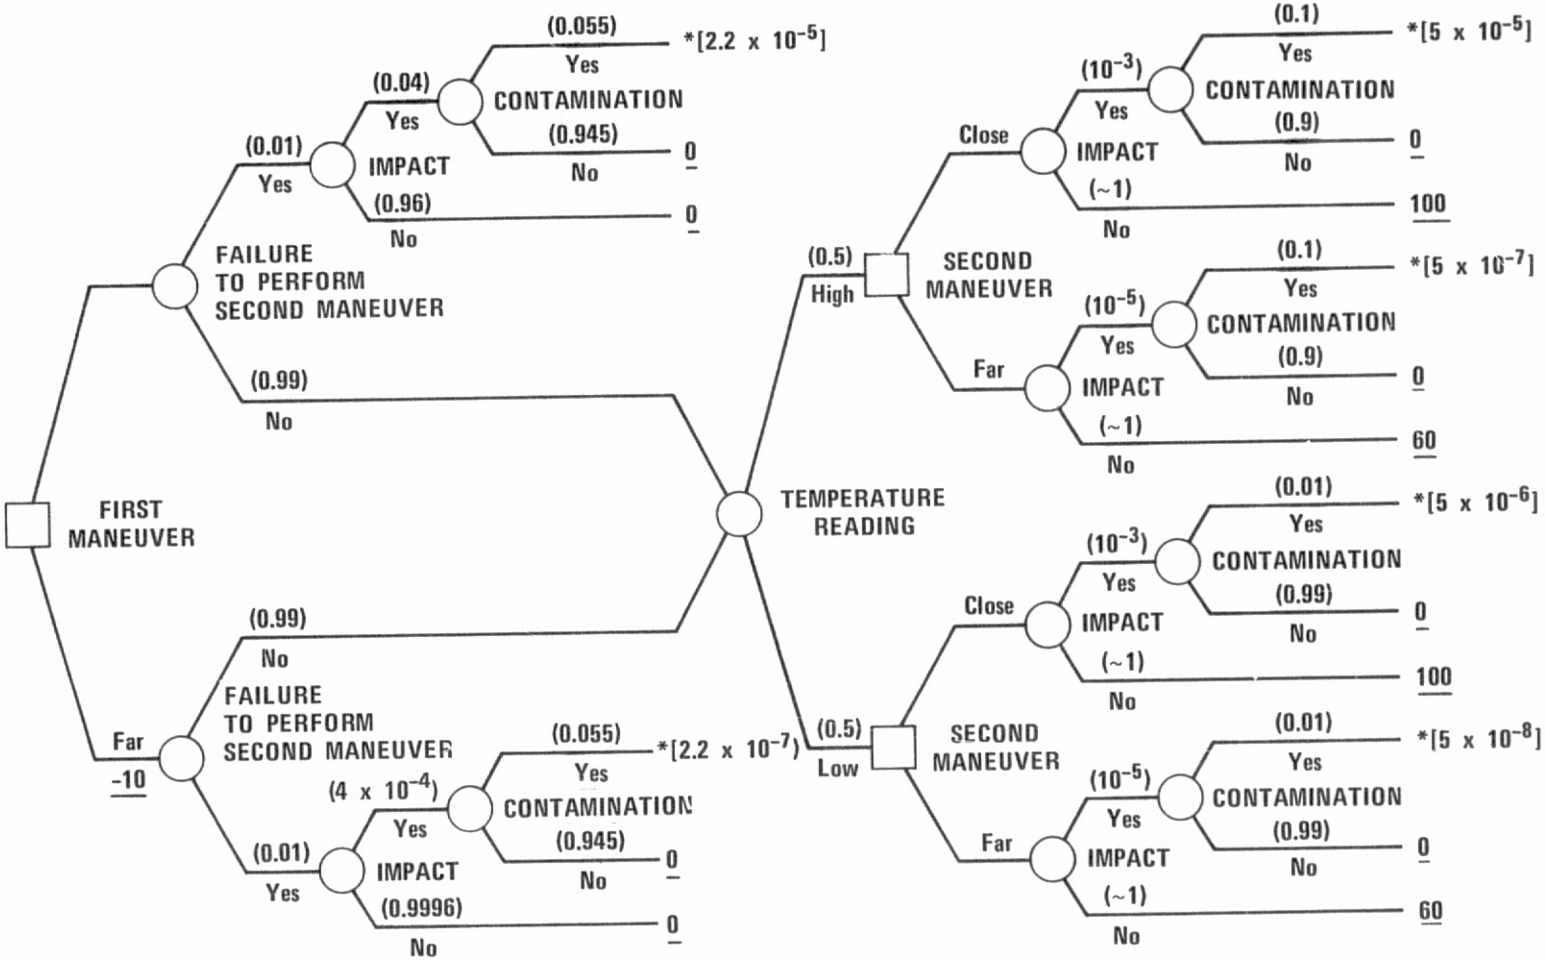

□ = Decision node; ○ = Chance node; * = Contamination; ( ) = Probability; [ ] = Probability of contamination; ___ = Expected value

NOTE: The consequences of the second maneuver are differentiated only by the results of the temperature reading and the -10 penalty associated with the first maneuver "far". Therefore, the consequences of the first maneuver when a second maneuver is possible have been coalesced and a single tree has been drawn for the temperature reading and the second maneuver.

Table 3.3

NEW MISSION STRATEGIES

Strategy Number	Decisions: First Maneuver	Second Maneuver: High Temp.	Second Maneuver: Low Temp.	Expected Value	Probability Contamination (10^{-5})
1	close	close	close	99	7.7
1′	close	close	far	79	7.2
2′	close	far	close	79	2.7
2	close	far	far	59	2.2
3	far	close	close	89	5.5
3′	far	close	far	69	5.0
4′	far	far	close	69	0.57
4	far	far	far	49	0.077

Note: All numbers have only two significant digits.

3.3.2.3 Determination of Optimal Strategy Using the Value Iteration Process

We shall not repeat the determination of optimal strategies using contamination penalties (see Section 3.3.1.3), but simply present the results.

Figure 3.10 shows only 4 strategies on the Pareto-optimal convex envelope. They are Strategies 1, 2′, 4′, and 4. The slopes of the convex envelope between these strategies are 4×10^5, 4.5×10^5, and 40×10^5, respectively. These results are confirmed in Table 3.4 and Figure 3.11, where the expected net value of each strategy is plotted versus the contamination penalty.

Note that Strategy 3, optimal according to the chance-constrained formulation, lies slightly below the POCE and therefore would not be selected with the value iteration approach. Furthermore, if Strategy 3 is the right choice for the basic mission, then the contamination penalty should be in the range of 4.7×10^5 to 7.3×10^5. The only optimal strategy in that contamination penalty range for the new mission is Strategy 4′.

FIGURE 3.10
NEW MISSION STRATEGIES

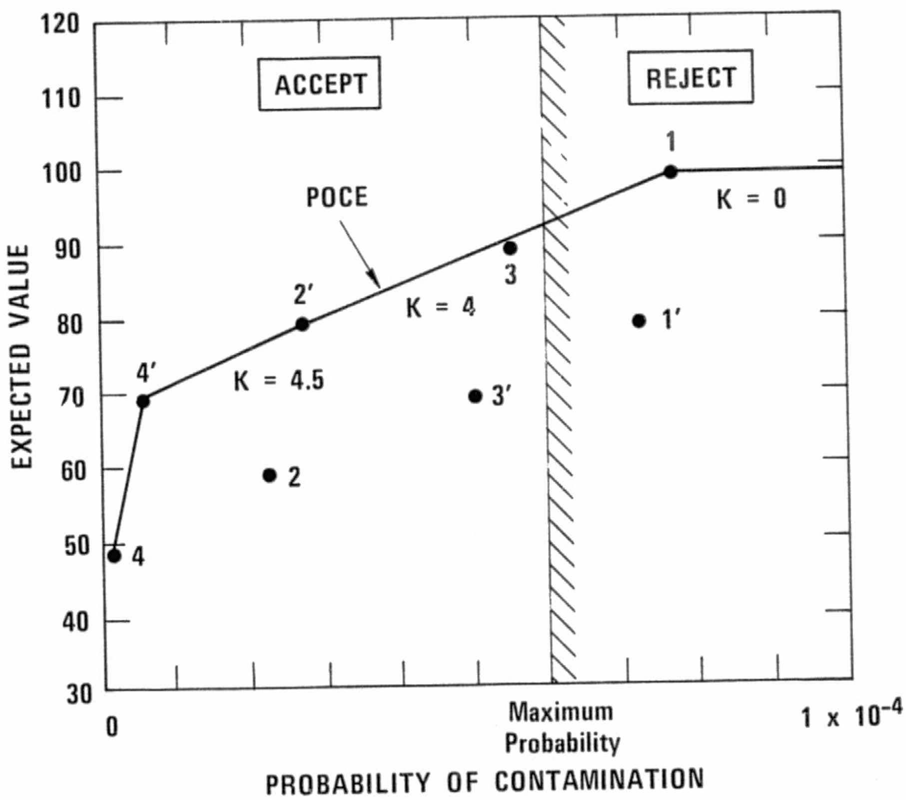

Table 3.4

CHOICES OF STRATEGY FOR THE NEW MISSION
AS A FUNCTION OF THE CONTAMINATION PENALTY K

Contamination Penalty (10^5) From	To	Optimal Strategy
0	4	1
4	4.5	2′
4.5	40	4′
40	infinity	4

Note, too, that in practice, the assessment of the expected value and the probability of contamination associated with a particular strategy are always subject to uncertainty. Therefore, a strategy lying near the POCE should not be rejected hastily; the value of reexamining this strategy and the neighboring strategies can be analyzed.

3.3.2.4 The Expected Value of Information

The chance-constrained formulation always suggests selecting Strategy 3 whether or not temperature information is available. Therefore, the chance-constrained formulation does not show any guidance decision-making value for the temperature measuring device.

Introducing a contamination penalty shows that if Strategy 3 is selected for the basic mission, then Strategy 4′ should be selected for the new mission. However, Strategy 4′ has an expected value of 69 only, whereas Strategy 3 has an expected value of 89. Is the expected value of the temperature information equal to -20? No; we must also take into account the probabilities of contamination associated with each strategy.

A look at Figure 3.11 will reveal that, in the range $K = 4.5 \times 10^5$ to 7.3×10^5, Strategy 4′ always has an expected net value including contamination penalty larger than Strategy 3. For example, if we believe that the contamination penalty should be $K = 5 \times 10^5$, the difference between the expected net values of Strategies 4′ and 3 is

FIGURE 3.11
EXPECTED NET VALUES FOR THE EIGHT STRATEGIES
OF THE NEW MISSION

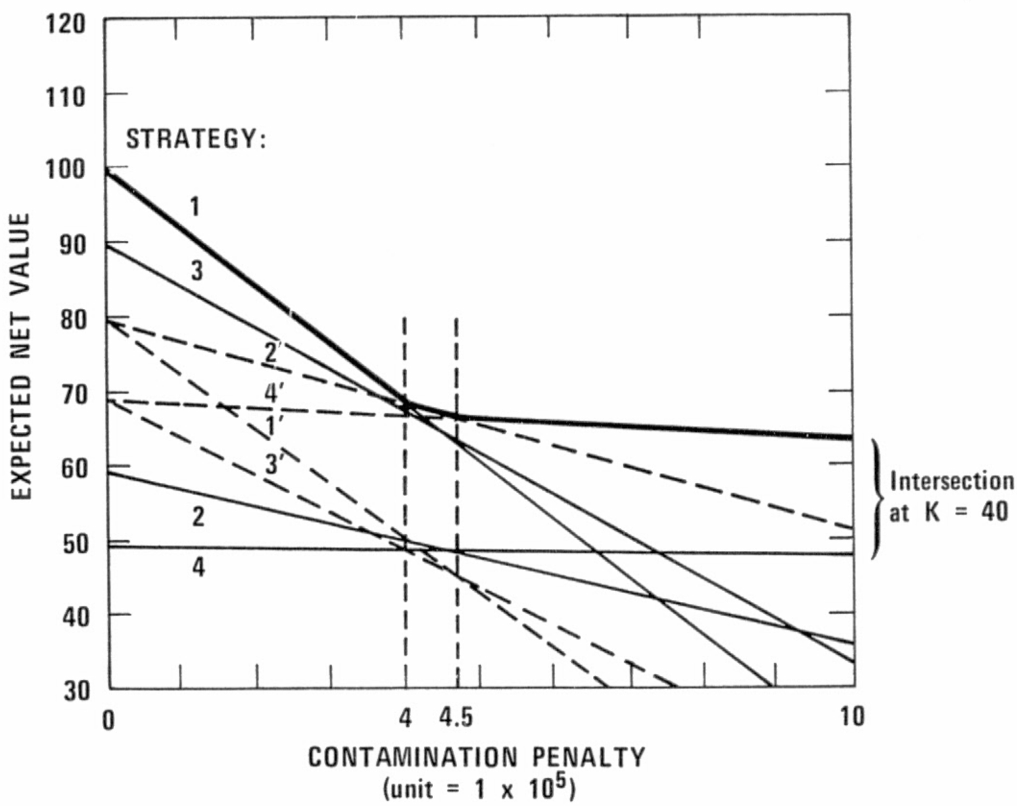

$$66.4 - 61.7 = 4.7 \quad .$$

In other words, the temperature information is expected to be worth 4.5 simply because it permits improving the guidance strategy. After a high temperature reading, we will not take the risk of flying by "close," which was inherent to Strategy 3. The savings are: (1) in scientific value, 60 - 100 = -40, and (2) in contamination penalty,

$$(1 \times 10^{-4} - 1 \times 10^{-6}) \times 5 \times 10^{5} = 49.5 \quad ,$$

and therefore (3) in total, 49.5 - 40 = 9.5. The high temperature reading having a probability of 0.5, the expected savings are 9.5 x 0.5 = 4.7, as we read in Figure 3.11.

Finally, we should note that determining the incremental value of the new mission (with the temperature sensor) over the basic mission could be pertinent for mission design decisions. If a choice had to be made between designs with and without the temperature sensor, carrying out the analysis above with dollar values would indicate the value of using the temperature sensor that could be expected solely from an improvement of the guidance decision. The sum of this expected value and the scientific value of surface temperature readings could then be compared to the cost of adding a temperature sensor to the basic mission.

4 A JUPITER ORBITER MISSION

4.1 Scope and Objective of the Application

The Jupiter Orbiter mission is a particularly appropriate application for the new methodology we proposed in Section 3. A Mariner or Pioneer spacecraft could be sent to orbit Jupiter and explore the Jovian system for a period of years.

It has been demonstrated (see [12]) that, using a succession of small propulsive changes of velocity and the gravitational assistance of the Galilean satellites, a great variety of trajectories can be obtained. Thus, a single spacecraft can therefore collect a wealth of information by performing particle and field experiments from a few to a hundred Jupiter radii, and imaging experiments during repeated encounters with the major Jovian satellites.

