

STRATEGIC DECISIONS GROUP

READINGS ON

The Principles
and
Applications of

DECISION ANALYSIS

EDITED BY

RONALD A. HOWARD &

JAMES E. MATHESON

VOLUME II: PROFESSIONAL COLLECTION

STRATEGIC DECISIONS GROUP

READINGS ON

The Principles
and
Applications of

DECISION
ANALYSIS

EDITED BY

RONALD A. HOWARD &

JAMES E. MATHESON

VOLUME II: PROFESSIONAL COLLECTION

EDITORS

Ronald A. Howard
James E. Matheson

PRODUCTION STAFF

Rebecca Sullivan, Coordinator
Paul Steinberg, Copy Editor
Leslie Fleming, Graphics

© 1984 by Strategic Decisions Group. Printed in the United States of America. All rights reserved. No part of this material may be reproduced or used without prior written permission from the authors and/or holders of copyright.

CONTENTS

VOLUME I: GENERAL COLLECTION

<i>Foreword</i>	vii
<i>What is Decision Analysis?</i>	viii
INTRODUCTION AND OVERVIEW	1
<i>Preface</i>	3
1. The Evolution of Decision Analysis	5
<i>R. A. Howard</i>	
2. An Introduction to Decision Analysis	17
<i>J. E. Matheson and R. A. Howard</i>	
3. Decision Analysis in Systems Engineering	57
<i>R. A. Howard</i>	
4. Decision Analysis: Applied Decision Theory	95
<i>R. A. Howard</i>	
5. A Tutorial Introduction to Decision Theory	115
<i>D. W. North</i>	
6. A Tutorial in Decision Analysis	129
<i>C. S. Staël von Holstein</i>	
7. The Science of Decision-Making	159
<i>R. A. Howard</i>	
8. An Assessment of Decision Analysis	177
<i>R. A. Howard</i>	
APPLICATIONS	203
<i>Preface</i>	205
<u>Investment and Strategic Planning</u>	
9. Decision Analysis Practice: Examples and Insights	209
<i>J. E. Matheson</i>	
10. Decision Analysis of a Facilities Investment and Expansion Problem	227
<i>C. S. Spetzler and R. M. Zamora</i>	
11. Strategic Planning in an Age of Uncertainty	259
<i>M. M. Menke</i>	
12. A Decision Analysis of a Petrochemical Expansion Study	269
<i>T. J. Braunstein</i>	
13. The Dangerous Quest for Certainty in Market Forecasting	285
<i>R. Abt, M. Borja, M. M. Menke, and J. P. Pezier</i>	
14. An Inside View: Analyzing Investment Strategy	299
<i>R. F. Egger and M. M. Menke</i>	
15. Managing the Corporate Business Portfolio	309
<i>J. E. Matheson</i>	

Research and Development

16. Overview of R&D Decision Analysis	327
<i>J. E. Matheson</i>	
17. Using Decision Analysis to Determine R&D's Value	331
<i>Article from <u>Inside R&D</u></i>	
18. Selecting Projects to Obtain a Balanced Research Portfolio	337
<i>D. L. Owen</i>	
19. Calling the Shots in R&D	363
<i>H. U. Balthasar, R. A. A. Boschi, and M. M. Menke</i>	
20. Quantifying and Forecasting Exploratory Research Success	375
<i>R. A. A. Boschi, H. U. Balthasar, and M. M. Menke</i>	
21. Evaluating Basic Research Strategies	385
<i>M. M. Menke, J. Gelzer, and J. P. Pezier</i>	

Social Policy

22. Social Decision Analysis	401
<i>R. A. Howard</i>	
23. The Decision to Seed Hurricanes	417
<i>R. A. Howard, J. E. Matheson, and D. W. North</i>	
24. Decision Analysis of the Synthetic Fuels Commercialization Program	435
<i>S. N. Tani</i>	
25. Decision Analysis of Space Projects: Voyager Mars	445
<i>J. E. Matheson and W. J. Roths</i>	

VOLUME II: PROFESSIONAL COLLECTION

<i>Foreword</i>	vii
<i>What is Decision Analysis?</i>	viii
HEALTH AND SAFETY	477
<i>Preface</i>	
479	
26. On Making Life and Death Decisions	481
<i>R. A. Howard</i>	
27. The Value of Life and Nuclear Design	507
<i>D. L. Owen, J. E. Matheson, and R. A. Howard</i>	
28. The Design of Hazardous Products	521
<i>D. L. Owen</i>	
29. On Being Environmentally Decisive	529
<i>P. A. Owen and S. B. Engle</i>	
30. On Fates Comparable to Death	545
<i>R. A. Howard</i>	

PROFESSIONAL PRACTICE	575
<i>Preface</i>	577
31. The Foundations of Decision Analysis	579
<i>R. A. Howard</i>	
32. The Difficulty of Assessing Uncertainty	591
<i>E. C. Capen</i>	
33. Probability Encoding in Decision Analysis	601
<i>C. S. Spetzler and C. S. Staël von Holstein</i>	
34. Risk Preference	627
<i>R. A. Howard</i>	
35. The Development of a Corporate Risk Policy for Capital Investment Decisions . . .	665
<i>C. S. Spetzler</i>	
36. The Used Car Buyer	689
<i>R. A. Howard</i>	
37. Influence Diagrams	719
<i>R. A. Howard and J. E. Matheson</i>	
38. The Use of Influence Diagrams in Structuring Complex Decision Problems	763
<i>D. L. Owen</i>	
TECHNICAL CONTRIBUTIONS	773
<i>Preface</i>	775
39. Information Value Theory	777
<i>R. A. Howard</i>	
40. Value of Information Lotteries	785
<i>R. A. Howard</i>	
41. The Economic Value of Analysis and Computation	795
<i>J. E. Matheson</i>	
42. Competitive Bidding in High-Risk Situations	805
<i>E. C. Capen, R. V. Clapp, and W. M. Campbell</i>	
43. Decision Analysis: Perspectives on Inference, Decision, and Experimentation . . .	821
<i>R. A. Howard</i>	
44. Bayesian Decision Models for Systems Engineering	835
<i>R. A. Howard</i>	
45. Proximal Decision Analysis	843
<i>R. A. Howard</i>	
46. Risk-Sensitive Markov Decision Processes	881
<i>R. A. Howard</i>	
CONTRIBUTIONS FROM PSYCHOLOGY	897
<i>Preface</i>	899
47. Judgment Under Uncertainty: Heuristics and Biases	901
<i>A. Tversky and D. Kahneman</i>	
48. Prospect Theory: An Analysis of Decision Under Risk	911
<i>D. Kahneman and A. Tversky</i>	
49. The Framing of Decisions and the Psychology of Choice	943
<i>A. Tversky and D. Kahneman</i>	
<i>List of Theses in Decision Analysis at Stanford University</i>	951

FOREWORD

Since the term "Decision Analysis" was coined in 1963 (see paper #4), both its theory and practice have developed profusely. Stanford University has been a center for the intellectual development of decision analysis and the catalyst for its extensive application. Consultants associated with Stanford, many of them graduates of the Engineering-Economic Systems Department, have accumulated hundreds of man-years of experience.

This collection is intended to portray the "Stanford School of Decision Analysis," as viewed by the editors. Because the Stanford decision analysis community has the broadest base of practical experience, we believe these papers represent the most successful methods of dealing with decision problems. We have not attempted to represent alternative approaches or to enter into any debate of their relative merits. We have, however, included a few papers from other fields, notably psychology, that have had, and are having, a significant impact on the practice of decision analysis.

In these two volumes, we have collected papers on both the theory and application of decision analysis. Although most of these readings have been published elsewhere, we have added a few unpublished papers to represent recent developments.*

The first volume is designed to be accessible to a general readership and contains introductory papers and descriptions of actual applications. Applications to corporate strategic decisions are necessarily disguised and underrepresented because of their proprietary nature.

The second volume is designed for the professional student of decision analysis. In addition to containing professional and technical papers, it contains some papers discussing recent developments in methodology for approaching health and safety problems. While papers in this volume use technical terminology, many of their ideas will be understandable to anyone.

* Where possible, we have indicated authors' current affiliations on the title page of each paper. Affiliation references appearing within the text are taken from the original publication and, therefore, may vary from those on the title pages.

WHAT IS DECISION ANALYSIS?



Reprinted by permission. © 1982 NEA, Inc.

When this nationally syndicated cartoon appeared in 1982, decision analysis had clearly become a common term. In common usage, however, the term has lost precision. By decision analysis, we mean a discipline comprising the philosophy, theory, methodology, and professional practice necessary to formalize the analysis of important decisions. Decision analysis includes procedures and methodology for assessing the real nature of a situation in which a decision might be made, for capturing the essence of that situation in a formal but transparent manner, for formally "solving" the decision problem, and for providing insight and motivation to the decision-makers and implementers.

Confusing the tools of decision analysis with decision analysis itself has contributed to the loss of precision. Because uncertainty is at the heart of most perplexing decision problems, decision analysts frequently use specialized tools, such as decision tree techniques, to evaluate uncertain situations. Unfortunately, many people, some of them educators, have confused decision analysis with decision trees. This is like confusing surgery with the scalpel. Although decision tree techniques are extremely useful in solving problems where uncertainty is critical, in a real decision analysis, most of the effort and creativity is focused on finding and formulating the correct problem and on interpreting the results rather than on performing computations.

HEALTH AND SAFETY

Preface

Important decisions are being made every day about our health and safety. These papers show the advantage of making these decisions consistently and present some recent theory that provides a sound basis for doing so.

"On Making Life and Death Decisions" develops a conceptual framework by considering how a person should value his own life. This paper shows that although life may be infinitely valuable in a moral sense, a person can rationally take on additional risks to his life or pay to remove them. The paper reveals that a monetary "value of life" is appropriate for an individual making choices involving a small probability of death. A numerical example is developed for a typical individual.

"The Value of Life and Nuclear Design" addresses the question of whether different "values of life" should be used in different safety decisions. It shows that using the same monetary value in all aspects of design produces the highest level of safety.

"The Design of Hazardous Products" addresses the design problem in a general setting, revealing that the designer of a hazardous product needs to know the small-risk value of life that has been assigned by the individual at risk. It also shows there is no rationale for situations where the individual is exposed involuntarily and does not bear the product cost.

"On Being Environmentally Decisive" demonstrates the insight gained from applying decision analysis to environmental issues within the corporate setting. Based on a real case, it describes a hypothetical example involving a company's decision to make a capital investment to reduce the exposure of workers to asbestos fibers. The paper shows how environmental issues can be incorporated in an economically-oriented investment decision analysis.

"On Fates Comparable to Death" extends the ideas for treating life risk to the risk of handicap or serious injury. The paper shows how small-risk values can be developed for each consequence for use in decision analysis.

ON MAKING LIFE AND DEATH DECISIONS

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted from Societal Risk Assessment: by R. C. Schwing and
W. A. Albers, Jr., General Motors Research Laboratories, 1980.
pp. 89-113.

Copyright © 1980. Plenum Publishing Corporation

ON MAKING LIFE AND DEATH DECISIONS

Ronald A. Howard

Stanford University, Stanford, California

ABSTRACT

Recent research has provided us with methods by which an individual can make decisions that involve risk to his life in a way that is consistent with his total preferences and with his current risk environment. These methods may ethically be used only by the individual himself or by an agent designated by the individual. In the absence of such delegation, anyone who imposes a risk on another is guilty of assault if the risk is large enough. Just as society has found ways to distinguish a "pat on the back" from physical battery, so must it now determine what risk may be placed upon another without his consent.

The research on hazardous decision making creates a framework for this exploration. The basic concept of this approach is that no one may impose on another a risk-of-death loss greater than a specified criterion value established by the experience of society. If anyone attempted to do so, he could be forbidden by injunction. The only way that an injunction could be avoided would be by showing evidence of insurance that would cover the damages to be paid by the imposer of the risk if the unfortunate outcome should occur. The methodological framework is used both to estimate the risk-of-death loss and the amount to be paid if death occurs, an amount that is likely to be much larger than present "economic" values of life. Evidence would be required both on the preferences of the individual-at-risk as revealed and corroborated by his behavior and on the magnitude of the risk as assessed by experts.

Such a system is likely to require revisions in the present legal codes. It is to be expected that when a logically and ethically based risk system is functioning, there will be an increased interest in purchasing the consent of people to imposed risk. Problems of securing the consent of contiguous property owners, for example, could be handled by interlocking options. People will also be more likely to be informed of the risk implied by using products or services. Thus risk would become an explicit part of purchasing decisions. The joining of logic and ethics in these new procedures offers hope for a more effective and humane treatment of risk issues in society.

References p. 106.

NOTATION

p :	probability of death
x :	required payment to undertake specified death risk
W :	present level of wealth
c :	constant annual consumption
l :	remaining length of life
\bar{l} :	expected remaining length of life
w :	worth numeraire
η :	consumption-lifetime trade-off
u :	risk preference function on worth
γ :	risk aversion coefficient
ρ :	risk tolerance; $1/\gamma$
ζ :	annuity factor; amount of annuity that \$1 will buy
i :	interest rate
p_{\max} :	maximum acceptable probability of death
$v(p)$:	life value in expected value sense when facing death with probability p
v_s :	small-risk life value
v_e :	economic life value; c/ζ
p_n :	probability of death in year n of life
q_n :	probability of death in year n of life given that individual was alive at beginning of year n .

INTRODUCTION

What risk may one impose on another? This question has achieved increasing importance as the sources of harm in our environment have increased. The spectrum of risk that one person imposes on another ranges from the relatively minor risks posed simply by existence up to the very serious risks represented by assault or attempted murder. Some of these risks society has chosen to ignore, while others have been treated as very serious matters requiring extensive social action. We shall examine both the ethical and practical questions of risk in society, propose measures for risk, suggest procedures for evaluating risk, and indicate how these procedures could be used in practice.

EFFICACY AND ETHICS

Social arrangements for any purpose may be judged in terms of both efficacy and ethics. Efficacy refers to what works in pursuing specific human goals; ethics refers to what actions are morally desirable in achieving those goals. For example, killing babies with genetic defects might be a very efficacious way of achieving the human goal of physical perfection, but it would be ethically unacceptable to most people. When we wish to judge any action or arrangement, we can think of examining it against standards of physical knowledge, ethics, and efficacy. For example, if someone threatened to bring the wrath of God against another, that threat would

LIFE AND DEATH DECISIONS

not be actionable in a court of law today because it is the present belief of a majority of our society that no one has such power. However, in the 12th century in Europe, such a threat may have been taken very seriously: the one who threatened might be condemned as a witch. Actions that seem physically feasible can then be subjected to the further tests of ethical acceptability and practical efficacy. Since there is much more discussion of efficacy than of ethics, our primary concern here will be the ethical one.

The ethical basis we shall use in our discussion is that every individual has a right to his own person. Or to put it in negative form, no one may initiate force against another without his consent. Of course, this allows for the use of force against the initiator of force in the sense of self-defense. Imposing a large risk of the use of force upon another is enjoined by the same principle. If the imposition is intended to be coercive, then the imposition is a threat. The robber who says "your money or your life" is thus violating the ethical principle even though you may avoid the use of force by surrendering your money.

Even when there is no intention to harm, the principle prohibits the imposition of a large risk on another. Thus, someone who is firing a gun in random directions may be restrained even though he has no intention of hurting anyone simply because he poses too great a threat to others.

While there might seem to be a wide variety of ethical principles from which to choose, the choice is not so large as one might think. In fact, the only other system with a claim to consistency (although a faulty claim, in my opinion) is that the king, czar, party, government, or church can do to any person whatever it likes. In such a system, of course, we don't have to worry about risk management; we simply ask the king-equivalent what to do.

Therefore, the ethic that shall guide us in this paper is that no one may impose a large risk on another without his consent. The remaining question, then, is how to measure risk and how to determine how large a risk may be imposed involuntarily.

It is important to distinguish this discussion of ethics from the usual discussion in terms of political and economic systems. The political system in many countries does incorporate ethical elements, such as the U.S. Constitution's Bill of Rights. However, it may also allow actions that many individuals consider unethical. Thus the political system technically contains both ethical judgments and other features based on the power possessed by various groups. As long as there exist two systems, political and economic, in the same society, then there exists the possibility of arbitrage, of people using political power to achieve what they cannot achieve economically or using economic power to achieve what they cannot achieve politically. For example, rent control is an action to transfer property ownership at least partially from owners to tenants. Environmentalists' objections against development can be attempted by some to raise their standards of living by political means at the expense of the standard of living of those not so well economically situated.

The main point is that unless political and economic systems have a common ethical basis, ethical conflict is bound to arise. The approach we take here is to follow ethical principles that preclude political and economic contradictions.

References p. 106.

MEASURING RISK: AN INDIVIDUAL DECISION MODEL

Recent research has shown one way life and death decisions can be made consistent with the non-coercive ethical principle [1], [2]. Naturally, then, this is a way for people to make their own risky decisions, *not* a way for other people to impose risky situations upon them. However, by seeing how an individual would view such an imposition by his own lights, we obtain a starting point for constructing a legal position regarding the imposition of risk.

The Black Pill — As a useful thought experiment, we imagine an individual faced with what we call the black pill question. He is offered the chance to take a pill that will kill him instantly and painlessly with a probability he assigns as p . If he takes the pill, he will receive x dollars. Should he accept? For example, should he accept a $p = 1/10,000$ incremental chance of death for a payment of $x = \$1000$? The choice is diagrammed in Fig. 1.

If the individual rejects the offer, he will continue his life with wealth W and face whatever future life lottery he presently faces. His future life lottery is the uncertain, dynamic set of prospects he foresees beginning with today. If, on the other hand, he accepts the proposition, his wealth will increase to $W + x$. If he lives after

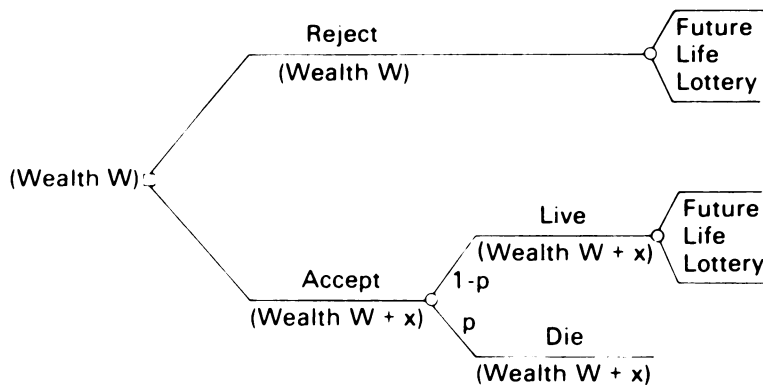


Fig. 1. The black pill decision tree.

taking the pill, he will begin his future life lottery with wealth $W + x$, presumably a more desirable situation. If he dies, he will leave $W + x$ in his estate, and, of course, have no opportunity to enjoy it. Clearly the value of this benefit might be different for different people, and could be included. But let us say, for the moment, that it has no value to him. Naturally, there would also be tax effects, but these too we shall ignore.

We have analyzed this question in quite general form [2], but here we present the simplest model we have used to answer it. We assume that everyone has a fundamental preference on both level of consumption and length of life. We begin by asking the individual how much consumption (measured in today's dollars) he

LIFE AND DEATH DECISIONS

expects to have at each year in the future. We then ask what constant level of consumption beyond bare survival over his lifetime would make him indifferent between this level and his present prospects. We call this the constant annual consumption for that individual. Now we give him choices between the different futures described by different constant annual consumptions c and different lifetimes ℓ , and find to what combinations he is indifferent. For the simple example, we shall assume that the indifference curves have the form

$$w(c, \ell) = c \left(\frac{\ell}{\bar{\ell}} \right)^\eta \quad \eta > 0, \quad (1)$$

where $w(c, \ell)$ is the worth numeraire associated with each indifference curve. The numeraire equals c when ℓ equals $\bar{\ell}$, the expected lifetime remaining.

Now we measure the risk preference on worth of the individual. For the example, we shall use the exponential form

$$u(w) = -e^{-\gamma w} = -e^{-w/\rho}, \quad (2)$$

where γ is the risk tolerance. With this structure and the assessment of the individual's joint probability distribution of c and ℓ , we can compute the utility of the individual for the case when he does not accept the black pill.

When he does take the pill, the probability p of dying immediately will transform his probability distribution on remaining life. The payment he receives, x , will increase his wealth. We assume that the individual will use the amount x to purchase an annuity over his remaining life at the prevailing interest rate i . In the calculation of annuity cost, we assume further that the seller of the annuity assigns the same probabilities on remaining life as those assigned by the individual. If we let ζ be the amount of annuity that one dollar will buy, then

$$\zeta = \frac{i}{1+i} \frac{1}{1 - \left\langle \left(\frac{1}{1+i} \right)^\ell \right\rangle}, \quad (3)$$

where $\langle \rangle$ denotes expectation.

When we set the utility of taking the pill equal to the utility of not taking it, we determine that at the point of indifference p and x must satisfy the equation.

$$p = \frac{\left\langle e^{-\gamma c \left(\frac{\ell}{\bar{\ell}} \right)^\eta} \right\rangle - \left\langle e^{-\gamma \left(c + \zeta x \right) \left(\frac{\ell}{\bar{\ell}} \right)^\eta} \right\rangle}{\left\langle 1 - e^{-\gamma \left(c + \zeta x \right) \left(\frac{\ell}{\bar{\ell}} \right)^\eta} \right\rangle} \quad (4)$$

By inverting this relationship computationally, we find for a given value of p what value of x will make the individual indifferent between taking the pill and not taking it.

References p. 106.

If we let x grow without limit in this equation, we find that p approaches a value p_{\max} given by

$$p_{\max} = \left\langle e^{-\gamma c \left(\frac{\rho}{\rho}\right) \eta} \right\rangle \quad (5)$$

Thus no amount of money could induce the individual to accept a death probability as large as p_{\max} .

To derive a measure of "life value", suppose that a risk-neutral observer examines the relationship between x and p . He could interpret $x = x(p)$ as the expected loss that the individual would incur from the risk if the individual valued his life at a number $v(p)$,

$$x(p) = pv(p) \quad (6)$$

and he could then determine this number from

$$v(p) = \frac{x(p)}{p} \quad (7)$$

Thus $v(p)$ is the value that the person is placing on his life in an expected value sense when he confronts a risk of magnitude p .

Of special interest in a safety context is the magnitude of this life value when the death risk is small. From a limiting analysis of the equation relating p and x , we find that as p approaches zero (and, of course, x also approaches zero), the ratio $v(p)$ approaches a value v_s given by

$$v_s = \frac{1 - \left\langle e^{-\gamma c \left(\frac{\rho}{\rho}\right) \eta} \right\rangle}{\zeta \gamma \left\langle \left(\frac{\rho}{\rho}\right) \eta e^{-\gamma c \left(\frac{\rho}{\rho}\right) \eta} \right\rangle} \quad (8)$$

We call this value the small-risk life value. It is the one number that an individual would need to keep in mind to make his safety decisions.

We shall be interested in comparing this small-risk value of life with an economic value of life comparable to that produced by other analyses. We shall define the economic life value, v_e , as the amount of money required to purchase an annuity paying the constant annual consumption c . Thus v_e is given by

$$v_e = \frac{c}{\zeta} \quad (9)$$

Illustrative Results — To illustrate the calculations implied by the model, let us consider a base case individual who is a 25-year-old male with a constant annual consumption of \$20,000 per year and a lifetime probability distribution given by a standard mortality table, Table 1. He chooses $\eta = 2$, which means that if he is sure to live his expected life (46.2 years), then a 1% decrease in his life would require a 2% increase in consumption for him to remain indifferent. From further questioning, we find that his risk tolerance is $\rho = \$6000$, which means roughly that he is indifferent between his present situation and equal chances of constant annual consumption of \$17,000 or \$26,000 for the remainder of his life. We also find that he faces a prevailing interest rate of 5% per year.

LIFE AND DEATH DECISIONS

The results of the calculation appear in the upper part of Fig. 2 where we show the amount x that he would have to be paid corresponding to each probability of death p . We observe that this amount increases proportionally to p until about

TABLE I
Life Table for White Males, U.S.
of 100,000 Born Alive. Number Dying During Age Interval

Age	Number Dying During Age Interval	Age	Number Dying During Age Interval	Age	Number Dying During Age Interval
0	2592	37	229	73	2775
1	149	38	251	74	2815
2	99	39	278	75	2841
3	78	40	306	76	2853
4	67	41	339	77	2855
5	60	42	376	78	2844
6	55	43	415	79	2821
7	52	44	458	80	2789
8	47	45	505	81	2738
9	43	46	556	82	2639
10	40	47	613	83	2482
11	40	48	681	84	2280
12	46	49	754	85	2096
13	56	50	835	86	1898
14	73	51	916	87	1693
15	90	52	995	88	1490
16	107	53	1071	89	1288
17	121	54	1144	90	1086
18	134	55	1216	91	888
19	143	56	1295	92	709
20	153	57	1383	93	548
21	162	58	1486	94	413
22	167	59	1598	95	300
23	163	60	1714	96	216
24	157	61	1827	97	152
25	149	62	1935	98	103
26	141	63	2039	99	70
27	137	64	2136	100	45
28	137	65	2231	101	29
29	141	66	2323	102	17
30	147	67	2409	103	11
31	154	68	2487	104	6
32	161	69	2559	105	3
33	170	70	2621	106	2
34	180	71	2678	107	1
35	194	72	2729	108	1
36	210				

References p. 106.

$p = 10^{-2}$, when it increases more rapidly and finally becomes infinite at $p_{\max} = 0.103$. No amount of money could induce this individual to play Russian roulette ($p = 1/6$).

The lower portion of Fig. 2 shows how the life value $v = v(p)$ depends on p . We observe that for small values of x , v is approximately equal to $v_s = \$2.43$ million, the small-risk life value of the individual. This means that for small probabilities of death (here less than 10^{-2}) the individual is acting as if his life were worth \$2.43 million in an expected value sense. Thus, if the individual faced the black pill problem with $p = 1/10,000$, the required compensation would be $v_s p = \$243$. He would accept any payment x greater than \$243 as an inducement to take the pill.

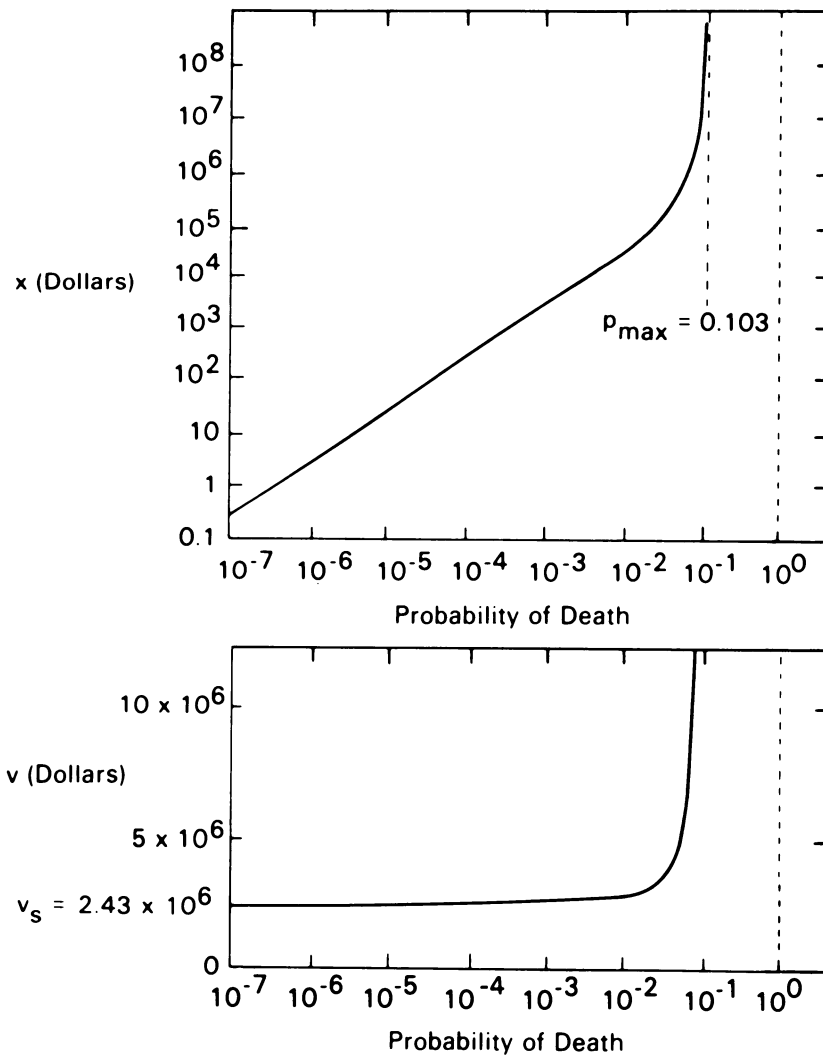


Fig. 2. Black pill results.

LIFE AND DEATH DECISIONS

The economic life value v_e for this individual is \$363,000. Such a number has sometimes been used for decision purposes. We observe that the small-risk value is about 6.7 times the economic value. If this model and the numbers used in it are representative, the economic values that have been used in the past considerably underestimate the individual's own value. This discrepancy has implications for both the efficacy and the ethics of risk decision-making in our society.

The White Pill — Our analysis up to this point has emphasized the question of what we must pay an individual to undertake an additional risk. However, more often we face the problem of spending resources to avoid risk or in other words increase safety. The same theoretical model serves to illuminate this problem with only a few small twists.

Suppose that an individual faces a hazard that will kill him with probability p ; for example, an operation. If he survives, he will live his normal life with whatever wealth he possesses. However, now someone arrives with a white pill that if taken will surely eliminate the death risk from this hazard. How much, x , would the individual be willing to pay for the white pill? Fig. 3 shows the relevant decision tree.

The unusual feature of the white pill question is that, of course, the amount x that he is willing to pay cannot exceed his wealth, no matter what death risk he faces. We assume that the individual can sell an annuity based on his lifetime distribution to pay the amount x for the purchase of the white pill. Since the most he can give up is his consumption beyond survival c , this means that in the white pill case the x versus p curve terminates on the economic life value of the individual

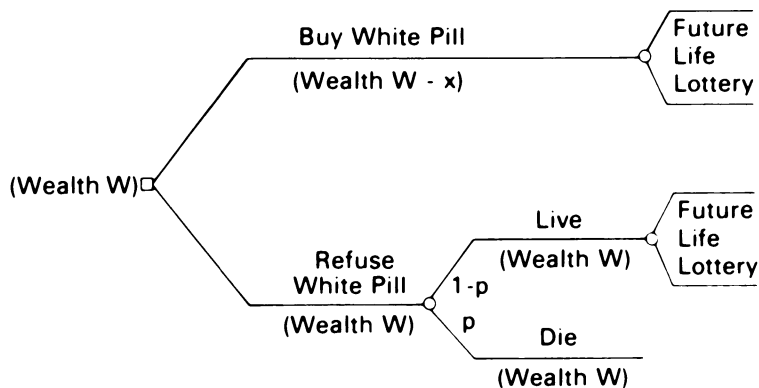


Fig. 3. The white pill decision tree.

when $p = 1$. The equation relating p and x ,

$$p = \frac{\left\langle e^{-\gamma(c - \zeta x)} \left(\frac{l}{l}\right)^\eta \right\rangle - \left\langle e^{-\gamma c} \left(\frac{l}{l}\right)^\eta \right\rangle}{1 - \left\langle e^{-\gamma c} \left(\frac{l}{l}\right)^\eta \right\rangle} \quad (10)$$

References p. 106.

confirms this observation, since at $x = \frac{c}{\zeta} = v_e$, $p = 1$.

Fig. 4 show the results for the base case individual. When $p = 1$, $x = \$363,000$, the economic value of his life. However, as p decreases, the x versus p curve becomes coincident with that of Fig. 2, and in particular implies the same small-risk value of \$2.43 million derived for the black pill case.

Table 2 shows how the small-risk value depends on the model variables. The first row shows the effect of changing annual consumption level from \$10,000 to \$30,000 while fixing the risk tolerance at 30% of consumption. We observe that the small-risk value is then proportional to consumption level. The second row shows that the effect of varying the interest rate i from 10% to 2.5% is to change the small-risk value from \$1.421 million to \$3.622 million, because the individual needs a higher cash payment to obtain the same increase in consumption. The third row shows the relative insensitivity to the consumption-lifetime trade-off ratio η , whereas the last row illustrates how the small-risk life value falls with age.




Fig. 4. White pill results.

TABLE 2
Sensitivity Analysis

	Variable			Small-Risk Value, v_s (\$ million)		
c:	10,000	20,000	30,000			
ρ :	3,000	6,000	9,000	1.215	2.430	3.645
i:	0.10	0.05	0.025	1.421	2.430	3.622
n:	3	2	1	2.418	2.430	2.541
Age:	35	25	15	2.157	2.430	2.671

Buying and Selling Hazards — Now that we have both the black pill and white pill results before us, we are in a position to make a few general observations. First, we see that the disparate results of the black and white pill cases for $p = 1$ show that we have answered a continual objection to analyses that place a finite value on life without regard to the distinction between accepting an additional risk and removing an existing risk. Since few people, if any, will sell their lives for any finite sum, all such analyses are doomed to failure. *However, the present model shows that it is perfectly consistent to refuse any finite offer for your life and yet be limited in what you can spend to save it.*

Of greater practical importance, however, is the result that for the wide range of hazardous decisions where we are buying and selling small hazards in our lives, the small-risk life value offers a simple and practical procedure to assure consistency.

To simplify the use of the small-risk life value and to emphasize the necessity that it be used only when the risk to life is small, we find it useful to define a unit for small risks to life. We shall use the term “micromort” to mean a one in one million chance of death, with symbol μmt . Then the small-risk life value can be conveniently expressed in dollars per micromort, or \$2.43 for the base case individual. With this terminology, it is easy to explain why an individual can set a value for a micromort that is valid up to, say, 1000 micromorts, but also why that price is inappropriate for larger risks.

The Value of Reducing Risk - We can use the base case individual's value of $\$2.43/\mu\text{mt}$ to see what he would be willing to pay annually to remove various hazards in his life. The first column of Table 3 shows U.S. accident statistics for 1966. The second column shows the number of micromorts/year each risk poses to the base-case individual if, as we now assume, he uses these statistics as his probability assignment to death from each risk. The final column shows what the base case individual would be willing to pay each year to eliminate each hazard, an amount obtained by multiplying the number of micromorts by the individual's value of a micromort. Note that he would be willing to pay \$900 just to eliminate the dangers of motor vehicles and falls. All other sources of accidents contribute collectively to an expected loss of less than \$500. This calculation is an important

References p. 106.

TABLE 3
U.S. Accident Death Statistics for 1966

Type of Accident	Total Annual Deaths*	Probability of Death in Micromorts/Year (1 Micromort [μmt] = 10^{-6})	Payment of Base Case Individual to Avoid Hazard @ \$2.43/ μmt
Motor vehicle	53,041	270	\$ 656.00
Falls	20,066	100	243.00
Fire and explosion	8,084	40	97.20
Drowning	5,687	28	68.00
Firearms	2,558	13	31.60
Poisoning (solids and liquids)	2,283	11	26.70
Machinery	2,070	10	24.30
Poisoning (gases and vapors)	1,648	8.2	19.90
Water transport	1,630	8.1	19.70
Aircraft	1,510	7.5	18.30
Inhalation and ingestion of food	1,464	7.3	17.70
Blow from falling or projected object or missile	1,459	7.3	17.70
Mechanical suffocation	1,263	6.3	15.30
Foreign body entering orifice other than mouth	1,131	5.7	13.90
Accident in therapeutic procedures	1,087	5.5	13.40
Railway accident (except motor vehicles)	1,027	5.1	12.40
Electric current	1,026	5.1	12.40
Other and unspecified	6,163	31.0	76.50
Total	113,563	580.	\$1,384.00

*U.S. Accident Statistics for 1966

starting point for determining whether feasible safety expenditures to modify these hazards would be worthwhile. It is clear that spending \$1000 to be free of motor vehicle accidents would not be a wise choice for the base case individual. There is a limit to the value of safety.

Continuing Risks - Hazard Modification — Many of the risks to life occur not at a single instant, as does the black pill, but rather over several years or even a lifetime. The risks of living with automobiles, of smoking, or of living near a power plant are of this type. We can use the previous formulation to analyze this situation after we deal with the concept of hazard as follows.

The lifetime mass function is defined by p_n , $n = 1, 2, 3, \dots$ where p_n is the

LIFE AND DEATH DECISIONS

probability that an individual will die in the n^{th} year of his life. Let q_n be the probability that the individual will die in the n^{th} year of his life given that he was alive at the beginning of that year. Then q_n for $n = 1, 2, 3, \dots$ is the hazard distribution or force of mortality. The lifetime mass function and the hazard distribution are related by the equations:

$$\begin{aligned} p_1 &= q_1 \\ p_n &= q_n \left[1 - \sum_{j=1}^{n-1} p_j \right] \quad n = 2, 3, 4, \dots \end{aligned} \quad (11)$$

and either may be constructed from the other.

We can now ask what present payment x would be required to induce an individual to accept a given modification of his hazard distribution, with the previous assumption that the payment will be converted into an annuity.

Consider first increasing the hazard in every remaining year of a person's life by adding 250 micromorts, a risk about equal to that posed by automobiles in American society. To induce the base case individual to accept such a hazard modification, which would decrease his life expectancy by 0.3 years, we would have to give him a lump sum of \$13,000, or an annuity paying \$700 per year.

If we doubled his hazard in every year, a risk considered by some the equivalent of heavy smoking, life expectancy would fall by 7.8 years and he would require a present payment of \$212,000 or an annuity of \$12,400.

If all benefits and costs associated with a general pattern of hazard modification are reduced to dollar terms, then the model can be used to determine the additional payment that the individual would demand or offer to be just indifferent to the modification.

Summary — We can use this model to evaluate how much an individual would have to be paid in money or its benefit equivalent in order to accept any given level of risk. Of special interest are those situations where the additional risk is small, for in this case the payment that the individual would require is equal to the probability of his death multiplied by a small-risk life value in dollars. This small-risk life value is likely to be constant over the range of risks involved in safety situations, for example, 1000 or less micromorts per year. The small-risk life value is typically several times larger than the economic value of life and is of the order of a few million dollars.

RISK ISSUES IN SOCIETY

Now that we have discussed both the ethical basis of imposing risk and a procedure by which an individual can make or delegate decisions that affect his chances of dying, we can proceed to an examination of the implications of those observations for various situations involving risk in our society. These situations include the treatment of risk in the marketplace and on the job, the imposition of excessive risk, and the creation of risky projects. We shall also discuss how these

References p. 106.

observations bear on the question of corporate liability.

Risk in the Marketplace and on the Job — Risks involved with the purchase and use of products or with holding a job are likely to represent 1000 or less micromorts per year, and hence, for practically all individuals, they will fall in the region of life value characterized by the small-risk value of life. This means that the individual who has assessed his chance of death can quickly calculate the expected death loss from the situation and balance it against other benefits and costs associated with the situation to make his decision.

However, we should also note that the designer of the product or the safety engineer of the job has already been making decisions that balance death against other considerations. In fact, if he is to be consistent, he should be using some small-risk life value in his design [3]. A logical next step would be to reveal this value to the purchaser of the product or the job applicant in a statement like, "We used a \$3 million small-risk value of life in designing this car (or coal mine)." Naturally, the individual will hope that the value used is at least his small-risk life value. Otherwise, he would be rightly concerned that the situation will not be safe enough for him. The small-risk value of life used in design could then become one of the features of the product or job that is advertised to the public. Companies that used too low a value would experience competitive pressure, based on safety concerns, to raise it; whereas, those who used too high a value would find their products overpriced relative to competition. A similar result would apply to jobs. Thus companies would be encouraged by the marketplace to balance safety and economics. As the standard of living increased, so would the level of safety.

A further step in this development would be for the companies to buy insurance that would pay the small-risk life value used in design to anyone killed as a result of the design. Since this amount would be listed on the product or in the job description, the product liability or safety liability of the company would be specified in advance. The estate of the person who bought a cheap hammer designed with a small-risk life value of \$10,000 would be able to collect only \$10,000 if the head came off and killed him. Of course, someone who wanted the hammer for use as a paperweight might still buy it.

The idea of describing products (or jobs) by the small-risk life value used in their design is only useful if the number can be believed. It would be fraud to post a number higher than was actually used, or, of course, to say that insurance paying this value in the event of death through product design is in force when it is not. To be fair, the insurance would pay off only if the product failed while being reasonably employed in its intended use. (If the purchaser of the cheap hammer commits suicide by hitting himself over the head, his estate has no claim.)

Recently one of America's largest automobile companies lost a multi-million dollar suit involving product design. Evidence presented at the trial showed that the value of life used in the design was an order of magnitude below those we have discussed. How many of the purchasers of the car would have bought it if they had known the design basis? Placing this number in view may be the most important single step that can be taken to insure the proper balance of safety with other considerations.

LIFE AND DEATH DECISIONS

Risk Imposition — The question of risk imposition can be addressed in terms of the legal procedures used when one person claims that another is imposing an unacceptable risk on him. Since nearly everyone is a potential danger to everyone else at some level, no absolute standard is possible. The question is at what level of risk different legal remedies may be imposed. We propose to measure the risk to an individual using the model described above. This model provides a monetary value of the risk in terms of the probability of death imposed and the preferences of the individual. To make the model operational, we must specify the source of the probabilities and the preferences.

Since there may be differences of opinion as to the probabilities of death imposed on one person by the actions of another, one function of the legal process would be to assign this probability in as objective and impartial a manner as possible. This may mean reviewing historical evidence, examining experimental findings, and ultimately considering the statements of experts. This procedure will not be easy, but it is necessary if serious concerns are to be separated from paranoia.

The individual's preferences, of course, are his alone. But to establish them in court, the individual will have to show that he acts consistently with his stated preferences. For example, it would be difficult for a circus performer who took large risks for money as part of his profession to then claim that no amount of money could compensate him for much smaller risks. The past practices and decisions of the individual would in most cases provide good evidence of his preferences.

In the majority of situations where this procedure will be implemented, the risk faced by the individual will be small enough that his preferences can be summarized by his small-risk value of life. In these cases, the procedure will reduce to the court's determining the probability of death and the small-risk life value. The product would then measure the extent of the risk imposed on the individual, a number we shall call his risk evaluation.

The risk evaluation would in turn indicate the kind of relief to which the individual is entitled. If the risk evaluation were very small, say less than 10¢ or perhaps 10^{-5} times the average annual income, then no relief would be provided under the principle that the law does not concern itself with trifles.

On the other hand, if the risk evaluation were greater than a serious level, say, \$10 or 10^{-3} times the average income, then the individual might be entitled to injunctive relief. That would mean that the imposer of the risk would be prohibited from imposing it. At this point, the risk imposer would either have to cease his activity, buy the right to impose the risk from the individual (for whatever he demands), or reduce the level of the risk evaluation below the serious level by making his activity considerably safer.

For risk evaluations in the intermediate region between trifling and serious, the medium range, a different remedy could be applied. This could be allowing the activity to continue only if the risk imposer buys insurance sufficient to pay damages if the activity actually kills the individual at risk. Moreover, the damages would not be the economic loss to the dead person's estate but rather his small-risk value of life, typically many times higher. This would mean that the risk imposer would always find it at least as desirable to pay the individual the risk evaluation in

References p. 106.

exchange for the right to impose the risk as he would to buy insurance. However, the insurance option does allow people to impose relatively small risks on others who may have unreasonable fears of certain kinds of activity.

Regardless of the level of risk involved, if the imposer buys the right to impose that risk on another individual from that individual, there would be no cause for litigation. This would include even cases of high risk, such as paying someone to play Russian roulette. However, in cases of high risk, the principle that every person has an inalienable right to his body would mean that the court would not allow the risk to be imposed in a situation where the seller of the right to impose risk changed his mind. The seller would, of course, remain liable for any damages he had promised to pay in the contract should he change his mind. Contracts involving the selling of rights to impose risk in the domain of safety, say, 10^{-3} chance of death or less, would not be subject to the alienability criterion, but would be considered as transfers of property.

On Creating Risky Projects — Much of the modern concern with risk arises from the building of what we might call risk-creating installations. These are installations that cause increased risks to the public, that is, to people who have not made any agreement to accept the increased risk. Such installations might be oil storage facilities, airports, or nuclear power plants. According to our preceding discussion, those creating such risks could proceed unencumbered only if the risk evaluations they created for those affected fell in the trifling range. If the risk evaluations fell into the medium range, then insurance would have to be bought that would pay the estate of anyone killed his small-risk life value. Of course, if no insurance company were willing to sell such a policy, the project could not proceed. Finally, if the risk evaluations fell in the serious level, the project could be prohibited regardless of insurance.

The entrepreneurs wishing to build risk-creating installations would be strongly encouraged by such a system to purchase in advance the risk rights from all individuals involved. But one immediately thinks of the problem of the holdout — someone who refuses to sell. When the risk evaluation is at the serious level, no one, including the government, could compel him to do so. For it is a violation of our basic principle regarding the initiation of force to use ideas like eminent domain to justify the initiation of force. How, then, can the practical entrepreneur proceed?

The basic idea that can solve this problem is the idea of risk options. The entrepreneur can buy from an individual an option to purchase his risk rights under certain conditions and at a specified price. For example, the entrepreneur might pay \$10 for the option to buy at some time in the next year the right to impose 100 micromorts per year at a price of \$100 per year. Then if the entrepreneur decides to build a 100 micromorts per year installation in the individual's area, he pays him \$100 per year. If he decides to build some place else, then the individual has received only the \$10 for the option. The entrepreneur can then buy options in several different areas knowing that he will in fact build in only one area. If the entrepreneur encounters holdouts in one area, he can move on to another.

LIFE AND DEATH DECISIONS

The option and the rights could both be negotiable to create a market in risk. Communities of people who were relatively more willing to accept risk for money would then be more likely to be those places where the risk-creating installation was built. Since risk-rights could be bought more cheaply in remote locations, such locations would also be favored for the building of the installations. Thus the risk market, like any market, would encourage the more efficient use of resources. However, we should remember that the efficiency in this case is not being achieved at the expense of ethical principle.

Liability — We must still consider the case where someone imposes a risk in the medium or serious range without having at least purchased insurance. Logically in this case if someone is killed as a result, the liability of the risk imposer should be at least the victim's small-risk life value, and more if the risk extended beyond the safety range. There would be an additional heavy penalty if the insurance was not bought after a finding that the risk was in the medium range. Deaths resulting from serious range risky activities would incur criminal penalties.

The lower portion of Fig. 2 shows how the minimum liability might depend on the prior death probability. The region where the value becomes infinite we might call the "murder" region.

However, to put this principle into operation when the risk imposer is a corporation will apparently require changes in corporate law. The reason is that today corporations have the same limited liability to third parties (like the victim of the risk) that they have to second parties, their knowing creditors. This means that if a corporation is so structured that its assets are insufficient to satisfy a claim, the victim's estate cannot reach beyond those assets to the stockholders for the settlement of the claim. While I have no objections to the limited liability to creditors because they entered into the credit arrangement with knowledge of the limited liability, I see no reason why this limit should extend to third parties. For example, if a group of individuals organized themselves into a corporation and then the actions of the corporation resulted in someone's death, the personal assets of those individuals would ordinarily not come into play, whereas if the individuals had organized as a partnership, their personal assets would be available to satisfy the judgment.

The limited liability to third parties is not a recent feature of corporate law. The corporate form is, after all, a human invention. At the time corporations were first allowed legal status, there was debate on this issue of third party liability. Unfortunately, from my point of view, limited liability to third parties was instituted as a feature of corporations. But there is no reason why this decision could not be reversed.

Suppose that corporations were treated like partnerships in matters of third party liability. This would mean that every stockholder would be liable for damages done by the corporation to third parties. The result would be increased care among corporations in controlling their effects on third parties, effects on their property as well as on their lives. A corporation that was careless in this regard would soon find that it had lost the favor of investors.

Consider a company that is in the business of constructing or operating nuclear

References p. 106.

reactors. At the moment its stockholders are protected not only by the Price-Anderson Act, but by the limitation of liability to third parties. If both of these limitations were removed, anyone investing in such a company would have to be quite sure that he was protected from a calamitous loss. This would mean, most probably, that such companies would have to buy insurance against all such losses.

If such insurance were unobtainable or prohibitively expensive, there would be good reason to question the economic viability of the industry. Furthermore, even if the insurance were available, it is likely that the insurance companies in their own self-interest would require that independent agencies certify the safety of production and operation. Thus, the ultimate effect of unlimited liability to third parties would be either prohibition of unsafe industries or considerable improvement in their safety.

CONCLUSION

Our present apparent impasse on many safety issues stems mainly from a reluctance to re-examine the ethical basis for risk management in our society. As long as these issues are stuck between the economic and the political system, we can expect little progress. Only by returning to more fundamental ethical considerations can the issues be clarified and ultimately resolved.

REFERENCES

1. R. A. Howard, "Life and Death Decision Analysis," *Second Lawrence Symposium on Systems and Decision Sciences, Berkeley, California, October 1978.*
2. R. A. Howard, "Life and Death Decision Analysis," *Research Report EES DA-79-2, Department of Engineering-Economic Systems, Stanford University, December 1979.*
3. R. A. Howard, J. E. Matheson and D. L. Owen, "The Value of Life and Nuclear Design," *Proceedings of the American Nuclear Society Topical Meeting on Probabilistic Analysis of Nuclear Reactor Safety, American Nuclear Society, May 8-10, 1978.*

LIFE AND DEATH DECISIONS

DISCUSSION

F. E. Burke (*University of Waterloo*)

I have one question of clarification. You had a slide in which you had computations where you compared various fairly modest risks. In one case an 8 year life expectancy change would require a present payment of \$212,000 or an annuity of \$12,400. That I could understand because my mental arithmetic was fast enough. But then you have a life expectancy change of 0.3 years which corresponds to a lump sum payment of \$13,000 or an annuity paying \$700/year. My mental arithmetic left me there and I wonder if you could help me out?

R. A. Howard

The payment is not proportional to the change in life expectancy because the effect is nonlinear. The changes in life expectancy and payments are computed from the same data but they are not obviously related to each other by any constants.

F. E. Burke

One is instant and the other is at the end of the expected life.

R. A. Howard

Unfortunately, there is a common belief that if I lose some expected life, the decrease always comes off the end. If I believed this, I could think there's no problem with smoking because it's going to hurt me after I'm too old to enjoy anything else. Of course, that's not true. The increased hazard is a change in the whole probability distribution of life.

E. V. Anderson (*Johnson and Higgins*)

One possible error is the use of the "fall" statistic to apply to a 20-year-old. About 90% of your fall deaths are of people aged 65 and over and you probably should take the deaths between 20 and 50 and use that as a basis.

R. A. Howard

Right, I think that's a very good observation. Everything in here should be interpreted from the individual's point of view. A 20-year-old should use a probability of falls that he believes describes his own risk. I am sure he would like your information in assigning his probability. I believe such modifications of general experience are proper. But the individual should realize that he may be biased. As we know, everybody thinks he's a better driver than everybody else.

J. H. Wiggins (*J. H. Wiggins and Company*)

Have you done anything that says, I'm not talking about gambling on me, but say on my wife, my children, my next door neighbor, the fellow down the street, or finally some person in Miami, Florida who I don't even know? In other words, how would this same kind of a thing deal with the case when it is a person other than myself?

R. A. Howard

That's what we call the value of a friend; it is discussed in the report referenced in the paper. The model shows that when you value your friend as yourself, you are willing to pay for him as yourself. As the degree of friendship goes down and down, of course, you will logically spend less for him than you will for yourself. But, only in the white pill case do I find this an ethically acceptable idea. You may not impose serious risks on others, but you can save people's lives so long as you don't affect them in any way in terms of coercing them. If you want to contribute to someone's medical plan, that's terrific.

J. H. Wiggins

We're doing Black Pill things all the time. All society is imposing serious risks on others.

R. A. Howard

Many people are, but I am not intentionally going to do anything like that, or to encourage it.

J. H. Wiggins

If you vote for a man who votes for certain legislation, then you are responsible?

R. A. Howard

I vote not because I support government coercion, but because I think it's o.k. for a slave to use any means at his disposal to secure his freedom. I am not responsible for the government any more than a slave is responsible for slavery.

M. E. Pate (*Massachusetts Institute of Technology*)

I have a practical question about the evaluation of public policy. I would like to know what kind of figure you would recommend for a group of people, with different ages, incomes, and risks. How would you aggregate their individual figures?

LIFE AND DEATH DECISIONS

R. A. Howard

That's precisely what I would not do. If we start a flying club of three or four people and we want to decide what kind of airplane to buy, then I can provide a small-risk life value to be used in a decision process I have agreed to. But I think it's very dangerous for people to say what other people's lives are worth. Because once you allow them that possibility, of course, you are open to their making the value as small as they like. So I don't like the idea of people setting a value on other people's lives. Ethically I think that each person may want to set a value for his own life; that's up to him. It's all right for him to make such decisions, but not for other people.

W. D. Rowe (*The American University*)

My question, Ron, is how many times a day do you have to make a calculation? For example, I get into my car and I'm driving along having made one calculation, and suddenly I see an accident and I decide I want to remake my calculations because of the imminence or the reminder of the reality factor. So there's a dynamic aspect here, isn't there?

R. A. Howard

Life value should change with age and changing circumstances, as we have seen. However, sudden changes would be unusual. You could spend your whole life making life and death decisions, but that's not what I am recommending. What I am saying is if you feel these issues are important this is a way to make such decisions.

A. S. Curran (*Dept. of Health, Westchester County*)

Getting back to the last question about putting a value on others' lives. In my position as Commissioner of Health, I frequently have to do that in a flip-flop way in that I am asking the taxpayer to pay a certain number of dollars so that he won't have trichloroethylene in his well or something. I am faced constantly with this type of analysis that has to be done, and then peoples' perceptions of what the risk really is. I think what you're talking about today can be very beneficial, but I think we do have to sometimes assume that responsibility.

R. A. Howard

I don't like the imposition of that responsibility on people who don't want it. There was a discussion earlier today about the freedom of the individual. It seems to me the ultimate freedom of the individual is to own and control his own life. So, I don't like the idea of health commissions making decisions about my life, but I guess a lot of other people must because we have a lot of such activity in our society.

B. Bruce-Briggs *Slave. (Laughter) (New Class Study)*

How would you handle a Typhoid Mary problem?

R. A. Howard

Typhoid Mary problem? I wouldn't eat in a restaurant that didn't inspect its workers and require that they be healthy. Typhoid Mary worked in a restaurant. Right? Look at the companies that have had problems with botulism in their canned goods. What has happened to them? Well, they have been in real business trouble. You can ruin your reputation when you take risks with what you put in the can. No restaurant is going to risk its reputation by not having health examinations for its employees. If you're in the restaurant business, you don't need a government regulator to tell you that having such examinations is a pretty smart idea, particularly when there is no corporate limited liability to third parties.

A. Curran

But then you're doing what you said you didn't want to do because if I'm taking the responsibility of saying I'm going to protect your health by sending in sanitarians to inspect that restaurant, I'm making some kind of decision for you. I'm assuming you want to be protected.

R. A. Howard

You heard me wrong. I don't want you to protect me. I want the owner to do it in his own self-interest. For one thing, he won't ask the taxpayers for money to do it.

M. G. Morgan *(Carnegie-Mellon University)*

What sort of plans do you have to use the technique to examine a significant sample of people so that we have some notions of how it would apply to different individuals in different walks of life?

R. A. Howard

That's a very good question. We have done, not what I would call experimentation, but rather class exercises with people of different ages and situations. There is quite a bit of divergence; some people have \$50,000,000 small-risk values and others have \$1,000,000 values. I have not done, nor am I likely to do, a sort of demographic study based on this model. Some other people, I understand, are interested in doing that. I wish them luck. I just hope that it won't be used for public policy decisions.

LIFE AND DEATH DECISIONS

E. A. C. Crouch (*Harvard University*)

Then how do you evaluate the small-risk life value to apply to products if each individual has his own?

R. A. Howard

Good point. You just display it on the product. In other words, if General Motors is going to use a \$1,000,000 value on life (I could have used other companies, of course), they just stamp it on the bumper of the car. A Mercedes Benz could have another number. As I said earlier, it becomes a product characteristic, just like color or how soft the seats are or anything else. Stating that a particular life value was used when it was not would be fraud. I would like the small-risk life value stamped on the product to be the indemnity paid by the manufacturer if someone is killed using it as a result of its design.

E. A. C. Crouch

Then you leave it up to the individual whether he should get that car or not?

R. A. Howard

Who else?

J. Huntsman (*Applied Decision Analysis*)

Ron, how can we be sure that the information that companies state is truthful? A sufficiently large company can lie about their information and there's no other information to go on.

R. A. Howard

Well, that would be fraud. Presumably at the time of trial all this would come out. The papers would be subpoenaed and so forth and so on.

J. H. Wiggins

We have found that when you ask people what they would do and then see what they really do, it's different. Can you use this method for real decisions?

R. A. Howard

I use this method but I can't tell you whether you ought to use it. I know I make my safety decisions this way.

D. McLean (*University of Maryland*)

I'm pursuing this small-risk life value number that's been put on cars. Do you propose that your analysis is limited to such personal product choices like cars or that you put the same sort of number on a highway?

R. A. Howard

Yes, highways, too. Of course the small-risk life value applies only to small risks.

I didn't have a chance to go into it, but where do you find death probabilities outside the safety range, that is, probabilities of death from one in one thousand up to one? They do exist in our society, mainly in medical problems. When you go to the doctor and he recommends an operation, you're often dealing with probabilities in that range and then you might use the more detailed model rather than simply the small-risk life value. But in the safety region, with death probabilities less than one in one thousand, then I'm quite content to use the small-risk life value.

A. Kneese (*Resources for the Future*)

I'm sorry to cut off this very interesting discussion. If somebody has a question that is just burning, I'll let him take one more question. Yes, sir.

M. Thompson (*Institute for Policy & Management Research*)

You say you use this method yourself.

R. A. Howard

That's correct.

M. Thompson

Well, honestly, I could not understand a word of it, so I can't use it for myself.

R. A. Howard

From this short presentation, I wouldn't expect anybody who hadn't heard it before to understand it. The paper will be clearer, and the report that is referenced in the paper even more explicit. But perhaps you're saying not that you didn't understand it, but that you disagreed with it.

M. Thompson

Yes, I did disagree . . .

R. A. Howard

There is hope on that issue, too.

THE VALUE OF LIFE AND NUCLEAR DESIGN

D. L. Owen
J. E. Matheson
R. A. Howard
Strategic Decisions Group

Presented at the American Nuclear Society Topical Meeting,
"Probabilistic Analysis of Nuclear Reactor Safety,"
May 8-10, 1978.

Copyright © 1978. D. L. Owen, J. E. Matheson, R. A. Howard.
All rights reserved.

Abstract

Using the decision analysis framework, an implicit value of life can be determined for design decisions that have death as a possible outcome. Our survey of the literature and our own calculations suggest a great inconsistency in the implied value of life. We show that by using a consistent, explicit value of life, the total expected number of deaths from all projects can be reduced without increasing total expenditures or reducing benefits. The explicit value of life directly affects design decisions. Some recent research indicates that the value that an individual places on his own life can be characterized by a few assessments of the individual's circumstances and his preferences.

I BACKGROUND FOR DECISION MAKING

The decision analysis methodology shows that assessing and combining three elements is essential to making a good decision. These elements are the decision maker's uncertainty about the outcomes, his values for the outcomes, and his attitude toward risk taking.(1) Consider a utility's hypothetical decision of whether or not to construct a new power plant. For simplicity, assume that the decision depends only on whether future customer demand for electricity is 'high' or 'low,' as shown in Figure 1. Four outcomes are possible: new plant/high demand, new plant/low demand, no new plant/high demand, no new plant/low demand.

Regardless of which decision the utility makes, the outcome is uncertain because customer demand is uncertain. Uncertainty is an

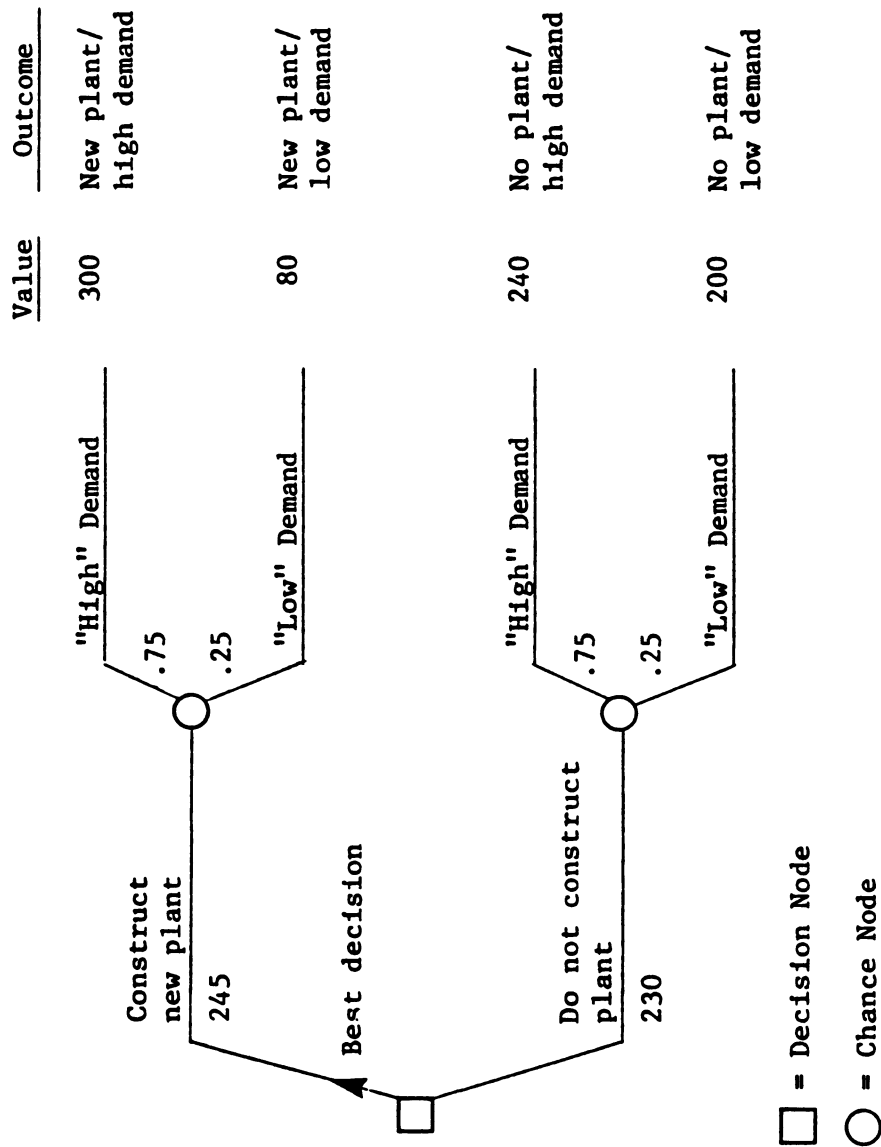


Figure 1: A Simplified Example of a Utility's Decision

Annual Value

Deaths	Property Damage	Project Cost	Project Benefits
--------	-----------------	--------------	------------------

1000	$\$1 \times 10^{10}$	$C + \$8.39 \times 10^6$	D
------	----------------------	--------------------------	---

0	0	$C + \$8.39 \times 10^6$	D
---	---	--------------------------	---

1000	$\$1 \times 10^{10}$	C	D
------	----------------------	---	---

0	0	C	D
---	---	---	---




Figure 2: Example of the Implicit Value of Life

important part of a decision, and probability is the language that the decision maker can use to describe uncertainty.(2) For example, the utility might believe there is a 75% chance of 'high' demand and a 25% chance of 'low' demand, independent of its construction decision.

A second element of the utility's decision is the values of the four outcomes. Value means the worth that the decision maker attaches to one outcome relative to another. A convenient measure of this worth is dollars. The utility could base the value of the new plant/low demand and no plant/high demand outcomes on the price at which it can buy or sell energy to neighboring utilities and could also include the subjective value of public or PUC reaction to these outcomes.

Finally, the utility's decision should depend on the company policy toward taking risks. Including the decision maker's risk attitude in the decision analysis is recognition that decision makers do not generally make decisions on an expected value basis, but evaluate each alternative at less than its expected value because the proposition is uncertain.(3) Since the purpose of this paper is illustrative, we shall ignore risk attitude for simplicity in our examples. In that case, Figure 1 shows the expected value of the 'construct new plant' decision to be 245, compared to 230 for the 'do not construct' decision. Consequently, building the plant is the utility's best decision.

As this example shows, decision analysis analytically combines the three elements provided by the decision maker to determine which decision is logically consistent with the decision maker's information and preferences. The determination of a 'design level of risk' for any part of the nuclear fuel cycle or for the entire fuel cycle is a decision. Thus, the design should depend on the three essential elements: the values of the outcomes, the uncertainty, and the risk attitude.

II NUCLEAR RISK ASSESSMENT AND THE VALUE OF LIFE

Government and industry have been engaged in quantifying the uncertainties in the safe operation of reactors, transportation of fuel, and storage of waste. However, deciding what risks are acceptable also requires specifying the values of the consequences.

There is, perhaps, a natural reluctance to place a value on the possible undesirable consequences. Some speakers at the 1977 International Conference on Reliability and Risk Assessment in Gatlinburg(4) tried to avoid the problem by suggesting that a probability of 10^{-6} of one reactor accident per year for 100 operating reactors is low enough to be acceptable without explicitly valuing the consequences of the accident. Unfortunately, this approach implies a value of the consequence and hides it from public view.

We will use the example of Figure 2 to demonstrate how specifying an 'acceptable' probability level implies a value of life. Suppose two reactor designs, A and B, are available to produce a fixed amount of power. A major accident produces the same consequences in both designs: 1,000 people killed and \$10 billion in property damage. If there are 100 operating reactors, the probability of occurrence of a major accident will be 10^{-6} and 10^{-5} for a year of operation for Designs A and B, respectively. However, 100 reactors of Design A cost \$100 million more than those of Design B (an additional \$8.39 million for forty years, amortized at 8%). The benefits D are assumed to be the same for both designs and to be obtained in both outcomes. The expected value on an annual basis from using Design A and Design B is:

	<u>Design A</u>	<u>Design B</u>
Deaths/year	1×10^{-3}	10×10^{-3}
Damage/year	$\$1 \times 10^4$	$\$10 \times 10^4$
Cost/year	$C + \$8.39 \times 10^6$	C
Benefits/year	D	D

If Design A is accepted in preference to Design B, then the logical implication is that the expected value of Design A must be more than that of Design B. Since the benefits are the same for both designs, the expected cost of Design A must be less than the expected cost of Design B. Letting V be the value of life,

$$1 \times 10^{-3}V + 8.39 \times 10^6 + 1 \times 10^4 + C < 10 \times 10^{-3}V + 10 \times 10^4 + C$$

or

$$\$922 \text{ million} < V.$$

For Design A to be preferred to Design B, the implied value of a life must exceed \$922 million.

The only difference between the two reactor designs in this example is the probability of a major accident. Making costly reactor design changes that reduce an already very low probability of a major accident implies a very high value of life. Of course, many design changes that reduce the probability of a major accident will also reduce the probability of minor accidents. Including this effect and others in a more realistic example could reduce the implied value of life. However, the point of this example is that every decision between alternative designs that affects the probability of death also implies a value of life.

III INCONSISTENCY IN THE VALUE OF LIFE

Because 'acceptable' probability levels for accidents in different industries have been defined without explicit consideration of the value of life, there is great inconsistency in the implied value of life due to different sources of risk. For example, D. Usher lists the implied value of life as \$34,000 to \$159,000 from the hazard pay to miners working underground, as \$161,000 from the hazard pay to test pilots, and as \$140,000 from the instructions to military pilots on when to crash land.(5) Linnerooth lists \$140,000 as the value of life explicitly used in the cost-benefit analysis of highways.(6) By comparison, our own calculations show that the proposed interim criteria for LWR Radwaste systems(7) imply a value of life of \$5 million, and the EPA's proposed Interim Primary Drinking Water Regulations(8) imply a value of life of \$2.5 million.

While we may disagree about the particular number that should be used for the value of life, we should agree that a consistent methodology for establishing the value is important. There is economic inefficiency in treating life as if it is worth \$5 million when setting radiation doses and \$.14 million when designing roads. By using a consistent value of life, the number of deaths could be reduced without reducing total benefits or spending more money.

As an illustration of the advantage of using a consistent value of life in design, consider two projects, both at the design stage (Figure 3). For Project I, a decision must be made regarding the total project cost C_1 , which can vary continuously over some range. There are only two possible outcomes for the project. With probability $P_1(C_1)$, the project results in benefits B_1 at cost C_1 . With probability $1-P_1(C_1)$, the benefits B_1 and cost C_1 still occur, but in addition N_1 lives are lost. The value of life associated with Project I is assumed to be V_1 .



As the cost of the project increases, the project is designed to include additional safety features, and the chance that no deaths result from the project increases. However, incremental safety is assumed to be increasingly expensive, so

$$\frac{d^2 P_1(C_1)}{dC_1^2} < 0 .$$

The second project, Project II, is similar to Project I. However, it has a chance of killing N_2 individuals, and the value of life for this project is taken as V_2 .

The expected value of Project I, given that the project budget is set to C_1 , is

Project I:



Project II:





Figure 3: Two similar Projects with the Death of Some Individuals as a Possible Outcome

$$(1) \quad \langle V|C_1 \rangle = P_1(C_1)[B_1 - C_1] + [1 - P_1(C_1)] [B_1 - C_1 - N_1 V_1]$$

where the symbol $\langle V|C_1 \rangle$ denotes expectation of the project value V conditioned on C_1 . The first order optimality condition for maximizing equation (1) is,

$$(2) \quad \frac{dP_1(C_1)}{dC_1} = \frac{1}{N_1 V_1} .$$

Similarly, for Project II,

$$(3) \quad \frac{dP_2(C_2)}{dC_2} = \frac{1}{N_2 V_2} .$$

Notice that the value of life explicitly appears in these two equations. Furthermore, since the second derivatives are negative, the cost of the projects increases as the value of life increases and as the number of possible deaths increases.

Let C_1^* and C_2^* satisfy equations (2) and (3) respectively. They are the optimum costs for each project.

Now, suppose we maintain the same total cost for the two projects, $C_1^* + C_2^*$, and ask how to distribute that cost between the two projects in order to minimize the total expected loss of life. The constrained minimization is:

$$\text{Min } N_1[1 - P_1(C_1)] + N_2[1 - P_2(C_2)]$$

$$\text{Subject to } C_1 + C_2 = C_1^* + C_2^*$$

As a solution, we find

$$(4) \quad N_1 \frac{dP_1(C_1)}{dC_1} = N_2 \frac{dP_2(C_2)}{dC_2}$$

Substituting (2) and (3) into (4) gives

$$(5) \quad \frac{1}{V_1} = \frac{1}{V_2} .$$

By assigning a consistent value of life for all projects, one minimizes the expected loss of life from all projects. If the value of life is

inconsistent among projects, then it is always possible to redistribute funds between the projects and reduce the number of expected deaths.