The planning of such a mission is, of course, complex. The number of possible guidance strategies seems to be practically limitless, and the variety of sometimes conflicting scientific objectives requires the assessment of many trade-offs.

From a planetary quarantine point of view, the problems seem unprecedented. The spacecraft will repeatedly fly by several celestial bodies, bodies that, although not well-known, are considered of potential biological interest. Moreover, there is no easy means of disposing of the spacecraft at the end of the period of controllable flight. A solution must be found that will guarantee that the planetary quarantine requirements will be met for the next few decades.

We have divided this application into two parts: first a tutorial analysis of the maneuver to insert the spacecraft into Jupiter orbit using the gravitational assistance of Ganymede; second, a survey of the complete mission to identify the major obstacles that must be cleared to satisfy the planetary quarantine requirements.

In the first part, the insertion maneuver is treated as a complete mission consisting of a single fly-by of Ganymede. The analysis does not pretend to solve this maneuver; rather, it demonstrates how realistic elements can be taken into account in such an analysis.

The second part is necessarily sketchy, given the early stage of planning of a JO mission. However, we have been able to point out major difficulties that will require planetary quarantine research as well as the attention of the mission planner.

4.2 The Insertion Maneuver into Jupiter Orbit

4.2.1 Problem Structure

4.2.1.1 The Fly-By Altitude Decision

The problem structure is depicted by the decision tree in Figure 4.1. We assume the spacecraft will use the gravitational assistance of Ganymede to enter into Jupiter orbit. A close fly-by of Ganymede would permit large fuel savings. However, Ganymede is a satellite of potential biological interest and the planetary quarantine requirements do conflict with fuel savings. We will limit our attention to this particular decision: the choice of a fly-by altitude H during the last maneuver before encounter with Ganymede (approximately four days before encounter). We assume that the other concomitant decisions (which are not crucial from a planetary quarantine point of view) are optimized independently.

4.2.1.2 Information for the Fly-By Altitude Decision

We assume there will be available two critical elements of information on which to base the fly-by altitude decision: one is the availability of the guidance system and the other, a surface temperature reading of Ganymede.

4.2.1.2.1 Availability of the Guidance System

The JO could be equipped with a high precision optical guidance system (using the spacecraft TV camera) in addition to a regular radio guidance system. For a given probability of impact, the optical and radio guidance systems permit a lower fly-by altitude than the radio guidance system alone.

We define three possible states of the guidance system with probabilities as indicated below:

FIGURE 4.1

DECISION TREE FOR INSERTION MANEUVER

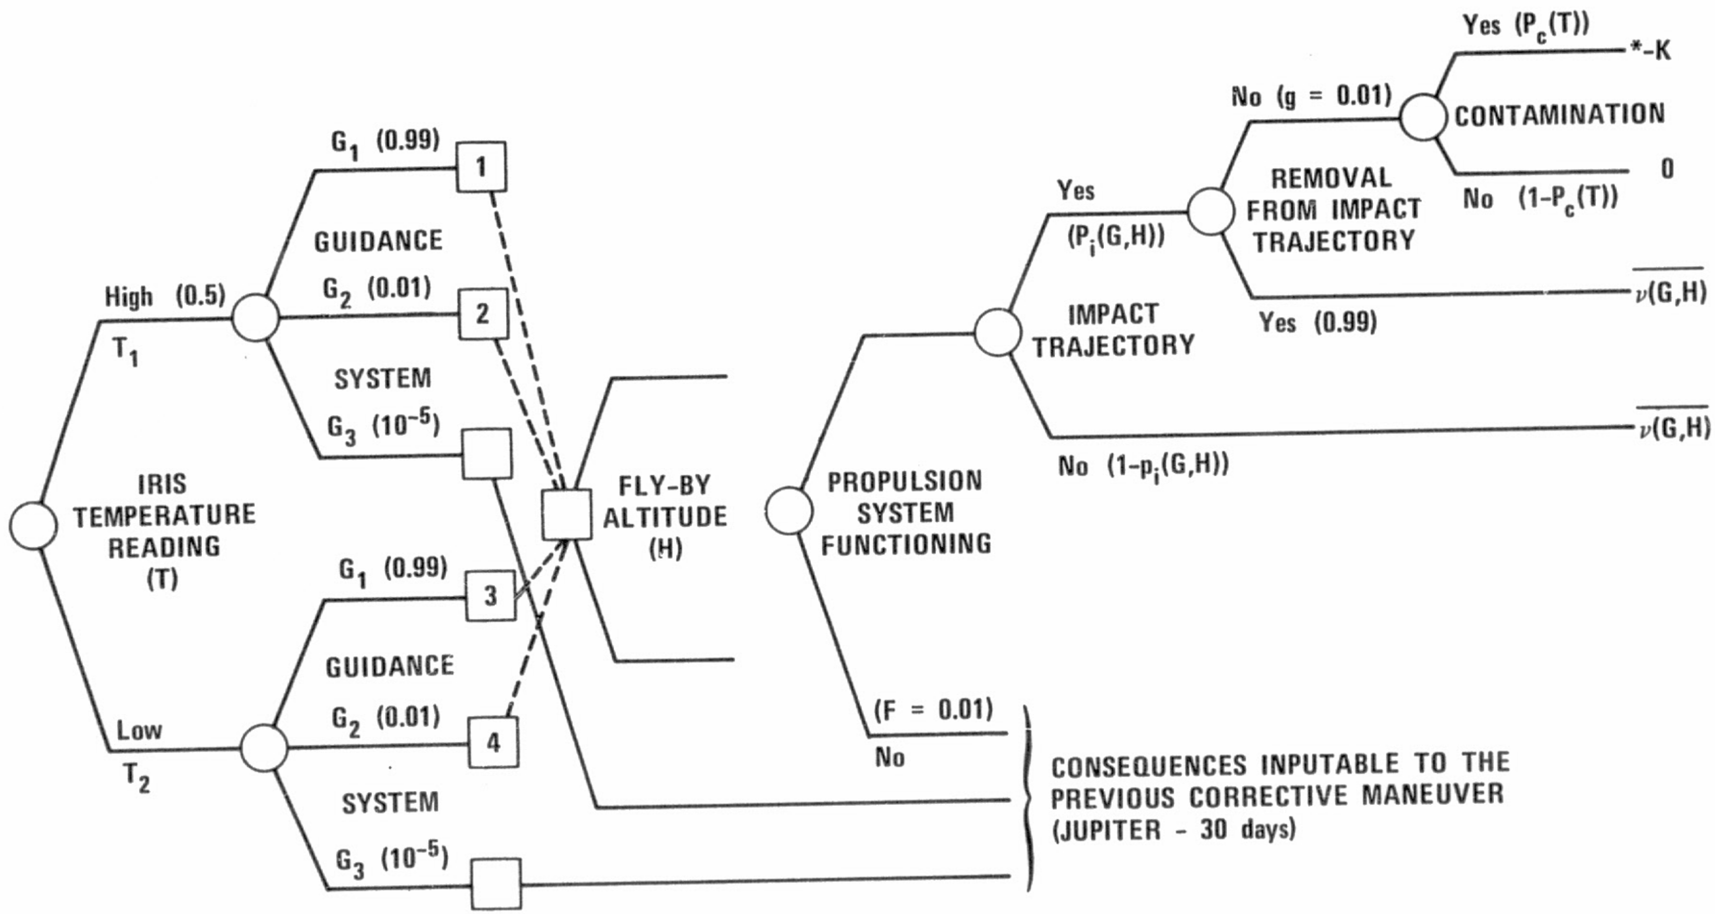

Decision Node; Chance Node; * Contamination; (Probability)

Symbol	Definition	Probability
G_1	Optical and radio systems function	0.99
G_2	Only radio system functions	0.01
G_3	No guidance system available	1×10^{-5}

The probabilities of impact given states G_1 and G_2 of the guidance system and fly-by altitude H are plotted in Figure 4.2 (from [13]).

4.2.1.2.2 Surface Temperature Reading

A temperature sensor may be part of the JO's scientific payload. This instrument will, among other things, provide precise surface temperature readings of planets, starting several days before encounter. The average surface temperature of Ganymede is believed to be on the order of -120°C, but it is not known with precision. A uniform temperature of -120°C would be a major obstacle to the development of life. However, some experts have tentatively assessed a probability of growth of terrestrial organisms on Ganymede as high as 0.1 (see [1]). We interpret this assessment as meaning that there is a small probability of finding localized surface temperatures as high as 0°C. Our major conclusions are not very sensitive to this assumption.

Figure 4.3 shows four curves pertinent to determining of the probability of growth p_g of terrestrial organisms on Ganymede. The two dotted lines are cumulative probability distributions. One represents an assessment of the prior probability that the maximum surface temperature T will not exceed the value indicated on the abscissa. For example, there is a probability of 0.6 that T will be less than -45°C. The second dotted line is a two-step discrete approximation of the first curve. Thus, assuming the temperature sensor to be well-calibrated and precise, we may simplify the description of the temperature information received by assuming two equally likely temperature readings T_1 = -25°C and T_2 = -75°C.

The two solid curves represent probabilities of growth. The first one is the probability of growth of a terrestrial organism on Ganymede at temperature T (approximated from [14]). The second curve indicates the probability of growth given that a maximum surface temperature of T has been observed. We assume that if T is larger than -100°C, then any temperature between -100°C and T can be found somewhere on Ganymede. Therefore, the second curve coincides with T as T increases from -100°C to 0°C and $p_g(T)$ increases simultaneously from 1×10^{-5} to 1. Then, for

FIGURE 4.2

PROBABILITY OF IMPACT ON GANYMEDE p_i (G,H)

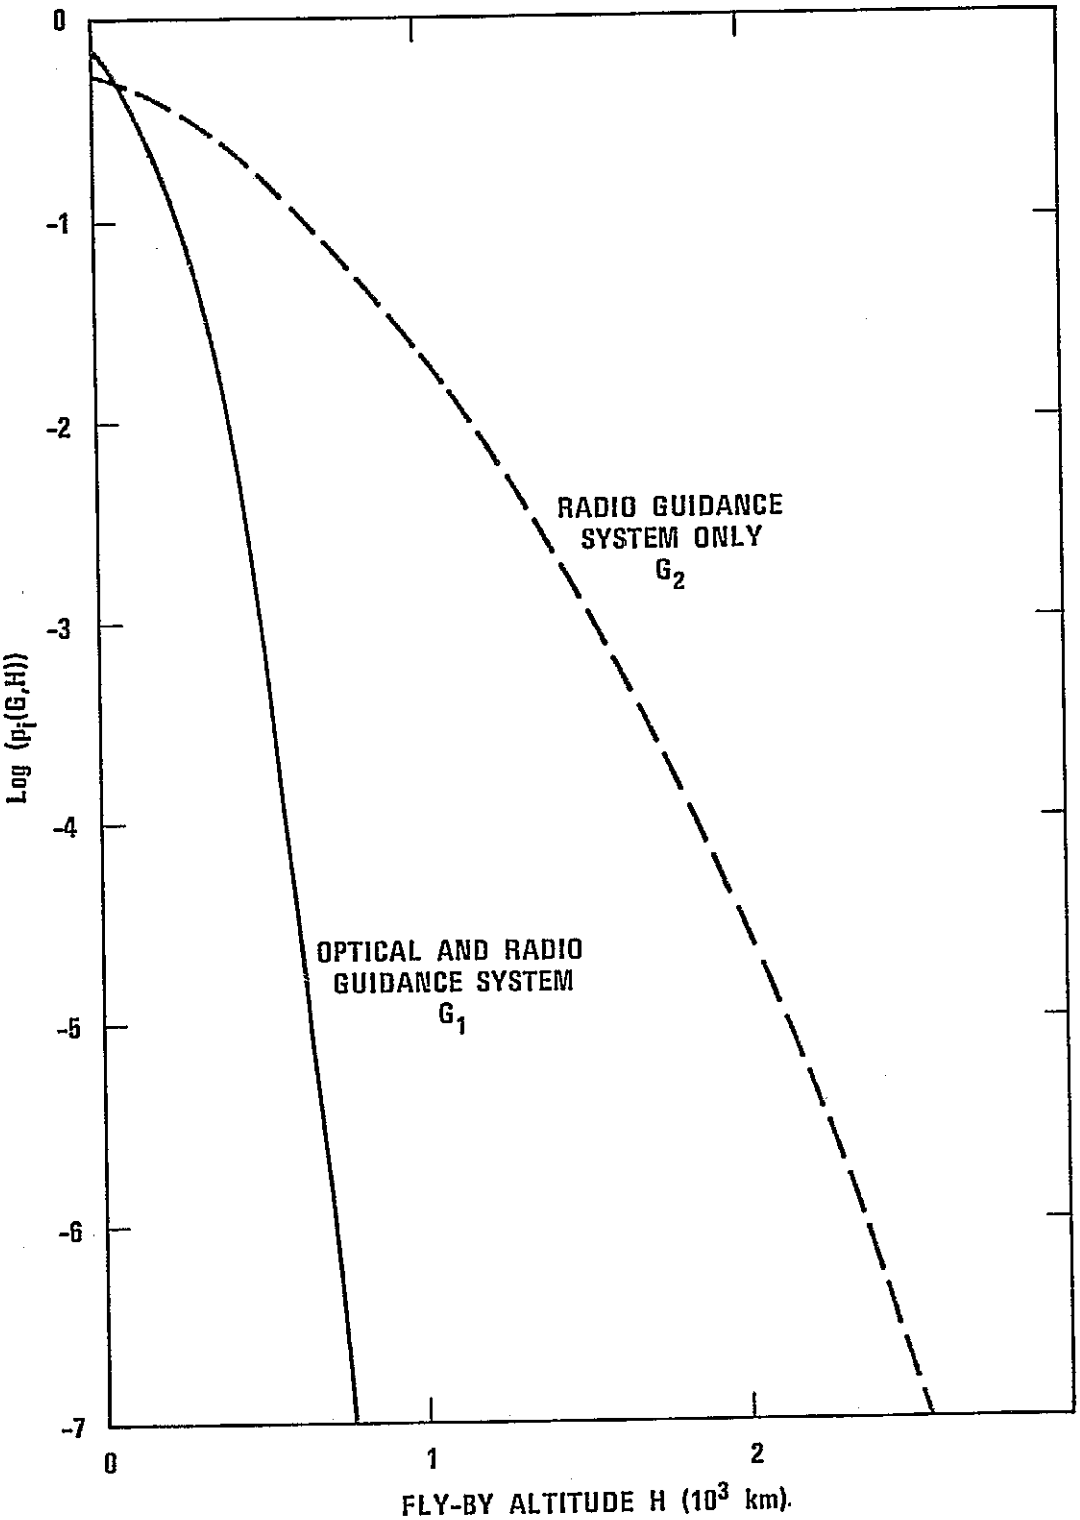

SOURCE: [13]