Unfortunately, most explicit life value calculations seem to be based on the individual's value to others rather than on his own values. One way to avoid this difficulty is to establish an individual's life value based on the individual's preferences between the length of his life and his level of consumption during it.(9) When this basic preference is augmented by his attitude toward risk, his ability to turn income into future consumption, and his remaining lifetime distribution, we can derive an asymptotic life value that the individual would use in an expected value sense. This value is asymptotic in the sense that it applies to situations that involve a small probability of death. Typically, this value turns out to be several times the economic value of life based on the present value of future earnings.

An example from Reference (9) is that of a 25-year old white male with an annual consumption of \$20,000 per year facing a 5% interest rate. Suppose he feels that a 1% decrease in his lifetime would require a 2% increase in consumption to make him indifferent. Suppose further that his risk attitude is described by a marginal willingness to accept a lottery that is equally likely to increase his annual consumption to \$26,000 or decrease it to \$17,000. Then we can calculate that his asymptotic life value is about \$2.4 million, while his economic life value is only about \$360,000, less than 1/6 as great. The \$2.4 million value would apply for life risks less than, for example, 1/1000. If this individual faced a 10^{-4} chance of dying, he should require compensation of $10^{-4} (2.4 \times 10^6) = \240 to undertake the risk.

IV WHERE IS AN EXPLICIT AND CONSISTENT VALUE OF LIFE NEEDED?

We have focused on one consequence of a major reactor accident, death. A complete value model for the consequences of a major accident must also include the value of non-fatal somatic health effects and a variety of genetic effects.(10)

A major reactor accident, itself, is only one of the outcomes of the many that may result from the operation of a nuclear plant. Other outcomes are energy, waste, and plutonium production without incident; sabotage of the reactor; and diversion of the reactor-produced plutonium. If the government faced the policy decision of whether to support development of nuclear energy generation in preference to coal, all of these nuclear plant operation outcomes and a similar set for coal plant operation would have to be evaluated.

For this high level decision, the value of life, or indeed the value of all external social costs (the cost of death, genetic damage, etc.) may not be particularly important to the decision. For example, using the illustrative data from "The Economic and Social Costs of Nuclear Power," Reference (11), the total external social cost from

nuclear power is .08 mills/kwh compared with the economic (internalized) cost of 24.9 mills/kwh. Since social cost is so small, the total cost of nuclear energy is not very sensitive to the value of life.

Where consistency in the value of life and other social values is very important is at the engineering and operating decision levels. Tradeoffs between improved emergency core cooling system design and increased cost depend directly on the value of life, as suggested by our last example. Designing against sabotage may increase plant personnel radiation exposure, and both the consequences of sabotage and the consequences of increased personnel exposure depend on the value of life. In addition, an explicit statement of the value of life is crucial since the different design criteria are set by different government agencies and implemented by different companies.

V SUMMARY

We have discussed the decision analysis approach to risk assessment and used it to show that specifying an 'acceptably low' probability for an outcome involving death implies a value of life. Our brief survey of the implicit values of life used in several instances shows a large inconsistency. The probability of death could be lowered without reducing benefits or increasing costs if a single value were used in all cases, as we demonstrated with a simple example. Finally, we suggested a value of life computation that could be used to determine the value of life to an individual affected by governmental decision making. The designers of a project must recognize the value of life, either explicitly or implicitly, when the project budget is specified.

REFERENCES

1. Readings in Decision Analysis, ed. by R. A. Howard, J. E. Matheson, and K. L. Miller, SRI International, Menlo Park, California. 1974.
2. C. S. Spetzler and C. S. Stael von Holstein, "Probability Encoding in Decision Analysis," in Readings in Decision Analysis.
3. R. A. Howard, "Risk Preference," in Readings in Decision Analysis.
4. "Nuclear Systems Reliability Engineering and Risk Assessment," Proceedings of the International Conference on Nuclear Systems Reliability and Risk Assessment, Gatlinburg, Tennessee, June 20-24, 1977. Ed. by J. B. Fussell and G. R. Burdick, SIAM, 1977.
5. D. Usher, "An Imputation to the Measure of Economic Growth Changes in Life Expectancy," in The measurement of Economic and Social Performance, ed. by Milton Moss, National Bureau of Economic

Research Studies in Income and Wealth, Vol. 38, Columbia University Press, 1973.

6. J. Linnerooth, "The Evaluation of Life-Saving: A Survey," joint IAEA/IIASA Research Report, International Atomic Energy Agency, Vienna, Austria.
7. "Cost-Benefit Analysis for Radwaste Systems for Light Water-Cooled Nuclear Power Reactors," draft Regulatory Guide 1.FF, Nuclear Regulator Commission, September, 1975.
8. "Interim Primary Drinking Water Regulations," Environmental Protection Agency, 40 CFR Part 141, Federal Register, Vol. 40, No. 158, August 14, 1975.
9. R. A. Howard, "Life and Death Decision Analysis," unpublished notes, Dept. of Engineering-Economic Systems, Stanford University, Stanford, California, May, 1977. A report is in preparation.
10. W. H. Cook, B. R. Judd, C. N. Smart, "Value in Nuclear Power Plant Safety Evaluation," SRI International Decision Analysis Group draft report.
11. S. M. Barrager, B. R. Judd, D. Warner North, "The Economics and Social Costs of Coal and Nuclear Electric Generation," report for the National Science Foundation by SRI International, March 1976.

THE DESIGN OF HAZARDOUS PRODUCTS

Daniel L. Owen
Strategic Decisions Group

The Design of Hazardous Products

DANIEL L. OWEN

Abstract—The kind of information required by the designers of hazardous products in order to provide individual consumers with their desired level of safety is considered. Normative consumer behavior with respect to multiple hazards, involuntary hazards, and public hazards is examined. Finally, a free market for safety is shown to be possible if corporate liability is properly arranged.

INTRODUCTION

Most research in the general area of safety has addressed the problem of how someone acting in the public interest should make safety decisions [1], [2]. Consequently, researchers have focused on two major components of that problem: the assessment of public attitude and the appropriate procedure for balancing costs and benefits of various actions affecting the public. Public attitude is important because from it comes the values, uncertainties, and attitudes toward risk-taking that are included in the cost-benefit equation.

At a recent General Motors conference on safety [2], Slovic *et al.* discussed psychological biases in public perception, including the disagreement between the statistical frequency of death and publicly perceived frequency of death for low-probability events [3]. Stan [4] and Schwing [5] have felt that by analyzing dangers that people currently face, one might gain insight into the public attitude toward safety.

With regard to the proper method of combining societal costs and benefits, there is an entire economic literature of welfare theory. The history of the application of welfare theory to safety issues is documented by Linnerooth [6]. In addition, Linnerooth discusses what she calls the "policy dilemma" resulting from a conflict between the economist's desire to reduce the total number of lives lost and the willingness-to-pay principle [7]. Her paper provides recommendations to the analyst or policymaker as to the appropriateness of the two formulations under various circumstances.

In this correspondence, we address a different problem from that addressed by most research in the safety area. Our concern is how an individual should make safety decisions about his own life and how he can communicate his desires to the product designers. Given an individual's information and preferences, there are certain normative implications for decisions affecting his own life that result from the theory of decision analysis [8]. A descriptive assessment of public attitude is irrelevant to this analysis because of our normative approach. Since our focus is individual decisionmaking, we also do not have the theoretical and practical difficulties of social cost-benefit analysis.

DESIGNING FOR SAFETY IN HAZARDOUS PRODUCTS

In general, design problems involve trade-offs between a large number of product attributes, such as size, output, efficiency, safety, and cost. We are concerned exclusively with the trade-off

of cost against safety. In order to address this issue, we use a simple design problem featuring only these two attributes. Though the results are developed by focusing on engineering design, the results are applicable to toxicity levels of various chemicals as well as other hazardous products.

Suppose we are approached by a client who wants us to design a hang glider for his use during the next year. Since he is very concerned about the possibility of death or serious injury from its use, he is not sure that the commercially available models are safe enough. After considerable effort, we are able to determine the feasible set of hang glider designs. Let $P(c)$ be our client's subjective probability assessment that he is not killed as a result of a hang glider accident, where c is the safety cost of the hang glider. Let the boundary of the feasible set of hang glider designs have the following properties over some range of c :

$$\frac{dP(c)}{dc} > 0$$

and

$$\frac{d^2P(c)}{dc^2} < 0.$$

These properties imply that increasing safety is increasingly more expensive, as shown in Fig. 1.

The design problem is to select a safety cost c for the product design, when the user will face death with probability $1 - P(c)$. This problem is presented in decision tree form in Fig. 2. The box and branches on the left represent the design decision about the safety cost of the product. Following the design decision, the user will use the product and live with probability $P(c)$ or die with probability $1 - P(c)$. The benefits B of this design are assumed to be received whether or not the user is killed and nonfatal injuries, pain, and suffering are not considered. Of course, these simplifications could be relaxed at the expense of some additional complexity in the model.

This figure shows that any selection of safety cost implies a particular balance between the client's resources and his chance of death. For example, a wealthy man could easily afford to reduce the risk of death to a very low level while a poor man could not, and for a given level of wealth we might expect Evel Knievel to accept a higher value of $1 - P(c)$ than you or I.

This author thinks most people would agree that this balancing of risk and cost is best done by the individual who is exposed to the risk. Professor Ronald A. Howard's work on life and death decision analysis shows how an individual can characterize his own balancing of cost against the risk of death by a single number—the small-risk value of life [9], [10]. Howard begins with the premise that in selecting a level of safety an individual is trading off consumption against lifetime. The individual could choose to live in relative safety without much consumption, having spent it all on reducing hazards, or have relatively more to consume by living dangerously. Howard shows how quantifying an individual's preferences for this trade-off leads to the computation of the small-risk value of life.

For example, Howard considers the case of a 25-year old male with a constant annual income of \$20 000, who is indifferent to a one percent decrease in lifetime coupled with a two percent increase in consumption. With a risk tolerance of \$6000 and a real interest rate of five percent, this individual's small-risk value of life is computed to be about \$2.4 million. The small-risk value

Manuscript received March 15, 1980; revised January 26, 1981 and July 7, 1981.

The author was with the Decision Analysis Group, SRI International, and the Decisions and Ethics Center, Stanford University, CA. He is now with Strategic Decisions Group, 3000 Sand Hill Road, Bldg. #3, Suite 150, Menlo Park, CA 94025.




Fig. 1. Properties of feasible set of designs.




Fig. 2. Designer must select safety cost c for design. O: safety outcome.

of life is the number of dollars an individual should attach to his life when he or a designated agent makes a decision that involves a very small possibility of the individual's death.

Returning to the design problem, as designers we could send our client to Professor Howard to have his small-risk value of life computed. When the client returns with his small-risk value of life V , then our design decision is described by Fig. 3. Since the value of life, as computed by Howard's method, incorporates the effects of the client's risk attitude, the value of the design to the client is its expected value:

$$P(c)[B - c] + [1 - P(c)][B - c - V].$$

The condition for maximizing this expected value is

$$\frac{dP(c)}{dc} = \frac{1}{V}. \tag{1}$$

If the probability of proper operation of this product $P(c)$ is interpreted as its safety, then (1) states that the safety cost of the product should be adjusted until the marginal safety is equal to the reciprocal of the client's value of life. Using a higher value of life increases safety, represented by $P(c)$, and since the second derivative is negative, also increases the cost. The properties that we have discussed above for the boundary of the feasible set of designs assure an optimum at some cost.

A particularly popular notion is that the whole question of value of life can be avoided by setting the chance of death from a product at some low level, say 10^{-6} . For example, this view was presented by some speakers at the Symposium on Nuclear and Non-Nuclear Energy Systems, Risk Assessment, and Governmental Decision Making held in Washington, DC, on February 5-7, 1979 [1]. However, for any particular value $P(c)$ there is a corresponding level of marginal safety $dP(c)/dc$ and a value of life implied by (1). This equation clearly shows that a safety decision made on the basis of cost or probability implies a value of life [11]. Studies have identified a large inconsistency in the implied value of life in past safety decisions [12], probably because they were made on the basis of cost or probability.




Fig. 3. Designer's problem with client-supplied value of life V . □: design decision.

THE IMPORTANCE OF CONSISTENCY

Impressed by our understanding of the trade-off between cost and safety, the client asks us to design a parachute for skydiving to be used during the same period as the hang glider. The safety decision for the simultaneous design of the two products is displayed in Fig. 4. The cost of product 1, c_1 , and the cost of product 2, c_2 , are specified during the design, and then the client faces the possibility of death from each of these products. Notice that each product has its own functional relationship between safety and cost. Initially we assume that the client may provide different value of life assignments for the two products. If death results from product 1, then it cannot also result from product 2. Hence the expected value of the two products is

$$B_1 - c_1 + B_2 - c_2 - (1 - P_1)V_1 - P_1(1 - P_2)V_2,$$

and the first-order condition for the optimum design is, for product 1,

$$\frac{dP_1(c_1)}{dc_1} = \frac{1}{V_1 - V_2(1 - P_2)} \approx \frac{1}{V_1} \tag{2}$$

and for product 2,

$$\frac{dP_2(c_2)}{dc_2} = \frac{1}{P_1V_2} \approx \frac{1}{V_2}. \tag{3}$$

The approximation results for small risks where $P_1 \approx P_2 \approx 1$.

Let C_1^* and C_2^* satisfy (2) and (3), respectively. Now suppose we maintain the same total cost for the two products $C_1^* + C_2^*$ and ask how we can minimize the aggregate probability of death. The constrained minimization is

$$\min 1 - P_1(c_1)P_2(c_2)$$

subject to

$$c_1 + c_2 = C_1^* + C_2^*.$$

As a solution we find

$$P_2 \frac{dP_1}{dc_1} = P_1 \frac{dP_2}{dc_2}. \tag{4}$$


Substituting (2) and (3) into (4) gives

$$V_1 = V_2. \tag{5}$$

By using a consistent value of life for all products, one minimizes the total chance of death from all products. If the value of life is inconsistent among products, then it is always possible to redistribute funds to reduce the chance of death without additional expenditure [11].

INVOLUNTARY RISK OF DEATH

Suppose an individual is forced to accept and pay for a product, which we call product 1, that he does not want. For example, some people would have preferred not to have paid for the seat-belt system that was mandatory on 1974 automobiles. Presumably, the individual does not want that product because



Design Problem
for Simultaneous Designs Min $1 - P_1(c_1) P_2(c_2)$
Subject to $c_1 + c_2 = c_1^* + c_2^*$

Fig. 4. Simultaneous design of two hazardous products.

TABLE I
RISK DESIGN CONDITIONS

	Voluntary Risk	Involuntary Risk
Exposed individual bears product cost	$\frac{dP}{dc} = \frac{1}{V}$	$\frac{dP}{dc} = \frac{1}{V}$
Exposed individual does not bear product cost	$\frac{dP}{dc} = \frac{1}{V}$	X

he believes it has a negative expected value. We will refer to products with a negative expected value that must be undertaken as involuntary. (Of course, the expected value should be taken considering all attributes of the product including fear, injury, death, etc. Those attributes could easily be added to our model.) Suppose there is also another product, product 2, that the individual voluntarily accepts.

If we are requested by this individual to design the two products for him, we find that the representation of Fig. 4 is still appropriate, (5) still applies, and the same value of life should be used for both products. Consequently, the involuntary nature of a product does not alter the normative design conditions. In obtaining this result we assumed that the individual who is exposed to the product pays its cost and determines its design.

A much preferable situation is for those who want an individual to use the product to compensate him so that he is indifferent between not having the product and having the product with compensation. Voluntary acceptance by the compensated individual requires that his expectation for the product be nonnegative, or a minimum compensation payment D given by

$$D = -B + c + [1 - P(c)]V.$$

The compensation D must equal the expected loss from the product. Those who pay the compensation want to design the product to minimize compensation. Hence their problem is

$$\min\{-B + c + [1 - P(c)]V\}$$

which has the familiar solution,

$$\frac{dP(c)}{dc} = \frac{1}{V}.$$

Notice that V , the small-risk value of life, is provided by the exposed individual.

A summary of these results in Table I shows that the design condition is the same in several very different situations. The first




Fig. 5. Group of individuals disagree about absolute level of safety, but agree about marginal cost of safety.

row corresponds roughly to democratic action. For example, three individuals may have the different beliefs shown in Fig. 5 about $P(c)$. Because individuals A and B believe a proposed product to be relatively safe, they vote to implement it. Individual C may see the expected value of it as negative because he believes it is relatively unsafe. The product gains a majority vote, and the product is forced upon C . Though he disagrees about the advisability of the product, individual C should agree with the others about its design according to (1) if he has the same value for his own life as A and B assign to theirs, agrees with the others about the marginal cost of safety dP/dc , and pays for his own product.

The first column of the matrix corresponds to a system in which the group accepts only those products to which no member objects. This unanimity is achieved through side payments. Under the assumptions above, those who require compensation and those who do not require it would still agree on the product design.

The lower right box of the matrix occurs when an individual is exposed to a risk but has no say in the design level of safety. In this situation, the exposed individual must rely on the altruism of those who make the safety decision. Hence, our conditions for agreement on the safety design can be summarized with the following.

Theorem: Given individuals who assign identical values to their own lives, V for use in the design of hazardous products, and who agree about the marginal safety per dollar on the boundary of the set of feasible designs $dP(c)/dc$, then there should be broad agreement on how to design the hazardous product among

- 1) those who think the product has a positive expectation and have to pay its cost,

- 2) those who think the product has a negative expectation and have to pay its cost,
- 3) those who have been compensated to be indifferent between not having the product and having the product with compensation,
- 4) those who pay compensation to get the product accepted.

An important point is that this result is derived by considering the value to the individual facing the risk. It is the exposed individual's subjective evaluation of probability (not necessarily official estimates) and his own value of life assessment that must be used in determining that expectation.

While there may be strong disagreement among individuals about the absolute level of safety $P(c)$ afforded by some particular design, agreement about the marginal safety $dP(c)/dc$ seems much more likely.¹ Agreement about the level of safety—that is, the probability of safe operation—is not required for agreement about the design of the product.

A MARKET FOR SAFETY

So far, our development has depended on perfect communication between the individual exposed to a hazardous product and the designer of the product. In large projects affecting many individuals or where a single product is designed for consumption by many individuals, direct communication with the product designer is impractical. Therefore we want to consider the design of hazardous products from the producer's side to determine under what conditions a market for safety might exist.

The design problem for the corporate designer is shown in Fig. 6. Following a decision about the safety cost c of the product, it will operate as intended with probability $P(c)$ and will kill the user with probability $1 - P(c)$. The company makes profit $\pi(c)$ in either event, but incurs a liability loss L if the user dies because the product fails. Selecting the product cost c on the basis of net profit maximization and assuming risk-neutral decisionmaking by the corporation gives as the design condition

$$\frac{dP}{dc} = -\frac{1}{L} \frac{d\pi(c)}{dc}. \quad (6)$$

Two conditions that would cause the product designer to select the same level of safety as in our previous examples, where he works directly with the client are

$$L = V \quad (6a)$$

and

$$\frac{d\pi(c)}{dc} = -1. \quad (6b)$$

Hence if the level of the liability award is equal to the customer's value of life and if each dollar of added safety decreases profits by one dollar, the company will design according to the level of safety desired by the client.

A possible way to implement the first condition is for the manufacturer to label his products with the value of life used in making his design safety decisions. For example, a power lawnmower manufacturer could have an "Evel Knievel" model selling at \$39.95 with a $V = \$100,000$ and a "Howard Hughes" model selling at \$229.95 with a $V = \$2,000,000$. Along with the purchase of either one comes a guarantee to pay V in the event the lawnmower kills its user due to design failure. Consequently, a market for safety is created that allows the consumer the freedom to choose his own level of safety.

This market for safety would also eliminate the need for the consumer to explicitly compute his own small-risk value of life. Successive purchases of hazardous products would lead the con-




Fig. 6. Corporate designer selects design cost that balances profits against possible liability losses. Corporate designer's problem: $\text{Max}_c \pi(c) - [1 - P(c)]L$.

sumer to a level of safety consistent with his own resources, just as he settles on a level of shoe quality consistent with his own resources.

Since profits are revenue less cost, condition (6b) requires that the corporate revenue be independent of the safety cost of the product. Furthermore condition (6b) will not be satisfied if revenues depend on the probability that the product operates as intended, since that probability depends on the cost. Given the postulated liability payment, consumers should purchase hazardous goods on the basis of the value of life used in design, ignoring the probability of death and the safety cost of the product. The value of life incorporates and balances both of these concerns. Thus if consumers make decisions consistent with their own value of life, condition (6b) is satisfied.

SAFETY AS A "PUBLIC" GOOD

In the above discussion on involuntary risk, we considered only products intended for individual use and did not address the design of products that may be hazardous to large numbers of people simultaneously. In the latter case, safety or $P(c)$ takes on the characteristics of a "public" good, because use of the safety by one individual does not detract from the use of it by another. Suppose a particular product has benefits B and can be purchased from the company at a price c , the safety cost. Then a designer authorized to act on behalf of a group of three individuals faces the decision shown in Fig. 7. Notice that we are not trying to solve the problem of how a group should make a decision. Furthermore this authorized decisionmaker is not a government administrator. Rather, the authorized decisionmaker is a designer for the group with perfect information about the beliefs of the individuals in the group. This section is included primarily to show that the previously obtained liability results hold for a group of individuals, each with a different value of life.

Solving as before, the resulting design condition is

$$\frac{dP(c)}{dc} = \frac{1}{V_1 + V_2 + V_3}. \quad (7)$$

The appropriate design condition for a hazardous product affecting many people is obtained by summing the values of life that each individual assigns to his own life.

Fortunately under the liability conditions discussed in the previous section, the company has an incentive to make decisions as if they were the decisionmaker acting on behalf of the individuals affected by the product. A company undertaking this hazardous product sees revenue less the safety cost c as profit. However if the company faces a liability L equal to the cumulative life value of the exposed individuals in the event of malfunction, then its expected profit is

$$\pi(c) - [1 - P(c)]L \quad (8)$$

or

$$\pi(c) - [1 - P(c)](V_1 + V_2 + V_3).$$

Maximization of expected profit leads to

$$-\frac{dP}{dc} \frac{dc}{d\pi} = \frac{1}{V_1 + V_2 + V_3}. \quad (9)$$

¹For example, the probability of a core meltdown in a nuclear power plant is difficult to assess and apparently subject to a great disparity of opinion. However, the reduction in meltdown probability resulting from the addition of an auxiliary feedwater pump is easier to assess and less controversial.




Fig. 7. Designer's problem for a product hazardous to three people. Designer's problem: $\max_c B - c - [1 - P(c)](V_1 + V_2 + V_3)$.

- [12] J. Linnerooth, "The evaluation of life-saving: A survey," Joint Int. Atomic Energy Agency/Int. Institute for Applied Systems Analysis, Res. Rep. IIASA RR-75-21, Laxenburg, Austria, 1975.
- [13] D. L. Owen, "The importance of a consistent value of life in project design," presented at ORSA/TIMS Nat. Meeting, Los Angeles, CA, Nov. 13-15, 1978.

Under our earlier assumption that $d\pi/dc = -1$, (9) and (7) are identical.

One important observation about this result is that each of the group members may assign a different value to his own life. The company designing the product would then agree to pay different liability payments for the death of these individuals.

SUMMARY

The information required by a designer of a hazardous product is the small-risk value of life, which is assigned by the individual exposed to the possibility of death. With this information the designer can equate marginal safety to the small-risk value of life to obtain the optimum design for that individual. An individual should strive to maintain a constant value of life or marginal safety for all hazardous products in order to minimize his total chance of death from all sources. Maintaining equal probabilities of death among hazardous products is not desirable because the individual could reduce his total chance of death without additional cost by redistribution of the expenditures for the products.

Even though an individual may dislike involuntary exposure to hazardous products, if he must bear the product cost, the design condition should be the same whether or not the risk is voluntary and whether or not compensation is paid.

Under certain liability and consumer behavioral assumptions, a market for safety appears possible. By tying the corporate liability to the consumer's value of life, the corporation has an incentive to design the product consistently with the consumer's preferences. Even products exposing large numbers of individuals to a chance of death would be designed as if the corporation were acting on behalf of the individuals exposed to the risk.

REFERENCES

- [1] "Risk assessment and governmental decision-making," in *Proc. Symp./Workshop Nucl. and Nonnucl. Energy Syst., Risk Assessment, and Governmental Decision Making*, Washington, DC, Feb. 5-7, 1979. McLean, VA: The Mitre Corp., Sept. 1979.
- [2] R. Schwing and W. A. Albers, Jr., Eds., "Societal risk assessment," in *Proc. Conf. at Gen. Motors Research Labs.*, Oct. 1979. New York: Plenum, 1980.
- [3] P. Slovic, B. Tuschoff, and S. Lichtenstein, "Societal risk assessment," in *Proc. Conf. at Gen. Motors Research Labs.*, Oct. 1979, R. Schwing and W. A. Albers, Jr., Eds. New York: Plenum, 1980.
- [4] C. Stan, "Societal benefit versus technological risk," *Sci.*, vol. 165, Sept. 1969.
- [5] R. Schwing, "Expenditures to reduce mortality risk and increase longevity," *General Motors Res. Pub.*, GMR-2353, Feb. 9, 1977.
- [6] J. Linnerooth, "Reevaluating the value of life," presented at Operations Res. Soc. of Amer./The Institute of Management Sciences Conf., Los Angeles, CA, Nov. 13-15, 1978.
- [7] E. J. Mishan, *Cost-Benefit Analysis*. London: Allen and Unwin, 1971.
- [8] *Readings in Decision Analysis*, R. A. Howard, J. E. Matheson, and K. L. Miller, Eds., SRI Int., Menlo Park, CA, 1974.
- [9] Ronald A. Howard, "Life and death decision analysis," Res. Rep. EESDA-79-2, Dept. of Engineering-Economic Systems, Stanford University, Stanford, CA, Dec. 1979.
- [10] Ronald A. Howard, "Life and death decision analysis," in *Proc. Second Lawrence Symp. Syst. Decision Sci.*, Oct. 1978, Western Periodicals, North Hollywood, CA, 1978.
- [11] R. A. Howard, J. E. Matheson, and D. L. Owen, "The value of life and nuclear design," in *Proc. Amer. Nuclear Soc. Conf. on Probabilistic Analysis of Nuclear Reactor Safety*, Los Angeles, CA, May 8-10, 1978.

ON BEING ENVIRONMENTALLY DECISIVE

Patricia A. Owen
SRI International

Steven B. Engle
Strategic Decisions Group

INTRODUCTION

An environmental health decision is one that changes the health and safety of workers or the public. Installing safety or emission-control equipment, changing production processes or work practices, putting warning labels on products, and setting priorities for health-related research are examples of environmental health decisions. In the past, many companies focused on complying with current government regulations. Companies relied on government standards because the hazardous nature of many substances had not been established, because technology was not available to detect low levels of hazardous substances, and because companies did not have an effective methodology for deciding whether they should adopt stricter standards.

However, the situation is rapidly changing. Awareness of the health impacts of many substances has increased greatly. Also, instrumentation for detecting hazardous substances and techniques for toxicologic and epidemiologic research have improved greatly. Moreover, concerns about potential liability have also led companies to direct greater attention towards environmental health problems. Finally, court awards and settlements amounting to millions of dollars are increasingly common.

Although environmental health decisions are more commonplace, industry is still limited in its ability to adequately handle these decisions. Health experts may not communicate their recommendations in terms that businessmen understand, and business managers may not understand complicated medical information generated by health scientists. In addition, because environmental health risks may be imposed involuntarily, management time is spent on value issues that are difficult to resolve.

Environmental health decisions are also uniquely difficult to analyze. The probabilities of detrimental health outcomes are often so small that it is difficult to have any confidence in whether there is justification for concern. There may be the potential for catastrophic outcomes or irreversible effects. In addition, there is often a long latent period between the exposure and the onset of the health problem, making it difficult to draw cause-and-effect conclusions.

Some companies have tried a new approach to increase confidence in their health and safety decisions: decision analysis. In the past, decision analysis has been routinely applied to forecast the range and likelihood of financial outcomes resulting from important decisions. Today, the same techniques are being used to forecast health and safety outcomes.

To demonstrate the insights that can be gained by applying decision analysis to environmental issues, we will describe an example of a company's decision to make a capital investment to reduce its workers' exposure to asbestos fibers. Although the example is hypothetical, it draws heavily from several actual analyses and is similar enough to them to provide general insights. We have constructed a hypothetical case because potential liability implications make the actual environmental health decisions we have analyzed too sensitive to discuss.

THE ASBESTOS RISK

In 1953, a major manufacturing company, Loftus Inc., installed sprayed-on asbestos insulation in the ceiling of its Odessa plant. Between 1953 and 1967, accumulated evidence showed that exposure to asbestos fibers caused serious side effects such as asbestosis and lung cancer. Loftus had taken a strong stance on safety. It was unclear whether vibration of the plant equipment and deterioration of the asbestos insulation were releasing asbestos fibers into the work environment. Also, since the Occupational Safety and Health Administration (OSHA) had not yet been established, no regulatory standards existed for asbestos exposure. However, the company voluntarily followed work practices to comply with the standard of 12 fibers per cubic centimeter (f/cc) recommended by the American Conference of Government Industrial Hygienists (ACGIH).

However, in 1967, increasing concern that workers might be exposed to asbestos fibers led the company to consider three options for handling the insulation. One possibility was to maintain the "status quo" and do nothing. However, this option might not achieve a satisfactory level of compliance if the proposed OSHA later promulgated regulations about exposure to asbestos. In addition, the potential liability under the status quo could be tremendous.

A second choice was to repair the insulation, which involved removing and replacing any insulation containing crocidolite asbestos fibers (thought to cause the most serious problem) and sealing any insulation containing chrysotile asbestos. The capital cost of such repairs was \$500,000 on a plant with an annual profit of \$10 million. In addition, inspection and resealing would add an annual cost of \$100,000. Repairing the insulation would be more likely to comply with OSHA regulations and would reduce (but not eliminate) potential liability.

The most expensive option was to remove and replace all asbestos insulation with other materials at a cost of \$5 million. This decision would allow Loftus to comply completely with any air quality standard or with a standard requiring the use of the best available material.

The medical department and the operating division of the company were divided about the asbestos decision. The medical staff was concerned about the risks to workers and wanted to replace all asbestos insulation. The staff was unhappy that its recommendations were being questioned because of a cost-conscious mentality prevailing in the company. The staff felt an ethical compulsion to replace the insulation, because a clear risk had been demonstrated.

The operating division, on the other hand, did not want to change the status quo. It argued that any money spent before the regulations were set might be wasted. In addition, it did not know how to evaluate the quality of the health department's recommendation. The operating division kept asking "How serious a problem is this?", without getting an answer it could understand.

Finally, Loftus tentatively decided to repair the asbestos insulation. This seemed a good compromise between the recommendations of the medical department and those of the operating division. Although it was a satisfactory decision, management was not confident that it was the best decision.

THE DECISION ANALYSIS APPROACH

Because she was uncomfortable about the decision, the president, Ann Loftus, decided to undertake a decision analysis. She assigned Bill Rowan, her staff assistant, to lead the analysis. Bill's first step was to structure the analysis by defining the decisions and the uncertainties facing the company. He reasoned that the decision about the asbestos insulation could be made immediately or postponed into the future. In addition, a decision to stick with the status quo now could be changed by a decision to repair or replace later. A "redesign" decision would be made if OSHA set a standard lower than the current ACGIH standard of 12 f/cc, or if the company accumulated operating experience that exposures were larger than expected. Bill decided to include in his analysis an option to redesign after five years.

Bill consulted the government relations department about whether OSHA standards would be set lower than those currently recommended by ACGIH. The department thought the OSHA standards would depend on the results of the research currently under way at an independent laboratory. The research was designed to investigate whether there is a threshold level (at 2 f/cc) below which there is no health effect. If a threshold was indicated, the department expected OSHA to promulgate a standard at 2 f/cc. On the other hand, if the threshold was disproven or if the research was inconclusive, OSHA would revise the standard, but probably less drastically. The government relations staff assessed only a 20 percent chance that the standard would be set at 2 f/cc in this situation. There remained an 80 percent chance that a standard of 5 f/cc would be set if the threshold effect was not confirmed.

The next step in Bill's analysis was to quantify the range and likelihood of possible worker exposures. Initially, the medical personnel were very skeptical about whether the exposures could be quantified, because there was so little information. Although no significant exposures had been detected thus far in the tests of personnel samplers worn by workers, some high, short-term exposures could have been missed, since exposures could fluctuate greatly from time to time and place to place. In addition, the scientists realized that the existence of other fibrous particulates in the air made the asbestos analysis difficult.

Continuous monitoring and complete analysis by electron microscope would have eliminated much of the uncertainty in exposures. To date, however, these more complete tests had not been undertaken because expected exposure levels did not warrant the added expense.

Bill consulted the industrial hygienist most familiar with the plant. The hygienist, who was not used to quantifying his judgment about exposures, was initially reluctant to give any estimates. Eventually, he estimated that if the status quo were followed, there would be about a 10 percent chance that workers would experience significant exposures to asbestos fibers within five years. By "significant," he meant a level of 5 or 10 f/cc. He also thought that exposures would be lower with repair than under the status quo. In addition, he estimated that if repairs were done, there would be a 5 percent chance that a significant release of asbestos would occur within five years.

Because the estimates of exposure were critical to the analysis, Bill attempted to verify them. First, he consulted a materials scientist about the chance that asbestos fibers would be released from the insulation. Then, he asked the hygienist to estimate the concentration of fibers that would be present in various plant locations if fibers were released. Finally, he consulted plant operating personnel about work practices and the time that workers spent in each location. The resulting exposure levels that were calculated confirmed the hygienist's original direct estimates of exposures. Seeing this result, the hygienist was willing to sign off on the exposure estimates.

Initially, Bill decided to focus on only one serious health effect: lung cancer. The health experts were uncertain about the appropriate dose-response relationship for predicting the fraction of the worker population that, given a particular exposure to asbestos, would develop lung cancer. The best data on the dose-response relationship was from laboratory experiments on animals at doses above the current standard. Because of limitations in these studies, a number of questions remained unanswered. For example, the latent period between the exposure and the onset of cancer, the existence of a threshold, and the rate at which the number of cases increases with exposure were not known.

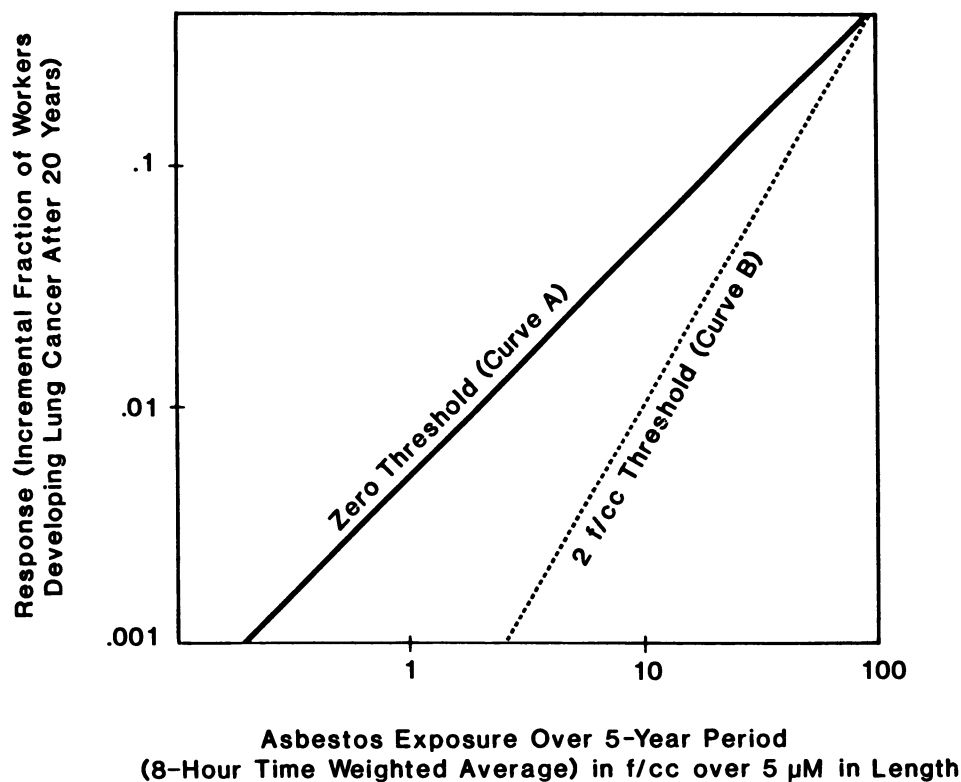
To model this uncertainty, Bill asked the medical staff to estimate the number of responses that would result from a purely hypothetical situation involving ideal laboratory or field conditions. For example, suppose that 10,000 workers received a specified constant dose for 5 years, and that a cellular examination for cancer was performed after a latent period of 20 years. With this specification of the dose and the methods for measuring the response, the medical staff felt confident in quantifying its judgments about response. Of course, where hard statistical data were available, he let the staff review the data before giving its estimates, or he incorporated the data explicitly in the analysis. In this way, the estimates were consistent with all the available hard and soft data.

Bill was surprised to find that experts with differing information and experience did not necessarily disagree about the number of workers that would respond to a given dose. In particular, the toxicologists who conducted experiments on laboratory animals and the epidemiologists who gathered data on human populations estimated the same dose-response relationship when questioned individually. However, a leading toxicologist hired by the company as a consultant disagreed with the head of the company's toxicology laboratory. Bill got them together to discuss their differences and eventually discovered that they disagreed because

they had different recollections of a particular piece of research relating to the body's defense mechanism. After reviewing the literature, they discovered that their estimates of the dose-response relationship were the same.

After careful consideration, the experts agreed that there was still uncertainty about the appropriate dose-response relationship. Thus, they decided to use two curves, shown in Figure 1, to represent the range of possible relationships. Curve A is a linear relationship between dose and response with a threshold at zero. This curve corresponds to the conservative assumption of most regulatory agencies that any exposure to a harmful substance results in some health effect. Curve B has a threshold of 2 f/cc; below this dose, there is no effect.

FIGURE 1: TWO CURVES WERE CHOSEN TO REPRESENT THE RANGE OF DOSE-RESPONSE RELATIONSHIPS.

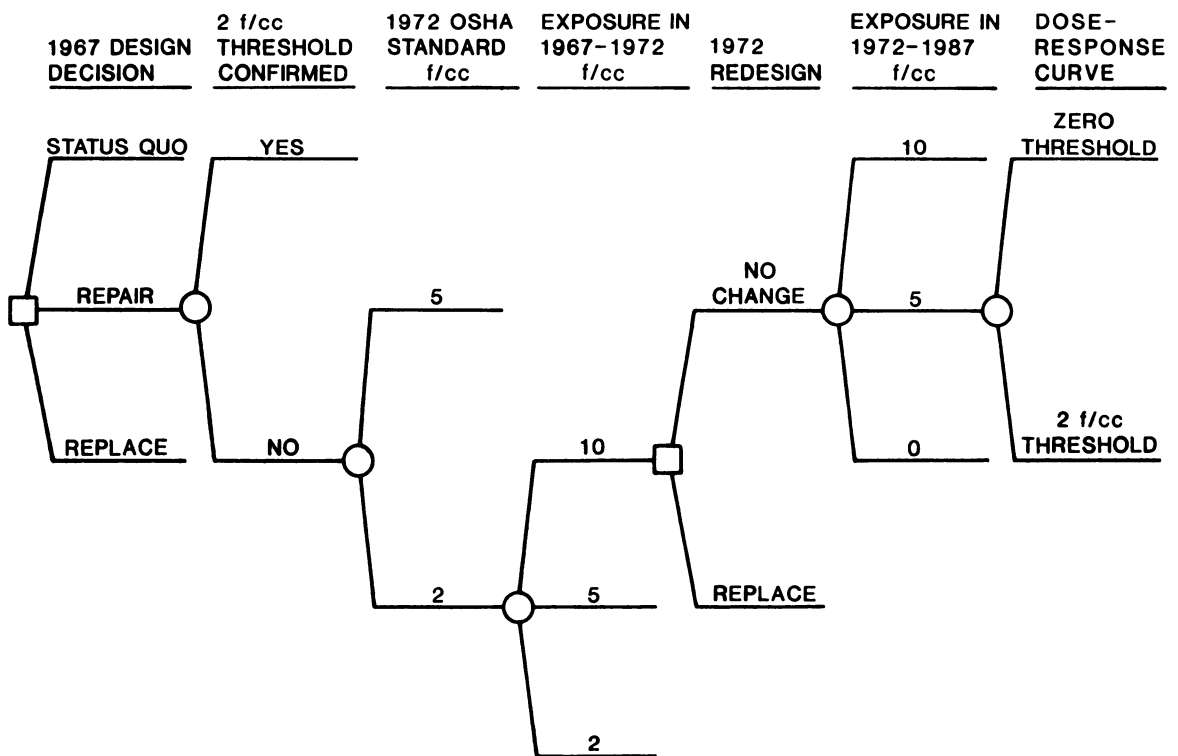


The medical staff believed that the estimate about which curve is appropriate would be influenced by the outcome of the research currently under way to investigate the existence of a threshold at 2 f/cc. If there was a threshold effect, then Curve B would certainly be appropriate. However, if the threshold effect was not confirmed or the research was inconclusive, they estimated a 90 percent chance for Curve A to be appropriate and a 10 percent chance for Curve B to be appropriate.

The Decision Tree

A decision tree (see Figure 2) shows all the possible scenarios that could occur, given the design and redesign decisions, and the uncertainties that make these decisions difficult. Located at the far left is the 1967 design decision. The options to maintain the status quo, to repair, or to replace the asbestos insulation are shown as branches at the square decision node. Then, within five years after this decision is made, the current research to investigate the threshold effect will be completed. The branches marked "yes" and "no" at the round node show the two possible outcomes of the research. Also, by this time, the newly proposed OSHA will have promulgated regulations about asbestos exposure, and the company will have accumulated some operating experience with exposures. In 1972, the company will have the option of redesigning if warranted by OSHA regulations or the company's operating experience. After the redesign decision, additional operating experience will accumulate. Finally, as shown on the far right, one of the two dose-response curves will be appropriate. Which curve will be judged most appropriate depends on the results of the current research.

FIGURE 2: THE DECISION TREE SHOWS ALL THE POSSIBLE SCENARIOS THAT COULD OCCUR.



Each scenario, or path through the decision tree, results in financial and health outcomes. For example, suppose that the company spends no money to repair or replace the insulation, but no exposures result from equipment vibration and deterioration of the insulation. This scenario might be called a "best case," because it results in a low capital cost and no health effects. In a different scenario, one at the other extreme, suppose that the company initially does not make any change in design, but high exposures accidentally occur and the conservative dose-response curve with a zero threshold turns out to be appropriate. At the same time, stringent OSHA standards are set, forcing the company to replace all of the insulation after five years. In this "worst case," there are both high capital costs and substantial health effects. Of course, these scenarios are only two of the many that could occur. All the possible scenarios are represented in the complete decision tree.

Results

The expected net present values and health effects for each design option are shown in Figure 3. The net present values are calculated from the discounted cash flows of the differences between the profits from sales of the product produced by the plant and the capital and operating costs of the immediate design option and any downstream redesign. The health effects are the expected number of cases of premature death due to lung cancer in the population of 2,000 workers. These expected values are averages of all the outcomes from the scenarios displayed in the decision tree, weighted by the probabilities of the scenarios occurring.

FIGURE 3: EXPECTED VALUES SHOW THE TRADE-OFFS MADE IN THE 1967 DESIGN DECISION.

<u>EXPECTED VALUES</u>		
<u>1967 DESIGN DECISION</u>	<u>NET PRESENT VALUE (\$ MILLIONS)</u>	<u>NUMBER OF DEATHS PER 2,000 WORKERS PER YEAR OF EXPOSURE</u>
STATUS QUO	104	.143
REPAIR	102.5	.078
REPLACE	99	0

On the basis of expected values, Loftus had to admit it faced a "problem." The number of incremental premature deaths that could result from exposure to asbestos fibers was significant. Also, the probability of death under the status quo or repair alternatives was nearly as large as many other risks to which individuals are exposed. For example, under

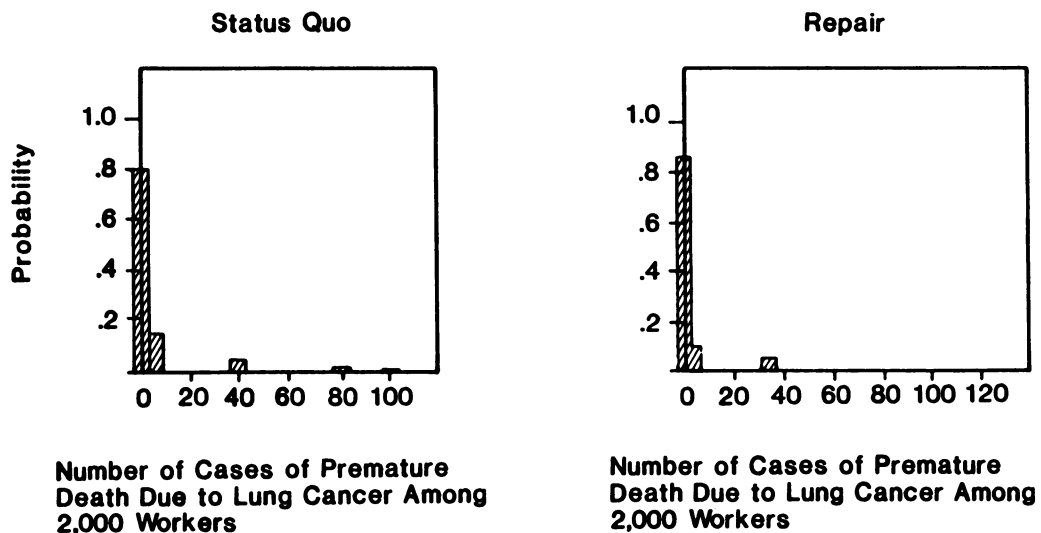
the status quo, .14 additional deaths from lung cancer would be expected per year of exposure among the 2,000 workers. By comparison, a general population of 2,000 people would experience .6 deaths from motor vehicle accidents and .2 deaths from falls over the same period.

However, Bill knew that management would be reluctant to approve the capital expenditure for replacing all the asbestos on the basis of a table of expected values. Expected values do not show the complete range of health effects that can occur. Also, there would undoubtedly be questions about the "value of life" that was being implied by a replace decision. Was the company being more conservative about asbestos than about other health and safety issues?

The company lawyers cautioned Bill about how to present the results to management. Because of liability implications, they felt that the company could not afford to put explicit dollar values on the health effects to compare them with the financial outcomes. In addition, they cautioned that small changes in the wording of the final report could make a big difference concerning liability.

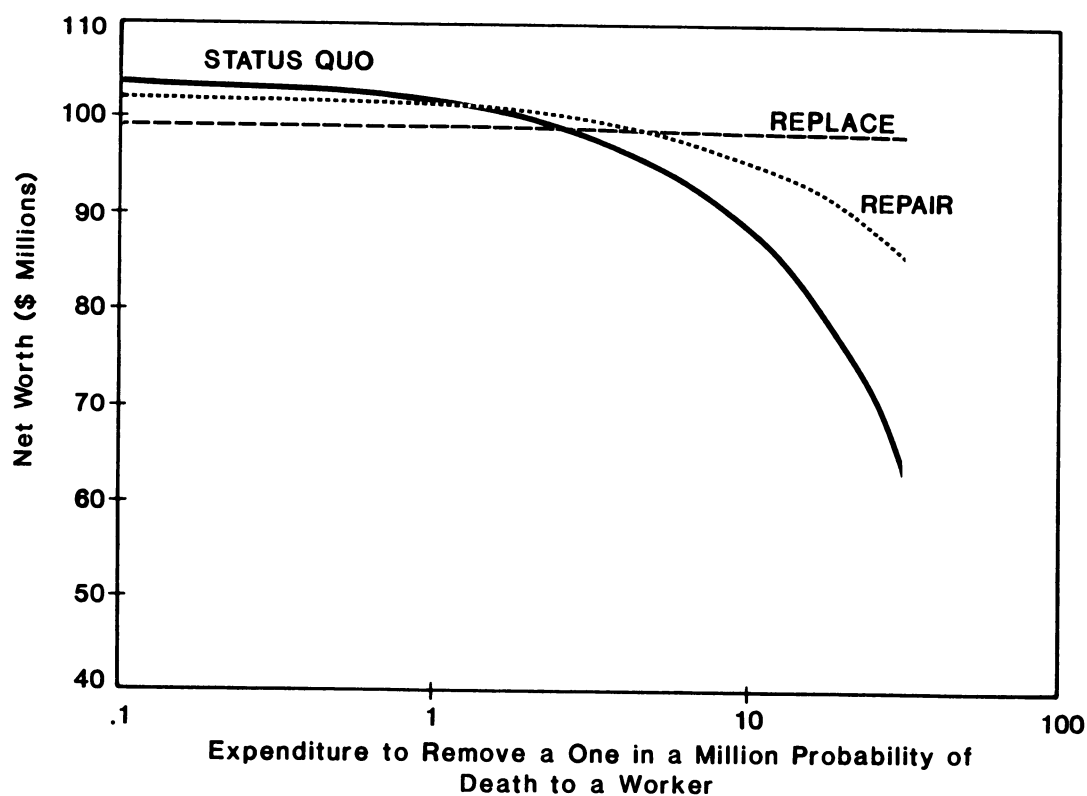
Bill decided to present the case for replacement in two ways. First, he used the decision tree to calculate a probability distribution on the number of cases of premature death that could occur under the status quo or repair alternatives (Figure 4). The probability distribution shows the range and likelihood of cases due to uncertainty about the exposures and the dose-response relationship. The results imply that although there is an 80 percent chance of no health effect, there could be as many as 100 cases of lung cancer under the status quo option. There is a 95 percent chance that there will be fewer than 20 cases under either the status quo or repair options. Of course, this also means there is a 5 percent chance of more than 20 cases.

FIGURE 4: ALTHOUGH THERE IS AN 85% CHANCE OF MINIMAL HEALTH EFFECT, THERE COULD BE AS MANY AS 100 CASES UNDER THE STATUS QUO OPTION.



Bill also calculated how much the company would be spending to reduce risks if it made either a repair or a replace decision. Having reviewed the literature on environmental health decisions, Bill was aware that many companies are willing to spend amounts in the range of \$1 to \$10 to eliminate a one-in-a-million chance of death to a worker. Figure 5 shows that Loftus could justify a replace decision if it was willing to spend at least \$5 to eliminate such a risk. On the other hand, if it was willing to spend less than \$1 to remove this risk, the status quo option would be preferred.

FIGURE 5: A REPLACE DECISION IMPLIES A WILLINGNESS TO SPEND MORE THAN \$5 TO ELIMINATE ONE CHANCE IN A MILLION OF PRE-MATURE DEATH TO A WORKER.



Given these results, management felt confident about deciding to replace all the asbestos insulation. This decision was generally consistent with the amount that other companies were spending on safety. In addition, a decision either to repair or stay with the status quo could put Loftus in a difficult ethical or legal position.

The results in Figure 5 shows that the decision is very sensitive to the willingness of the company to spend money to reduce risks. Management did not want to find itself being accused of undervaluing life and death outcomes. In addition, the possibility that as many as 100 cases of lung cancer could occur at some later date under the status quo or repair options was unacceptable. The corporate aversion to this catastrophe made replacement look even better.

OTHER EXAMPLES OF ENVIRONMENTAL HEALTH DECISIONS

The asbestos example is typical of the type of environmental health decisions that companies have been making with the aid of decision analysis. Some other examples are the following.

- A petrochemical company established priorities for further research relating to a newly discovered toxic effect. Limited research from an independent toxicology laboratory indicated that a broad class of the company's products had possible side effects on workers and customers. The company considered what further research should be undertaken. Possible programs included long-term toxicology studies, clinical studies of exposed workers, and surveys of customers to determine exposure levels. The trade-offs between the accuracy of the information produced by each study and the cost of conducting each study were evaluated.
- A consumer products company determined the level of risk associated with the use of one of its products. The only data available were preliminary results from a short-term toxicology program. The analysis showed how the opinions of toxicologists could be quantified and combined with the limited data to determine the level of risk. One result was a table showing the expected number of cases and the maximum number of cases of various health effects that could be occurring among customers. The probabilities of death and injury from using the product were compared with risks resulting from using other substitutes.
- A major chemical company decided whether to install new equipment that would reduce worker exposures to a toxic substance. The options included upgrading the current technology or changing to an entirely new process. This decision was difficult because of uncertainty about future changes in regulatory requirements and because of the effectiveness of the new process. The analysis showed that doing nothing or installing the new process were the only viable options. Upgrading the technology, which had been originally recommended to management, cost nearly as much as replacement, but accomplished little in terms of risk reduction.

IMPLICATIONS FOR MANAGEMENT

Our experience in applying decision analysis to situations like those described above has led us to some general conclusions. We have found that applying decision analysis can improve the evaluation of environmental health decisions. One of the most important benefits is the integration of the wide variety of information and data relevant to these problems. In addition to health and safety personnel, financial, marketing, and operating officers can be consulted to generate the inputs for this type of analysis. Where differences of opinion arise, examining them explicitly allows them to be resolved more quickly. This process can be very important in achieving consensus in the organization and in increasing management's comfort with the decision.

However, applying any methodology to environmental health decisions is subject to a number of problems. These include technical difficulties in doing quantitative analyses, ethical questions, and liability implications.

Technical Difficulties

A typical criticism of quantitative analysis is that the uncertainty about the health effects is too great to be quantified. Sometimes, the only data available are the results of laboratory experiments on animals at very high doses. The extrapolation of impacts resulting from high doses to those resulting from low doses and the translation of effects in animals to those in humans are very questionable approaches.

The lack of information, however, is precisely the reason that decision analysis is necessary. Decision analysis is most useful when there is a great deal of uncertainty. In such a situation, the best information available is probably the experience of the company's most trusted staff and consultants. Decision analysis provides techniques for capturing the judgment of these experts and combining it with any available hard statistical data.

A second difficulty results from the complexity of environmental health problems. First, there may be dependencies between the risks associated with company actions and those to which the workers expose themselves. For example, workers who smoke cigarettes heavily are more likely to develop lung cancer from exposure to asbestos than workers who do not smoke. Second, there also may be latent effects that occur long after the exposure. Finally, when the exposure ends, some effects are reversible while others are not. With effort, an experienced analyst can include these factors in the evaluation.

Including the possibility of a wide range of health effects in the analysis reduces the chance of making an incorrect decision. In many studies we have done, the range of health effects from a particular product or process is not certain. Limited data indicating one possible effect leave unanswered the question of whether there are other effects. For example, evidence of carcinogenicity may or may not influence the likelihood that there are reproductive or neurological effects. Management should be aware that other impacts are possible and should carefully consider whether those impacts should be included in the analysis.

Another difficulty surfaces because management often wants to know how the risks associated with its product or process compare with risks in other industries or with those to which workers or the public are exposed. This comparative-risk approach, however, may not lead to the correct conclusion. First, there is the philosophical problem that current risk levels may have been set illogically and that they may not be consistent from one situation to another. Even more importantly, the conclusions often vary, depending on the units in which the comparison is made. For example, if there are many individuals involved, the probability of death or illness per person may look very small, while the expected number of

cases may appear large. Similarly, if we express the results in terms of a percentage change in the background rate of death or illness, a radically different impression of the significance of the impact may be created.

A final technical difficulty concerns biases in expert judgment. For example, we have observed a "conservative bias" while obtaining information from health specialists. Because of a genuine concern for preserving life that arises from their training and orientation, health scientists may unconsciously bias the information they provide for the decision analysis. This conservative bias is analogous to the motivational bias that an optimistic sales manager may have about next year's forecast of sales. As a result of this conservative bias, we have sometimes found health specialists more concerned that management makes what the specialists consider the "right decision" than that management makes the decision on the basis of the best information. Skilled interviewers can use assessment techniques to identify and greatly mitigate this bias and others.

Ethical Questions

Environmental health decisions, and the analysis of those decisions, raise difficult ethical questions. It is important to realize that these questions exist even if they are not explicitly addressed. Unfortunately, there are no widely agreed upon answers to these questions. However, management can take a consistent ethical position and be aware of the potential implications of these issues.

Ethical judgments arise as a result of how the analysis is conducted and what factors are included. For example, is it ethical to trade off costs, public image, and liability in analyses of corporate environmental health decisions? In addition, is it ethical to exclude some option or some potential outcome from an analysis because of possible liability implications?

The valuation or weighting of health outcomes is another ethical issue arising in analyses. What if the health effect does not impair the individual's function in his job or personal life, but changes only his likelihood of developing other problems? Or what if a toxic substance causes cellular damage to an organ but does not affect the organ's function? What value should be put on these effects?

The decision itself also has ethical implications. While these implications are not necessarily a part of the decision analysis, the analysis process often brings these issues into the open. In our experience, we have heard senior management debate the ethics of the following actions.

- Protect a worker by denying him a particular job, when his own actions (such as smoking) or physical condition leave him more susceptible to certain risks.
- Pay workers higher salaries for accepting risky jobs.
- Allow workers in some plants to be exposed to different probabilities of death and injury than workers' in other plants (for example, plants in foreign countries).
- Undertake research prohibited in one country in another place where human experimentation of the type needed is allowed.
- Restrict sales to customers known to use the product in a safe way.
- Inform workers or the public about risks when, although the company is meeting current regulatory requirements, new but uncorroborated information not reflected in the regulations becomes available.
- Warn end users of risks that are not advertised by downstream packagers or distributors of products.
- Sell similar products with different levels of safety at different prices.
- Decide whether to undertake further research that may reveal unrecognized hazards.

Although the above list is not exhaustive, management should be aware that many subtle ethical issues do exist.

Liability Implications

Quantifying the probability of a health effect may increase the company's chance of being liable for knowing about that effect. Quantification may be interpreted as an admission that there is a finite risk and may put the company in a worse liability position than if the analysis were not done. Although this possibility contradicts common sense, it is an unfortunate fact in today's judicial system.

However, in terms of liability, decision analysis has an important advantage over conventional analysis. Because the output of the decision analysis is a probability distribution showing the range and likelihood of health effects, the uncertainty in the environmental health decision is apparent. Conventional analysis, which focuses on a single case, does not emphasize the uncertainty surrounding such problems. Instead, the results of a single-case analysis make it seem that the company is certain of the problem's severity. An analysis that presents only the best, most likely, or even the worst case can be easily criticized.

In addition to this advantage, there are several things that can be done to ensure that the decision analysis does not aggravate the company's liability position. First, it goes without saying that to protect the company's liability position, the analysis must not have obvious technical mistakes.

Second, trade-offs of health and financial outcomes can be evaluated using a wide range of weighting factors on the health consequences. In many cases, the best option is obvious regardless of the values chosen. If the decision is sensitive to the values, then, as in the asbestos example, other considerations, such as liability, may recommend a particular option.

Third, small changes in the wording of reports and presentations may make a big difference concerning the company's potential liability. Because of the uncertainties involved, the analysis can be described as preliminary or the data, as illustrative. Value-laden words such as "large," "insignificant," or "serious" can be avoided in describing health effects, and notes can be included if forthcoming data may change the conclusions.

Fourth, the company may be able to improve its liability position by undertaking the decision analysis jointly with other companies through a trade or professional association. This action may prevent one company from being in a worse liability position than others. It also spreads the cost of the analysis and may increase the perceived objectivity of the study.

Finally, legal counsel should be consulted before doing any analysis. This does not mean that the lawyers should dictate whether, or how, an analysis is done. However, their inputs are just as important as those of the operating division or the medical staff. In doing the analysis and in presenting it, companies should never place themselves in an untenable position.

CONCLUSION

As we have discussed, environmental health decisions are fraught with technical, ethical, and legal difficulties. But because management faces these strategic decisions with increasing frequency, and because most senior managements are not equipped by experience or training to deal with the medical aspects of environmental health decisions effectively, management needs an analytic framework to arrive at the best decision. Our experience with decision analysis suggests that it is one effective tool for increasing management's confidence in its decision.

ON FATES COMPARABLE TO DEATH

Ronald A. Howard
Strategic Decisions Group

Reprinted from Management Science, Volume 30, No. 4, April, 1984.

Copyright © 1984. The Institute of Management Sciences.

INTRODUCTION

Two fallacies have often impeded the studying of decision-making about hazards to human life. The first fallacy--the altruistic fallacy--is that the value a person places on his life is primarily related to such measures as the life insurance he carries, court judgments, and discounted future earnings--even though these measure mainly the value of his life to others rather than to himself. The effect of this fallacy is usually to place value on life too low for people to accept in making their own decisions.

The second fallacy is the incremental fallacy. Simply stated, it is the belief that because a person is willing to incur additional risk for money, one can infer a cash payment that the person should accept for being killed. Clearly, this is absurd, but this fallacy is often why discussions about placing values on lives have seemed both silly and frightening to the public.

The incremental fallacy is also present in other spheres of life. Thus the price a person would exact to accept a given incremental risk should rise as the total risk the person faces increases. No conceptualization of life decision-making is complete until it has captured this phenomenon.

This paper extends previous work concerned with risks to life [1, 2, 3, 4, 5] to include the possibility of living with various degrees of disability. We shall investigate not only risks of disability themselves, but the interaction between death and disability risks that such events as medical operations might pose.

The author thanks the referees for several helpful suggestions in both the presentation and the citation of references.

THE ETHICAL FOUNDATION AND ITS IMPLICATIONS

Before proceeding, we must be clear on the ethical framework of the analysis. Our ethical assumption is that each person, and only that person, has the right to make or to delegate decisions about risks to his life or well-being. We note in passing that several present social arrangements do not meet this ethic and are the source of continuing controversy; for example, consider the development of the "right-to-die" movement.

Once we accept this ethic and deal with an individual who is rational in that he desires to follow the axioms of probabilistic logic, we can develop a procedure that will allow the individual to make consistent rational choices about the hazards in his life. In particular, the procedure focuses on what the person's life is worth to himself, using his preferences for different life states. Furthermore, it shows that an individual is consistent in being willing to trade risk of life or disability for money at low risk levels and yet refusing to do so at high risk levels.

SUMMARY OF PREVIOUS WORK

We shall now summarize briefly the model we will use and extend in this paper. Readers are referred to other sources [2] and [4] for a more complete description. In the simplest model investigated, we assume the individual has a fundamental preference on consumption and length of life and ignore the question of legacies. Consumption is defined as the constant level of consumption beyond bare survival over the remainder of life, measured in real dollars, that would be equivalent to current expectations; this is called the constant annual consumption c . When we combine c with the remaining years of life ℓ , we have the fundamental descriptors

of life quality and quantity. We assume that the individual can trade between c and ℓ to develop a worth numeraire $w(c, \ell)$, and we use the simple model:

$$w(c, \ell) = c \left(\frac{\ell}{\bar{\ell}} \right)^\eta \quad \eta \geq 0 .$$

Here η is the parameter governing the tradeoff, and $\bar{\ell}$ is the expected lifetime remaining, a useful benchmark, when $\ell = \bar{\ell}$, $w(c, \ell) = c$.

In fact, the emphasis of this model is on percentage changes in c and ℓ rather than on their magnitudes. When ℓ is small, this property may not be appropriate, and other models may be useful [2].

Now we must specify the risk preference of the individual on the numeraire w . We use the exponential form with risk tolerance ρ ,

$$u(w) = -e^{-w/\rho} .$$

To find the payments that will compensate for risk, we use those payments to modify c in this model. We assume that any payment received will be used to buy an annuity at the prevailing interest rate i from a seller who agrees with the buyer's probability distribution on life. Payments to be made can be financed in a complementary manner. For example, if the individual receives one dollar, then he will be able to increase his annual consumption by

$$\zeta = \frac{1}{1+i} \frac{1}{1 - \left\langle \left(\frac{1}{1+i} \right)^\ell \right\rangle}$$

where $\langle \rangle$ denotes expectation.

One result is immediately derivable from this model. If the individual is offered larger and larger amounts of money to face a probability of imminent death p , there is a maximum value of p ,

$$P_{\max} = \left\langle e^{-\frac{c}{\rho} \left(\frac{\ell}{\ell}\right)^{\eta}} \right\rangle$$

No amount of money, however, would induce the individual to accept a risk of death as large as P_{\max} .

The Micromort Value

While curves can be, and have been, derived to determine the payments $x(p)$ required for risks of death $p < P_{\max}$, the case of small incremental risk is of special interest. We define v_s , the small-risk value of life, as

$$v_s = \lim_{p \rightarrow 0} \frac{x(p)}{p},$$

which is readily computable from the model. If a person had calculated his small-risk life value, then it would suffice for all decisions in the safety range, say for $p < 1/1000$. To determine the death risk to him, in dollars, that should be compared with other costs and benefits, a person would simply multiply p by v_s .

However, since even the small-risk life value may lead some to the incremental fallacy, it is better to use v_s in the form of the value of a micromort [μm], a 10^{-6} chance of death. Then $v_{\mu\text{m}} = 10^{-6} v_s$. Now, by keeping in mind a micromort value $v_{\mu\text{m}}$ when confronted with a risk

in the safety range, any person can simply compute the number of micromorts in the risk, multiply by $v_{\mu m}$, and establish the death risk in dollars. Although this change is cosmetic only, we should remember the size of the cosmetic industry.

As we have said, the model described is the simplest one we have analyzed that both possesses the desired qualities of finite p_{\max} and a small-risk life value and is rich enough to suit many tastes. However, to represent a wide range of value functions and risk preferences, more general formulations are available [2].

We should note that the concept of developing a worth numeraire on attributes and then assessing a risk preference upon it is as general a procedure as using a multi-attribute utility function.

Numerical Results

To obtain a feeling for the model, it is useful to summarize previous numerical results. Consider a base-case individual--a 25-year old male with a \$20,000 per year constant annual consumption. He chooses $\eta = 2$, which means he would forgo 2% of his consumption over his remaining life to have it be 1% longer. His risk tolerance is \$6000, which means, for example, that he is roughly indifferent between his present situation and a lottery with a 2/3 chance of \$32,000 per year and a 1/3 chance of \$14,000 per year. The prevailing interest rate i is 5%.

How much would this individual have to be paid to accept a probability of death p ? Figure 1 shows the answer. Note that no amount of money will induce him to accept a risk of death greater than $p_{\max} = 0.103$. In the safety region, the curve becomes a straight line corresponding to a small-risk life value of \$2.43 million.

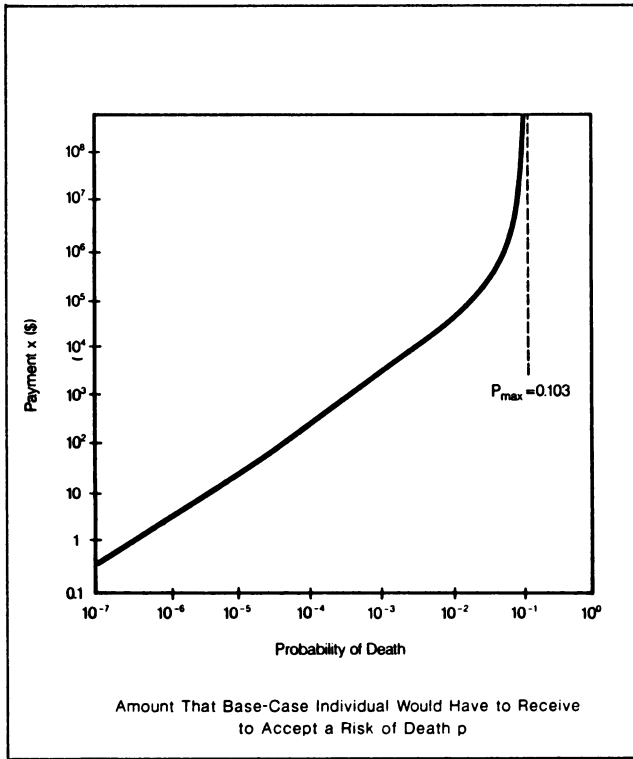


Figure 1




Figure 2

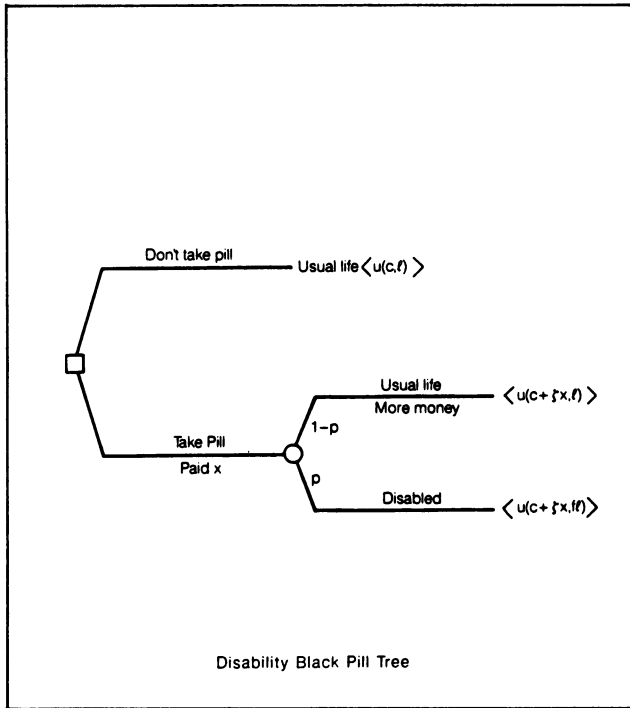


Figure 3

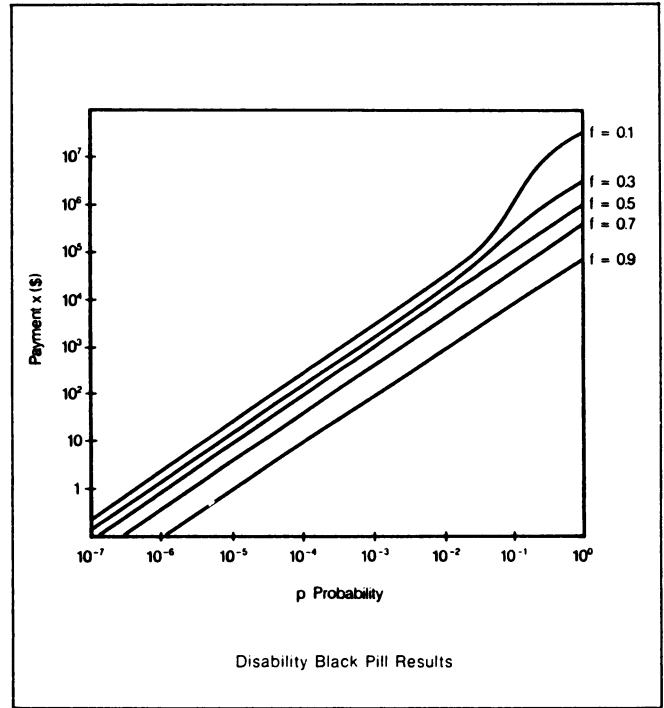


Figure 4

If the individual can avoid a risk of death p by paying money, then we obtain the curve of Figure 2. Even to avoid certain death, he would not be able to pay more than the value of the annuity represented by his constant annual consumption c/ζ ; we call this the economic value of his life, $v_e = \$363,000$. The straight-line portion of the curve in the safety region again corresponds to a small-risk life value of \$2.43 million.