FIGURE 4.3

REVISION OF p_g AS A FUNCTION OF TEMPERATURE READING

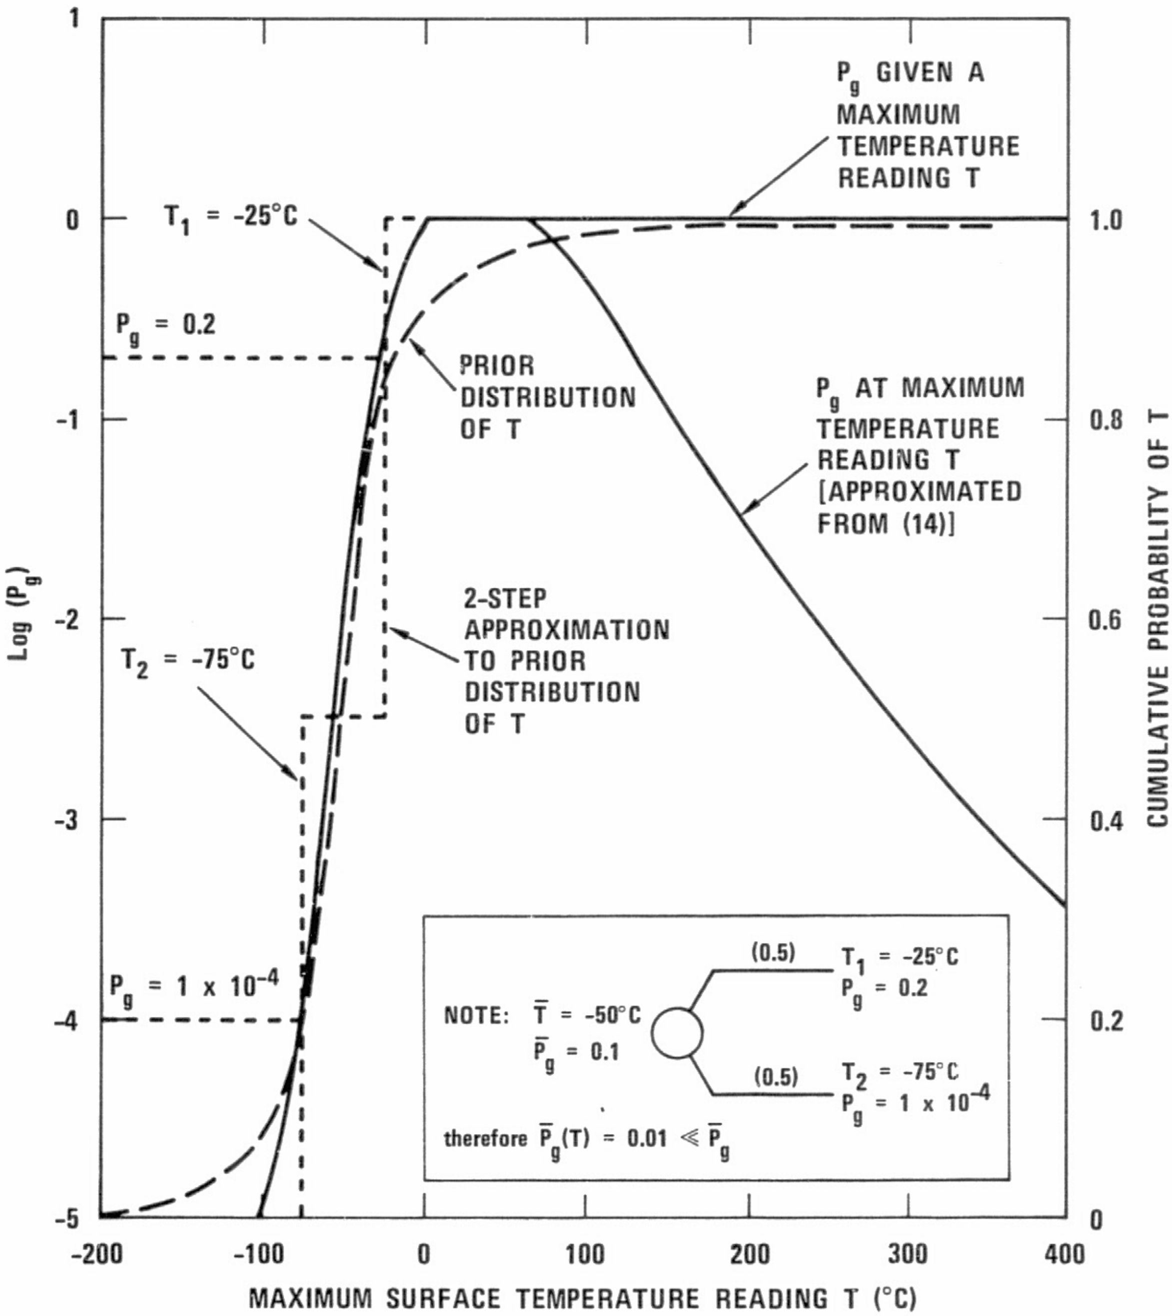

T above 65°C, $P_g(T)$ begins to decrease, but the probability of growth somewhere on Ganymede remains equal to 1. For the temperature readings T_1 and T_2, the probabilities of growth are 0.2 and 1×10^{-4}, respectively.

The probability of growth is one of the factors required to compute the probability of contamination given impact of the spacecraft. A careful contamination analysis also includes a description of the spacecraft bioload by type of organisms and location on the spacecraft and a model for the release and propagation of organisms following an impact. We shall assume that a contamination analysis indicates probabilities of contamination given impact of 1, and 0.01 given probabilities of growth of 0.2 and 1×10^{-4}, respectively.

The information from the temperature sensor and its consequences on contamination probabilities are summarized below:

Temperature Reading Probability	Temperature (T)	Probability of Growth (p_g)	Probability of Contamination Given Impact (p_c)
0.5	-25°C	0.2	1.0
0.5	-75°C	1×10^{-4}	0.01

The combination of the two likely states of the guidance system (G_1 and G_2) and the two possible temperature readings form four states of nature that will condition the fly-by altitude decision.

4.2.1.3 Outcomes

The consequences of a fly-by altitude decision H can be decomposed as shown in Figure 4.1. First, the insertion maneuver may not be executed because of a failure of the main propulsion system. The probability of failure is f = 0.01 and is assumed independent of the state of the guidance system. The consequences of a failure depend on the last maneuver executed and are of no direct concern for the decision at hand. (We shall consider the consequences of a failure in some detail in Section 4.3.) If the propulsion system does not fail, the consequences can be divided between impact trajectories [probability $p_i(G,H)$] and nonimpact trajectories.

Orbit determination procedures will permit detecting an impact trajectory in time to execute an emergency maneuver. The emergency maneuver will fail with probability $q = 0.01$, given that the main propulsion system functioned shortly before for the insertion maneuver. Impact is avoided when the emergency maneuver succeeds.

4.2.2 Values

Most of the values associated with the outcomes of the insertion maneuver are indirect values except, of course, for the impact outcome.

We have assumed that except for the choice of the fly-by altitude H, all other concomitant decisions were solved independently, e.g., we have assumed that an optimal balance had been reached between orbit parameters and fuel consumption by analyzing the subsequent phases of the mission.

We shall therefore consider the expected value of the mission as a function of the fly-by altitude decision as given. For the sake of simplicity, we have adopted the expected value function shown in Figure 4.4 and Table 4.1. A value of 100 has been assigned to a nominal fly-by ($T = T_1$, $G = G_1$, $H = 1000$ km, and so forth), and a value of zero has been assigned to an impact (exclusive of contamination penalty). All the other values are relative to this scale. For example, a fly-by at an altitude of 2000 km reduces the amount of fuel available for the rest of the mission by 25 percent. We have assumed that, as a result, the value of the mission would be reduced by 20 percent.

Failure of the optical guidance system has two effects: (1) reduction of the scientific value to be expected from a given flight and (2) increase in the fuel requirements to control a given flight, hence the parameters assigned to the value function given failure of the optical system.

Finally, if the spacecraft is successfully removed from an impact trajectory by an emergency maneuver, we assume that the expected value of the mission given a fly-by aiming point decision is the same as if the spacecraft had never been on an impact trajectory. The analytical results are not sensitive to this assumption except in the extreme case of a decision to fly by very low, leading to a high probability of impact.

FIGURE 4.4

INSERTION MANEUVER VALUE AS A FUNCTION OF FLY-BY ALTITUDE

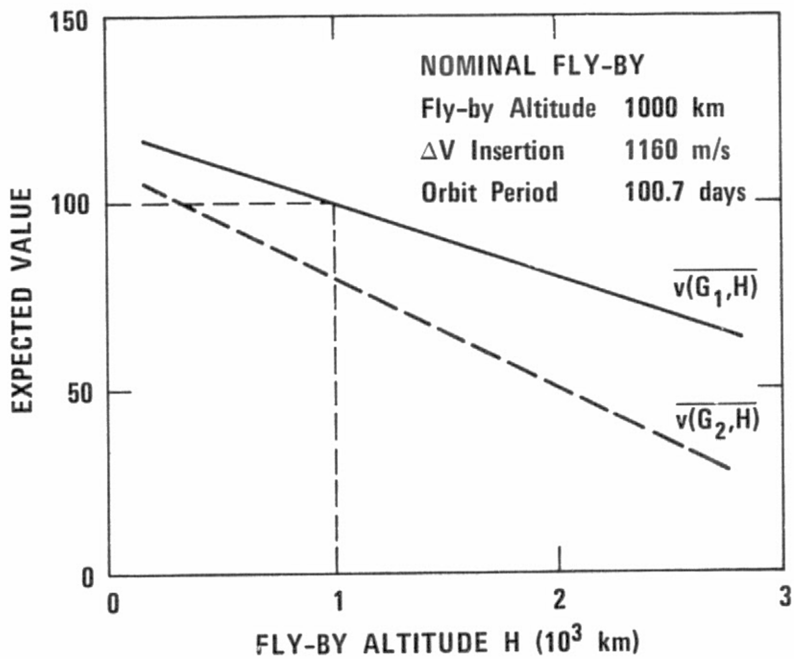

Table 4.1

VALUE FUNCTION PARAMETERS FOR THE INSERTION MANEUVER

Guidance System	Value at $H = 10^3$ km	Slope (per 10^3 km)
Optical (G_1)	100	-20
Radio only (G_2)	80	-30

4.2.3 Analysis

4.2.3.1 Formulation

In analyzing the fly-by altitude decision, we are concerned only by the sequences of an execution of this decision. These consequences are found on the upper right quadrant of Figure 4.1, on the branches following a reliable performance of the propulsion system. The expected value of the mission given the execution of the fly-by maneuver is

$$(1 - p_i(G,H)q)\overline{v(G,H)} - Kp_i(G,H)qp_c(T) \quad , \tag{4-1}$$

where K is the contamination penalty. The prior probability of contamination is

$$\sum_{T,G} r(T,G)p_i(G,H)qp_c(T) \quad , \tag{4-2}$$

where r(T,G) is the probability of temperature reading T and state of the guidance system G.

The optimization of the fly-by altitude decision subject to a probability constraint Q can therefore be written as

$$\begin{aligned} &\underset{\text{over } H}{\text{maximize}} \quad (1 - p_i q)\bar{v} - Kp_i qp_c \\ &\text{subject to} \quad \sum_{T,G} rp_i qp_c \le Q \quad . \end{aligned} \tag{4-3}$$

4.2.3.2 Choice of an Initial Contamination Penalty K

Consider a likely state of nature such as the state where the optical guidance system is available and the surface temperature reading is high, i.e., $T = T_1$ and $G = G_1$. Let us find a contamination penalty K that would lead under these circumstances to an optimal fly-by with a probability of contamination of 1×10^{-5}.

The probability of contamination due to the insertion maneuver is $p_i(G_1,H)qp_c(T_1)$. If it is limited to 1×10^{-5}, the probability of impact should be

$$p_i(G_1,H) = \frac{1 \times 10^{-5}}{qp_c(T_1)} = \frac{1 \times 10^{-5}}{0.01 \times 1} = 1 \times 10^{-3} \tag{4-4}$$

and the fly-by altitude decision should be H = 500 km (see Figure 4.2).

With this small probability of impact, the expected value of the mission, Eq. (4-1), reduces approximately to

$$\overline{v(G_1,H)} - Kp_i(G_1,H)qp_c(T_1) \quad . \tag{4-5}$$

The expected value is maximum when the differential of Eq. (4-5) with respect to H at H = 500 km is null, that is, when

$$\frac{\Delta\overline{v(G_1,H)}}{\Delta H} - K\frac{\Delta p_i(G_1,H)}{\Delta H}qp_c(T_1) = 0 \quad . \tag{4-6}$$

But we know, according to Table 4.1, that

$$\left.\frac{\Delta\overline{v(G_1,H)}}{\Delta H}\right|_{H = 500 \text{ km}} = -20/1000 \text{ km} \tag{4-7}$$

and, also according to Figure 4.2, that

$$\left.\frac{\Delta p_i(G_1,H)}{\Delta H}\right|_{H = 500 \text{ km}} = -2 \times 10^{-3}/100 \text{ km} \quad . \tag{4-8}$$

We find, therefore, that the initial value of the contamination penalty K can be chosen as

$$K = \frac{(\Delta\bar{v}/\Delta H)}{(\Delta p_i/\Delta H) q p_c} = \frac{-2}{-2 \times 10^{-3} \times 10^{-2}} = 10^5 \quad . \qquad (4\text{-}9)$$

4.2.3.3 <u>The Value Iteration Process</u>

4.2.3.3.1 <u>$K = 10^5$</u>

We could again represent the strategy set in an expected value versus probability of contamination plane. However, since there are four likely states of nature and the fly-by altitude decision is continuous, each strategy would be difficult to identify. Figure 4.5 is a simpler picture representing the expected value and the expected contamination penalty as functions of the fly-by altitude H for the value $K = 10^5$.

As can be seen in Figure 4.5, the expected contamination penalty is negligible for high fly-by altitudes, and increases sharply when the fly-by altitude is lowered. It is therefore easy to determine the fly-by altitudes leading to the maximum expected net values. These values are represented in Figure 4.5 and Table 4.2.

When the temperature reading is low (T_2) and only the radio guidance system is available (G_2), the planetary quarantine requirement is not binding; that is, the fly-by altitude is not affected by the contamination penalty. In this case, an emergency maneuver to remove the spacecraft from impact trajectory becomes likely, and the optimization of the fly-by altitude decision depends critically on the value assigned to a successful removal from impact. For simplicity, we have assumed this value to be equal to $\overline{v(G,H)}$, the expected value of the fly-by when the spacecraft is not on an impact trajectory; a more careful value assignment would be needed in this particular case.

In the three other states of nature, the contamination penalty is binding and leads to sharp optima.

The prior probability of contamination for the insertion maneuver is found to be approximately 1×10^{-5}, half the mission allocation. A smaller contamination penalty should be tried next.