Thus, while there is a considerable difference between buying and selling risks for large probabilities of death, the treatment is symmetric in the safety region. The individual is willing to buy and sell micromorts at a price of \$2.43. If someone wanted him to take on a risk of $\frac{1}{10,000}$ probability of death, he would evaluate the death risk at \$243 (100 μm). Of course, if someone wanted to buy 1,000,000 μm , he would refuse. This micromort value will increase as the individual takes on more and more risk; however, for most of us who are both buying and selling very small risks all the time, it should be relatively constant and hence a useful guide to safety decisions. Naturally, the death risk cost computed in this way must be combined with other costs and benefits to arrive at a decision.

Continuing Hazards

While we have discussed risk as if it will occur only in the present year, many of life's activities, like smoking, for example, imply a change in the risk of death in future years [2] and [4]. Knowing the changes in risks in future years will enable us to derive the corresponding lifetime distribution and to use the model to compare this situation with any other.

A serious hazard, like smoking, which might double the chances of dying in any year, would reduce life expectancy from 46.2 to 38.4 years. To be consistent, the base-case individual would demand a lump-sum payment of

\$212,000 or an annuity of \$12,400 before assuming such a risk. If smoking is not worth that much to him, he should reconsider whether to pursue it.

HEALTH STATES

Now let us show how to use this same structure with a slightly more complex description of life. We shall retain constant annual consumption as the economic indicator of life, but consider more possibilities for the non-economic aspects of life. Suppose we define a number of health states $i = 1, 2, \dots, n$ for the individual that he might occupy at any time in his life. Such states might be defined in many ways; we are concerned here, however, with the concept. The states might be characterized by the potential mobility of the person, by the amount of pain he feels, or by the senses he has available. For any life, the transitions through the states would be probabilistic, both as to the succession of states and as to the time spent in each. A semi-Markov model would be a fairly simple probabilistic model to describe this process, provided that the Markovian assumption was justified in this particular case [6]. For related approaches, see references [7] and [8].

To complete the description of the system, we need to extend the earlier value model, which depended only on c and ℓ . As a simple extension, let us retain the original form for $w(c, \ell)$, but let ℓ now be a weighted sum of the years spent in each state i , ℓ_i ,

$$\ell = \sum_i f_i \ell_i .$$

For the state of full health and capacities, the weight would be one; however, for less desirable states, it would be less than one.

For example, suppose a person faced a normal lifetime except that from this point on he would be blind. If he said that living his remaining life blind would be equivalent to living only 30 percent of that life sighted, then the weighting for the blind years would be 0.3. While more complicated schemes can be developed to represent preferences, let us see how valuable this simple model can be.

One question that immediately arises is whether the weights can be negative; that is, are some life states so bad that living them would be worse than being dead? Many people feel that there are, seeing total paralysis as one such state. Since we have assumed that each person has the right to make and delegate decisions about his own life, any person could choose suicide at any time and, thus, weight that life state as zero. The restriction to non-negative weights is, therefore, not a problem for those who have suicide as an option.

DISABILITY

We shall focus our analysis on the case of disability: where the person faces the possibility of spending the rest of his life with a serious health impairment, like blindness, paralysis, or severe pain. In each case, we imagine assessing the worth function by asking the person the following question: "Suppose that instead of living the rest of your life in the state of health you expect, you would live it in this way (a specified disabled state, like blindness or total paralysis). However, you have a choice--you can continue to live with the state of health you expect, but for a reduced time. What fraction of your remaining years lived in this state would be just equivalent to living all your remaining years in the disabled state?" The fraction f that the person answers will be used as

the weighting factor in the value model. We can then see how the person would make various decisions involving a risk of disability or a risk of either death or disability. We shall assume that the fraction f is assessed independent of ℓ , although there is only modest difficulty in making it a function of ℓ . For the moment, we shall further assume that the person's income will be unaffected by his disability and, hence, concentrate on the qualitative aspects of the disabled state.

The Disability Black Pill

We proceed now in a manner analogous to our earlier work on death risk. Suppose the person is offered an amount of money x to assume a risk p of being disabled to a level f . For example, someone could be offered \$1000 to take a $\frac{1}{10,000}$ chance of becoming totally paralyzed, an outcome he regards as equivalent to $f = 0.1$. The choice is diagrammed in Figure 3. For simplicity, we can imagine the risk is contained in a pill and that everyone agrees on the probability p . If the person refuses to take the pill, then he lives his normal life, receiving any windfalls and calamities that may be in store for him--his future life lottery. This has a utility of $\langle u(c, \ell) \rangle$, obtained by multiplying the utility of each constant annual consumption level c and remaining life ℓ by the joint probability distribution on these quantities. If the person takes the pill, then with probability $1-p$ he will live his future life lottery as before, only with more money. He converts the payment x into the constant annual consumption $c + \zeta x$, with expected utility

$$\langle u(c + \zeta x, \ell) \rangle .$$

However, if he becomes disabled, which he will with probability p , then he transforms the payment in the same way, but has his remaining years multiplied by f so far as his value function is concerned. Therefore, he has expected utility

$$\langle u(c + \zeta x, fl) \rangle .$$

Setting the expected utility of the two alternatives equal, we find

$$\langle u(c, l) \rangle = p \langle u(c + \zeta x, fl) \rangle + (1-p) \langle u(c + \zeta x, l) \rangle$$

or

$$p = \frac{\langle u(c + \zeta x, l) \rangle - \langle u(c, l) \rangle}{\langle u(c + \zeta x, l) \rangle - \langle u(c + \zeta x, fl) \rangle} .$$

Base-Case Results

This equation allows us to calculate the risk of a disability level f that would be assumed for a payment x once all the other parameters of the model have been specified. With the results shown in Figure 4, we evaluate the equation for the base-case individual described earlier. We see that the curves begin to diverge from straight lines on these scales only for p close to one and f small, that is, when there are severe risks of serious disabilities. Within the safety region ($p = 10^{-3}$ or less), they are all straight lines. This means that in this region, we can define a small-risk value of disability level f , $v_{sd}(f)$ that

can be used by the individual in the expected value sense to compute x as

$$x = p v_{sd}(f) ,$$

where $v_{sd}(f) = \lim_{\substack{x \rightarrow 0 \\ p \rightarrow 0}} \frac{1}{dp/dx}$. If we define

$$v_{\mu d}(f) = 10^{-6} v_{sd}(f) ,$$

then we have the value of a 10^{-6} chance of a disability level f in consonance with our definition of a micromort.

The results for the base-case individual are listed in Table I and plotted in Figure 5. Notice that the microdisability value approaches the micromort value of \$2.43 when f is small, and that it approaches zero as f approaches 1 .

Table I
Small Risk Value of Disability at Level f for Base-Case Individual

f	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
$v_{sd}(f)$ [M\$]	2.333	2.066	1.696	1.297	0.927	0.620	0.383	0.210	0.087
$v_{\mu d}(f)$ [\$]	2.33	2.07	1.70	1.30	.93	.62	.38	.21	.09




Figure 5

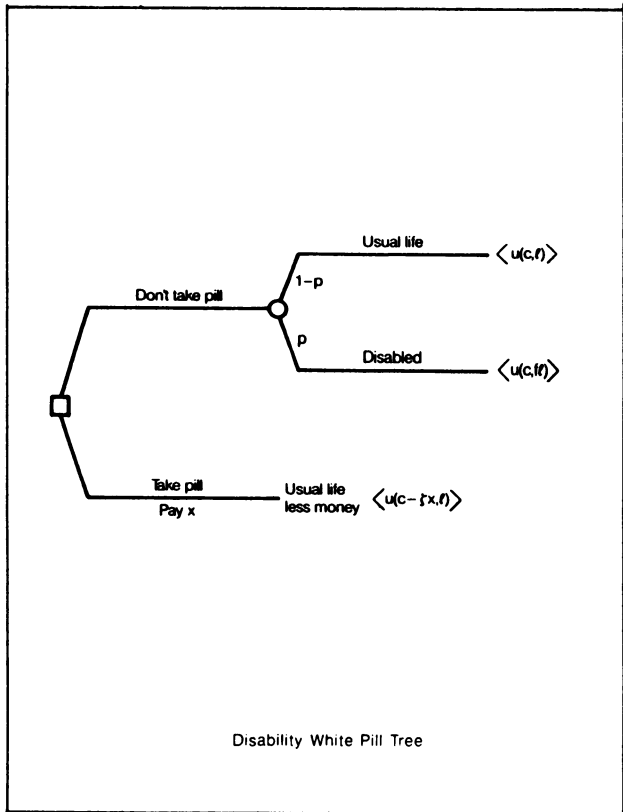


Figure 6

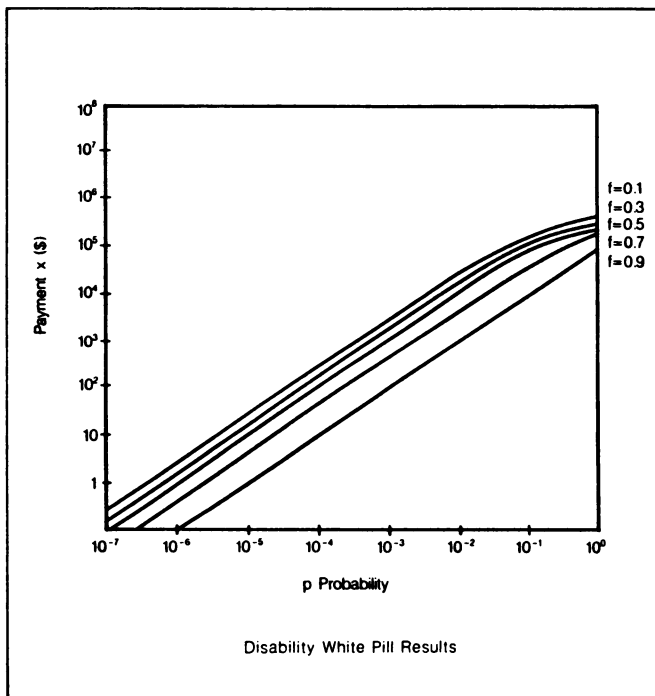


Figure 7

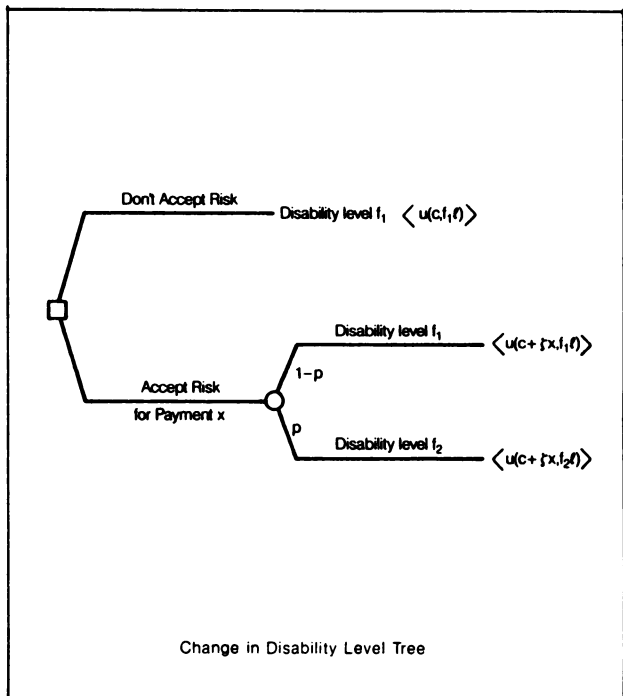


Figure 8

The Disability White Pill

We can also think of selling risks of being disabled. Suppose that a person facing a risk p of being disabled to level f could avoid that risk for a payment x . How much should he pay? We call this the disability white pill question and diagram it as shown in Figure 6. Note that if the person pays x , he finances the payment by reducing his constant annual consumption to $c - \zeta x$. Of course, the most he could pay would be

$$x = \frac{c}{\zeta} = v_e ,$$

the economic value of his life. Equating the utility of the two alternatives, we find

$$\langle u(c - \zeta x, l) \rangle = p \langle u(c, fl) \rangle + (1-p) \langle u(c, l) \rangle$$

and

$$p = \frac{\langle u(c, l) \rangle - \langle u(c - \zeta x, l) \rangle}{\langle u(c, l) \rangle - \langle u(c, fl) \rangle} .$$

Base-Case Results

Figure 7 shows how much the base-case individual would pay to avoid a probability p of a disability level f . Note that the curves are again straight lines in this plot until they approach $p = 1$. In this region, they bend downward as the individual encounters the finiteness of his resources. However, for values of p in the safety region, they are straight lines and, in fact, the same straight lines as in Figure 4. This means that the individual will buy and sell small risks of being disabled for the microdisability values in Table I and Figure 5. He thus treats small risks of disability in the same way as small risks of death.

CHANGES IN DISABILITY LEVEL

We can also consider the risks of changing from one level of disability to another. In particular, we shall consider the case of a person who is currently disabled becoming more disabled, even to the point of death. The situation is diagrammed in Figure 8: a person who is currently disabled at level f_1 is offered x to assume a probability p of being disabled to a more restricted level f_2 , including even the possibility of death, $f_2 = 0$. To find the x to which he is indifferent, we equate the utility of the two alternatives,

$$\langle u(c, f_1, \ell) \rangle = p \langle u(c + \zeta x, f_2, \ell) \rangle + (1-p) \langle u(c + \zeta x, f_1, \ell) \rangle$$

and find

$$p = \frac{\langle u(c + \zeta x, f_1, \ell) \rangle - \langle u(c, f_1, \ell) \rangle}{\langle u(c + \zeta x, f_1, \ell) \rangle - \langle u(c + \zeta x, f_2, \ell) \rangle}$$

where

$$1 \geq f_1 > f_2 \geq 0 .$$

If we let $v_{sd}(f_1, f_2)$ be the small-risk value of disability from level f_1 to level f_2 defined by

$$v_{sd}(f_1, f_2) = \lim_{\substack{x \rightarrow 0 \\ p \rightarrow 0}} \frac{1}{dp/dx}$$

then in the safety region we can compute the required payment x from

$$x = pv_{sd}(f_1, f_2)$$

or from the equivalent microdisability value of this change

$$v_{ud}(f_1, f_2) = 10^{-6} v_{sd}(f_1, f_2) .$$

The results for the base-case individual are shown in Table II. Notice that the first row of the table reproduces the small-risk values of Table I for the case of risk of disability to a healthy individual. In general, the value decreases as f_1 decreases for a fixed f_2 and increases as f_2 decreases for a fixed f_1 . The first phenomenon results from contemplating the same prospect from worsening initial states, and the second one results from contemplating worsening final states from the same initial state.

Table II
 Small-Risk Value of Changing from Disability
 Level f_1 to Disability Level f_2 , $f_1 > f_2$
 (Thousands of Dollars)

From f_1 \ To f_2	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0
1.0	87	210	383	620	927	1297	1696	2066	2333	2431
0.9		96	231	416	656	944	1256	1547	1752	1829
0.8			107	254	444	673	920	1149	1314	1374
0.7				120	277	464	667	855	991	1041
0.6					136	300	477	641	759	802
0.5						156	325	482	594	636
0.4							184	354	477	522
0.3								225	386	446
0.2									291	398
0.1										372

COMBINED HAZARDS OF DISABILITY AND DEATH

Deciding to participate in certain activities, like driving a motor vehicle or operating a chain saw, entails hazards ranging from minor disabilities through death. Suppose there are n such hazards, and let p_i be the probability that the individual will be disabled to level f_i , $i = 1, 2, \dots, n$, where $f_i = 0$ represents death. Then, the total probability of one of these outcomes is

$$p = \sum_{i=1}^n p_i,$$

and the conditional probability of disability to level i given that some hazard occurs is $q_i = p_i/p$. The amount x that the individual would have to be paid to take on this combined hazard can be determined from Figure 9 by equating the expected utility of the two alternatives:

$$\langle u(c, l) \rangle = (1-p) \langle u(c + \zeta x, l) \rangle + p \sum_{i=1}^n q_i \langle u(c + \zeta x, f_i l) \rangle$$

or

$$p = \frac{\langle u(c + \zeta x, l) \rangle - \langle u(c, l) \rangle}{\langle u(c + \zeta x, l) \rangle - \sum_{i=1}^n q_i \langle u(c + \zeta x, f_i l) \rangle}.$$

We can now let v_{sh} be the small-risk value of the combined hazard, defined by

$$v_{sh} = \lim_{\substack{p \rightarrow 0 \\ x \rightarrow 0}} \frac{1}{dp/dx}$$

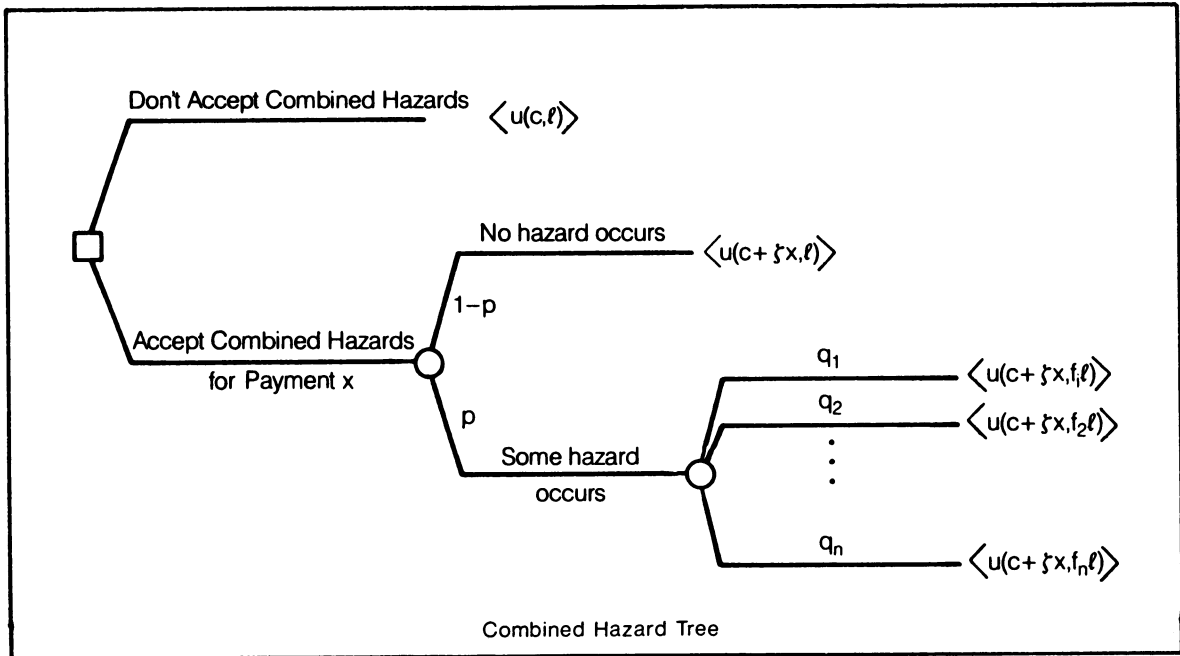


Figure 9

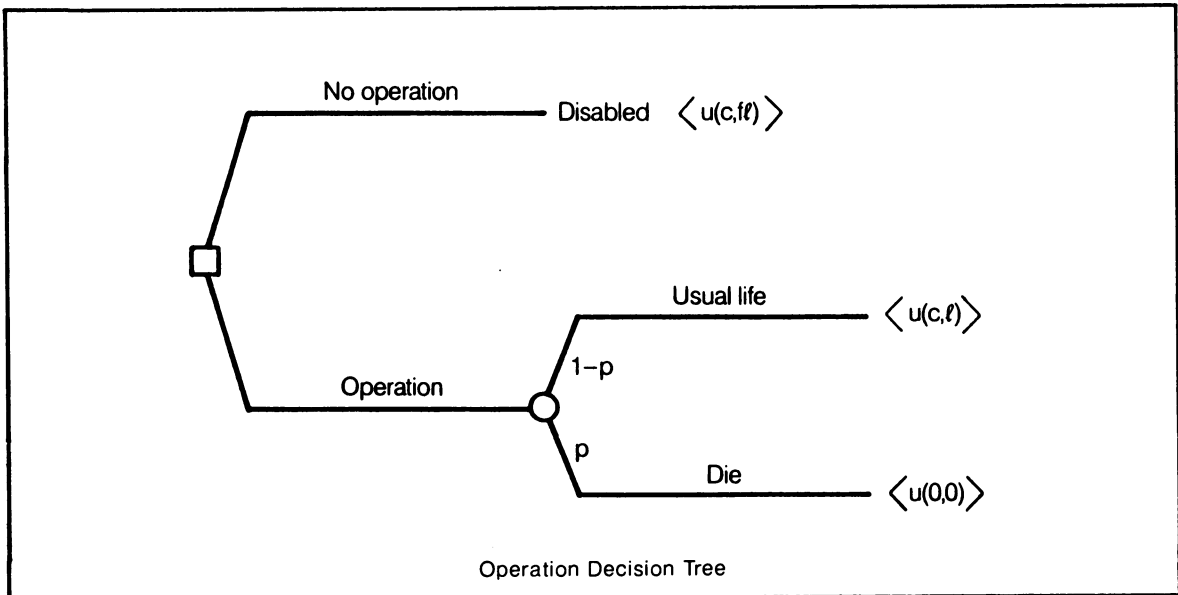


Figure 10




Figure 11

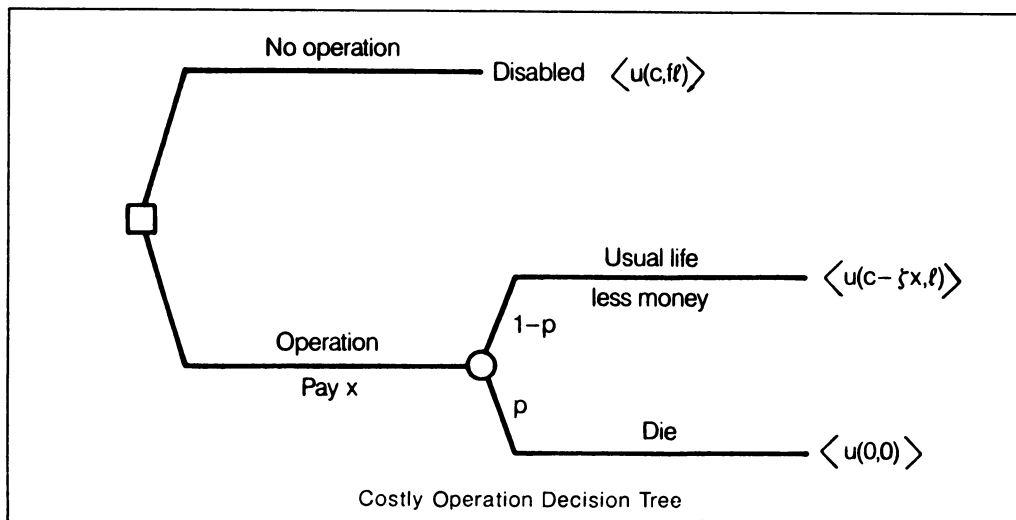


Figure 12

and compute x using $x = pv_{sh}$ when the probability p is in the safety region. Equivalently, we can use

$$v_{\mu h} = 10^{-6} v_{sh}$$

as the microhazard value of a 10^{-6} chance of combined hazard.

To illustrate these computations, suppose that the base-case individual confronts a situation where he has probabilities $p_1 = 0.00005$ of death, $p_2 = 0.00025$ of disability level $f_2 = 0.3$ (perhaps from losing a leg), and $p_3 = 0.0007$ of disability level $f_3 = 0.9$ (perhaps from losing a finger). We find that the microhazard value is \$0.606 and since $p = 1000 \times 10^{-6}$, that the required payment to accept this combined hazard would be $x = \$606$.

An equivalent and more convenient way to compute x in the safety region is to add the amounts he would have to be paid to assume each of the individual hazards,

$$x = \sum_{i=1}^n p_i v_{sd}(f_i)$$

where $v_{sd}(0)$ is equal to v_s , the small-risk value of life. If p_i is measured in microprobability (10^{-6}) units, then

$$x = \sum_{i=1}^n p_i v_{ud}(f_i) ,$$

and using the values in Table I, we find

$$\begin{aligned} x &= 50 (2.430) + 250 (1.696) + 700 (0.087) \\ &= 121.5 + 424 + 60.9 \\ &= 606. \end{aligned}$$

We see that the major contributor to the payment is the \$424 required by the threat of disability to level 0.3.

The payment to accept combined hazards in the safety region is thus easily computed by adding the payments required to accept each of them individually.

Disability or Death

Suppose an individual disabled at level f faces an operation that may cure him at some risk p of death. How large a risk of death could he tolerate as a function of his disability level? The choice is diagrammed in Figure 10. Equating expected utilities produces

$$p = \frac{\langle u(c, l) \rangle - \langle u(c, fl) \rangle}{\langle u(c, l) \rangle - \langle u(0, 0) \rangle} .$$

We can, therefore, determine the maximum death risk p that would be associated with each level of disability f . The results for the base-case individual appear in Figure 11. Note that a severely handicapped individual could tolerate very high death probabilities. This corresponds to the fact that the riskiest treatments are reserved for patients without much to lose.

We can easily add the possibility that the operation will be costly as well as dangerous. Figure 12 diagrams the case where the operation costs the patient an amount x . We note that even if $p = 0$, the operation may not be desirable. If

$$\langle u(c, fl) \rangle < \langle u(c - \zeta x, l) \rangle ,$$

then the operation costs so much that the patient considers himself worse off than he is now even if it succeeds. The maximum payment for the operation is obtained by setting the two utilities equal

$$\begin{aligned} \langle u(c, fl) \rangle &= \langle u(c - \zeta x, l) \rangle \\ w(c, fl) &= w(c - \zeta x, l) \\ c \left(\frac{fl}{l} \right)^\eta &= (c - \zeta x) \left(\frac{l}{l} \right)^\eta \end{aligned}$$

or

$$\begin{aligned} x &= \frac{c}{\zeta} [1 - f^\eta] \\ &= v_e [1 - f^\eta] . \end{aligned}$$

If $f = 0$, the individual would pay his whole economic value v_e , which is \$363,048. As f increases, the amount he would pay falls as shown in Table III..

Table III
Maximum Payment for Operation as Function of Disability Level

Disability Level	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Payment[K\$]	363	359	349	330	305	272	232	185	131	69	0

Provided we limit the payment as discussed, we can then use the tree of Figure 12 to write

$$p = \frac{\langle u(c - \zeta x, l) \rangle - \langle u(c, fl) \rangle}{\langle u(c - \zeta x, l) \rangle - \langle u(0, 0) \rangle} ,$$

and compute the maximum tolerable probability of death for an individual at disability level f facing an operation that costs x .

We have done this for the base-case individual for the cost $x = \$100,000$ --the dashed line in Figure 11. For seriously disabled people, the cost is not important, but for moderately disabled people, it is. No operation would be considered at a disability level of 0.85. At $f = 0.7$ only a 0.10 rather than a 0.16 probability of death could be tolerated.

Compensation

We can use the model even without risk to determine the amount of compensation that an individual would require to be indifferent to a given disability level. Thus, to determine the amount x he would have to be paid for being placed at disability level f , we equate

$$\langle u(c + \zeta x, fl) \rangle = \langle u(c, l) \rangle$$

$$w(c + \zeta x, fl) = w(c, l)$$

$$(c + \zeta x) \left(\frac{fl}{l} \right)^\eta = c \left(\frac{l}{l} \right)^\eta$$

or

$$x = \frac{c}{\zeta} \left(\frac{1}{f^\eta} - 1 \right) .$$

The annual consumption increase required would be

$$\zeta x = c \left(\frac{1}{f^\eta} - 1 \right) .$$

Table IV shows the results for the base-case individual. Given that his constant annual consumption is \$20,000, that amount would have to be doubled to make him indifferent to a 0.7 level of disability.

Table IV

Compensation Required by the Base-Case Individual for a Disability Level f

Disability Level f	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Lump-Sum Compensation x [K\$]	35,942	8,713	3,671	1,906	1,089	645	378	204	85	0
Annual Payment z [K\$]	1,980	480	202	105	60	36	21	11	5	0

Income Effects of Disability

Up to this point, we have assumed that the effect of disability is to decrease the desirability of the life experience, but not to change either the life expectancy or the income received by the person. Let us now relax this income assumption, since disability will usually depress income. A simple way to investigate this effect is to assume that a person's income will be reduced by the same factor f that appears in his value function. (Naturally, we could make the reduction factor different from f if necessary.) The net effect of this assumption is that in every expression we have seen where we have multiplied ℓ by f we now also multiply c by f .

Consider, for example, the effect of this change on the microdisability value $v_{\mu d}(f)$ of this change. Figure 13 shows that when this income effect is incorporated, microdisability values for the base-case individual increase substantially at every level of f except towards $f = 0$ and $f = 1$ where they are bounded respectively by the micromort value and by zero. At $f = 0.5$, the increase is from \$0.927 M to \$1.484 M. The prospect of losing both health and income is significantly more worrisome than losing health alone.




Figure 13




Figure 14

As we would expect, including the income effect makes the base-case individual more willing to accept a potentially fatal operation that could cure him. Figure 14 shows an increase in the tolerable death risk at all levels except $f = 0$ and $f = 1$. At $f = 0.5$, for example, the tolerable death risk for a free operation increases from 0.38 to 0.61.

Just as we have been able to explore the income effect rather easily within the model, so too could we incorporate the combined effects of disability on health state, income, and life expectancy.

CONCLUSION

Analyzing fates comparable to death has required only a relatively straightforward extension of the models used earlier to analyze risks of death. The idea of a small-risk value that can be used to evaluate safety decisions is directly applicable to the case of disability. Using the combined results of these analyses, an individual can evaluate safety situations where he faces both risks of death and disability.

The concept of micromort, microdisability, and microhazard values are more important than the details of the model used to derive them. Even without an analytic model, an individual could directly assign such values and then use them for decisions in the safety region. While he would not be assured of consistency with underlying preferences, he would be assured of consistency across different hazardous situations.

REFERENCES

1. R. A. Howard, "Life and Death Decision Analysis," Second Lawrence Symposium on Systems and Decision Sciences, Berkeley, California, October 1978.
2. R. A. Howard, "Life and Death Decision Analysis," Research Report EES DA-79-2, Department of Engineering-Economic Systems, Stanford University, December 1979.
3. R. A. Howard, J. E. Matheson and D. L. Owen, "The Value of Life and Nuclear Design," Proceedings of the American Nuclear Society Topical Meeting on Probabilistic Analysis of Nuclear Reactor Safety, American Nuclear Society, May 8-10, 1978.
4. R. A. Howard, "On Making Life and Death Decisions," Societal Risk Assessment, How Safe Is Safe Enough?, Edited by Richard C. Schwing and Walter A. Albers, Jr., General Motors Research Laboratories, (New York: Plenum Press, 1980), pp. 89-113.
5. R. A. Howard, "The Risks of Benefit-Cost-Risk Analysis," Risk/Benefit Analysis in Water Resources Planning and Management, Edited by Yacov Y. Haimes, (New York: Plenum Press, 1981).
6. R. A. Howard, Dynamic Probabilistic Systems, Volumes I and II, (New York: John F. Wiley and Sons, Inc., 1971).
7. S. Fanshel and J. W. Bush, "A Health Status Index and Its Application to Health Services Outcomes," Operations Research 18: 1021-1066.
8. D. Shepard and R. Zeckhauser, "Where Now for Saving Lives," Law and Contemporary Problems, Fall 1976, 5-45.
9. S. E. Bodily, "Analysis of Risks to Life and Limb," Operations Research 28 (1980): 156-175.
10. M. Jones-Lee, "The Value of Changes in the Probability of Death or Injury," Journal of Political Economics 82 (1974): 835-849.
11. M. Jones-Lee, The Value of Life (Chicago: University of Chicago Press, 1976).
12. R. L. Keeney, "Evaluating Alternatives Involving Potential Fatalities," Operations Research 28 (1980): 188-205.
13. J. Linnerooth, "The Evaluation of Public Programs Affecting Population Mortality," unpublished Ph.D. thesis, Department of Economics, University of Maryland, College Park, 1977.

PROFESSIONAL PRACTICE

Preface

These papers focus on issues and methodology that arise in the professional application of decision analysis.

"The Foundations of Decision Analysis" shows the steps taken, the models constructed, and the computations made in performing a decision analysis. It discusses the role that concepts such as stochastic dominance and value of clairvoyance play in professional practice.

"The Difficulty of Assessing Uncertainty" reports on an experiment on how well engineers assess uncertainty. The results illustrate several of the universal biases in probability assessment.

"Probability Encoding in Decision Analysis" prescribes a procedure for obtaining probability assessment that avoids common biases. The paper summarizes the results of psychological experimentation and their implications for encoding procedures.

"Risk Preference" presents the use of utility functions to capture risk-taking attitude and describes both theory and assessment procedures. The paper shows that the "delta property" axiom makes it possible to use simpler assessment procedures.

"The Development of a Corporate Risk Policy for Capital Investment Decisions" describes an early experimental effort to determine quantitatively the risk attitude of a major industrial corporation.

"The Used Car Buyer" is an extensive example of decision tree analysis requiring probability revision using Bayes' rule. The paper provides a step-by-step solution of a sequential decision problem and emphasizes qualitative interpretation and quantitative evaluation of results.

"Influence Diagrams" is a previously unpublished paper that has been circulating among students of decision analysis for several years. It defines influence diagrams, a promising new concept for treating decision problems that may supersede decision trees in both structuring and evaluating decision situations, and develops in detail an application to screening chemicals.

"The Use of Influence Diagrams in Structuring Complex Decision Problems" emphasizes the need for communication with a decision-maker to capture the probabilistic structure of a problem. It illustrates a sequential process of building the influence diagram in a hierarchy beginning with the value attributes.

THE FOUNDATIONS OF DECISION ANALYSIS

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from IEEE Transactions on
Systems Science and Cybernetics, Vol. SSC-4, No. 3,
September 1968. pp. 1-9

Copyright © 1968 by Institute of Electrical and Electronics Engineers.

The Foundations of Decision Analysis

RONALD A. HOWARD, SENIOR MEMBER, IEEE

Abstract—Decision analysis has emerged from theory to practice to form a discipline for balancing the many factors that bear upon a decision. Unusual features of the discipline are the treatment of uncertainty through subjective probability and of attitude toward risk through utility theory. Capturing the structure of problem relationships occupies a central position; the process can be visualized in a graphical problem space. These features are combined with other preference measures to produce a useful conceptual model for analyzing decisions, the decision analysis cycle. In its three phases—deterministic, probabilistic, and informational—the cycle progressively determines the importance of variables in deterministic, probabilistic, and economic environments. The ability to assign an economic value to the complete or partial elimination of uncertainty through experimentation is a particularly important characteristic. Recent applications in business and government indicate that the increased logical scope afforded by decision analysis offers new opportunities for rationality to those who wish it.

INTRODUCTION

DECISION analysis is a term that describes a combination of philosophy, methodology, practice, and application useful in the formal introduction of logic and preferences to the decisions of the world. There was a time less than a decade ago when suggesting that decision theory had practical application evoked only doubtful comment from decision makers. The past five years have shown not only that decision theory has important practical application, but also that it can form the basis for a new professional discipline, the discipline of decision analysis. Many of the professional aspects of the field have already been described in the literature (see Howard [1]). Here we shall concentrate on the rationale and methodology of decision analysis.

In discussing the rationale and philosophy of decision analysis, we shall focus on those concepts that are most unfamiliar to the intuitive decision maker. These concepts are generally concerned with the measurement of uncertainty and with the decision maker's reaction to it. In providing a methodology for decision analysis, we shall be concerned primarily with developing a procedural form that will be broad enough to cover the important areas of application.

THE RATIONALE OF DECISION ANALYSIS

The problem of the decision maker is to select a course of action in a world that is perceived as uncertain, complex, and dynamic. To follow a course of action is to make an

irreversible allocation of resources, an act that we call making a decision. Perhaps the resource whose allocation is least reversible is time, but other resources may vie for this characteristic.

Although the development of a theory of decision that comprises uncertainty, complexity, and dynamic effects is a formidable task, such a theory would not be complete, for it often turns out that what is most perplexing to the decision maker is not the mystery of his environment, but rather the specification of his own preferences. Thus we shall discuss the rationale of decision analysis by commenting on the three topics of uncertainty, structure, and preference.

Our primary interest in the topic of uncertainty is the philosophical basis for the treatment of uncertainty according to the mathematical laws of probability. The topic structure includes the complex and dynamic interactions that may exist among the many facets of a decision problem. Finally, we shall discuss under preference not only the difficulty of assigning values, but also the necessity for a value language that will be useful in a dynamic and uncertain environment.

Uncertainty

The problem of describing uncertainty has tormented philosophers for centuries. Pascal and Fermat laid the mathematical foundations of probability over three hundred years ago, and its development continues today. It might seem obvious that this theory would be the natural medium for thinking about uncertainty. However, the obvious was not proved until the present century, when it was shown that reasonable axioms for a theory of uncertainty led directly to the mathematical theory of probability.

Subjective Probability: While virtually everyone agrees on the proper use of the probability calculus, there is considerable disagreement on the interpretation of its results. Many users of probability theory consider probability to be a physical characteristic of an object as its weight, volume, or hardness. For example, they would say that a coin "has" a probability of falling heads on any toss and that to measure this probability would merely require a large number of tosses. This view of probability is called the objective interpretation.

Another group considers probability as a measure of the state of knowledge about phenomena, rather than about the phenomena themselves. This group would say that one "assigns" a probability of heads on the next toss of a coin based on all the knowledge that he has about the coin. A coin would be "fair" if, on the basis of all available evidence, there is no reason for asserting that the coin is more

Manuscript received July 2, 1968. This research was partially supported by the National Science Foundation under Grant NSF-GK-1683 and by the Office of Naval Research under Contracts ONR N00014-67-A-0112-0008 and ONR N00014-67-A-0112-0010.

The author is with the Department of Engineering-Economic Systems, Stanford University, Stanford, Calif.

likely to fall heads than tails. This view is called the subjective interpretation.

The distinction between the interpretations might seem small, but it is the key to the power of decision analysis. The objectivist requires repeatability of phenomena under essentially unchanged situations to make what he would consider to be meaningful inferences. The subjectivist can accept any amount of data, including none, and still apply logic to the decision. The objectivist was able to survive and even flourish, when the main problems of inference arose in areas such as agriculture that provide large amounts of cheap data. Today, when decisions regarding space programs must be based on a single launch of a one hundred million dollar rocket, the ability of the subjectivist to apply logic to one-of-a-kind situations has become indispensable.

These examples might lead one to believe that the subjective view of probability is modern; in fact, it was clearly held and understood by Bayes and Laplace two hundred years ago. The objectivist view is associated primarily with the founding of the British school of statistics in the early 1900's. It is the feeling of many, including decision analysts, that the creation of the field of statistics through the advent of the objective interpretation was a heresy in the development of the treatment of uncertainty. While objectivists are definitely in the majority at present, their ranks seem to be diminishing.

Subjective Probability Notation: Since the decision analyst necessarily holds the subjective viewpoint, he prefers a notation for probability that reveals that it is an assignment based on a certain set of information. Such a notation is constructed as follows: Let A be an event and S be the state of information on which the probability of the occurrence of A is to be assigned. Then $\{A|S\}$ is the symbol for the probability of A given S . If x is a random variable, then the probability density or mass function of x assigned on the basis of S is $\{x|S\}$. The expectation of x based on S is written $\langle x|S \rangle$ and is defined by

$$\langle x|S \rangle = \int_x x \{x|S\}$$

where \int_x is a general summation on x to be interpreted as a summation or integration as appropriate. The n th moment of x based on S would then be

$$\langle x^n|S \rangle = \int_x x^n \{x|S\}.$$

The variance of x is written $\sigma^2\langle x|S \rangle$ and defined by

$$\sigma^2\langle x|S \rangle = \langle x^2|S \rangle - \langle x|S \rangle^2.$$

One very special state of information is the total knowledge available at the beginning of the problem under consideration, the total prior experience denoted by \mathcal{E} . Then $\{x|\mathcal{E}\}$ would be called the prior density function on x , or the "prior" for short. The quantities $\langle x|\mathcal{E} \rangle$ and $\sigma^2\langle x|\mathcal{E} \rangle$ would then be the prior mean and variance.




Fig. 1. Problem space.

Although this notation often seems strange, it provides a mathematical language for uncertainty that describes precisely both the quantities on which the probability assignment is to be made and the state of information to be used in the assignment. The subjective view thus induces not only care in the interpretation of probability but also precision in its written expression.

Structure

The primary function of the decision analyst is to capture the relationships among the many variables in a decision problem, a process called structuring. The complexity of structure required will differ from problem to problem: from a "back-of-the-envelope" decision tree to a system of interconnected programs that tax the largest computers.

The Problem Space: A diagram like Fig. 1 is an aid in visualization. This diagram, the problem space, permits characterizing decision problems by their underlying structure. The dimensions of the problem space are degrees of uncertainty, time dependence, and complexity. Degree of uncertainty can range from the deterministic situations, where all variables are known, to the highly probabilistic situations, where little information is available about any problem variables. The time dependence can range from static to dynamic; complexity is measured in terms of the number of variables required.

Each corner of the problem space corresponds to certain mathematical models. Corner 1 is the deterministic static one-variable decision problem, such as that of finding the largest rectangular area that can be fenced with a given length of fencing. The models of elementary calculus, developed over 300 years ago, would be appropriate. Corner 2, the deterministic dynamic single-variable decision problem, would arise in elementary automatic control applications. The mathematical models of differential equations and transform calculus would be relevant; they were developed over 100 years ago. Corner 3 represents the probabilistic static single-variable problem, such as whether or not to buy life insurance. Three-hundred-year-old elementary probability would be quite helpful in reaching a decision. Corner 4 introduces complexity in the form of the deterministic static, but many-variable problem. Decision problems like assigning customers to warehouses or men to jobs provide an illustration.

One-hundred-year-old matrix algebra and 20-year-old linear optimization techniques would be very useful.

Corner 5 combines the two factors of uncertainty and dynamism in the uncertain dynamic, but single-variable problem, such as simple inventory control. Here the theory of stochastic processes and queuing models developed over the last 50 years would be most relevant. Corner 6 corresponds to the probabilistic static multivariable problem. Decision problems like bidding on new product introduction might have such an underlying structure. The mathematics of joint probability distributions would be especially helpful. Corner 7 refers to the deterministic dynamic multivariable decision problem, such as the complicated control problems posed by a space vehicle or a steel mill. Although probabilistic elements may be present, they are usually treated as perturbations of the deterministic model. The modern theory of control developed in the past three decades applies successfully to these problems.

Finally, corner 8 is the most complex corner, describing problems involving uncertainty, dynamism, and complexity. In a sense, all decision problems could be located here because they all involve the three factors to some degree. However, this corner is used to indicate problems where the three elements are indispensable to a meaningful analysis. Problems like electrical power system planning or business mergers are particular examples. Useful models might be Markov processes and their derivatives.

The extent to which formal models are available varies considerably over the problem space. Near the origin there are usually several alternative models for the problem; near corner 8 it is more a matter of patching together approximations to obtain a useful representation. As technology advances, more realistic models of uncertain, dynamic, and complex processes will be developed. However, it will continue to be the job of the decision analyst to be the engineer who matches technology to the requirements of the problem. His product is the embodiment of logic.

Preference

The problem of preference measurement is to determine in quantitative terms just what the decision maker wants.

Value: The first step is to assign a single value v to each possible outcome of the decision problem. If the problem is concerned with the allocation of monetary resources, then it is logical to measure this value in monetary terms. In business organizations, some form of profit may be appropriate. But the need for monetary values as a precedent for monetary allocation applies even if the outcome involves the loss of life or limb. As decision analysis is increasingly used in problems of social significance, a monetary value may have to be assigned to such outcomes as a cultured life or an ignorant life. Though these assignments may be very difficult, there is no rational alternative.

Time Preference: However, even in dynamic world, the preference question would not be resolved until the decision maker had stated his preference for outcomes

that are distributed in time: a preference called time preference. The importance of time preference is revealed when the analyst studies problems like the development of the national parklands or management of an individual's investment portfolio.

The phenomenon of time preference could be described as the greed-impatience tradeoff. It is characteristic of individuals and organizations that they want more now. However, the alternatives provided often give them a choice between more later or less now. Examples would be the choice between hydroelectric and gas turbine electricity production or, in general, the choice between investment in capital goods and consumer goods.

While the problem of preference is complicated, it is usually treated in decision analysis by the specification of a discount or interest rate and the rule that the alternative with the highest discounted, or present, value is to be preferred. Even within this framework, selecting the appropriate interest rate is not easy; it involves the nature of the interaction between the organization and its financial environment.

Risk Preference: The most unusual and challenging preference problem concerns preference for risk. The existence of the phenomenon is established by noting that few people are willing to bet double or nothing on next year's salary, even though the proposition is fair. Most people and organizations are averse to risk: they are willing to engage in uncertain propositions only if the expected value of the proposition is positive and relatively large. The description of this type of preference requires a set of concepts that are unusual, but logical.

To be specific in describing the concepts, it is necessary to define the technical term "lottery." A lottery is a set of prizes or prospects, one and only one of which will be received. Associated with each prize is a probability; the sum of all the probabilities is one. In many cases the prizes will each correspond to the amount of some commodity, such as money, that will be received. In these cases we can think of the lottery as a random variable described by either a probability mass or probability density function.

Utility theory: The most common structure for encoding risk preference requires that the individual subscribe to a set of axioms concerning lotteries. The first is that he must be willing to provide a transitive rank ordering of all prizes in any lottery. That is, if the prizes in a lottery are A , B , and C , he must be able to say in what order he prefers the prizes; further, if he prefers A to B and B to C , then he must prefer A to C .

The second axiom is that if he says he prefers A to B to C then there must exist a value of p such that he is indifferent between receiving B for certain and participating in a lottery that produces A with probability p and C with probability $1 - p$. When the appropriate value of p has been found, we would say that B is the certain equivalent of the lottery on A and C .

The third axiom is that if he prefers prize A to prize B and if he is presented with two lotteries, each offering A




Fig. 2. Utility curve.

and *B* with different probabilities, then he must prefer the lottery that yields *A* with the higher probability.

These axioms are the most significant ones. However, two others are necessary for completeness. One is that a certain equivalent of a lottery may be substituted for the lottery in any situation without changing the preferences of the decision maker; we might call this a "did you really mean it?" axiom. The other is that a lottery whose prizes are themselves lotteries is equivalent to a lottery that produces the same ultimate prizes with probability computed according to the laws of probability; this could be termed a "no fun in gambling" axiom.

Mathematical arguments reveal that an individual who subscribes to these axioms can encode his risk preference in terms of a function on the prizes of the lotteries, a function called a utility function. The utility function has two important properties: first, that the utility of any lottery is the expected utility of its prizes; second, that if one lottery is preferred to another by the individual, then its utility will be higher.

Thus the utility function assigns to any lottery a real number; the lotteries will be preferred in the order of these numbers. However, the actual magnitude of the utility is not important, because the preferences revealed by the utility function are unchanged if the utility function is modified by multiplication by a positive constant or by addition of any constant. Thus the utility function serves as a risk preference thermometer that can be used for ranking lotteries according to the risk preference of an individual.

In problems of professional interest the lottery prizes are usually measured in a commodity such as money. In this case the utility function can be represented by a curve that shows the utility to be assigned to any amount of the commodity. Such a utility curve appears as Fig. 2. The curve $\langle u|v \rangle$ shows the utility *u* assigned by some individual to amounts of money *v* between 0 and 100 dollars. Be-

cause of the invariance to linear transformation, the scale of measurement can be selected arbitrarily; this curve assigns a utility of 0 to 0 dollars and a utility of 1 to 100 dollars.

The two lotteries below the curve show how it is used. The expected value of a lottery *L* is defined in our notation by

$$\langle v|L \rangle = \int v \{v|L \rangle$$

Lottery *L*₁ has an expected value of 38 dollars; *L*₂, an expected value of 36 dollars. Someone who was indifferent to risk would prefer *L*₁. However, to determine the preference of the individual with the utility function in Fig. 2, we first determine the utility of each prize in each lottery from the utility curve and then find the expected value of the utility. The expected utility of a lottery is given by

$$\langle u|L \rangle = \int \langle u|v \rangle \{v|L \rangle$$

Since the expected utility of lottery *L*₁ is 0.44, while that of lottery *L*₂ is 0.51, the individual would prefer lottery *L*₂, in spite of its lower expected value. We would describe individuals whose utility curves are concave downwards as risk averse.

The certain equivalent: Although this calculation serves to determine the individual's preference, it gives us no feeling about the strength of the preference. The magnitude of the utility can be no help because we see that if we added 10 to all utility numbers, we would derive exactly the same preference ordering but with much smaller percentage difference in utility numbers. To measure strength of preference, it is helpful to return to the concept of certain equivalent.

To evaluate a lottery in a single meaningful monetary number, we can ask what amount of money received for certain would have the same utility as the lottery. The certain equivalent of a lottery *L*, denoted by $\bar{z}|L \rangle$, is thus the amount of money shown by the utility curve to have the same utility as the lottery. The certain equivalent is mathematically defined by the equation

$$\langle u|\bar{z}|L \rangle = \langle u|L \rangle$$

Thus from the curve we see that the utility of 0.44 for lottery *L*₁ corresponds to a certain equivalent of 28 dollars, while the utility of 0.51 for lottery *L*₂ would mean a certain equivalent of 34 dollars. The individual would be just indifferent between receiving either 28 dollars for certain or lottery *L*₁ and between receiving 34 dollars for certain or lottery *L*₂. It would be slightly inaccurate, but intuitively satisfying, to say that lottery *L*₂ is worth 6 dollars more to the individual than is lottery *L*₁.

Exponential utility curves: In some cases the individual is willing to subscribe to a sixth axiom: that if all prizes in a lottery are increased by any amount Δ , the certain equivalent of the lottery will also increase by Δ . The axiom is persuasive, since the increment Δ will be received with certainty regardless of the outcome of the lottery. However,

the axiom is very powerful, for someone who subscribes to it must have a utility curve that is linear or exponential in form; that is, $\langle u|v \in \mathcal{E} \rangle$ is proportional either to v or to $e^{-\gamma v}$. Furthermore, the curve is completely described by the constant γ called the risk aversion coefficient. Although few individuals may in fact wish to be governed by this axiom, the exponential utility curve is very useful in analyses, as we shall see.

Stochastic dominance: There is one important case in which risk preference need not be measured at all. That is the case in which the choice between two alternatives would be clear to a rational man regardless of his risk preference; it is called the case of stochastic dominance. Lottery L_1 stochastically dominates lottery L_2 if the probability of receiving a monetary return in excess of c is higher for L_1 than for L_2 for any value of c ; that is,

$$\{v > c|L_1\} > \{v > c|L_2\}, \quad -\infty < c < \infty.$$

If one lottery stochastically dominates all others, then it will be preferred by the individual regardless of his attitude toward risk; there is no need to use the utility function.

Joint Time-Risk Preference: Individuals often have to choose between monetary rewards that are not only uncertain, but distributed over time. In these situations time and risk preference must be jointly encoded. The description of joint time-risk preference is a problem that admits many solutions. Here we shall employ the idea of reducing any time stream of value to a present value using the time preference measure and then applying the utility function to determine which lottery on present values is most desirable.

THE METHODOLOGY OF DECISION ANALYSIS

With this background we can go on to a discussion of how a decision problem can be progressively analyzed using decision analysis principles. The procedure is best explained in terms of a diagram like that in Fig. 3. Here we view the decision analysis procedure as divided into three major phases, the deterministic, probabilistic, and informational phases. The deterministic phase establishes the deterministic relationships among the variables of the problem. The probabilistic phase introduces uncertainty and risk preference. Finally, the informational phase determines the economic value of gathering more information. Following these phases, a decision is required on whether to act or to gather new information. If additional information is obtained, e.g., through market testing or building a pilot plant, then this information must be incorporated into the structure and probability assignments of the problem; the cycle is then repeated.

The decision analysis cycle is a convenient conceptual model rather than an inevitable method for analyzing decision problems. With this point in mind, we shall now examine the steps required in each phase.

The Deterministic Phase

The first step in the deterministic phase is to construct a deterministic model of the decision problem.

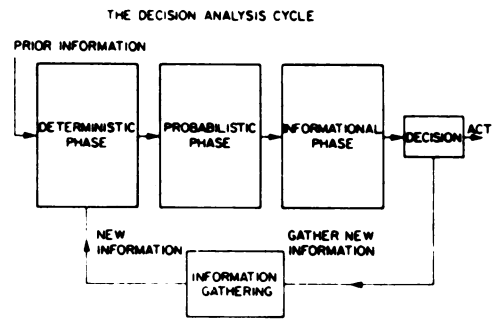


Fig. 3. Decision analysis cycle.




Fig. 4. Deterministic model.

The Deterministic Model: Fig. 4 is an abstract representation of the model. The model relates the important variables in the problem that are not under the control of the decision maker and the variables that are under his control to the production of value in time. These variables are called the state variables s_i and decision variables d_i . We can visualize the state variables as a set of knobs on the model that are set by a disinterested nature; the decision variables are knobs set by the decision maker. Fig. 4 shows that the values developed over time $v(0), v(1), v(2), \dots$ are operated upon by the time preference specification to produce a present value reading v that we may regard as appearing on a value meter. Thus any setting of the state- and decision-variable knobs will produce a value reading. The deterministic model will generally be realized in the form of a computer program.

Deterministic Sensitivity: Fig. 5 shows the first analytical step in the deterministic phase, the measurement of deterministic sensitivity. In the representation of Fig. 5 the time preference measure is shown incorporated into the deterministic model to produce a single present value reading. The analysis begins by assigning each state variable a nominal value and a range that might correspond to the 10- and 90-percent point on its marginal cumulative probability distribution. Decision variables would also be assigned nominal values and ranges to reflect initial feelings about what the best decision might be.

With all variables but one set to their nominal values, that one variable would be swept across its range to deter-

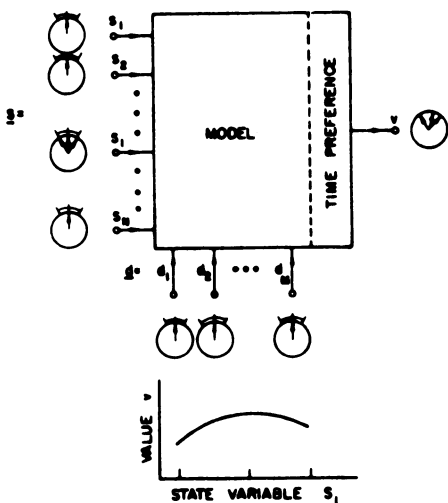


Fig. 5. Deterministic sensitivity.

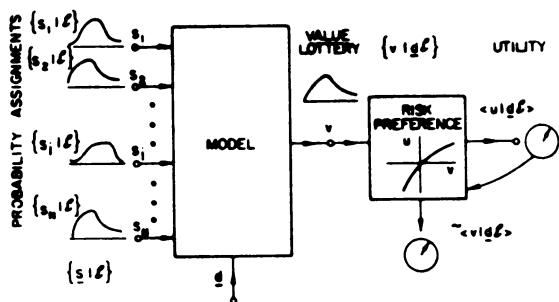


Fig. 6. Value lottery.

mine the effect on the value reading. The figure shows the measurement for the i th state variable s_i . State or decision variables that showed high sensitivity would be retained in the further analyses of the model. A variable could show a high deterministic sensitivity because of its wide range, crucial nature, or a combination of these effects.

In some problems this one-at-a-time type of sensitivity analysis will not be sufficient: the joint sensitivity of variables will have to be measured by sweeping more than one variable at a time over their ranges. Because the number of possibilities for joint sensitivity increases combinatorially with the number of variables, the analyst must use judgment in determining where joint sensitivity measurements will be required.

The net effect of the deterministic sensitivity analysis will be to determine the state variables and decision variables that have a major effect on value. The next step will be to introduce the current state of knowledge on uncertainty in the state variables and determine which decision would be best, given the uncertainty; this is done in the probabilistic phase.

The Probabilistic Phase

The probabilistic phase requires assignment of probability distributions on the state variables.

The Value Lottery: Fig. 6 shows this assignment as a marginal probability distribution $\{s_i|\mathcal{E}\}$ on each state variable. Since the state variables will generally be joint y related, the complete description of the state of knowledge about them would be the joint probability distribution $\{s_1, s_2, \dots, s_N|\mathcal{E}\} = \{\mathbf{s}|\mathcal{E}\}$, but the marginal distributions shown will serve as a pictorial representation. The settings of the decision variables are summarized by the decision vector $\mathbf{d} = [d_1, d_2, \dots, d_M]$. For any setting \mathbf{d} the joint distribution $\{\mathbf{s}|\mathcal{E}\}$ on the state variables will imply a probability distribution on the value, $\{v|\mathbf{d}\mathcal{E}\}$, a distribution we call the value lottery. The decision problem then reduces to finding the setting \mathbf{d} that produces the most desirable value lottery.

The determination of the value lottery corresponding to any decision vector \mathbf{d} will be performed by analytical or simulation methods, as appropriate. Efficient search procedures are helpful in establishing the best setting for \mathbf{d} .

Risk Preference: There remains the question of which value lottery is best. Perhaps the question will be easily resolved by the observation that one setting of \mathbf{d} produces a value lottery that stochastically dominates the lotteries produced by all other settings. But if not, then it will be necessary to encode the risk preference of the decision maker in a utility curve. This curve will allow each value lottery and hence each setting of \mathbf{d} to be rated by its utility. The setting that produces the highest utility $\langle u|\mathbf{d}\mathcal{E}\rangle$ would then be judged the best. To gain intuitive meaning, the utility of each lottery could be returned to the utility curve to show the certain equivalent value $\sim\langle r|\mathbf{d}\mathcal{E}\rangle$ implied by the decision setting \mathbf{d} .

This procedure establishes the setting of the decision variables $\mathbf{d}(\mathcal{E})$, that is most desirable to the decision maker in view of his state of knowledge regarding uncertainties and his risk preferences,

$$\mathbf{d}(\mathcal{E}) = \max_{\mathbf{d}}^{-1} \langle u|\mathbf{d}\mathcal{E}\rangle = \max_{\mathbf{d}}^{-1} \sim\langle r|\mathbf{d}\mathcal{E}\rangle.$$

Furthermore, it shows the utility $\langle u|\mathcal{E}\rangle$ and certain equivalent $\sim\langle r|\mathcal{E}\rangle$ of the best decision,

$$\begin{aligned} \langle u|\mathcal{E}\rangle &= \langle u|\mathbf{d} = \mathbf{d}(\mathcal{E})\mathcal{E}\rangle \\ \sim\langle r|\mathcal{E}\rangle &= \sim\langle r|\mathbf{d} = \mathbf{d}(\mathcal{E})\mathcal{E}\rangle. \end{aligned}$$

In a sense, this step completes the solution of the decision problem. However, since decision analysis is more engineering than mathematics, the procedure does not stop here, but rather continues to the measurement of another kind of sensitivity, stochastic sensitivity.

Stochastic Sensitivity: The idea behind stochastic sensitivity is the desire to measure the effect of a variable on the result of the decision problem not in the deterministic environment where all other variables are set to their nominal values, but in the probabilistic environment where all other variables are governed by their appropriate probability distributions. As Fig. 7 shows, if the i th state variable s_i were known, the other state variables would be governed by the conditional distribution $\{\mathbf{s}|\mathbf{s}_i, \mathcal{E}\}$ obtained by dividing $\{\mathbf{s}|\mathcal{E}\}$ by $\{s_i|\mathcal{E}\}$. Thus the specification of any




Fig. 7. Stochastic sensitivity.




Fig. 8. Clairvoyance.

value for s_i would imply some joint probability distribution of the remaining state variables, and in turn a value lottery $\langle v|s_i, \mathbf{d}, \mathcal{E} \rangle$ for the given setting of the decision vector. The risk preference encoding would describe this value lottery by a certain equivalent $\tilde{v}(v|s_i, \mathbf{d}, \mathcal{E})$.

Suppose now that the decision vector \mathbf{d} is adjusted to the value $\mathbf{d}(s_i, \mathcal{E})$ that produces the highest certain equivalent for this value of s_i , $\max_{\mathbf{d}} \tilde{v}(v|s_i, \mathbf{d}, \mathcal{E})$, that is,

$$\mathbf{d}(s_i, \mathcal{E}) = \max_{\mathbf{d}}^{-1} \langle u|s_i, \mathbf{d}, \mathcal{E} \rangle = \max_{\mathbf{d}}^{-1} \tilde{v}(v|s_i, \mathbf{d}, \mathcal{E}).$$

If this procedure is repeated for the various values of s_i within its range, the plot of $\max_{\mathbf{d}} \tilde{v}(v|s_i, \mathbf{d}, \mathcal{E})$ will show the stochastic sensitivity of the variable s_i .

Stochastic sensitivity shows how the certain equivalent of the decision problem depends on a particular state variable when all other state variables are uncertain. Stochastic sensitivity can be measured in a different sense if, rather than choosing the best decision variable setting \mathbf{d} for each s_i , the setting $\mathbf{d}(\mathcal{E})$ that was best for $\{s|\mathcal{E}\}$ is used throughout. This technique measures the stochastic sensitivity to the i th state variables under the original decision rule rather than under a decision rule adjusted to take advantage of knowledge of s_i . Stochastic sensitivity to a decision variable d_i can be measured by using the probability assignment $\{s|\mathcal{E}\}$ for the state variables and then seeing how the certain equivalent changes with d_i , either with other decision variables fixed or continually optimized.

The problems of joint sensitivity measurement arise just as they did in the case of deterministic sensitivity. However, here the cost of joint sensitivity measurement is even greater than before because of the need to develop lotteries on value rather than single numbers.

Stochastic sensitivity can provide important additional insight into problem relationships. It can show the need for further structure to allow available information to be encoded more effectively. It might reveal that variables originally thought to be of vital importance on the basis of deterministic analyses are relatively unimportant in the probabilistic environment. At a minimum, it yields a useful measurement of the robustness of the indicated decision.

The Informational Phase

The probabilistic phase of the analysis provides further insight into the importance of uncertainty in state variables, but it stops short of what we would really like to know, namely, what is the worth in monetary terms of the various forms of uncertainty remaining in the problem? The informational phase covers this last step of measuring economic sensitivity and hence indicates what sort of additional information could be economically gathered.

Clairvoyance: A useful concept in discussing the informational phase will be the clairvoyant. The clairvoyant is an individual who can tell us the precise value of any uncertain variable. Clearly, such help would be valuable, but how valuable?

Fig. 8 illustrates the case where we have engaged the clairvoyant to tell us the value of the i th state variable s_i at a cost k_{s_i} . Knowing s_i will have two effects on the result. First, the probability assignments on the other state variables will be governed by $\{s_i|s_i, \mathcal{E}\}$. Second, whatever present value v is produced will have to be reduced by the clairvoyant's charge k_{s_i} to a net present value v' . Once s_i is reported, the best setting $\mathbf{d}(s_i, k_{s_i}, \mathcal{E})$ of the decision vector will be the setting that produces a net present value lottery having the highest utility. Thus

$$\mathbf{d}(s_i, k_{s_i}, \mathcal{E}) = \max_{\mathbf{d}}^{-1} \langle u|s_i, k_{s_i}, \mathbf{d}, \mathcal{E} \rangle = \max_{\mathbf{d}}^{-1} \int_{s_i} \langle u|s_i, k_{s_i}, \mathbf{d}, \mathcal{E} \rangle \{s_i|s_i, \mathcal{E}\}$$

and

$$\max_{\mathbf{d}} \langle u|s_i, k_{s_i}, \mathbf{d}, \mathcal{E} \rangle = \langle u|s_i, k_{s_i}, \mathbf{d}(s_i, k_{s_i}, \mathcal{E}), \mathcal{E} \rangle.$$

Therefore, if we knew that the clairvoyant would report a particular value of s_i , the utility of the resulting lottery would be $\langle u|s_i, k_{s_i}, \mathbf{d}(s_i, k_{s_i}, \mathcal{E}), \mathcal{E} \rangle$. However, we are not sure that he will report that value; indeed, if we were sure, there would be no point in employing him. Consequently, we must weight the utility we shall derive if he reports a value of s_i by the probability that he will report that value in order to determine the utility $\langle u|k_{s_i}, \mathcal{E} \rangle$ of the lottery we enter by engaging him. The probability we assign to his reporting any value of s_i is, of course, just $\{s_i|\mathcal{E}\}$ since he is assumed competent and trustworthy. Therefore,

$$\langle u|k_{s_i}, \mathcal{E} \rangle = \int_{s_i} \langle u|s_i, k_{s_i}, \mathbf{d}(s_i, k_{s_i}, \mathcal{E}), \mathcal{E} \rangle \{s_i|\mathcal{E}\}.$$

If the cost of the clairvoyant k_{s_i} were equal to zero, we would expect this utility $\langle u|k_{s_i} = 0 \mathcal{E} \rangle$ to be greater than the utility $\langle u|\mathcal{E} \rangle$ of the best lottery without clairvoyance. However, as the cost of the clairvoyant increases, his service will become progressively less desirable until the utility of the lottery with clairvoyance is just equal to the utility of the best lottery without clairvoyance. The value of k_{s_i} that satisfies the equation

$$\langle u|k_{s_i} \mathcal{E} \rangle = \langle u|\mathcal{E} \rangle$$

is called the value of clairvoyance about the variable s_i .

The value of clairvoyance on a variable is an important quantity because it represents the largest amount that one should pay to eliminate completely uncertainty regarding the variable. Since most real information gathering opportunities provide less than perfect information, they should never be employed when their cost exceeds the cost of clairvoyance.

Notice that the actual availability of a clairvoyant is irrelevant to this argument. The clairvoyant in decision analysis plays exactly the same role as the Carnot engine in thermodynamics: a conceptual reference against which to compare the performance of physically realizable alternatives.

As with sensitivity measurement, the value of simultaneous clairvoyance on several variables can also be calculated with somewhat more difficulty. In the preceding argument, s_i would be replaced by a subset of state variables, but the nature of the calculations remains the same. Even if the state variables are independent, the value of clairvoyance on several of them can differ from the sum of the values of clairvoyance on each separately. (See [4], [5].)

The value of clairvoyance on any state variable or set of state variables will depend on the prior distribution $\{s|\mathcal{E}\}$. It is clear that some prior distribution will maximize the value of clairvoyance; we might call this the maximum value of clairvoyance. It is the value of clairvoyance to a decision maker who had the most unfortunate initial state of information as far as purchasing clairvoyance is concerned. The calculation is useful because it shows the most that anyone should pay for clairvoyance regardless of his state of information. Of course, the calculation is predicated on a given time and risk preference.

Experimentation: The real-world approximation to clairvoyance is some form of experimentation. An important question in guiding the gathering of additional information is, therefore, the value of a given experiment. The calculation follows almost the same form as the computation of the value of clairvoyance.

Fig. 9 illustrates the nature of the calculation. Suppose that the experiment costs $k_{\mathcal{E}}$ and that after it was conducted, it produced the data D . Knowledge of D would change the probability distribution on s to $\{s|D\mathcal{E}\}$, which is related to the prior distribution $\{s|\mathcal{E}\}$ by Bayes' equation,

$$\{s|D\mathcal{E}\} = \frac{\{D|s\mathcal{E}\} \{s|\mathcal{E}\}}{\{D|\mathcal{E}\}}$$

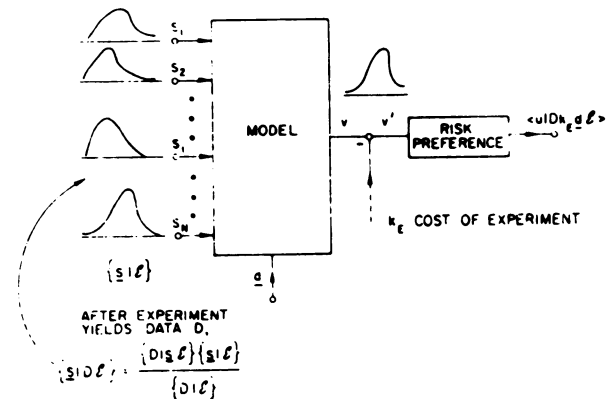


Fig. 9. Experimentation.

The new quantities $\{D|s\mathcal{E}\}$ and $\{D|\mathcal{E}\}$ are interesting in themselves. The quantity $\{D|s\mathcal{E}\}$ is the probability of observing the particular data D for any setting of the state variables; it is called the likelihood function. The quantity $\{D|\mathcal{E}\}$ is the probability of observing D assigned before the experiment is performed; it is related to the likelihood function and the prior by

$$\{D|\mathcal{E}\} = \int_s \{D|s\mathcal{E}\} \{s|\mathcal{E}\}$$

and is called the preposterior distribution.

Once D is known, the best setting $d(Dk_{\mathcal{E}}\mathcal{E})$ of the decision vector will be the setting that produces the net present value lottery of highest utility,

$$\begin{aligned} d(Dk_{\mathcal{E}}\mathcal{E}) &= \max_d \langle u|Dk_{\mathcal{E}}d\mathcal{E} \rangle \\ &= \max_d \int_s \langle u|sk_{\mathcal{E}}d\mathcal{E} \rangle \{s|D\mathcal{E}\}. \end{aligned}$$

The utility of this lottery will be

$$\max_d \langle u|Dk_{\mathcal{E}}d\mathcal{E} \rangle = \langle u|Dk_{\mathcal{E}}d(Dk_{\mathcal{E}}\mathcal{E}) \rangle.$$

However, this utility will be received conditional on the reporting of D . The probability that D will be reported by the experiment is the preposterior probability $\{D|\mathcal{E}\}$. Therefore, the overall utility of the experiment at a cost $k_{\mathcal{E}}$, $\langle u|k_{\mathcal{E}}\mathcal{E} \rangle$, will be just

$$\langle u|k_{\mathcal{E}}\mathcal{E} \rangle = \int_D \langle u|Dk_{\mathcal{E}}d(Dk_{\mathcal{E}}\mathcal{E}) \rangle \{D|\mathcal{E}\}.$$

The number $k_{\mathcal{E}}$ that satisfies the equation

$$\langle u|k_{\mathcal{E}}\mathcal{E} \rangle = \langle u|\mathcal{E} \rangle$$

and thus makes the utility of the best lottery with the experiment equal to the utility of the best lottery without the experiment is the value of the experiment.

Comparing this calculation with the one for the value of clairvoyance shows that we can interpret clairvoyance as a very special kind of experiment: one that completely eliminates uncertainty in one or several state variables.