FIGURE 4.5

DETERMINATION OF OPTIMAL GANYMEDE FLY-BY ALTITUDE

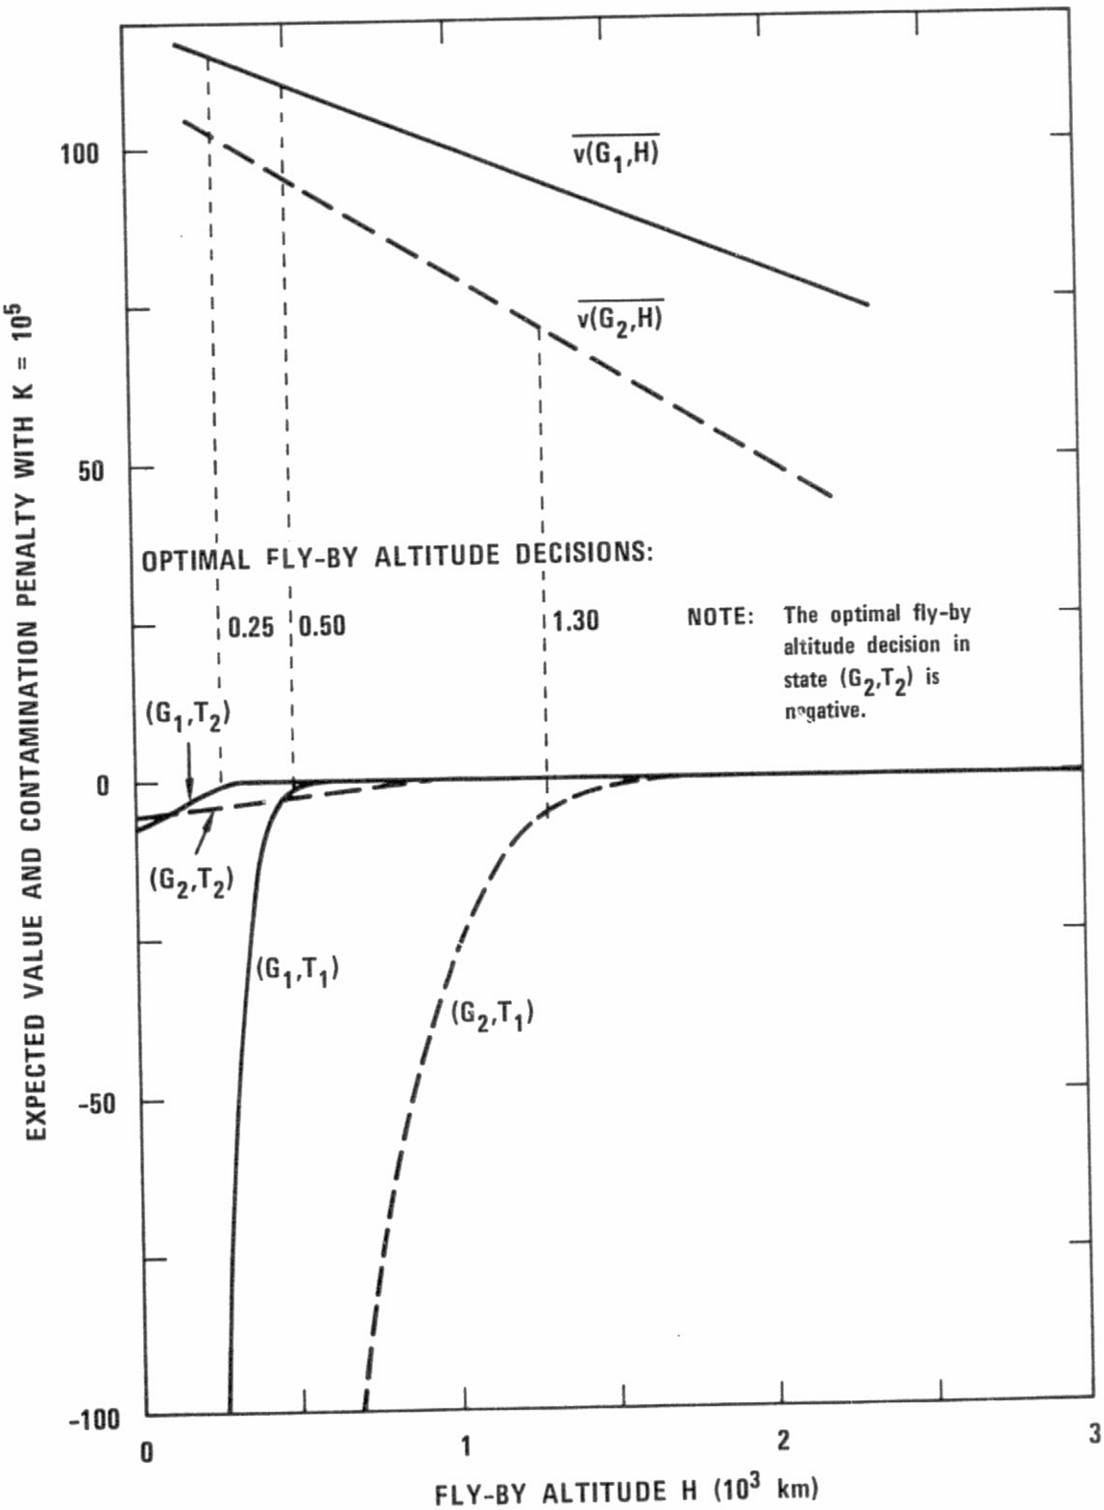

Table 4.2

OPTIMAL GANYMEDE FLY-BY ALTITUDE DECISION
FIRST ITERATION: CONTAMINATION PENALTY $K = 10^5$

State of Nature	Temperature Reading T	State of Guidance System G	Probability of State of Nature $r\ (T,G)$	Ganymede Fly-by Altitude Decision H (km)	Probability of Impact p_i	Probability of Contamination Given Impact p_c	Probability of Contamination: Conditional $p_i q p_c$ (10^{-5})	Probability of Contamination: Total $r p_i q p_c$ (10^{-5})
1	T_1	G_1	0.495	500	1×10^{-3}	1.0	1.0	0.495
2	T_1	G_2	0.005	1300	5×10^{-3}	1.0	5.0	0.025
3	T_2	G_1	0.495	250	0.1	0.01	1.0	0.495
4	T_2	G_2	0.005	negative	1.0*	0.01	10.0*	0.050*

Prior probability of contamination $p(C) \simeq 1 \times 10^{-5}$

*These numbers depend critically on the value assigned to the outcome of a successful emergency maneuver to remove the spacecraft from impact trajectory. More realistic values would require further refinements of the current model.

4.2.3.3.2 $K = 5 \times 10^4$

The value iteration process will normally be carried out using a standard computerized search method. In this illustration, we use the following heuristic to guess at the next trial value of K: since $\log(p_i(H))$ is approximately linear for most values of H (see Figure 4.2) and K is inversely proportional to $\Delta p_i/\Delta H$ [see Eq. (4-9)], to double the probability of contamination, we cut the contamination penalty in half.

The new results are represented in Table 4.3. The prior probability of contamination has jumped to 6×10^{-5}, far in excess of the mission allocation $Q = 2 \times 10^{-5}$. The largest contribution to the probability of contamination comes from the third state of nature, $T = T_2$ and $G = G_1$ (high temperature reading and optical guidance system functioning); the new contamination penalty $K = 5 \times 10^4$ is no longer binding. Note also that the probabilities of contamination conditional upon each state of nature exceed the mission allocation in all cases except the first one ($T = T_1$, $G = G_1$).

The contamination penalty must be revised upwards to reduce the prior probability of contamination.

4.2.3.3.3 $K = 7.5 \times 10^4$

We shall rely again on a heuristic to guess at the next trial value of K. The value $K = 5 \times 10^4$ leads to an excessive probability of contamination because it is not binding for the third state of nature ($T = T_2$, $G = G_1$); the smallest value of K that is still binding can be obtained as follows: From Figure 4.2 we can compute

$$\frac{\Delta p_i(G_1,H)}{\Delta H} = -0.28/100 \text{ km at } H = 0$$

and therefore, from Eq. (4-7), we compute

$$K = \frac{-2}{-0.28 \times 10^{-2} \times 10^{-2}} = 7.14 \times 10^4 \quad .$$

With a slightly larger value of K, $K = 7.5 \times 10^4$, we obtain the results in Table 4.4, which satisfy the mission allocation. The optimal fly-by altitudes are sharply defined when the temperature reading is high ($T = T_1$). When the temperature reading is low ($T = T_2$) and the

Table 4.3

OPTIMAL GANYMEDE FLY-BY ALTITUDE DECISION
SECOND ITERATION: CONTAMINATION PENALTY $K = 5 \times 10^4$

State of Nature	Temperature Reading T	State of Guidance System G	Probability of State of Nature r (T,G)	Ganymede Fly-by Altitude Decision H (km)	Probability of Impact p_i	Probability of Contamination Given Impact p_c	Probability of Contamination: Conditional $p_i q p_c$ (10^{-5})	Probability of Contamination: Total $r p_i q p_c$ (10^{-5})
1	T_1	G_1	0.495	475	2×10^{-3}	1.0	2.0	0.99
2	T_1	G_2	0.005	1240	7×10^{-3}	1.0	7.0	0.03
3	T_2	G_1	0.495	negative	1.0*	0.01	10.0*	4.95*
4	T_2	G_2	0.005	negative	1.0*	0.01	10.0*	0.05*

Prior probability of contamination $p(C) \simeq 6 \times 10^{-5}$

*These numbers depend critically on the value assigned to the outcome of a successful emergency maneuver to remove the spacecraft from impact trajectory. More realistic values would require further refinements of the current model.

Table 4.4

OPTIMAL GANYMEDE FLY-BY ALTITUDE DECISION
THIRD ITERATION: CONTAMINATION PENALTY $K = 7.5 \times 10^4$

State of Nature	Temperature Reading T	State of Guidance System G	Probability of State of Nature r (T,G)	Ganymede Fly-by Altitude Decision H(km)	Probability of Impact p_i	Probability of Contamination Given Impact p_c	Probability of Contamination	
							Conditional $p_i q p_c$ (10^{-5})	Total $r p_i q p_c$ (10^{-5})
1	T_1	G_1	0.495	487	1.5×10^{-3}	1.0	1.5	0.74
2	T_1	G_2	0.005	1260	6×10^{-3}	1.0	6.0	0.03
3	T_2	G_1	0.495	175	0.25	0.01	2.5	1.24
4	T_2	G_2	0.005	negative	1.0*	0.01	10.0*	0.05*

Prior probability of contamination $p(C) \simeq 2 \times 10^{-5}$

*These numbers depend critically on the value assigned to the outcome of a successful emergency maneuver to remove the spacecraft from impact trajectory. More realistic values would require further refinements of the current model.

optical guidance system is available ($G = G_1$), the optimal fly-by altitude is very low (175 km) and corresponds to a flat optimum. When the temperature reading is low ($T = T_2$) and only the radio guidance system is available ($G = G_2$), the fly-by altitude can be chosen arbitrarily low; with the value assignments of this illustration, it is even preferable to aim for an impact and plan to resort to an emergency maneuver.

The probabilities of contamination conditional upon each state of nature are found in the penultimate column of Table 4.4. These numbers should have been the probability suballocations for each decision if a suballocation procedure had been used.

4.3 Sketch of a JO Planetary Quarantine Analysis

4.3.1 Structure of the Mission

The planetary quarantine of the JO mission can be divided into three phases: the insertion maneuver, the exploration of the Jovian system with multiple satellite encounters, and the disposal of the spacecraft (see Figure 4.6). Each phase raises specific planetary quarantine problems.

The critical decisions for the insertion maneuver are the choice of a satellite for gravitational assistance and the choice of a fly-by altitude. We analyzed the second problem in Section 4.2, assuming that Ganymede had been selected for gravitational assistance, a logical but not necessary choice. Regardless of the choice of satellite, large quantities of fuel can be saved by flying by at a low altitude. The major trade-off is therefore between fuel conservation and risk of contamination inherent in a close fly-by.

There is a large number of promising strategies for exploring the Jovian system, and the planning is still in a rapidly evolving stage. However, two planetary quarantine issues can already be perceived in all the strategies proposed thus far.

One is a minor issue: Some important changes of velocity will be desired to modify rapidly the spacecraft orbit characteristics. To conserve fuel and save time, these velocity changes will require close fly-bys of some of the Galilean satellites. These problems are comparable in nature to the choice of fly-by altitude for the insertion maneuver but do not have the same magnitude.