Once the value of the experiment has been computed, it can be compared with its real-world cost. Experiments whose value exceeds their cost are profitable alternatives for the decision maker; others are not. Determining the profitability of various information gathering plans shows which, if any, should be pursued before the primary decision is made.

The Decision Analysis Cycle

This discussion of the decision analysis cycle has indicated most, but not all, of the types of analyses that may be useful. For example, determining sensitivity of the best decision and its present value to the discount rate representing time preference would be an obvious test to perform. In some decision problems, particularly those requiring the consensus of several interested parties, it may be wise to measure risk sensitivity. This would involve seeing how the best decision and its certain equivalent value change as the risk aversion coefficient is increased. Fortunately, it often happens that the same policy remains best for a range of risk coefficients that includes those of all participants. In these cases, there is no point in argument over just what attitude toward risk should govern the decision.

Division of Effort: The total effort devoted to the cycle is not typically equally divided among the phases. Because of the need for a detailed understanding of fundamental problem relationships, the deterministic phase requires about 60 percent of total effort. The probabilistic phase might receive 25 percent; the informational phase, the remaining 15 percent. As the analysis progresses through the phases, the nature of the work changes from the construction and tuning of the model to the development of insight by exercising it.

Computational Demands. The difficulty of exercising the model changes from phase to phase. For example, a computer run to establish stochastic sensitivity might require ten times as much time as a run to measure deterministic sensitivity. Similarly, an economic sensitivity run in the informational phase might require ten times as much computation as the measurement of stochastic sensitivity. Thus we see the need for the continued screening of variables to assure that only important factors are retained in each phase of the analysis. To think of performing a decision analysis by including all possibly relevant variables in each phase would be very unrealistic.

The Model Sequence: Typically, a decision analysis is performed not with one, but with a sequence of progressively more realistic models. The first model in the sequence we call the pilot model; it is an extremely simplified representation of the problem, useful only for determining the most important relationships. Its aeronautical counterpart would be the wind tunnel model of a new airplane. It looks

very little like the desired final product, but it is indispensable in achieving that goal. Perhaps 20 percent of total effort might be devoted to construction and testing of the pilot model.

The next model in the sequence is called the prototype model. It is a quite detailed representation of the problem that may, however, still be lacking a few important attributes. Its aeronautical analogy would be the first flying model of a new airplane. While it will generally have bugs that must be eliminated, it does demonstrate overall appearance and performance of the final version. Because of the need for verisimilitude of the prototype model, it might require 60 percent of the total effort.

The final model in the sequence is the production model; it is as accurate a representation of reality as will be produced in the decision analysis. Like the production airplane, it should function well even though it may retain features that are treated in a less-than-ideal way. Perhaps 20 percent of the total effort might be devoted to this final stage of model development. When completed, the production model should be able to withstand the test of any good engineering design: additional modeling resources could be utilized with equal effectiveness in any part of the model.

It would be unrealistic to expect the decision analysis of any large problem to employ all the phases, sensitivity analyses, and models that we have discussed. However, having the concepts and nomenclature necessary to depict these steps is a powerful aid in the planning and execution of a decision analysis. The future should bring continual refinements in the theory and application of the methodology.

CONCLUSION

The last few years have seen decision analysis grow from a theorist's toy to an important ally of the decision maker. Significant applications have ranged from the desirability of kidney transplants through electric power system planning to the development of policies for space exploration. No one can say when the limits of this revolution will be reached. Whether the limits even exist depends more on man's psychology than on his intellect.

REFERENCES

- [1] R. A. Howard, "Decision analysis: applied decision theory," *Proc. 4th Internat'l Conf. on Operational Research* (Boston, Mass., 1966).
- [2] —, "Bayesian decision models for system engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-1, pp. 36-40, November 1965.
- [3] —, "Dynamic inference," *Operations Res.*, vol. 13, pp. 712-733, September-October 1965.
- [4] —, "Information value theory," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 22-26, August 1966.
- [5] —, "Value of information lotteries," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-3, pp. 54-60, June 1967.

Reprinted from IEEE TRANSACTIONS ON SYSTEMS SCIENCE AND CYBERNETICS
Vol. SSC-4, No. 3, September 1968
pp. 211-219

Copyright 1968, and reprinted by permission of the copyright owner
PRINTED IN THE U.S.A.

THE DIFFICULTY OF ASSESSING UNCERTAINTY

E. C. Capen
ARCO Exploration Company



The Difficulty of Assessing Uncertainty

E. C. Capen,
ARCO Exploration Company

Introduction

The good old days were a long time ago. Now, though we must harness new technology and harsh climates to help provide needed energy supplies, we are also faced with the complex problem of satisfying not altogether consistent governments, the consumer, our banker, and someone's time schedule. Judging from the delays, massive capital overruns, and relatively low return this industry has experienced lately, it would seem that we have been missing something. At least one explanation is that we have not learned to deal with uncertainty successfully.

Some recent testing of SPE-AIME members and others gives rise to some possible conclusions:

1. A large number of technical people have little idea of what to do when uncertainty crosses their path. They are attempting to solve 1976 problems with 1956 methods.

2. Having no good quantitative idea of uncertainty, there is an almost universal tendency for people to understate it. Thus, they overestimate the precision of their own knowledge and contribute to decisions that later become subject to unwelcome surprises.

A solution to this problem involves some better understanding of how to treat uncertainties and a realization that our desire for preciseness in such an unpredictable world may be leading us astray.

Handling Uncertainty

Our schooling trained us well to handle the certainties of the world. The principles of mathematics and physics work. In Newton's day, force equaled mass times ac-

celeration, and it still does. The physicists, when they found somewhat erratic behavior on the atomic and molecular level, were able to solve many problems using statistical mechanics. The extremely large number of items they dealt with allowed these probabilistic methods to predict behavior accurately.

So we have a dilemma. Our training teaches us to handle situations in which we can accurately predict the variables. If we cannot, then we know methods that will save us in the presence of large numbers. Many of our problems, however, have a one-time-only characteristic, and the variables almost defy prediction.

You may embark on a new project whose technology differs from that used on other projects. Or perhaps your task is to perform a familiar project in a harsh environment. Try to estimate the total cost and completion time. Hard! You cannot foresee everything. And, for some reason, that which you cannot foretell seems to bring forth more ill than good. Hence, the predictions we make are often very optimistic. Even though we see the whole process unfolding and see estimate after estimate turn out optimistic, our next estimate more than likely will be optimistic also.

What happens? Is there some deep psychological phenomenon that prevents our doing better? Because we are paid to know, do we find it difficult to admit we do not know? Or can we obtain salvation through knowledge? As we were trained to handle certainty, can we also find a better way to estimate our uncertainty?

I think so, but it will take some special effort — just as it did when we first learned whatever specialty that

What do you do when uncertainty crosses your path? Though it seems that we have been taught how to deal with a determinate world, recent testing indicates that many have not learned to handle uncertainty successfully. This paper describes the results of that testing and suggests a better way to treat the unknown.

got us into the business. As one of the Society's Distinguished Lecturers for 1974-75, I had a unique opportunity to collect information on the way our membership treats uncertainty. I do not claim that what you are about to read will set the scientific or business communities to quaking (others have noticed similar phenomena before¹). But there are lessons that should help to improve our perceptions of uncertainty and, we hope, increase our economic efficiency by giving us better information on which to base decisions.

SPE-AIME Experiment

The experiment went like this. Each person put ranges around the answers to 10 questions, ranges that described his personal uncertainty. The questions were the following:

1. In what year was St. Augustine (now in Florida) established as a European settlement?
2. How many autos were registered in California in 1972?
3. What is the air distance from San Francisco to Hong Kong in miles?
4. How far is it from Los Angeles to New Orleans via major highways in miles?
5. What was the census estimate of U.S. population in 1900?
6. What is the span length of the Golden Gate Bridge in feet?
7. What is the area of Canada in square miles?
8. How long is the Amazon River in miles?
9. How many earth years does it take the planet Pluto to revolve around the sun?
10. The English epic poem "Beowulf" was composed in what year?

For some, the task was to put a 90-percent range around each answer. The person would think up a range such that he was 90-percent sure the range would encompass the true value. For example, in one section a gentleman put a range of 1500 to 1550 on Question 1. He was 90-percent sure that St. Augustine was established after 1500, but before 1550. In his view, there was only a 5-percent chance that the settlement came into being after 1550. If he were to apply such ranges for many questions, we would expect to find about 10 percent of the true answers outside of his intervals.

Other groups were asked to use 98-percent ranges — virtual certainty that their range would encompass the true value. I also asked for ranges of 80, 50, and 30 percent. The 30-percent interval would supposedly allow 70 percent of the true answers to fall outside the range.

Most sections used a single probability range. However, a few groups were divided in two, with each half using different intervals, usually 30 and 90 percent. I shall refer to these ranges as probability intervals.

You may want to test your skill on the test, too. The answers are in the Appendix. Use a 90-percent interval so you can compare with results given later.

Results and Conclusions

My testing turned up traits that should be of interest. [From this point on, the people referred to are the 1,200+ people at the local section meetings who answered the questions sufficiently to be counted. There

were a significant number (350 or so) at the meetings who either had no idea of how to describe uncertainty or thought it chic not to play the game.]

1. People who are uncertain about answers to a question have almost no idea of the degree of their uncertainty. They cannot differentiate between a 30- and a 98-percent probability interval.

2. The more people know about a subject, the more likely they are to construct a large probability interval (that is, one that has a high chance of catching the truth), regardless of what kind of interval they have been asked to use. The converse seems to hold also: the less known, the smaller the chance that the interval will surround the truth.

3. People tend to be a lot prouder of their answers than they should be.

4. Even when people have been previously told that probability ranges tend to be too small, they cannot bring themselves to get their ranges wide enough, though they do somewhat better.

5. Simultaneously putting two ranges on the answers greatly improved performance, but still fell short of the goal.

Such conclusions come from the following observations. Looking at the data collected on each of the sections, we find that the average number of "missed" questions was close to 68 percent. We could adopt the following hypothesis:

SPE-AIME sections will miss an average 68 percent of the questions, no matter what probability ranges they are asked for.

Mathematical statisticians have invented a way to test such hypotheses with what they call confidence intervals. They recognize, for instance, that the Hobbs Petroleum Section average of 6.26 misses out of 10 questions is subject to error. Slightly different questions, a different night, a longer or shorter bar — all kinds of things could conspire to change that number. By accounting for the variability of responses within the Hobbs chapter and the number of data points that make up the average, these statistical experts can put a range around the 6.26 much like the ranges the members were asked to use. Except that (unlike the members) when the statistician says he is using a 95-percent range, he really is!

For Hobbs, that range comes out to be 5.45 to 7.07. Since that range includes 6.8, or 68-percent misses, the statistician will agree that, based on his data, he would not quarrel with the hypothesis as it applies to Hobbs.

Table 1 shows all the 95-percent ranges and Fig. 1 illustrates how these ranges compare with the 68-percent hypothesis. You will see a portion of the Los Angeles Basin Section whose confidence interval (5.24 to 6.68) does not include 6.8. There are three possible explanations:

1. The group has a bit more skill at handling such a problem than most.

2. Being part of an audience that was asked to use two different ranges, there was a more conscious effort on their part to use a wider range.

3. The statistics are misleading, and the group is not different from the others. We expect this to happen about 5 percent of the time. (Our testing mechanism

TABLE 1 — SUMMARY OF 95-PERCENT RANGES

SPE-AIME Section	Number of Usable Responses	Requested Range (percent)	Expected Number of Misses	Actual Number Average Misses	95-Percent Confidence Interval
Hobbs Petroleum	34	98	0.2	6.26	5.45 to 7.07
Oklahoma City	111	98	0.2	7.00	6.64 to 7.36
Los Angeles Basin (1)	28	90	1	5.96	5.24 to 6.68
San Francisco	61	90	1	6.41	5.89 to 6.93
Oxnard	26	90	1	7.38	6.64 to 8.12
Long Beach (1)	28	90	1	6.04	5.20 to 6.88
New York	29	90	1	6.52	5.76 to 7.28
Bridgeport/Charleston (1)	16	90	1	7.63	6.89 to 8.37
Anchorage	63	90	1	6.54	6.00 to 7.08
Bartlesville	44	90	1	6.30	5.61 to 6.99
Lafayette	79	90	1	6.51	6.03 to 6.99
Shreveport	41	90	1	6.83	6.18 to 7.48
Vernal	13	80	2	7.23	6.30 to 8.16
Denver	129	80	2	6.46	6.12 to 6.80
Cody	42	80	2	7.31	6.74 to 7.88
Columbus	27	50	5	6.96	6.47 to 7.45
Lansing	30	50	5	6.83	6.16 to 7.50
Chicago	41	50	5	6.54	5.97 to 7.11
Tulsa	53	50	5	6.79	6.33 to 7.25
Los Angeles Basin (2)	27	30	7	7.00	6.26 to 7.74
Long Beach (2)	28	30	7	7.39	6.80 to 7.98
Bridgeport/Charleston (2)	15	30	7	7.82	6.97 to 8.67

was a 95-percent confidence interval.)

Likewise, the Bridgeport/Charleston (W. Va.) sections had ranges that did not encompass 6.8. In their defense, the meal service had been poor, the public address system had disappeared, and there were more than the normal misunderstandings. Even so, their lower limits of 6.87 and 6.97 just barely missed the 6.8 target.

One group of highly quantitative people also took the test. I mention this group because of the large number of members it includes and because it provides evidence that the more quantitative people may do a little better in estimating uncertainty — but still not as well as they would like. (See Table 2.)

The 68 percent would not be expected to hold on all kinds of questions or all kinds of people. In fact, it is clear that the number would have been higher had it not been for relatively easy questions such as Questions 1 and 4. Most people know St. Augustine was a Spanish community and, therefore, had to be established between 1492 and 1776. By making the range a bit more narrow than that, they could be reasonably sure of bracketing the true answer. Even so, more than one-third of the members missed that one — regardless of their instructions on range.

Based on a sample of the 1,200+ quizzes, here are the average misses for each question:

Question	Average Misses (percent)
1	39
2	67
3	60
4	50
5	69
6	68
7	76
8	69
9	74
10	85

Questions such as Questions 9 and 10 were difficult,

and we found 80 percent or so misses — again regardless of the requested probability of a miss.

People who have no idea of the answer to a question will apparently try to fake it rather than use a range that truly reflects their lack of knowledge. This trait may be as universal a part of human nature as laughter; certainly it is not peculiar to SPE-AIME members.

Is the Problem Costly?

Why should anyone get excited about such results? Because, I think, similar behavior on the job can cost industry a bundle. Our membership at various levels of




Fig. 1 — The 95-percent confidence intervals of SPE-AIME sections. Average number of misses on 10-question quiz.

TABLE 2 — COMPARISON OF RESULTS

Section	Number of Usable Responses	Requested Range (percent)	Expected Number of Misses	Actual Number Average Misses	95-Percent Confidence Interval
Atlantic Richfield R&D SPE-AIME Section	52	98	0.2	4.52	3.84 to 5.20
(Hobbs and Oklahoma City)	145	98	0.2	6.83	6.50 to 7.16

management is responsible for all sorts of daily estimates that ultimately work their way into investment decisions. To the extent that the success of the investments relies on those estimates, business can be in trouble. If one's range so seldom encompasses the truth on tough questions, then the more common single-point estimates have little chance of being very close. Even those beloved "what-ifs" cannot be of much help since such questions would only be expected to test "reasonable" ranges. This research seems to indicate that most of us have little idea of what is a reasonable range.

Other Experiments

Earlier, I mentioned that we might be able to practice this business of estimating uncertainty and improve our track record. Experience with the SPE-AIME sections says that the practice may have to be substantial. Having established the 68-percent norm during the early part of my tour, I was able to do some other experimenting later.

One section had the benefit of knowing ahead of time what all the other sections had done. They knew before they started that no matter what range I had asked for, the membership always responded with about 68-percent misses, or a 32-percent probability interval. This group of 143 knew, then, that the tendency was to give much too tight a range and that they should be very careful not to fall into the same trap. (See Table 3.) It would seem that my warning had some effect. The mere telling of the experience of others is not, however, enough to shock most people into an acceptable performance.

Menke, Skov, and others from Stanford Research Institute's (SRI) Decision Analysis Group have experimented along similar lines. (and, in fact, their work gave me the idea for these tests). They say that if groups repeatedly take quizzes such as those described here, they are able to improve. Initially, people gave 50-percent ranges even though 98-percent ranges had been asked for. After several such tests (different each time, of course), the participants were able to reach a 70-percent range, but could never quite break that barrier. Their results show, apparently, that many intelligent men and women (they dealt largely with business executives) can never admit all their uncertainty. SRI made sure that some of their tests were built from subject matter familiar to the executives, such as questions extracted from their own company's annual report. Therefore, the phenomenon we are describing must have very little to do with the type of question.

Value of Feedback

For several years now we have asked our exploration people for 80-percent ranges on reserves before drilling an exploratory well. But we recognized that the act of putting down a 10-percent point and a 90-percent point would not in itself be sufficient. We also asked them

to see what their 80-percent range told them about other points on the distribution curve. If one is willing to assume a certain form of probability distribution, then the 80-percent range also specifies every other point. Hence, the explorationist can essentially put himself into a feedback loop. He puts two points into a simple time-share computer program, and out pop all the others. He now may check the 90-percent point, the 50-percent point, or any other. He well may find some that do not fit his notions — for example, his 80-percent range does not yield a 40-percent range that suits him. So he compromises one or the other until he gets the fit he likes.

All that is design and theory. In practice, most people throw in the 80-percent range and just accept whatever comes out. Based on the recent testing with SPE-AIME groups, I would have to guess that the 80-percent range constructed without feedback is actually much more narrow — perhaps 50 percent. It would take a lot of data, which we do not have, to measure the range. Almanacs and encyclopedias cannot help much here.

My estimate of 50 percent comes from the following judgment. It must be more difficult to put ranges on exploration variables than to put them on questions such as when St. Augustine was founded. On the other hand, it should be easier for a geologist to conceive of his vocational uncertainties than for him to handle Beowulf-type questions. Since the audiences' average ranges on those two questions were about 40 and 85 percent, respectively, I chose 50 percent.

The feedback process, if used, can be of benefit. The following experiment was performed with some sections. I asked the members to write down two ranges simultaneously. That forced some sort of feedback. And since both ranges could not have 68-percent misses, it seemed logical to expect that such a ploy would yield better results — which, in fact, was what happened. (See Table 4.)

By having to use two ranges, the members were able to greatly improve their 90-percent range compared with those who worked with only one interval. The 50-percent range, however, was shoved in the other direction. I would guess that the best strategy for one faced with an uncertainty problem would be to consider whole distributions (that is, many ranges), continually playing one against the others. That scheme should result in even better definition of one's uncertainty.

Even then, studies suggest that people may come up short. I once saw the results of a full-scale risk analysis, including a probability distribution of project cost. A few months later the same people did another risk analysis on the very same project. Amazingly, the cost distributions did not even overlap. Changes had taken place on that project in the space of a few months that moved the results far beyond those contemplated when the experts were laying out their original ranges. People tend to build into their ranges those events that they can

TABLE 3 — KNOWLEDGE OF PREVIOUS RESULTS

Section	Number of Usable Responses	Requested Range (percent)	Expected Number of Misses	Actual Number Average Misses	95-Percent Confidence Interval
New Orleans	143	90	1	5.46	5.08 to 5.84

TABLE 4 — RESULTS USING FEEDBACK PROCESS

Section	Number of Usable Responses	Requested Range (percent)	Expected Number of Misses	Actual Number Average Misses	95-Percent Confidence Interval
Bay City	26	90	1	5.04	3.99 to 6.09
Bay City	26	50	5	8.31	7.67 to 8.95
Houston	98	90	1	4.05	3.63 to 4.47
Houston	98	50	5	7.32	6.94 to 7.70

see as possibilities. But since much of our uncertainty comes from events we do not foresee, we end up with ranges that tend to be much too narrow.

Are the Tests Valid?

There may be those who still feel that the kinds of questions I used cannot be used as indicators of what one does in his own specialty. I know of several arguments to counter that view, but no proof.

The less one knows about a subject, the wider should be his range. An English scholar might have a 90 percent range of A.D. 700 to 730 for the "Beowulf" question. The typical engineer might recognize his limitations in the area and put A.D. 500 to 1500. Both ranges can be 90-percent ranges because the degree of uncertainty is a very personal thing. One's knowledge, or lack of it, should not affect his ability to use 90-percent ranges. So the type of question should not matter.

I mentioned earlier that SRI's use of material from a company's own annual report did not change the results. Regardless of whether one is an expert, the ranges generally come in too narrow.

Another criticism of these questions has been that they test one's memory of events already past rather than the ability to predict the future. Conceptually, is there any difference regarding the uncertainty? There may be more uncertainty associated with, for instance, the timing of an event yet to take place. But it seems that the difference is only one of degree when compared with recalling a date in history from an obscure and seldom-used brain cell. In either case, one does not know for sure and must resort to probability (likely a nontechnical variety) to express himself.

Bean Counting

You may find a third argument even more compelling. We asked groups of people to estimate the number of beans in a jar. Not only were they asked for their best-guess single number but also for a 90-percent range. The players were mostly professional people with technical training, and most had or were working part time on advanced degrees. Since we built in a reward system (money), the estimators were trying to do a good job, at least with their best guess. The following table gives their results. The jar contained 951 beans.

Best Guess	90-Percent Range
217	180 to 250
218	200 to 246
250	225 to 275
375	200 to 500

385	280 to 475
390	370 to 410
450	400 to 500
500	150 to 780
626	500 to 700
735	468 to 1,152
750	500 to 1,500
795	730 to 840
800	750 to 850
960	710 to 1,210
1,000	900 to 1,100
1,026	700 to 1,800
1,053	953 to 1,170
1,070	700 to 1,300
1,080	700 to 1,400
1,152	952 to 1,352
1,200	500 to 3,600
1,200	1,000 to 1,500
1,201	1,000 to 1,400
1,300	500 to 2,000
1,300	600 to 2,000
1,400	1,200 to 1,600
1,500	400 to 1,800
1,500	800 to 2,000
1,600	1,350 to 1,950
1,681	1,440 to 2,000
1,850	1,400 to 2,200
4,655	4,000 to 5,000
5,000	2,000 to 15,000

The experiment provides added insight because everyone could see the beans. No one had to test his memory of geography or history or his company's performance reports. The jar was somewhat square in cross-section so as not to introduce any tricks in estimating volume, though no one was allowed to use a ruler. Still, the requested 90-percent ranges turned out to be more like 36-percent ranges because only 12 of 33 included the true value. After our testing, Elmer Dougherty of the U. of Southern California tried the same experiment and privately reported very similar results. We then asked some of our exploration people to go through the exercise, and they too repeated the earlier performances of others.

Interestingly, we have three more bean estimates made by people using a computer model (Monte Carlo simulation) to get ranges. They estimated their uncertainty on the components (length, width, height, and packing density) to get an over-all range. All included the true value of 951. Equally competent people not using the simulation approach could not do as well.

<u>Best Guess</u>	<u>90-Percent Range</u>
1,120	650 to 1,900
1,125	425 to 3,000
1,200	680 to 2,300

This experiment provides evidence that even a simple approach to probability modeling usually will be a lot better than what one dreams up in his head when it comes to assessing uncertainty.

Still More Experiments

Few people give in easily when confronted with this kind of material. They complain that I am testing groups and it was the "other guys" who caused the problems we see reflected in the data. Or they did not know my game was a serious one. Or they had no real incentive to do well, as they normally have on the job. Or that while they admit to having missed cost estimates, project completion times, producing rates, inflation rates, crude oil prices, etc., now and then, those were caused by external circumstances and certainly nothing they could have been responsible for. (Who ever said that we should only estimate that part of uncertainty for which we have responsibility?)

To counter such talk, I have engaged in other testing. One group had money riding on their ability to properly assess probability ranges. I asked them for 80-percent ranges and even agreed to pay them if, individually, they got between 60 and 90 percent. If they did not, they had to pay me. The group was so convinced the game was in their favor that they agreed to pay for the privilege of playing! And it was not sight unseen, either. They had already taken the test before the wager (same 10 questions given to SPE-AIME sections). They lost. But the point is that before getting their results, they did not feel that the questions were in some way beyond their capabilities.

At the SPE-AIME Fall Technical Conference and Exhibition in Dallas, I needed to save time while presenting this paper but I still needed to illustrate the point. I used a color slide of some beans spread about in an elliptical shape. It was the easiest test yet; the audience could clearly see every single bean. We used a 12-ft screen so the images would be large even for those in the rear. Still, only about one-third of the several hundred present came up with a 90-percent range that encompassed the true value.

As early as 1906, Cooke² did some testing of meteorological questions to see how well he could assess uncertainty. Since then, others^{3,4} have examined the problem and noticed similar results. Lichtenstein *et al.*⁵ have an extensive bibliography.

Don Wood of Atlantic Richfield Co. has been using a true/false test to study the phenomenon. The subject answers a question with true or false and then states the probability he thinks he is correct. Most people are far too sure of themselves. On those questions they say they have a 90-percent chance of answering correctly, the average score is about 65 percent.

To illustrate his findings, Wood describes the results on one of his test questions: "The deepest exploratory well in the United States is deeper than 31,000 ft." Several knowledgeable oil men have said the statement is false and that they are 100-percent sure of their answer.

Other oil men have said true, also believing they are 100-percent sure of being correct. Two petroleum engineers argued about another of Wood's questions: "John Wayne never won an academy award." Each was 100-percent sure of his answer, but one said true and the other said false. By the way, an Oklahoma wildcat has gone deeper than 31,000 ft and "True Grit" won an Oscar for the actor.

Where this paper reports results on how SPE-AIME groups act, Wood gives a test that has enough questions so that an individual can calibrate himself apart from any group. The grade one receives after taking the test may be loosely defined as the probability he knows what he is doing. It comes from a chi-square goodness-of-fit test on binomial data. Typical scores have been smaller than 1×10^{-5} , or less than 1 chance in 100,000.

Every test we have performed points in the same direction, as have most of the tests performed by others. The average smart, competent engineer is going to have a tough time coming up with reasonable probabilities for his analyses.

What Can We Do?

First, think of a range of uncertainty without putting any probability on that range. Since our sample showed that people tend to use the same range no matter what kind of range they were asked for, it seems plausible that a range such as we obtained during the tour would be forthcoming.

Having written it down, we arbitrarily assign some relatively small probability to the range encompassing the truth, say 40 percent. Decide on the form of the error. For example, in estimating project completion time, one may feel his uncertainty is symmetrical (± 6 months). (See Fig. 2.)

If the uncertainty is best expressed as symmetrical, then get some normal probability paper like that illustrated in Fig. 3. Plot the low end of your range at the 30-percent point and the high end at the 70-percent point. Note that $70 - 30 = 40$. Your range has a 40-percent chance of encompassing the truth. Connect the points with a straight line and extend the line all the way across the paper. By reading the ordinates at the 5-percent and 95-percent points, you have your 90-percent range ($95 - 5 = 90$). Our ± 6 months has been converted to $\pm 1\frac{1}{2}$ years. If that range seems uncomfortably large, good! Remember that if you are like most people, your natural tendency is to make such ranges too narrow. To repeat an earlier idea, uncertainty comes about because of what we do not know. Ranges constructed using what we do know are likely to be too small. (Bias, either pessimism or optimism, may be a problem too, but we have not addressed it here.)

You may feel the uncertainties are asymmetrical with a long tail region to the right, such as in estimating reserves (see Fig. 4). One cannot have less than 0 bbl, though with small probability he can have very large numbers.

In such cases, use log-probability paper as in Fig. 5. Say the range is 3 to 6 million bbl. Again, go through the ritual of plotting the low and high, drawing the line, and checking to see how comfortable you are with a 90-percent range. This time our range has been con-

verted from 3 to 6 to something like 1.4 to 12. Discomfort is a good sign.

Because they fit so much of the world so well, the normal and lognormal distributions are logical choices for describing uncertainty. Do not worry a great deal about this apparent straight jacket. A realistic range (that is, wide) is often more important than the form of the distribution anyway.

Nor is there anything particularly holy about defining your original range as 40 percent. I could have used 50 or 30 percent. I am just proposing a simple way to get started in this business of defining the degree of your uncertainty and at the same time paying homage to the finding that people tend to overestimate the extent of their knowledge.

If each bean counter had plotted his range on log-probability paper as a 40-percent range and graphically determined his 90-percent range, 25 of the ranges (or 76 percent) would have included the true value of 951. Using such a technique, the group would have achieved a significant improvement in their ability to set ranges. After all, 76 percent is not that far from their target of 90 percent.

As you begin to keep records of your probability statements and compare them with actual outcomes, you will begin to build your own rules for making estimates. And, ultimately, your own tested rule is going to work better for you than anything others design.

The Value of Training

Winkler and Murphy⁶ reported on some meteorologists who showed little or no bias in assessing probability. Training through years of almost immediate feedback on their predictions very likely accounts for this rare but enviable behavior. The oil business seldom allows such feedback. We may not find the answers to our predictions for several years, and by then we have been retired, promoted, banished, or worse.




Fig. 2 — Estimating with symmetrical uncertainty.




Fig. 3 — Determining range, symmetrical uncertainty.

But since training in this area appears to be vital, I urge you to set up a program for yourself. Every month make some predictions about the future, predictions whose outcome will become known during the following few weeks. Assign probabilities to your predictions, and religiously check your results. Find out what happens when you are 90-percent sure, 70-percent sure, etc. Example:

1. The next holiday weekend will see more highway deaths recorded than the similar period last year.

True 60 percent

2. The Cincinnati Reds will lead their division on July 4.

True 70 percent

3. XXX Corp. common stock will close above \$Z before Sept. 1.

False 50 percent

To find out how well you are doing, consult some binomial probability tables (or a friendly expert). Say you had 20 statements to which you assigned a 70-percent chance of being right. You would have expected to get 14 of them right. What if you only got 10 right? Is that good? The tables show a probability of 4.8 percent of getting 10 or less right under conditions when you expect to get 14 right out of 20. It would be long odds (1 in 20) to claim, therefore, that you had learned to set the probabilities correctly. Better practice some more. Ask your stockbroker to do likewise.

Does a Better Range Lead to a Better Mean?

One might be tempted to argue that improving our understanding of uncertainty would not in itself improve the estimate of the mean, best guess, or whatever people tend to use for making their decision. But look, for example, at the Alyeska Pipeline and the 1969 cost estimate of \$900 million. Most everyone associated




Fig. 4 — Estimating with asymmetrical uncertainty.




Fig. 5 — Determining range, asymmetrical uncertainty.

with the project knew that it could not cost much less than \$900 million. If everything had gone off without a hitch (roughly equivalent in probability of occurrence to all the molecules congregating on one side of a room), it might have come in for around \$800 million.

What kind of things could happen to drive the cost in the other direction?

1. Labor problems such as jurisdictional disputes and the lack of an adequate supply of necessary skills in such a harsh environment.

2. Weather.

3. Shortages of equipment and supplies resulting from the unique nature of the project and remoteness of the site.

4. Design problems. An axiom of engineering: All doth not work that man designeth.

5. Economy of scale in full retreat. Some projects are so large that they are most difficult to manage effectively.

6. Bureaucratic delays brought about by masses of government regulations.

(Note that the list does not include the large cost increase brought about by government inflationary policies and the oil embargo, nor does it include the problems caused by so-called environmentalists. Reasonably intelligent forecasters might have missed those events back in early 1969.)

An analysis of these six items would have led one to imagine some chance for a pipeline costing as much as \$3 billion giving the following range.

Rock-bottom cost	\$0.8 billion
Best estimate	\$0.9 billion
High-side cost	\$3.0 billion

How long could such a "best guess" survive in such a range? Merely writing down the numbers exposes the best guess to sharp criticism and doubtless would force it to a higher and more realistic level. Though the new best guess would still have been far below present cost estimates of almost \$8 billion, it nevertheless would have been very useful. Crude prices, we remember, were much lower then.

It seems logical, then, to expect that quite a number of projects would benefit similarly from a better range analysis. Consider the bean counters mentioned earlier. What if all those whose best guesses were less than 500 had known that there was a chance the truth might be up around 1,000? Is it not likely that they would have moved those best guesses up somewhat?

The Payoff

The payoff for having a better grasp on uncertainty should be quite a sum. In recent years both industry and government could have been more cautious in their estimates and perhaps achieved a better return for their investments.

The *Oil and Gas Journal* of Oct. 9, 1967, quoted management at the Great Canadian Oil Sands plant dedication: "Operating in the northland offers no unusual problems — in fact, it has some advantages." *Business Week*, Jan. 5, 1974, quoted the GCOS President: "We're the proud owners of a \$90 million loss. This is the cost of being a pioneer."

Most tax payers remember the many government

programs that ended up costing much more than original estimates (TFX, C5A, Interstate Highway Program, BART, and the Dallas-Fort Worth Regional Airport, for example). There has been a long history of cost underestimates for all kinds of projects because of not adequately accounting for future unknowns.

The whole planning and budget process stands at the mercy of supposedly expert estimates. It may be that we have gotten ourselves into trouble by looking for "the answer" (never attainable) when we should have concentrated on realistically setting our uncertainties. If the ranges are adequate, then at least the plan can cope with possible events of the future.

A better view of our uncertainties should have a significant effect on our success as risk takers and ultimately on profits.

Acknowledgments

I would like to thank the many SPE-AIME members who took the test and made this project so enjoyable. Also, I extend thanks to friends in Atlantic Richfield Co. who gave many helping hands.

References

1. Tversky, A. and Kahneman, D.: "Judgment Under Uncertainty: Heuristics and Biases." *Science* (1974) **185**, 1124-1131.
2. Cooke, W. E.: Forecasts and Verifications in Western Australia." *Monthly Weather Review* (1906) **34**, 23-24.
3. Fischhoff, B. and Beyth, R.: "I Knew It Would Happen" — Remembered Probabilities of Once-Future Things." *Organizational Behavior and Human Performance* (1975) **13**, 1-16.
4. Lichtenstein, S. and Fischhoff, B.: "Do Those Who Know More Also Know More About How Much They Know?." *Bull., Oregon Research Institute, Eugene* (1976) **16**, 1.
5. Lichtenstein, S., Fischhoff, B., and Phillips, L. D.: "Calibration of Probabilities: The State of the Art." *Bull., Oregon Research Institute, Eugene*.
6. Winkler, R. L. and Murphy, A. H.: "Evaluation of Subjective Precipitation Probability Forecasts." *Proc., First National Conference on Statistical Meteorology, American Meteorological Society* (May 27-29, 1968) 148-157.

APPENDIX

Answers to the ten questions used in the quiz.

1. A.D. 1565.
2. 12.8 million (10.3 million autos).
3. 6,904 miles.
4. 1,901 miles.
5. 76.2 million people.
6. 4,200 ft.
7. 3.85 million sq miles.
8. 3,900 miles.
9. 248.4 years.
10. A.D. 700 to 730.

The answers to the questions came from the *Official Associated Press Almanac*, 1974 edition. Any source can be in error, and thus I discovered after the testing that I had been led astray on Question 2. The source said automobiles, but in checking other sources, I am now sure they meant motor vehicles. Strangely, the "new" answer does not affect our results very much. Most of those who missed that one were so far off that they were beyond help.

JPT

Original manuscript received in Society of Petroleum Engineers office July 27, 1975. Paper accepted for publication Jan. 8, 1976. Revised manuscript received April 19, 1976. Paper (SPE 5579) was first presented at the SPE-AIME 50th Annual Fall Technical Conference and Exhibition, held in Dallas, Sept. 28-Oct. 1, 1975. © Copyright 1976 American Institute of Mining, Metallurgical, and Petroleum Engineers, Inc.