FIGURE 4.6

JUPITER ORBITER MISSION PHASES

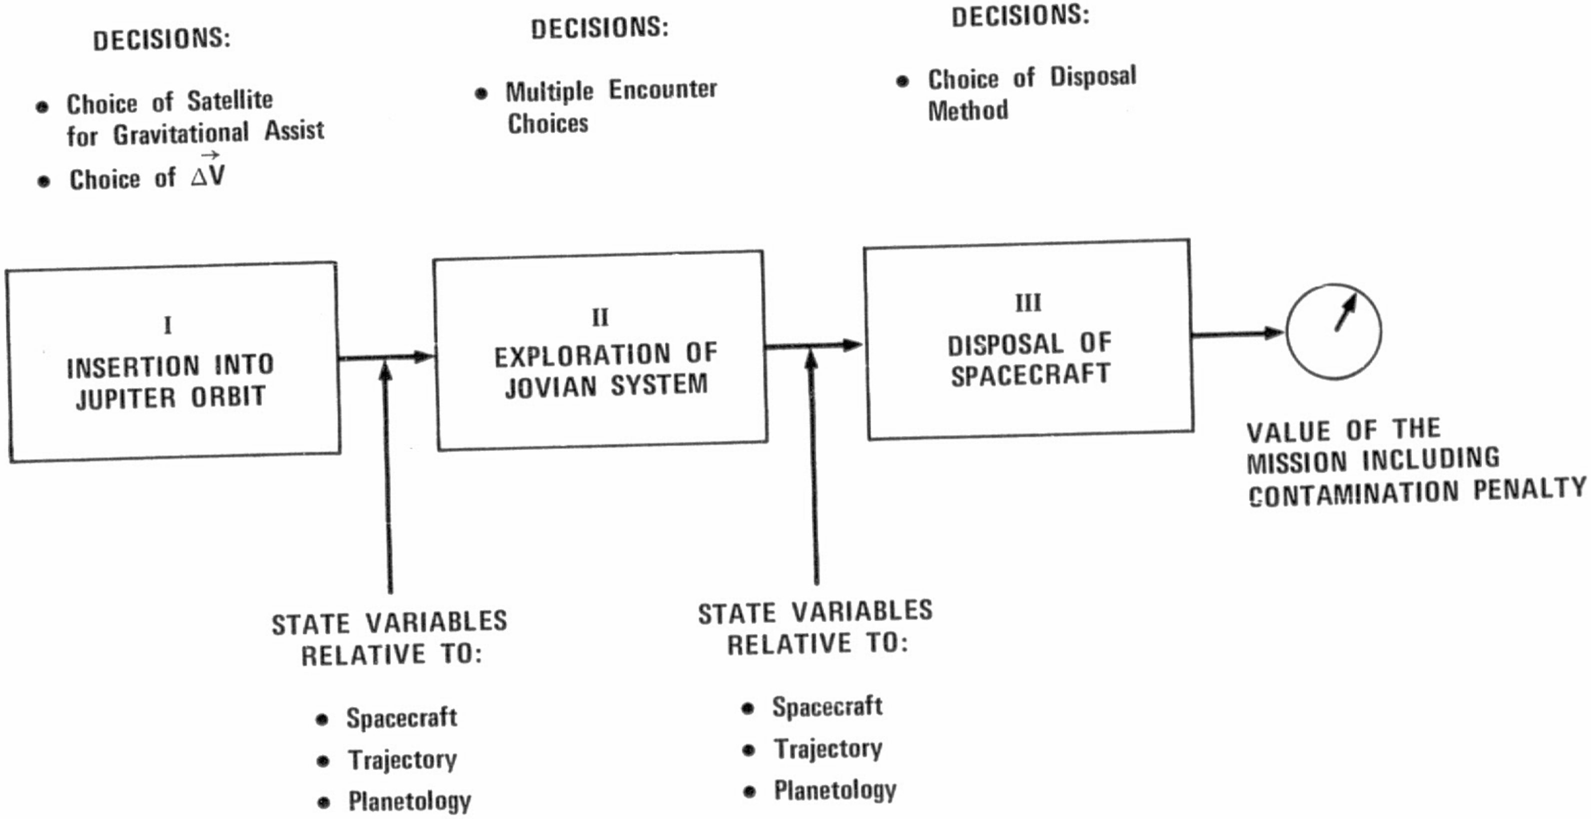

The other issue is critical: If control of the spacecraft trajectory is lost during the exploration phase, the spacecraft is very likely to impact one of the major bodies of the Jovian system sometime during the planetary quarantine period [15]. Given the current design estimates of spacecraft reliability and the probabilities of contamination if the spacecraft impacts on the Galilean satellites, the prior probability of contamination of the exploration phase exceeds by far the contemplated mission allocation. Hence, the need for an emergency disposal alternative.

The spacecraft disposal problem appears again in the third phase of the mission, but this time under standard circumstances, i.e., with complete control of the spacecraft. A standard disposal alternative must be found to guarantee that the probability of contamination will be kept below the mission allocation for the duration of the planetary quarantine period.

We shall review successively the third and the second phase of the JO mission.

4.3.2 The Third Phase: Disposal of the Spacecraft

Many alternatives are available for disposing of the spacecraft while it still remains under control. An attempt to rank some of these alternatives according to their expected net values and their fuel requirements is illustrated in Figure 4.7. For example, the alternative of injecting the spacecraft in an orbit that will have a negligible chance of impacting a body of biological interest may have the highest scientific value and the lowest contamination penalty. However, it may be one of the most fuel-demanding alternatives and may therefore require an early termination of the exploration phase. At the other extreme, crashing into one of the Galilean satellites may be the most easily achievable outcome, but it is the least preferred from a planetary quarantine point of view.

Between these extremes there is a large number of alternatives with different values and fuel requirements; for example, crashing on a polar cap of a Galilean satellite, injecting into a circular orbit within Io's orbit, crashing on Jupiter or on a dead satellite, and so forth.

Mathematically, the selection of a standard disposal alternative can be formulated as follows:

FIGURE 4.7

DETERMINATION OF OPTIMAL STANDARD DISPOSAL ALTERNATIVE

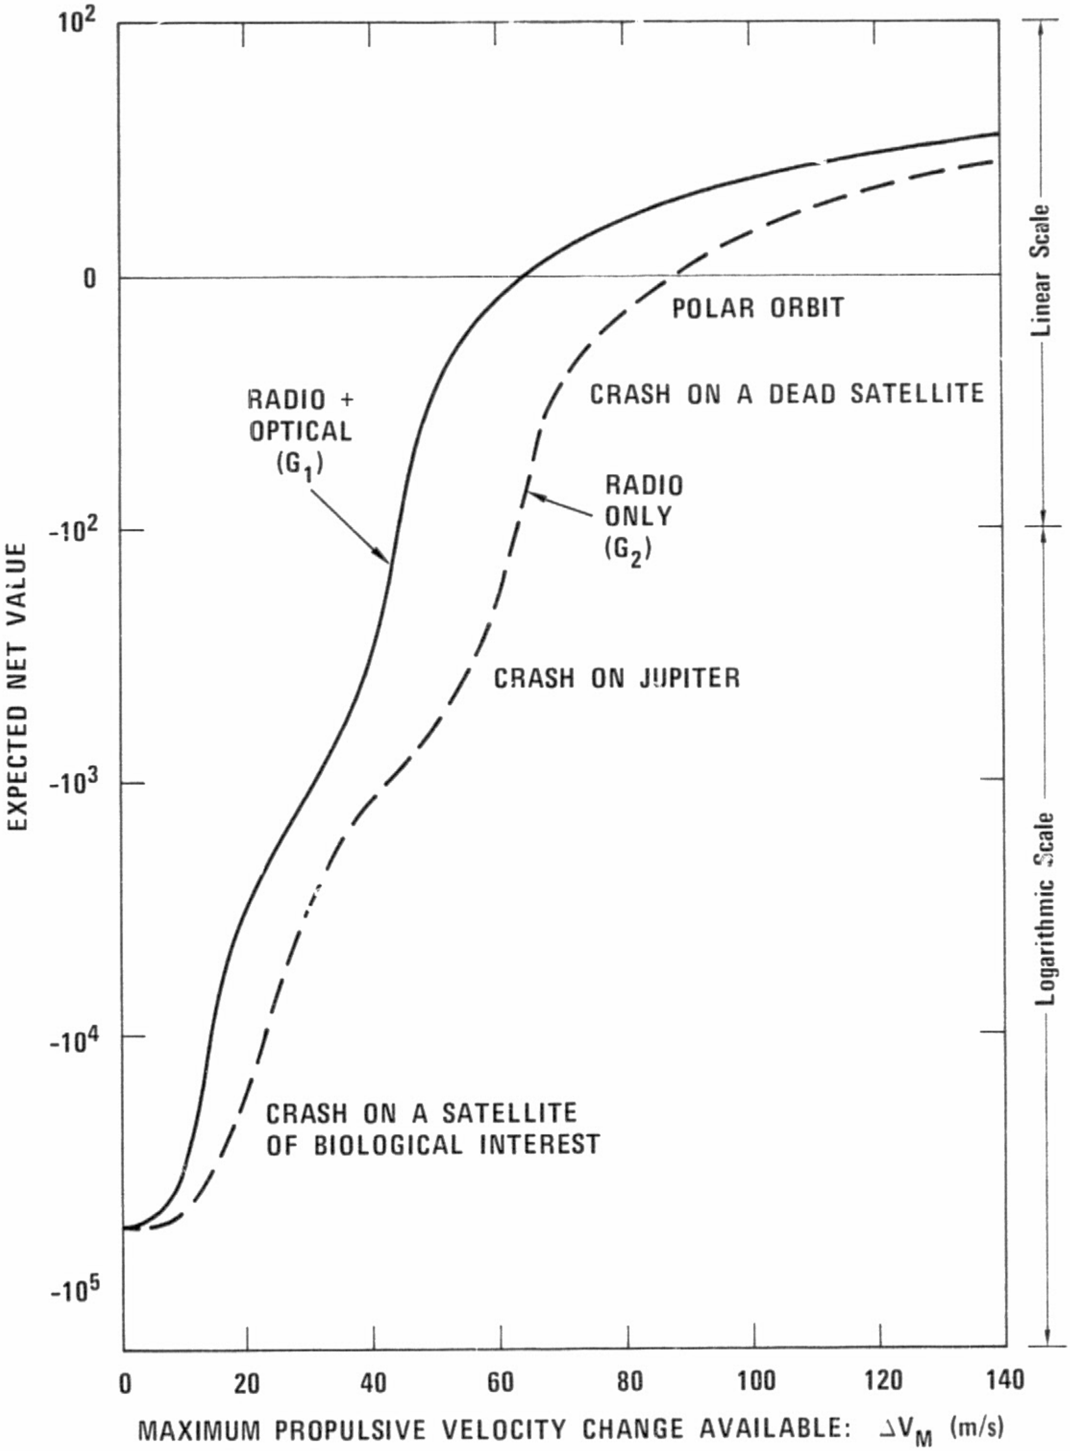

Let d = a disposal decision

$\overline{v(d)}$ = the expected scientific value of disposal decision d

$p_j(d)$ = the probability of contaminating planet j given the disposal decision d

K_j = the contamination penalty for planet j

$\Delta V(d)$ = the propulsive change of velocity requirement for disposal decision d

ΔV_M = the propulsive change of velocity still available for disposal.

Then, the problem is to maximize over all disposal alternatives the expected net value, the maximization being constrained by the ΔV availability, i.e.,

$$\underset{d}{\text{maximize}} \quad \overline{v(d)} - \sum_{j=1}^{m} K_j p_j(d) \tag{4-10}$$

$$\text{subject to} \quad \Delta V(d) \le \Delta V_M \quad .$$

Alternatively, the fuel constraint can be introduced into the objective function by pricing the ΔV availability. Letting K_V be the price of a ΔV unit, the alternative formulation is

$$\text{maximize} \quad \overline{v(d)} - \sum_{j=1}^{m} K_j p_j(d) - K_V \Delta V(d) \quad . \tag{4-11}$$

The price K_V will be chosen in an iterative manner so that the ΔV constraint for the entire mission will be met.

We shall denote by $\overline{v_3}$ the expected net value for the optimal disposal decision. According to Eq. (4-11), $\overline{v_3}$ is a function of ΔV_M [Figure 4.7 shows two illustrative plots of $\overline{v_3}$ (ΔV_M) conditional upon two states of the guidance system]. According to Eq. (4-11), $\overline{v_3}$ is a function of K_V.

C.2

4.3.3 The Second Phase: Multiple Encounters

4.3.3.1 The Probability of Contamination

The structure of the retargeting decision problem faced at each encounter with a Jovian satellite is depicted in Figure 4.8. If the main propulsion system fails (probability f), the spacecraft is likely to crash onto a Galilean satellite sometime during the planetary quarantine period. The corresponding probability of contamination is represented by $\overline{p_c}$. If the main propulsion system does not fail, the retargeting maneuver will result in a probability of immediate impact p_i. An emergency maneuver has a probability q of removing the spacecraft from an impact trajectory. The probability of contamination for one encounter is therefore

$$p(C) = \overline{fp_c} + (1 - f)p_i q p_c \quad . \tag{4-12}$$

We shall simplify and extend this result to a sequence of n encounters by assuming that:

(1) All the probabilities of contamination given impact are approximately equal.

(2) The probabilities f, p_i, and q are small.

The probability of contamination for n encounters is then approximately

$$p(C) \simeq n(f + p_i q)p_c \tag{4-13}$$

where $(f + p_i q)$ is the total probability of impact at each encounter and p_c is still the probability of contamination given impact.

A limitation of p(C) to 1×10^{-5} will impose strict constraints on the probability of impact, i.e., on the spacecraft's reliability (p,q) and on the guidance strategy (n, p_i).

4.3.3.2 Determination of the Probability of Contamination Given Impact, p_c

The assessment of the probability of contamination given impact depends on the spacecraft bioload, the release and transportation mechanisms on the impacted planet of the terrestrial organisms, and the planet's environment. Simplified models of contamination given impact

FIGURE 4.8

PROBABILITY OF CONTAMINATION FOR ONE ENCOUNTER

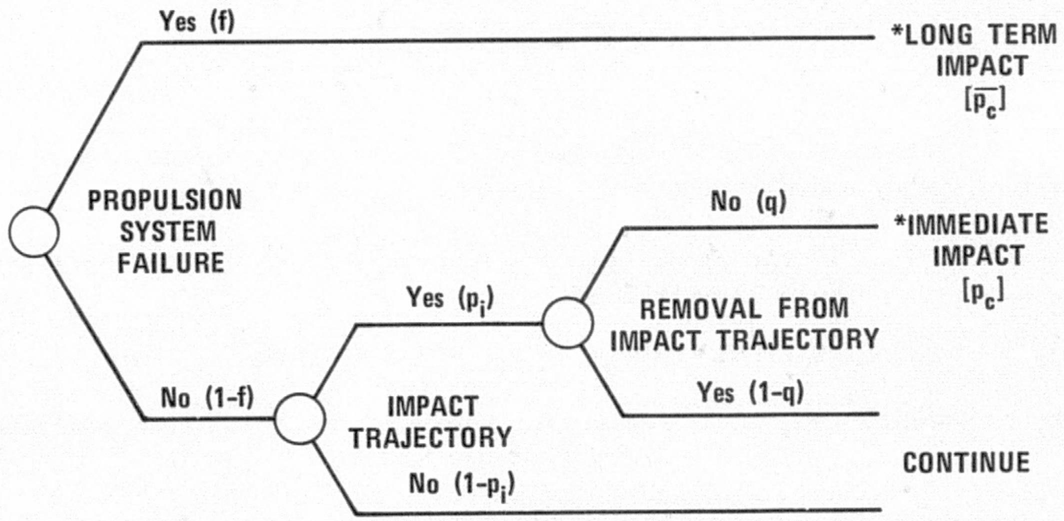

$$p(C) = f\overline{p_c} + (1-f)p_i q p_c$$

☐ = Decision node; ○ = Chance node; * = Contamination;

( ) = Probability; [ ] = Probability of contamination summarize these factors by two quantities: the expected number N of viable organisms released on the planet and the probability of growth p_g of each of these organisms. The probability of contamination given impact can then be computed by taking the product of these quantities (see a discussion of the Sagan-Coleman formula in [16]), that is,

$$p_c = N \times p_g \quad . \tag{4-14}$$

This equation assumes, among other things, that each organism has an independent chance of survival p_g. A different result would be obtained with the opposite assumption of complete dependence: All organisms will survive and proliferate with probability p_g or die. The determination of p_c may be extremely sensitive to the choice of assumption as illustrated in Figure 4.9.

For example, assuming that the probability of growth on Jupiter is 1×10^{-7} and that the spacecraft bioload is 1×10^5, the probability of contamination given impact may be as high as 1×10^{-2} or as low as 1×10^{-7}. The probability of contamination given impact on a Galilean satellite, assuming that $p_g = 0.1$, is, of course, much less sensitive to the dependence assumptions: p_c will always be between 0.1 and 1. If p_g for the Galilean satellites is revised downwards, sensitivity to the dependence assumptions may be significant, as with Jupiter.

4.3.3.3 Constraints on Spacecraft Reliability and Guidance Strategy

Based on Eq. (4-14), the probability of impact at each maneuver has been plotted against the ratio $p(C)/p_c$ in Figure 4.10. Each plot in Figure 4.10 corresponds to a given number of encounters. Thus, if $p(C) = 1 \times 10^{-5}$, $p_c = 0.5$, and $(f + p_i q) = 1 \times 10^{-5}$, we can read in Figure 4.10 that no more than two encounters are permissible. However, current estimates of the spacecraft reliability are no better than $f = 1 \times 10^{-2}$ to 1×10^{-4}. With these values, no encounter is permissible. Of course, the mission allocation may be raised, the probabilities of contamination given impact may be revised downward, and the spacecraft reliability may be improved. However, considerable modification of these parameters would be required to permit a reasonable number of encounters.

4.3.3.4 Desirable Characteristics of an Emergency Disposal System

Assuming that the reliability of the main propulsion system cannot be greatly improved, an emergency disposal system must be designed. The characterization of such a system is beyond the scope of this project, but we can describe two of its main specifications: the probability that

FIGURE 4.9

DETERMINATION OF THE PROBABILITY OF CONTAMINATION GIVEN IMPACT

(A) INDEPENDENCE ASSUMPTION

$$\phi_g = p_g$$

(B) A PARTIAL DEPENDENCE ASSUMPTION

$$\phi_g = \begin{cases} p_g^{1/2} & \text{with prob. } p_g^{1/2} \\ 0 & \text{-------------------- } (1-p_g^{1/2}) \end{cases}$$

(C) COMPLETE DEPENDENCE ASSUMPTION

$$\phi_g = \begin{cases} 1 & \text{with prob. } p_g^{1/2} \\ 0 & \text{with prob. } (1-p_g) \end{cases}$$

ϕ_g = long run frequency of survival

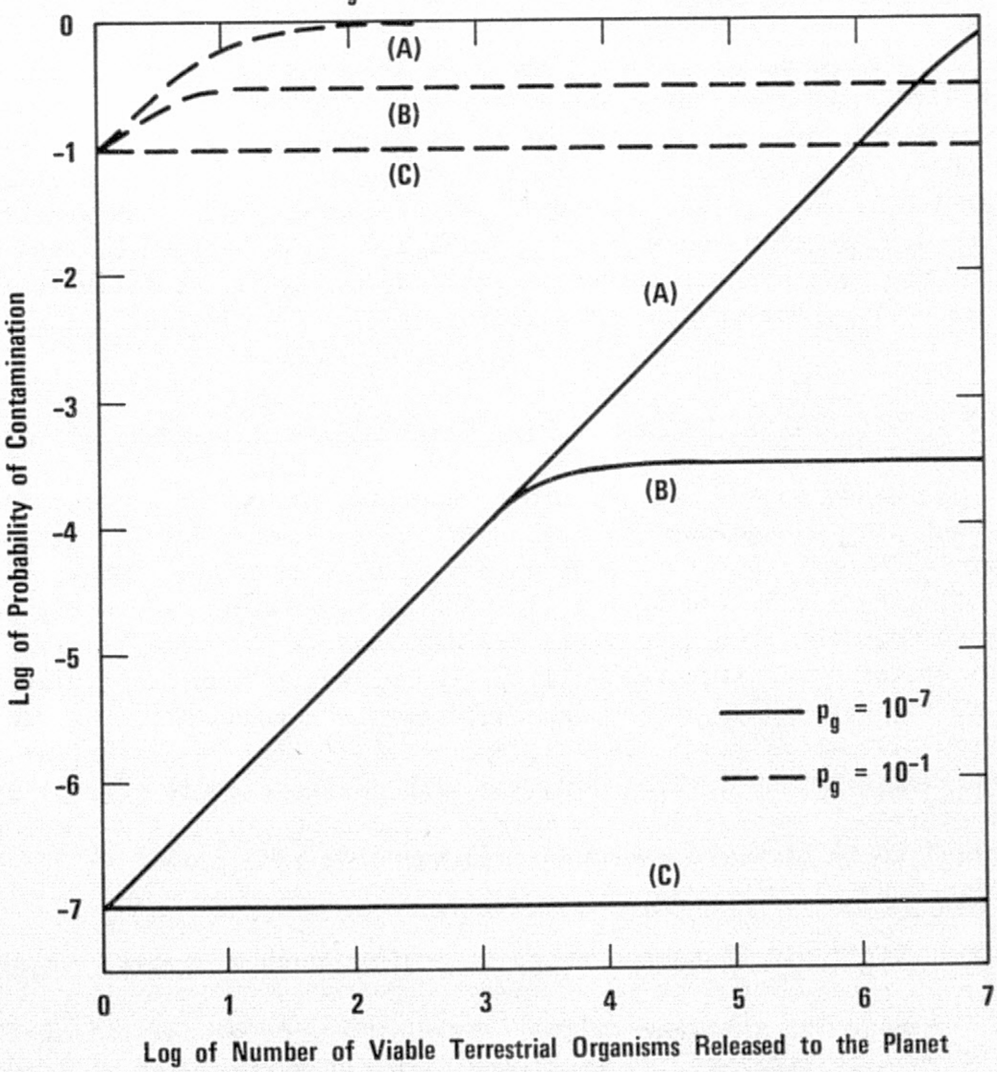

FIGURE 4.10
FEASIBLE PROBABILITIES OF IMPACT PER ENCOUNTER FOR MULTIPLE-ENCOUNTERS

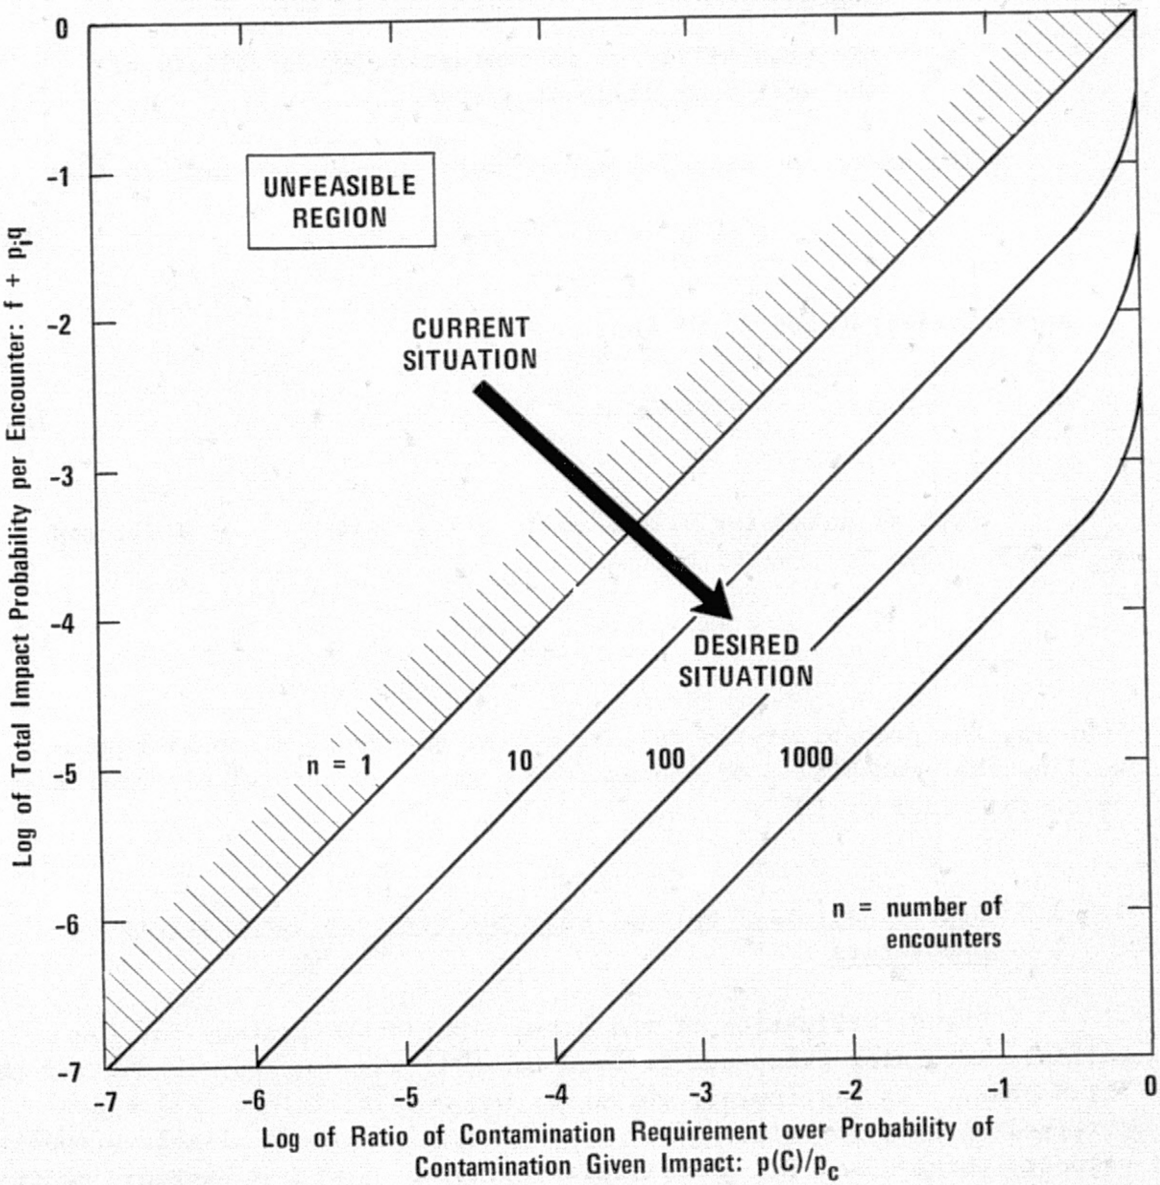

this system will fail f' and the probability of contamination given that it functions successfully p'_c (see Figure 4.11).

We shall also use the following notations:

Q = the probability of contamination of the multiple encounter phase

f'' = the probability that the main propulsion system will fail sometime during the multiple-encounter phase

$\overline{p_c}$ = the probability of contamination given failure of the emergency disposal system.

The emergency disposal system must then be designed so that

$$f''[f'\overline{p_c} + (1 - f')p'_c] < Q \tag{4-15}$$

or approximately (with $f' \ll 1$)

$$f''(f'\overline{p_c} + p'_c) < Q \quad . \tag{4-16}$$

Typical parameter values are: $Q = 2 \times 10^{-5}$, $f'' = 0.01$, and $\overline{p_c} = 0.5$. Then Eq. (4-16) becomes

$$f' + 2p'_c < 4 \times 10^{-4} \quad . \tag{4-17}$$

That is, the probability of failure of the emergency disposal system as well as the probability of contamination given successful disposal must be on the order of 10^{-4}.

4.3.3.5 <u>Determination of Optimal Guidance Strategy for Multiple Encounters</u>

The determination of the optimal guidance strategy for the multiple-encounter phase can be formulated in the same way as that of the third phase. In particular, there are two possibilities: (1) a constrained optimum formulation with a constraint on the available propulsive velocity change, and (2) a Lagrangian formulation with a Lagrange multiplier K_V. The expected net value of the optimal disposal strategy, $\overline{v_3}$,

FIGURE 4.11

REQUIREMENTS FOR AN EMERGENCY DISPOSAL SYSTEM

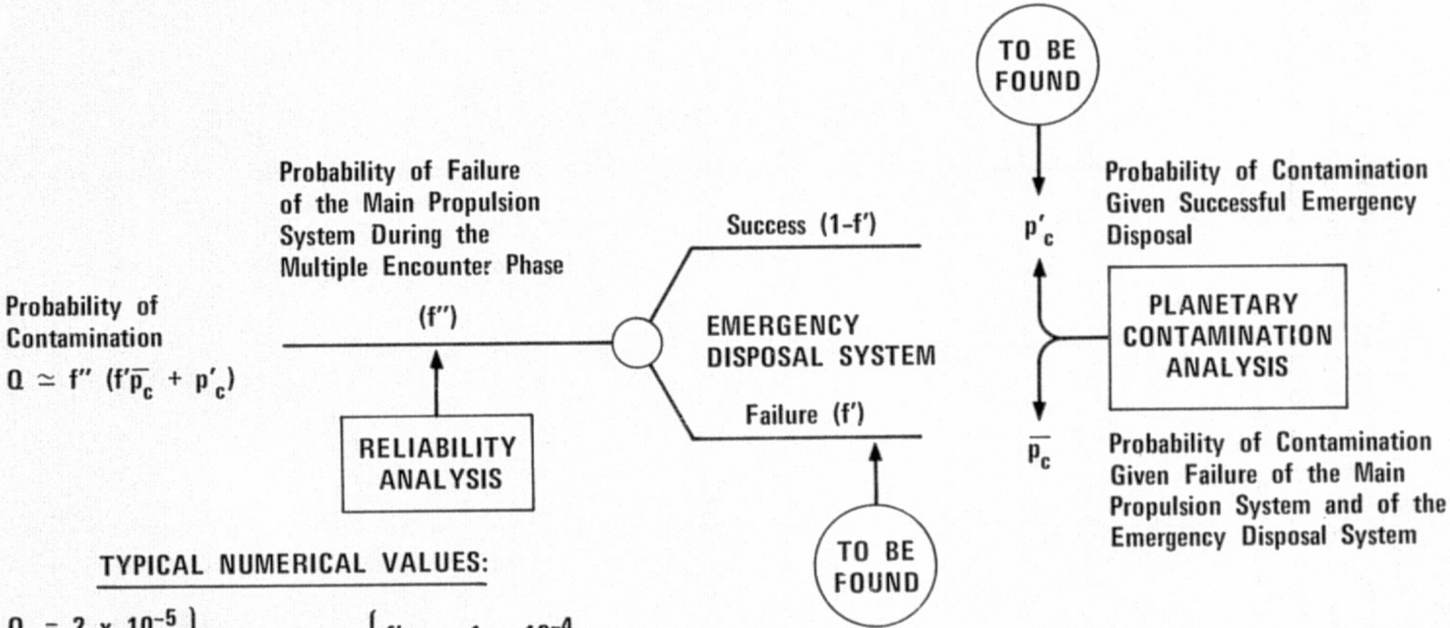

TYPICAL NUMERICAL VALUES:

$$\left.\begin{aligned} Q &= 2 \times 10^{-5} \\ f'' &= 0.1 \\ \bar{p}_c &= 0.5 \end{aligned}\right\} \text{ and therefore } \left\{\begin{aligned} f' &\approx 1 \times 10^{-4} \\ p'_c &\approx 1 \times 10^{-4} \end{aligned}\right.$$

is assumed to be known. The only difficulty is in the dimensionality of the problem and the large number of probability and value assessments that must be made.

The expected net value of the optimal strategy for the second and third phases of the JO mission can be denoted by $\overline{v_2}$. This is the value function that we have directly assessed in Table 4.1 and Figure 4.4 to analyze the insertion maneuver.

Appendix

SOME SHORTCOMINGS OF THE CHANCE-CONSTRAINED FORMULATION OF PLANETARY QUARANTINE

Appendix

SOME SHORTCOMINGS OF THE CHANCE-CONSTRAINED FORMULATION OF PLANETARY QUARANTINE

A.1 Irreconcilability of Chance-Constraints and Expected Utility Theory

The selection of space exploration programs subject to current planetary quarantine requirements can be summarized as follows: Find the space exploration program having the maximum expected value provided that the probability of planetary contamination be less than or equal to a given number Q.

In mathematical formulation,

$$\underset{i \in I}{\text{maximize}} \int_{x \in X} v(x)\, dF_i(x) \tag{A-1}$$

$$\text{subject to} \int_{x \in C} dF_i(x) \le Q \tag{A-2}$$

where x is a complete description of a space program outcome; $F_i(x)$ is the cumulative distribution of outcomes corresponding to program i; C is the set of outcomes entailing contamination, a subset of the set of all possible outcomes X; and v(x) is a value function assessed by the decision maker over the set of possible outcomes.

Problems of the form (A-1) and (A-2) are known in mathematical programming as chance-constrained problems. A prospect of outcomes x with probability distribution F(x) is often referred to as a lottery.

An alternative formulation is to represent <u>all</u> the decision makers' preferences, including the risk and consequences of contamination by a utility function u(x), and to maximize the unconstrained expected utility, i.e.,

PRECEDING PAGE BLANK NOT FILMED

$$\underset{i \in I}{\text{maximize}} \int_{x \in X} u(x)\ dF_i(x) \quad . \tag{A-3}$$

However, we shall prove now that a chance-constrained formulation with a positive probability constraint is inconsistent with the existence of a utility function.

Consider two space exploration programs (or lotteries) L_1 and L_2 (see Figure A.1) having expected values V_1 and V_2 [computed from Eq. (A-1)] and probabilities of contamination q_1 and q_2 [computed from Eq.(A-2)] such that

$$V_1 > V_2 \tag{A-4}$$

$$q_1 > Q > q_2 \quad . \tag{A-5}$$

According to the chance-constrained formulation, the first program has a larger expected value than the second but must be rejected in favor of the second because its probability of contamination q_1 exceeds the chance constraint.

The expected utility formulation will lead to the same preference ordering between L_1 and L_2 if and only if the expected utility U_2 of L_2, computed from Eq. (A-3), is larger than the expected utility U_1 of L_1, namely,

$$U_2 > U_1 \quad . \tag{A-6}$$

Let us now construct a new program L_3 in the following manner: First, consider a random event E; if the random event occurs (probability p), we proceed with program L_1, whereas if the random event does not occur [probability (1-p)], we proceed with program L_2. The expected value and probability of contamination of program L_3 are therefore

$$V_3 = pV_1 + (1 - p)V_2 \tag{A-7}$$

$$q_3 = p\, q_1 + (1 - p)q_2 \quad . \tag{A-8}$$

FIGURE A.1

CHANCE-CONSTRAINED PREFERENCE ORDERING OF SPACE EXPLORATION PROGRAMS

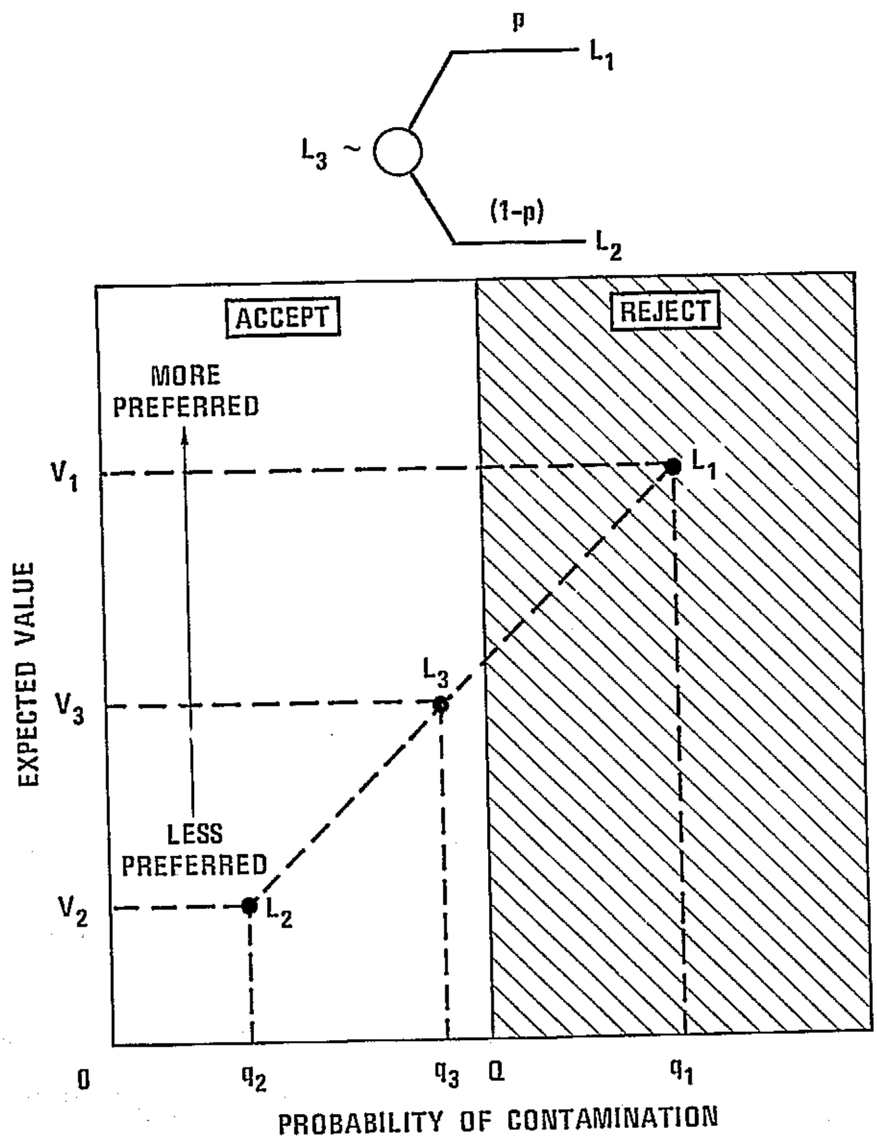

Now, since $q_1 > Q > q_2$, it is always possible to find $p > 0$ so that $q_3 < Q$, i.e., so that program L_3 meets the chance-constraint. [For example, choosing $p = (1/2)(Q - q_2)/(q_1 - q_2)$ will give $q_3 = (1/2)(Q + q_2) < Q$.] The availability of a random event E with the appropriate probability p is assumed for the sake of the argument.

From Eq. (A-4), (A-7), and $p > 0$ we can see that

$$V_3 > V_2 \tag{A-9}$$

and, therefore, program L_3 will be preferred to program L_2, according to the chance-constrained formulation.

However, the expected utility of program L_3 is

$$U_3 = pU_1 + (1 - p)U_2 \quad , \tag{A-10}$$

and since $U_2 > U_1$, a positive p implies

$$U_3 < U_2 \quad . \tag{A-11}$$

That is, program L_2 should be preferred to program L_3 on an expected utility basis. The preference orderings of space exploration programs resulting from a chance-constrained formulation and from expected utility theory may not only differ, but may differ in a way that cannot be resolved by a mere respecification of a utility function.

When the chance-constraint Q is nil, i.e., when even the slightest chance of contamination is rejected, utility theory can explain the preference ordering of chance-constrained programs. It suffices to assign an infinite negative utility to contamination outcomes.

A.2 The Violation of Two Axioms of Utility Theory by Chance-Constraints

Expected utility theory is founded on a set of intuitive and widely accepted axioms (for example, see [10], Section 2.5). Chance-constraints may also seem intuitive at first, but must be recognized as an arbitrary decision procedure. The irreconciliability of the two approaches is not surprising in view of the fact that chance-constraints violate the two following axioms of utility theory:

(1) Continuity--A decision maker preferring alternative L_3 to L_2 and L_2 to L_1 will be indifferent to the choice between having alternative L_2 or having the possibility to play some lottery involving just L_1 and L_3.

(2) Monotonicity--If two lotteries involve the same two alternatives, a decision maker will prefer the lottery offering the most preferred alternative with the higher probability.

We shall now demonstrate that chance-constraints violate these two axioms.

A.2.1 Axiom 1: Continuity

As before, we denote by V_i and q_i the expected value and probability of contamination associated with alternative L_i; Q is a positive chance-constraint.

Consider three alternatives L_1, L_2, and L_3 (see Figure A.2) such that

$$q_1 > Q > q_3 \text{ and } q_2 \tag{A-12}$$

$$V_3 > V_2 \tag{A-13}$$

$$V_1 > V_2 + (V_2 - V_3)\left(\frac{q_1 - Q}{Q - q_3}\right) \quad . \tag{A-14}$$

(Note: $V_1 > V_2$ would do, a fortiori, and would simplify the demonstration.) Then L_3 is preferred to L_2, which is preferred to L_1. The characteristics of a lottery between L_1 with probability p and L_3 with probability (1 - p) are

$$V = pV_1 + (1 - p)V_3 \tag{A-15}$$

$$q = pq_1 + (1 - p)q_3 \quad . \tag{A-16}$$

There are two possibilities: If $q > Q$, the lottery is unacceptable, and L_2 is preferred to the lottery; if $q \leq Q$, the lottery is acceptable. Given inequalities (A-12) and (A-13), we see that

FIGURE A.2

VIOLATION OF THE AXIOM OF CONTINUITY BY CHANCE CONSTRAINTS

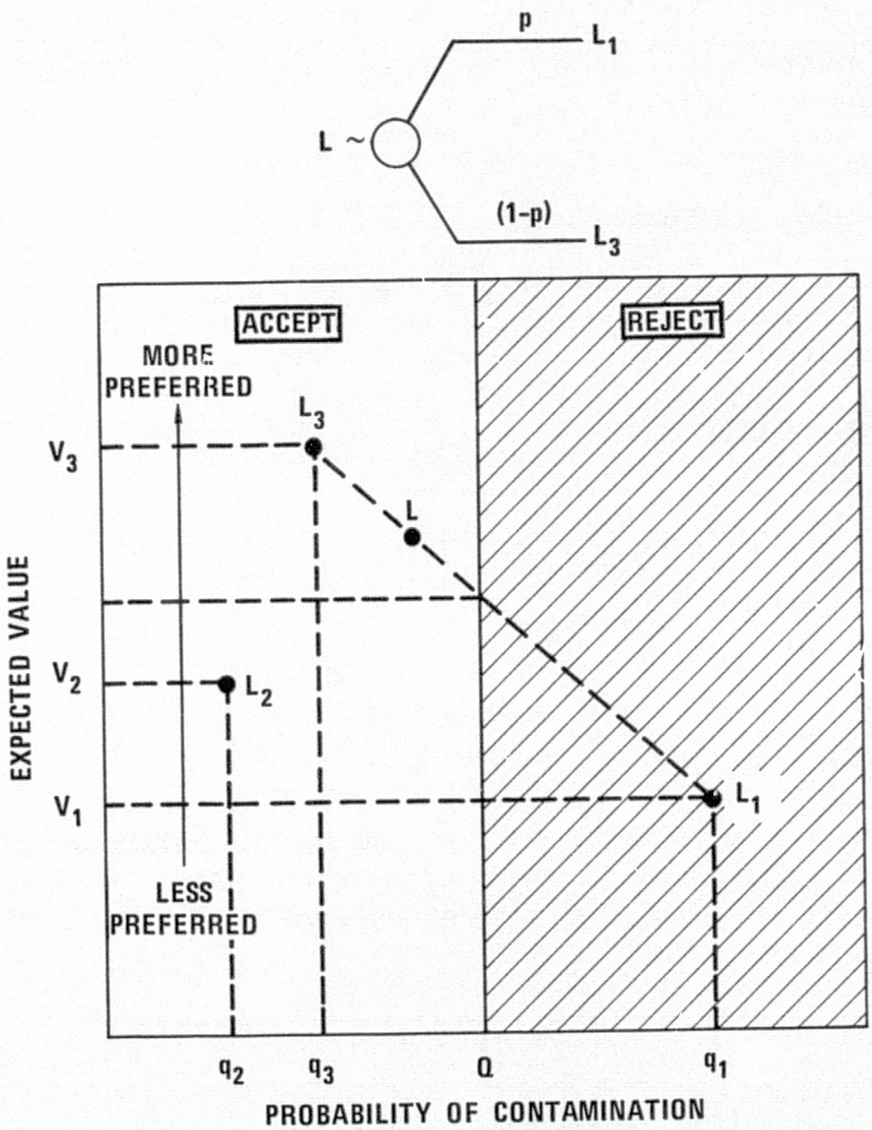

NOTE: In this illustration L_3 is preferred to L_2, which is preferred to L_1. However, any lottery L between L_1 and L_3 is either more desirable than L_2 or less desirable than L_2; No lottery L is equally desirable as L_2.

$$(V_2 - V_3)\left(\frac{q_1 - Q}{Q - q_3}\right) \geq (V_2 - V_3)\left(\frac{q_1 - q}{q - q_3}\right) \tag{A-17}$$

and, therefore, from (A-14) and (A-17)

$$V_1 > V_2 + (V_2 - V_3)\left(\frac{q_1 - q}{q - q_3}\right) \quad . \tag{A-18}$$

Substituting (A-16) for q in (A-18) and the right side of (A-18) for V_1 in (A-15), we obtain, after simplification,

$$V > V_2 \tag{A-19}$$

That is, the lottery is preferred to L_2.

There is no lottery between L_1 and L_3 such that the decision maker would be indifferent between having L_2 or having the possibility to play the lottery.

A.2.2 Axiom 2: Monotonicity

Consider two alternatives L_1 and L_2 such that

$$V_1 > V_2 \tag{A-20}$$

$$q_1 > Q > q_2 \quad , \tag{A-21}$$

i.e., L_1 is unacceptable and L_2 is preferred to L_1.

The characteristics of a lottery between L_1 with probability p and L_2 with probability (1 - p) are

$$V = pV_1 + (1 - p)V_2 \tag{A-22}$$

$$q = pq_1 + (1 - p)q_2 \quad . \tag{A-23}$$

Consequently, V and q increase linearly with p. It is therefore possible to find two lotteries L' and L'' with $p' > p''$ such that

$$V_1 > V' > V'' > V_2 \tag{A-24}$$

$$q_1 > Q > q' > q'' > q_2 \quad , \tag{A-25}$$

i.e., L' is preferred to L'', although L_2 is preferred to L_1 and L'' offers the greater probability of obtaining L_2.

A.3 Chance-Constraints and the Negative Value of Information

The intuitive appeal of a set of axioms is always debatable; it may be more enlightening to explore some of the consequences of rejecting the axioms. Thus we shall show now that rejecting the axioms of expected utility theory and adopting instead chance-constraints can lead to a distasteful consequence: If a space exploration program is selected using chance-constraints, information pertinent to the program may have a negative expected value.

Consider the fly-by mission depicted in Figure A.3. This mission has a structure similar to program L_3, described above, that is, two states of nature may prevail. In this instance, free water may be present on the target planet with probability 0.4 or be absent with probability 0.6. If free water is present, the probability of contamination given impact is 1.0, whereas it is only 0.01 otherwise. The only decision is to launch or to forego the mission. If the mission is flown, the probability of impact is 2×10^{-4}. The mission allocation is assumed to be 1×10^{-4}. We shall analyze the effect of receiving perfect information about the presence or absence of free water prior to the launching decision (1) when chance-constraints are used to make the decision (2) when the decision is based on expected utilities.

A.3.1 Chance-Constrained Formulation

For simplicity, a scientific value of 100 (in millions of dollars, if you will) has been assigned to a successful mission. A scientific value of zero has been assigned to the other two possibilities: crash or no launch. The analysis could have been carried out with more refined values distinguishing between these last two outcomes or conditional upon the presence of free water, or both.

Figure A.3 (a) represents the analysis of the launching decision when only the prior probability about the presence of free water is

FIGURE A.3

CHANCE-CONSTRAINED FORMULATION OF A SINGLE-DECISION FLY-BY MISSION WITH ONE UNCERTAIN STATE VARIABLE

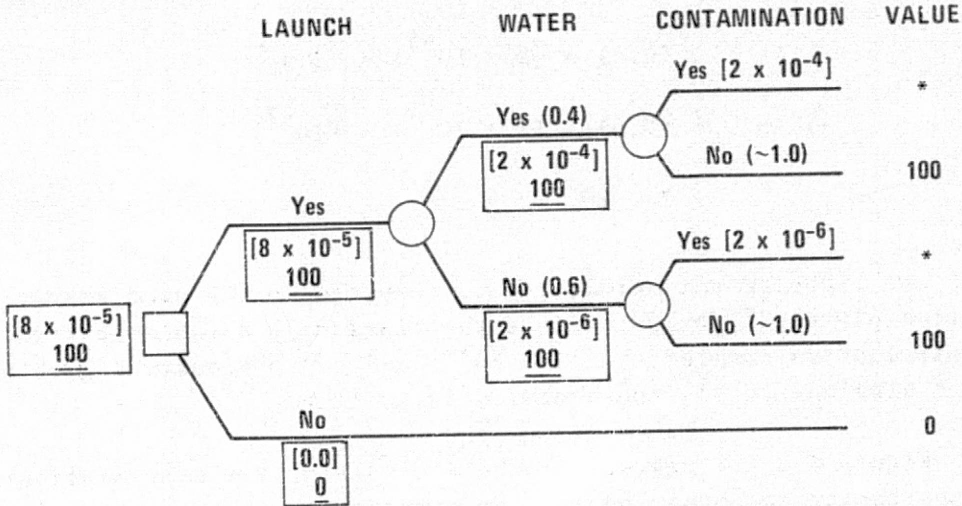

(a) WITHOUT INFORMATION ABOUT PRESENCE OF WATER

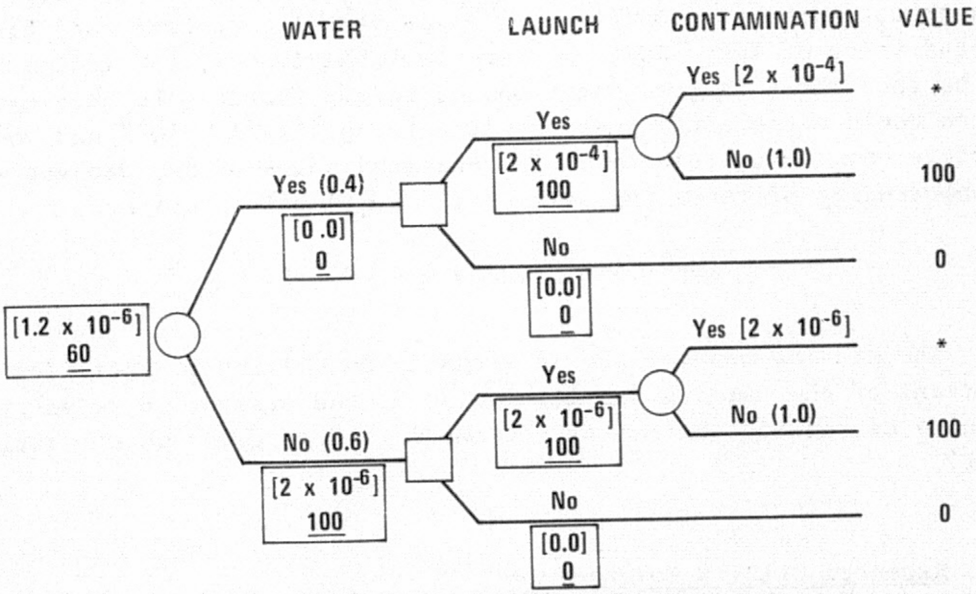

(b) WITH INFORMATION ABOUT PRESENCE OF WATER

□ = Decision node; ○ = Chance node; * = Contamination; ( ) = Probability;

[ ] = Probability of contamination; ____ = Expected value

□ Boxes contain the results of solving the tree

NOTE: Information about the presence or absence of water diminishes the expected value of the mission.

known. The probability of contamination associated with the mission is equal to

Pr. (water) $\times$ Pr. (Impact) $\times$ Pr. (Contamination given impact and water) + Pr. (No water) $\times$ Pr. (Impact) $\times$ Pr. (Contamination given impact and no water) =

$$(0.4) \times (2 \times 10^{-4}) \times (1.)$$

$$+ (0.6) \times (2 \times 10^{-4}) \times (10^{-2})$$

$$= 8 \times 10^{-5} \quad .$$

The expected scientific value, computed in the same manner, is approximately equal to 100. Since the mission is acceptable (probability of contamination smaller than 1×10^{-4}) and has obviously a greater value than no mission at all, the decision is to launch.

Figure A.3 (b) represents the analysis of the same mission with the opportunity to obtain perfect information about the eventual presence of free water on the target planet prior to launch. With probability 0.6, the indication will be that there is no free water. A rapid inspection shows that, under this condition, the mission is acceptable and has an expected value of 100. However, with probability 0.4, the indication will be that there is free water on the target planet. In this case, the mission would cause contamination with probability 2×10^{-4} and would therefore have to be rejected. The expected value of the project with the opportunity of receiving perfect information is therefore

$$100 \times 0.6 + 0 \times 0.4 = 60 \quad .$$

The mission manager should actually be willing to sacrifice up to 40 percent of the total scientific value of the mission to reject the opportunity of knowing whether or not there is free water on the target planet!

A.3.2 Expected Utility Formulation

The aberration of a negative value of information cannot emerge from an expected utility formulation. Figure A.4 illustrates the same mission with an extremely negative value of -750000 assigned to the contamination outcomes in order to reproduce the same strategies as with the chance-constrained formulation.

FIGURE A.4

EXPECTED VALUE FORMULATION OF A SINGLE-DECISION FLY-BY MISSION WITH ONE UNCERTAIN STATE VARIABLE

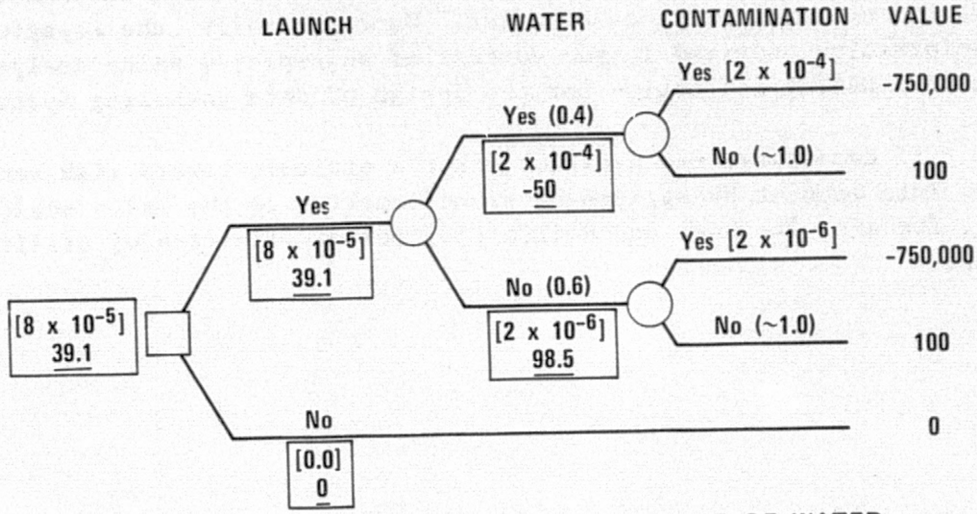

(a) WITHOUT INFORMATION ABOUT PRESENCE OF WATER

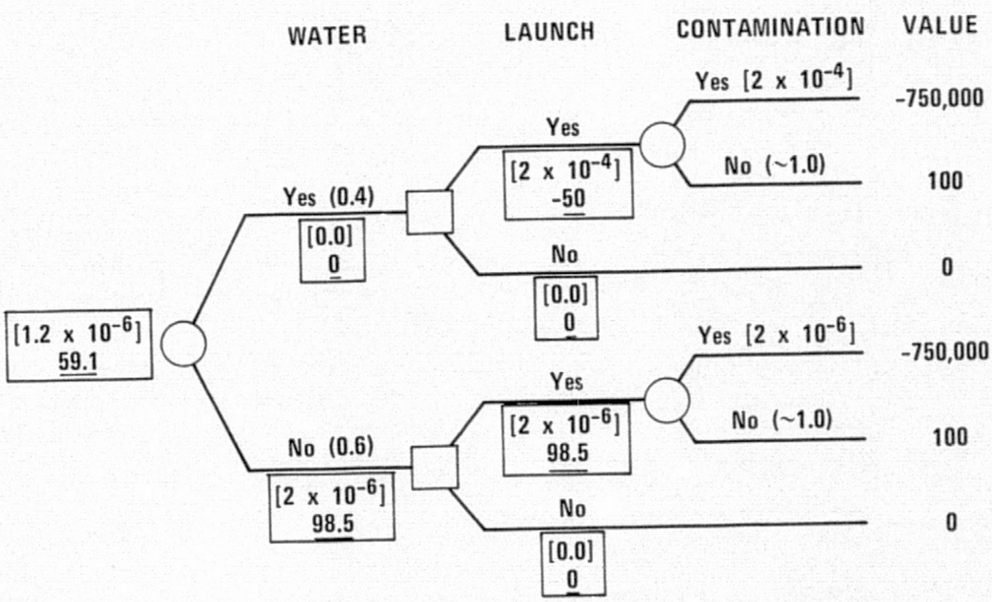

(b) WITH INFORMATION ABOUT PRESENCE OF WATER

☐ = Decision node; ◯ = Chance node; * = Contamination; ( ) = Probability;

[ ] = Probability of contamination; ___ = Expected value;

☐ Boxes contain the results of solving the tree

NOTE: Information about the presence or absence of water raises the expected value of the mission (positive value of information).

The analysis shows that without prior information about the presence of water the best decision is to launch; the expected value is only 39.1 because the contamination event has been penalized. On the other hand, with perfect information, the expected value of the project is 59.1. Therefore, a positive value of 20 (in millions of dollars, if you will) can be derived from the opportunity of receiving perfect information about the eventual presence of water. More generally, the expected value of information, derived in the context of an expected value analysis, will be a useful indication for the design of data gathering systems.

If necessary, the decision maker's attitude toward risk can be taken into account by applying a transformation to the value scale. (See, for example, Luce and Raiffa [10] for a definition of utility scales.)

REFERENCES

1. "Planetary Quarantine Specification Sheets," report prepared for the Chief, Planetary Quarantine Program, NASA Headquarters by Exotech Systems, Inc., Gaithersburg, Maryland, tentative specification p. II-7 (Revised January 1975).

2. B. R. Judd, D. W. North, and J. P. Pezier, "Assessment of the Probability of Contaminating Mars," Final Report, SRI Project MSU-2788, NASA Contract NASW-2535, SRI, Menlo Park, California (April 1974).

3. R. W. Davies and M. G. Comuntzis, "The Sterilization of Space Vehicles to Prevent Extraterrestrial Biological Contamination," Proceedings Tenth International Astronautical Congress, London, 1959, pp. 495-504 (Springer Verlag, Vienna, 1960).

4. J. O. Light, "The Non-Contamination Constraint and its Application to Capsule Delivery Techniques," Jet Propulsion Laboratory TM 312-615, Pasadena, California (19 November 1965).

5. R. G. Chamberlain, "Adaptive Allocation of Planetary Quarantine Probabilities," Jet Propulsion Laboratory TM 312-876, Pasadena, California (28 February 1968).

6. C. Sagan and S. Coleman, "Spacecraft Sterilization Standards and Contamination of Mars," Astronautics and Aeronautics, Vol. 3, p. 22 (May 1965).

7. J. E. Matheson and W. J. Roths, "Decision Analysis of Space Projects: Voyager Mars," paper presented at National Symposium, Saturn/Apollo and Beyond, American Astronautical Society (11 June 1967).

8. J. E. Matheson and A. B. Pollard, "A New Approach to Space Project Planning: Decision Analysis of Voyager," Final Report, SRI Project 6152, prepared for General Electric Company Missile and Space Division by SRI, Menlo Park, California (December 1967).

9. R. A. Howard, "Dynamic Programming," Management Science, Vol. 12, p. 317 (January 1966).

10. R. D. Luce and H. Raiffa, Games and Decisions, John Wiley & Sons, Inc. (1957).

11. L. S. Lasdon, Optimization Theory for Large Systems, MacMillan Co. (1970).

12. J. C. Beckman, J. R. Hyde, and S. I. Rasool, "Exploring Jupiter and its Satellites with an Orbiter," Astronautics and Aeronautics, Vol. 12, p. 24 (September 1974).

13. W. F. Brady (Jet Propulsion Laboratory), private communication, 10 January 1975.

14. A. Hoffman (Jet Propulsion Laboratory), private communication, 31 January 1975.

15. A. L. Friedlander, "Jupiter Orbiter Lifetime--The Hazard of Galilean Satellite Collision," Report No. SAI 1-12-194-T2, NASA Contract NASA-2613, Science Applications, Inc., Rolling Meadows, Illinois (14 February 1975).

16. J. M. Harrison and D. W. North, "Probabilistic Models of Planetary Contamination," Final Report, SRI Project MSU-2274, NASA Contract NASW-2451, SRI, Menlo Park, California (April 1973).