PROBABILITY ENCODING IN DECISION ANALYSIS

Carl S. Spetzler and Carl-Axel S. Staël von Holstein
written while associated with
Decision Analysis Group
Stanford Research Institute

Presented at the ORSA-TIMS-AIEE 1972 Joint
National Meeting, Atlantic City, New Jersey,
8-10 November 1972.

All rights reserved by Carl S. Spetzler and
Carl-Axel S. Staël von Holstein, 1972.

1. INTRODUCTION

Probability encoding plays an important role in the application of decision analysis, since it is the process of extracting and quantifying individual judgment about uncertain quantities. This paper is intended as a start in disseminating probability encoding methodology. It summarizes the probability encoding methods currently used by the Decision Analysis Group at Stanford Research Institute. These methods are based on several years of experience with probability encoding in decision analysis applications, as well as on evidence from experiments.

There is a vast literature that relates to probability encoding. The annotated bibliography by Staël von Holstein [8] covers the items that are most relevant to this paper. Some encoding techniques are summarized in [7]. The last twenty years have seen a flood of psychological experiments dealing with various aspects of man as an "intuitive statistician" or "processor of probabilistic information"; many of the experiments provide relevant insights. Two recent overviews of the field are provided by Peterson and Beach [5] and Rapoport and Wallsten [6]. However, the psychological studies have restricted usefulness for probability encoding in practical situations for three reasons. Most studies deal with binary probability distributions (an event either occurs or does not occur) rather than continuous distributions. Moreover, they are based on laboratory experiments rather than actual decision situations. Finally, while the studies show how well (or poorly) subjects perform in various tasks, they do not develop procedures for improving performance.

During our research, we have collaborated with Professors Daniel Kahneman and Amos Tversky of Hebrew University, Jerusalem; the material in Section 4 is based on their work. We have benefitted from many discussions of the subject with our colleagues in the Decision Analysis Group, in particular with Dr. James E. Matheson, and we are grateful for their valuable comments. The paper has further benefitted from a careful review by Dr. Michael M. Menke.

2. THE DECISION ANALYSIS FRAMEWORK

Probability encoding is primarily done in the context of a decision problem. A brief overview of decision analysis is given below to provide a frame of reference. More extensive discussions of decision analysis are found in Howard [1], [2] and Staël von Holstein [9]. A second, but not necessarily less important, reason for encoding probabilities is that they provide a clear means for communication about uncertainty.

Decision analysis procedures usually involve three phases--the deterministic, probabilistic, and informational phases. The deterministic phase accomplishes the basic structuring of the problem by defining relevant variables, characterizing their relationship in formal models, and assigning values to possible outcomes. The importance of the different variables is measured through sensitivity analysis.

Uncertainty is explicitly incorporated in the probabilistic phase by assigning probability distributions to the important variables. These distributions are obtained by encoding the judgment of knowledgeable people. They are transformed in the model to exhibit the uncertainty in the final outcome, which again is represented by a probability distribution. After the decision maker's attitude toward risk has been evaluated and taken into account, the best alternative in the face of uncertainty is then established.

The informational phase determines the economic value of information by calculating the worth of reducing uncertainty in each of the important variables in the problem. The value of additional information can then be compared with the cost of obtaining it. If the gathering of additional information is profitable, the three phases are repeated again. The analysis is completed when further analysis or information gathering is no longer profitable.

Throughout, the analysis is focused on the decision and the decision maker. Expanding the analysis is considered of value only if it helps the decision maker choose between available alternatives.

3. MODELING AND ENCODING

The personal interpretation of probability represents a cornerstone in the decision analysis philosophy. Probability represents an encoding of information. Since various people are likely to have different information, two persons can make different probability assignments to the same uncertain quantity. We have found an interview process to be the most effective way of encoding a probability distribution.

The decision maker is the person (or group of persons) who has the responsibility for the decision under consideration. It follows that a decision analysis must be based on the decision maker's beliefs and preferences. He may be willing to designate some other person or persons as his expert(s) for encoding the uncertainty in a particular variable if he feels that the expert has a more relevant information base. The decision maker can then either accept the expert's information as his input to the analysis or modify it to incorporate his own judgment.

Definition of Decision and State Variables

A decision analysis model includes two kinds of input variables: decision variables and state variables. The two must be carefully distinguished from one another because while the decision maker can choose the values of the decision variables, the values of the state variables are beyond his control. Thus it is only meaningful to discuss encoding with respect to state variables. Some variables, such as price, may at first seem difficult to classify as decision or state variables. This difficulty, however, may be resolved by further structuring of the problem: e.g., into a controllable price strategy and the uncertain market response. A similar problem can arise when variables interact. For example, development time, program cost, and product performance are closely related in new product decisions. One or two can be selected as decision variable(s) and the other(s) become(s) a state variable(s). The problem must be structured carefully according to which variables are best considered decision variables and which are state variables. Often this separation is most easily achieved by redefinition of the variables.

There is always a choice between encoding the uncertainty in an important variable or modeling the problem further. At one extreme, it is conceivable that the final worth or profit contribution of a project

could be encoded directly, thus bypassing a need for examination of the underlying variables. Generally, a distribution for final worth is more easily reached, or provides more confidence, if a model is constructed that relates final worth to other variables. The modeling effort tends to be most effective and most economical if it starts with a gross model that is successively refined. The model should be refined only while the cost of each addition provides at least comparable improvement in information. This test depends on how the information bears on the decision at hand.

The choice between additional modeling and encoding may need to be reconsidered during the encoding process, since the subject may reveal biases during the interview that often can be treated by further structuring of the problem.

Some Encoding Principles

The following list of principles should be used in defining and structuring any variable whose uncertainty is to be encoded: Violating them invariably leads to problems in the probability encoding. It serves as a checklist before the actual encoding takes place. These principles are:

- The uncertain quantity should be important to the decision, as determined by a sensitivity analysis.
- The quantity should be defined for the subject as an unambiguous state variable. If the subject believes the outcome of the quantity can be affected to some extent by his decision, then the problem needs restructuring to eliminate this effect.
- The level of detail required from the encoding process depends on the importance of the quantity and should be determined by sensitivity analysis before the interview. It may sometimes be sufficient to elicit only a few points on the distribution.
- The quantity should be well structured. The subject may think of the quantity as conditional on other quantities; accordingly, conditionalities should consciously be considered and brought into the structure because our minds deal ineffectively in combining uncertain quantities. Mental acrobatics should be minimized.
- The quantity should be clearly defined. A good test of this quality is to ask whether a clairvoyant could reveal the value of the quantity by specifying a single number without requesting

clarification. To cite an example, it is not meaningful to ask for the "price of wheat in 1974," because the clairvoyant would need to know the quantity, kind of wheat, at what date, at which Exchange, and the buying or selling price. However, "the closing price of durum wheat on June 30, 1975 at the Chicago Commodity Exchange" would be a well-defined quantity.

The quantity should be described by the analyst on a scale that is meaningful to the subject. For example, in the oil industry, the subject--depending on his occupation--may think in terms of gallons, barrels, or tank cars. The wrong choice of scale may cause the subject to spend more effort on fitting his answers to the scale than on evaluating his uncertainty. It is important, therefore, to choose a unit with which the subject is comfortable; after the encoding, the scale can be changed to fit the analysis. As a rule, let the subject choose the scale if there is no obvious scale.

4. MODES OF JUDGMENT*

Relevance for Probability Encoding

People perceive and assess uncertainty in a manner similar to the way they perceive and assess distance. They use intuitive assessment procedures that are often based on cues of limited reliability and validity. At the same time the procedures (we will use mode of judgment as a synonym) generally produce reasonable answers. For example, an automobile driver is generally able to estimate distance accurately enough to avoid accidents, and a business executive is generally able to evaluate uncertainties well enough to make his enterprise profitable. On the other hand, a particular mode of judgment may lead to answers that are systematically biased, sometimes with severe consequences.

To pursue the analogy with estimation of distance, people are known to overestimate the distance of a remote object when visibility is poor and to underestimate the distance when the sky is clear. Thus, they exhibit a regular systematic bias. This is because we normally use the haze as a cue to distance. This cue has some validity, because more distant objects are usually seen through more haze. At the same time, this mode of judgment may lead to predictable errors. Three features of this example are worth noting: (1) People are not generally aware of the cues on which their judgments are based. Few people know that they use haze to judge distances, although research shows that virtually everybody does. (2) It is difficult to control the cues we use; the object seen through haze looks more distant, even when we know why. (3) People can be made aware of the bias, and can make a conscious attempt to control its effects, as the captain of a ship does when navigating in a mist.

An analogous problem exists in the assessment of uncertain quantities. Here too, one relies on certain modes of judgment that may introduce systematic biases. Here too, modifying impressions and intuitions

* Much of the material in this and the next section is based on private communications with Daniel Kahneman and Amos Tversky. The analogy between judgment of distance and judgment of uncertainty is due to them.

is exceedingly difficult, but it is possible to learn to recognize the conditions under which such impressions are likely to lead us astray.

We will now briefly categorize biases that may be encountered in probability encoding. In the subsequent sections we will discuss some modes of judgment that are often used in responding to questions about uncertain quantities.

Biases in Probability Encoding

For the purpose of this discussion the subject is assumed to have an underlying stable knowledge regarding the quantity under investigation. This knowledge may be changed through receiving new information. The task of the analyst is to elicit from the subject a probability distribution that describes the underlying knowledge. Conscious or subconscious discrepancies between the subject's responses and an accurate description of his underlying knowledge are termed biases. Biases may take many forms. One is a shift of the whole distribution upward or downward relative to the basic judgment; this is called displacement bias. A change in the shape of the distribution compared with the underlying judgment is called variability bias. Discrepancies may be a mixture of both kinds of bias. Variability bias frequently takes the form of a central bias, which means that the distribution is tighter (has less spread) than is justified by the subject's actual state of information. Biases are illustrated in Figure 1 in the form of three density functions, where A represents the underlying state of information and distributions B and C, respectively, represent the effects of central bias and displacement bias.




FIGURE 1 EXAMPLES OF VARIABILITY AND DISPLACEMENT BIASES

The sources of biases can be cognitive or motivational. Motivational biases are either conscious or subconscious adjustments in the subject's responses motivated by his perceived system of personal rewards for various responses; he may want to influence the decision. He may also want to bias his response because he perceives his performance will be evaluated by the outcome: For example, a sales manager may consciously give a low prediction of sales because he thinks he will look better if the actual sales exceed his forecast. Finally, the subject may suppress the full range of uncertainty that he actually believes to be present because he feels that someone in his position is expected to know what will happen in his area of expertise.

Even when a subject is honest--in the sense that he lacks motivational biases--he may still have cognitive biases. Cognitive biases are either conscious or subconscious adjustments in the subject's responses systematically introduced by the way the subject is intellectually processing his perceptions. For example, a response may be biased towards the most recent piece of information simply because that information is the easiest to recall. Cognitive biases depend on the judgment mode used; they will be discussed further in the next section.

Basic Modes of Judgment

A bias results from the use of a mode of judgment. An important responsibility of the interviewer is to try to elicit what modes of judgment may be used by the subject and then try to adapt the interview to minimize biases. In this section, we will define five different modes of judgment and give examples of how they might operate.

Availability

Probability assignments are based on information that the subject recalls or visualizes. The probability of a breakdown in a production process may be assigned by recalling past breakdowns. Availability refers to the ease with which relevant information is recalled or visualized [10]. It is easy to recall information that made a strong impression at the time it was first presented. Past results and present business plans also become easily available. Recent information is more available than old information and is often given too much weight. For example, a piece of recent news regarding a competitor may influence a sales forecast much more than should be allowed on the basis of past experience with such news. Some events may become overly available because

of their potentially disastrous consequences (e.g., an accident with a nuclear reactor) and are thus assigned probabilities that are too high.

Availability will be an important judgment mode in most probability encoding sessions. It can also be introduced deliberately by the interviewer. For instance, if the interviewer believes that the subject has a central bias, he can ask the subject to make up scenarios for extreme outcomes, which thereby become more available and help counteract the central bias.

Adjustment and Anchoring

The most readily available piece of information often forms an initial basis for formulating responses; subsequent responses then represent adjustments from this basis. For example, the current business plan is often used as an available starting point. Likewise, when predicting this year's sales, the subject may use last year's sales as a starting point. He may use the recent years with the biggest and smallest sales as the bases for formulating judgment about the extreme values for this year's sales. The initial response many times serves as a basis for later responses, especially if the first question concerns a likely value for the uncertain quantity.

The subject's adjustment from such a basis is often insufficient. We then say that the response is anchored on the basis; the result is likely to be a central bias. Anchoring thus occurs when some information has become overly available at the beginning of the procedure. It results from a failure to process information regarding other points on the distribution independently from the point under consideration.

Representativeness

Representativeness means that the probability of an event or a sample is evaluated by the degree to which it is representative of, or similar to, major characteristics of the process or population from which it originated [3], [4]. We can then say that probability judgments are being reduced to judgments of similarity. For example, people tend to assign roughly the same distribution to the average of a sequence of uncertain quantities (e.g., the average production volume for a group of machines) as to each individual quantity forming the average when they usually should assign a much tighter distribution to the average. The main characteristic of the average value is the population from which the individual quantities were sampled; information about that population

therefore has a much greater influence on the distribution of the average than has the number of quantities making up the average.

There is a tendency to disregard general information and base probability assignments on what appears to be a specific fact. For example, a company had to decide whether to introduce a new product that was considered to have a high demand potential. The product was test-marketed with a slightly unfavorable outcome, and the revised judgment of the market was then a low demand. This revision was made in spite of past experience with similar market tests that had been less than accurate in predicting the final market size and in contrast to the strong prior judgment indicating a high demand. This is a case of focusing on information that relates to an individual hypothesis and of ignoring general information, which perhaps should carry the main weight in the probability assignment.

Biases can sometimes be explained by different modes of judgment. For example, the fact that people attach too little weight to general information can also be explained by availability. That is, the market test information in the example above was more recent than the general information, and therefore more available.

The biases resulting from representativeness can often be reduced or eliminated by further structuring of the problem. In the marketing example, it is easy to encode the prior probabilities for various levels of demand and encode the probability distribution for the test result conditional on the demand. A simple application of probability calculus will then provide the posterior probabilities of demand level given the outcome of the market test.

Unstated Assumptions

A subject's responses are typically conditional on various unstated assumptions; consequently, the resulting probability distribution does not properly reflect his total uncertainty. This means that the subject may not have considered such possibilities as future price controls, major strikes, currency devaluation, war, and so on, when expressing his judgment. He does not hold himself responsible for considering such events. One result is that he may be less surprised than might be expected when the revealed value of an uncertain quantity falls outside the range of his distribution. He justifies this because of a drastic change in some condition that he did not feel responsible for incorporating into his judgment.

The subject cannot be held responsible for all of the unstated assumptions; rather, he is responsible for stating the assumptions he is using so they can be built into the model and so that the most appropriate expert (who may or may not be the current subject) can assign their probabilities.

Coherence

People sometimes assign probabilities to an event based on the ease with which they can fabricate a plausible scenario that would lead to the occurrence of the event. The event is considered unlikely if no reasonable scenario can be found; it is judged likely if many scenarios can be composed that could make the event occur or if one scenario is particularly coherent. The credibility of a scenario to a subject seems to depend more on the coherence with which its author has spun the tale than on its intrinsically "logical" probability of occurrence. For example, the probability assigned to the event that sales would exceed a high volume may depend on how well market researchers have put together scenarios that would lead to that volume; for instance, scenarios on what markets might be penetrated and what the penetration rate might be with a reasonable marketing effort. Arguments in court are another example of evaluation based on the coherence of the sequence of evidence (as presented by the prosecution as well as the defense). It is thus important that the discussion of possible outcomes for an uncertain quantity be well-balanced, since the discussion and generation of arguments may affect the probability assignments.

5. ENCODING METHODOLOGY

Encoding Methods and Response Modes

Most encoding methods are based on questions for which the answers can be represented as points on a cumulative distribution function. We classify encoding methods as follows:

- P-methods ask questions on the probability scale with the values fixed.
- V-methods ask questions on the value scale with the probabilities fixed.
- PV-methods ask questions to be answered on both scales jointly; the subject essentially describes points on the cumulative distribution.

The encoding procedure consists of a set of questions that requires response either directly or indirectly through choices between simple bets. In the direct response mode, the subject is asked questions that require numbers as answers. The answers can be given in the form of probabilities (or equivalently in the form of odds) or values.

In the indirect response mode, the subject is asked to choose between two or more bets (or alternatives). The bets are adjusted until he is indifferent; this indifference can then be translated into a probability or value assignment. With a reference process, one bet is defined with respect to the uncertain quantity and the other with respect to the reference process.

The choice can also be made between events defined on the value scale for the uncertain quantity, where each event represents a set of possible outcomes for the uncertain quantity (e.g., sales being less than or equal to 2,000 units or sales being greater than 2,000 units). We can say that this response mode makes use of internal events.

Specific Techniques

Each probability encoding technique can be classified according to the encoding method and response mode used. The techniques which we have found most useful are given in Table 1.

Table 1

CLASSIFICATION OF PROBABILITY ENCODING TECHNIQUES

Encoding Method	Response Mode		
	Direct	Indirect	
		External Reference Process	Internal Events
Probability (value fixed)	Cumulative probability	Wheel	Odds
Value (probability fixed)	Fractiles	Wheel; fixed probability events	Interval technique
Probability--Value (neither fixed)	Drawing graph; Parametric description	--	--

The probability wheel is useful with most subjects. As an external reference process, it can be used as a P-method or a V-method, but the former is the method generally preferred. The probability wheel is a disk with two sectors, one blue and the other red, with a fixed pointer in the center of the disk. The disk is spun, finally stopping with the pointer either in the blue or the red sector (see Figure 2). A simple adjustment changes the relative size of the two sectors and thereby also the probabilities of the pointer indicating either sector when the disk stops spinning. The subject is asked whether he would prefer to bet either on an event relating to the uncertain quantity, e.g., that next year's production will not exceed x units, or on the pointer ending up in the red sector. The amount of red in the wheel is then varied until the expert becomes indifferent. When indifference has been obtained, the relative amount of red is assigned as the probability of the event. This is a P-method since the event (value) is fixed and the probability is determined through the process.

One advantage of the probability wheel is that the probability can be varied continuously from zero to one. It is only useful, however, for probabilities in the range from 0.1 to 0.9 because it is difficult for the subject to discriminate between sizes of small sectors. One alternative to the probability wheel (with the same restricted usefulness) is a horizontal bar with a marker (defining two events, to the left and to the right of the marker); another is an urn with, say, 1000 balls of




FIGURE 2 A PROBABILITY WHEEL

two colors (a ball is supposed to be drawn at random from the urn and the reference event is "the ball drawn is red"; the composition of the urn can then be varied). We prefer to use the probability wheel because it is easier to visualize the chance process than in the case of the bar or the urn.

Other reference processes may be useful, particularly when reference has to be made to low-probability events. For example, the event "ten heads in a row with an unbiased coin" has a probability close to $1/1000$. An event that some subjects might identify with is a royal flush which has a probability of roughly $1/65,000$. Typical for the reference processes mentioned in this paragraph is that they concern events with known probabilities and therefore only work as V-methods.

The interval technique is an example of the internal events response mode and is a V-method. An interval is split into two parts, and the subject is asked to choose which part he would prefer to bet on, or which part he considers most likely. The dividing point is changed to reduce the size of the part considered most likely (and thereby to increase the size of the other part), and the subject is asked to choose between the two new parts. The procedure of changing the dividing point is continued

until indifference is reached, and the subintervals are then assigned equal probabilities. Starting from an interval covering all possible outcomes and then splitting into two parts will first give the median, then the quartiles, and so on. The method does not seem to be very meaningful after the quartiles have been obtained because each question depends on earlier responses, thus errors may be compounded. The interval technique can also be based on splitting the interval into three parts.

A P-method with the internal events response mode asks the expert to assign the relative likelihoods (or odds) to two well-defined events. For example, the expert may first be asked whether he considers next year's sales more likely to be above or below 5,000 units. The next question is then: how much more likely is it? This method is used primarily for uncertain quantities with only a few possible outcomes.

In the direct response mode one asks for the probability level (cumulative probability) at a given value (e.g., what is the probability that next year's sales will be less than or equal to 3,000 units?), or asks for the value (fractile) corresponding to a probability (e.g., what is the level of sales that corresponds to a 10-percent probability?). The probability response can be given as an absolute number, 0.20; as a percentage, 20 percent; or can be expressed in a fractional way as "one in five" or "two in ten." The last way is particularly useful for small probabilities because the subject can discriminate more easily between "one in 100" and "one in 1000" than between the absolute numbers 0.01 and 0.001. Expressing a probability in the fractional form is closely related to expressing it in terms of odds, in particular for probabilities close to zero.

The direct response mode can also be used in a free format (making it a PV-method) where the subject either draws a picture of a density function or a cumulative distribution, or states a series of pairs of numbers (value and probability). The distribution can also be described in parametric form, e.g., a beta distribution with parameters 2 and 7.

Verbal encoding makes use of verbal descriptors for events (e.g., high, medium, and low production cost) in the first phase of the encoding. The descriptors are those that the subject is accustomed to. The interpretation of the descriptors is encoded in a second phase. This method might have some use for quantities that have no ordinal value scale. It can be viewed as a PV-method.

6. THE INTERVIEW PROCESS

While the structure of the interview process is still evolving, the following approach has been found quite effective. The process is divided into five phases.

- Motivating--Rapport with the subject is established and possible motivational biases are explored.
- Structuring--The structure of the uncertain quantity is defined.
- Conditioning--The subject is conditioned to think fundamentally about his judgment and to avoid cognitive biases.
- Encoding--This is the actual quantification of judgment in probabilistic terms.
- Verifying--The responses obtained in the encoding are checked for consistency.

Motivating

This phase has two purposes. The first is to introduce the subject to the encoding task. This may entail an explanation of the importance and purpose of probability encoding in decision analysis, as well as a discussion of the difference between deterministic (single number) and probabilistic (probability distribution) predictions.

The second purpose is to explore whether any motivational biases might operate. The interviewer and the subject should have an open discussion on what payoffs might be associated with the probability assignment as well as on possible misuses of the same information. The subject may be aware of misuses of single-number predictions, e.g., that they often are interpreted as "commitments." It should be pointed out that no commitment is inherent in a probability distribution. In fact, the distribution represents the complete judgment of the subject.

Structuring

The next step in the encoding process is to define and structure clearly the uncertain quantity. This quantity is assumed to be important to the decision. It should be defined as an unambiguous state variable and the definition should pass the clairvoyant test, i.e., a clairvoyant should be able to specify the outcome without asking additional questions for clarification. The structure may have to be expanded so that the subject does not have to model the problem further before making each judgment. It is also important to choose a scale that is meaningful to the subject.

The subject should be required to think the problem through carefully before the actual encoding session begins. He should decide what background information might be relevant (or irrelevant) to the problem. Otherwise, only the readily available information will be used initially, and new information may later rise to the surface in the course of the session and invalidate all prior answers. Even if it does not, however, the resulting distribution may be highly biased with respect to the subject's underlying judgment.

Conditioning

The aim of this phase is to head off biases that otherwise might surface during the encoding and to condition the subject to think fundamentally about his judgment. Basically, the phase should be directed toward finding out how the subject goes about making his probability assignments. This will reveal what information seems to be most available, what (if any) anchors are being used, what assumptions are made, and so on. The interviewer should thus watch out for (and make use of) the modes of judgment discussed in Section 4. The following are some suggestions for a checklist that we have found useful in many applications.

The subject can be asked to specify the most important bases for his judgment. These may often be values from the current business plan or results from previous years. Such values could then be expected to act as anchors and often lead to a central bias.

Asking the subject what he is taking into account will show what information becomes most easily available. The interviewer can also make use of availability if he suspects the subject to have a central bias. He can then ask the subject to compose scenarios that would produce extreme outcomes.

Uncertain quantities sometimes represent averages, such as average productivity or average reliability. The interviewer should then try to determine whether the distribution assigned by the subject really is a distribution for the average or a distribution for an individual unit. (The reason is that people often have difficulty in discriminating between the two situations.) If the latter is the case, it is probably best to use the resulting distribution and restructure the model. Representativeness may come into play in another situation when one is concerned with revising a probability assignment in the light of new information. The best way to handle such a situation is often to ask for the probability distribution of the quantity without the new information and for the probability of the information conditional on the outcome of the quantity; it is then a matter of applying probability calculus to obtain the distribution for the quantity given the information.

It is important to specify all assumptions (conditionalities) that will underlie the probability distribution, as well as those factors that the subject is supposed to integrate into his judgment. The structure may sometimes be changed because some conditionalities have been made explicit. The encoding may then be made conditional on different sets of assumptions, and the probabilities that the various assumptions will hold are then encoded separately (from the current subject or from someone else).

When a subject is assigning a probability to the occurrence (or non-occurrence) of some event (e.g., that a product will be successful in the market), he may base his assignment on whether he can generate plausible scenarios leading to the occurrence of the event in question. Asking him to state the basis for his probability assignment may reveal that the coherence of such scenarios has been an important factor. The interviewer may then want to generate more scenarios that would or would not lead to the occurrence of the event. For example, simply devising a scenario that implies the opposite outcome might considerably change the first probability assignment.

Encoding

The procedure outlined for this phase of the interview process is suggested as a guideline. It is primarily based on the use of the probability wheel as the encoding technique. Responses may often indicate a need to return to the tasks in the previous three phases.

Begin by asking for extreme values for the uncertain quantity. Then ask for scenarios that might lead to outcomes outside of the extremes.

Also ask for the probabilities (or odds) of outcomes outside of the extremes. This deliberate use of availability is to counteract the central bias that is otherwise likely to occur. Eliciting the scenarios makes them available to the subject, and he is then likely to assign higher probabilities to extreme outcomes; this has the effect of increasing the variability in his assigned distribution.

Next turn to the wheel. Take a set of values and encode the corresponding probability levels. Do not choose the first value in a way that may seem significant to the subject, otherwise you will cause him to anchor on that value. In particular, do not begin by asking for a likely value and then encode the corresponding probability level. Make the first few choices easy for the subject so that he will be comfortable with the task. Plot all responses as points on a cumulative distribution and number them sequentially. An example is shown in Figure 3. This will point out any inconsistencies and will also show the gaps in the distribution that need one or more additional points.

The interval technique can be used next to generate values for the median and the quartiles. The order of the questions and of the different types of questions should be determined by the situation. The length of the encoding session depends on the ease with which the subject can answer the questions and on the convergence to coherent




FIGURE 3 GRAPHICAL REPRESENTATION OF RESPONSES

responses. The length of the session also depends on the importance of the variable. The interviewer should be aware of attention shifts (for example, the shift of attention between the encoding process and the actual problem), changes in the subject's modeling of the situation, and the appearance of new information.

Each response will lead to a point on a cumulative distribution. The importance of the variable for the decision problem at hand determines the number of points to encode. After enough points have been encoded, a curve should be fitted to the points. An example is shown in Figure 4.




FIGURE 4 EXAMPLE OF A CURVE FITTED TO RESPONSES

Verifying

The judgments are tested in the last phase to see if the subject really believes them. If needed, there may be iteration through some of the stages mentioned above.

A graphical representation of the responses as points on a cumulative distribution and an interpretation of this distribution (perhaps in terms of a density function) provides an important test and feedback. An examination of the distribution itself cannot show whether or not the distribution agrees with the subject's judgment. However, it can show

implications of the subject's responses and thereby provide feedback. For instance, the plotted distribution may turn out to be bimodal whereas the subject may state that he believes the distribution to be unimodal. If some responses are not consistent with the subject's final judgment, they will have to be modified.

A second part of the verification process is based on a sequence of pairs of bets. Each pair is chosen so that the two bets would be equally attractive if the curve from the preceding phase is consistent with the subject's judgment.

There should be a number, say three to five, of such indifference responses before the process is ended. This provides the subject and the interviewer with confidence that the curve represents the subject's judgment.

Other Methods

It should be clear that the encoding methods discussed in this paper stress the interaction between interviewer and subject. We find it dangerous to have the subject assign a probability distribution without the help of an analyst. This is true even for subjects that are well trained in probability or statistics. The main reason is that it is difficult to avoid serious biases without an analyst present.

Questionnaires have dubious value for the same reasons, namely that they eliminate the interaction between analyst and subject. Questionnaires can be used as a first approximation to the encoding process, but the subjects interviewed should preferably be experienced in probability encoding.

An interactive computer interview can make use of iterative methods, such as the interval method, and thereby avoid some pitfalls with a direct response mode, but the personal interaction is still missing. We do not recommend using such a program unless the subject has been through a number of actual interviews regarding similar uncertain quantities. It should be avoided for new quantities even though the subject may have had long experience with the computer interview. An example of an interactive computer program is the Probability Encoding Program (PEP) developed by the Decision Analysis Group at Stanford Research Institute.

There are procedures that ask the subject for the parameter of a named distribution; e.g., a normal distribution or a beta distribution. Our experience indicates that subjects will give such parameters, but

they usually do not understand the full implications. We consider the choice of named distributions a modeling consideration and believe that it should not be made part of the encoding process.

7. SUMMARY AND CONCLUSION

In this paper we have presented a general methodology for probability encoding. This has been done in the spirit of a checklist or rules to remember rather than a cookbook with well-structured recipes on many different procedures. The ultimate form of the procedure used for a particular uncertain quantity will depend on the quantity, its importance for a decision, the subject, and the interviewer. The interviewer is an important factor since his function is perceiving problems or biases that the subject might have and adapting the encoding procedures accordingly.

The methodology differs considerably from what is generally described in the literature. The following are some important points that we want to stress in this summary. First, the pre-encoding stage is essential and may even take longer than the quantification of probabilities. The purpose is to establish rapport between the interviewer and the subject, to make sure the problem is well defined, and to ensure that possible sources for bias have been detected. Second, probability assignments should be inferred from choices among bets that require only an ordinal judgement rather than from answers to direct questions. Third, reference processes, such as the probability wheel, appear to provide an effective encoding technique for most subjects.

There are many special topics related to probability encoding that may be relevant in particular situations. They include encoding of discrete distributions, encoding of probabilities for rare events, accuracy and calibration of probability assignments, the use of multiple experts. These topics are presently being researched by the SRI Decision Analysis Group and discussion will be deferred to another paper.

REFERENCES

1. Howard, R. A., "Decision Analysis: Applied Decision Theory," in D. B. Hertz and J. Melese (eds.), Proceedings of the Fourth International Conference on Operational Research, Wiley, New York, 1966 (1968), pp. 55-71.
2. _____, "The Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4 (1968), 211-219.
3. Kahneman, D., and Tversky, A., "Subjective Probability: A Judgment of Representativeness," Cognitive Psychology, Vol. 3 (1972), 430-454.
4. _____, "On Prediction and Judgment," Oregon Research Institute Research Bulletin, Vol. 12 (1972), No. 4.
5. Peterson, C. R., and Beach, L. R., "Man as an Intuitive Statistician," Psychological Bulletin, Vol. 68 (1967), 29-46.
6. Rapoport, A., and Wallsten, T. S., "Individual Decision Behavior," Annual Review of Psychology, Vol. 23 (1972), 131-176.
7. Staël von Holstein, C.-A.S., Assessment and Evaluation of Subjective Probability Distributions, The Economic Research Institute at the Stockholm School of Economics, Stockholm, 1970.
8. _____, "A Bibliography on Encoding of Subjective Probability Distributions," Unpublished manuscript, Stanford Research Institute, 1972.
9. _____, "A Tutorial on Decision Analysis," Unpublished manuscript, Stanford Research Institute, 1972.
10. Tversky, A., and Kahneman, D., "Availability: A Heuristic for Judging Frequency and Probability," Oregon Research Institute Research Bulletin, Vol. 11 (1972), No. 6.

RISK PREFERENCE

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

RISK PREFERENCE

Ronald A. Howard

THE PHENOMENON OF RISK PREFERENCE

One of the most difficult choices a person or an organization can face is how to choose between propositions whose outcomes are uncertain. For example, which proposition would you prefer, one that pays \$500 if a coin falls Heads or one that pays \$10,000 if a pair of dice come up "snake eyes" (both show a one)? In the corporate case the payoffs may be measured in millions of dollars and the probabilities may result from the evaluation of marketing or research programs, but the same problem remains. Our purpose is to develop a theory of risk preference and to show how it may be applied to make choices under uncertainty on a consistent, logical basis.

A LANGUAGE OF RISK PREFERENCE

To discuss risk preference we need a language.

A LOTTERY

We shall describe any uncertain proposition as a lottery. "Lottery" is a technical term, not a gambling term. When we participate in a lottery we shall receive exactly one of a specified set of prizes or prospects. Each prize has associated with it the probability that we shall receive that prize. Thus the proposition described above of tossing a coin for \$500 is a lottery with prizes \$0 and \$500, each to be received with probability 1/2. A graphical description of this lottery appears as Figure 1.

There is no need for the prizes in a lottery to be in any way commensurate. For example, Figure 2 shows a lottery where the prizes are the Hope diamond, a ton of cheese, and a case of pneumonia to be received with probability p_A , p_B , and p_C where, of course, $p_A + p_B + p_C = 1$. We might also choose to describe this lottery in the notation $(p_A, A; p_B, B;$




FIGURE 1 A COIN-TOSSING LOTTERY.

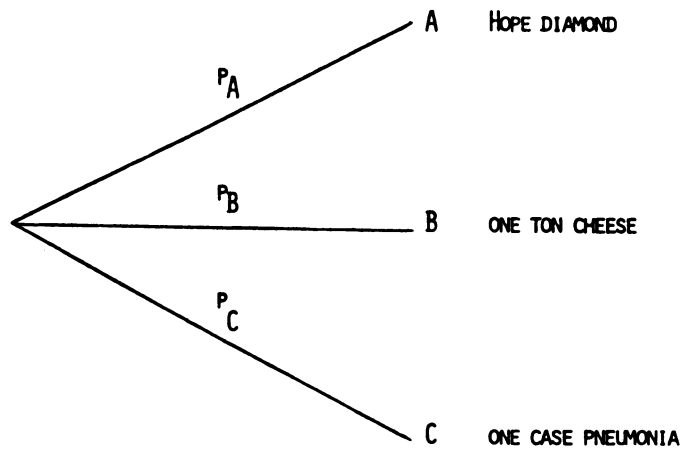


FIGURE 2 AN UNUSUAL LOTTERY.

p_C, C). Depending on the specification of the probabilities, this lottery could be very favorable or very unfavorable.

LOTTERIES ON A COMMODITY

Though it is interesting to speculate about such unusual lotteries, the lotteries of greatest practical interest are those whose prizes are all measured in terms of some commodity like time or, most usually, money. In this case a lottery is simply a random variable. If the lottery is described by a discrete random variable x then we have a choice of two representations shown in Figure 3, the tree or the probability mass function. Either representation shows that the payoffs $-\$10$, $\$20$, and $\$50$ will be received with probabilities 0.1, 0.3, and 0.6.

If the lottery is described by a continuous random variable x , then the tree representation does not work and we can specify the lottery only by indicating the probability density function or cumulative probability distribution of the random variable. For example, if the payoff is given by a normal density function with mean m and standard deviation σ , then we could use either the density function or cumulative distribution of Figure 4 as its description.

EXPECTATION

One measure of a commodity lottery x that might be suggested as a basis for comparing lotteries is the expectation \bar{x} computed by multiplying the amount of each prize by its probability and summing over all prizes. This calculation is performed in Figure 5 for the two lotteries we originally posed. Although the lottery using the dice has expectation $\$27.78$ higher than that using the coin, it is not at all clear that one would or should prefer it. Indeed, from a behavioral point of view many people would prefer the coin lottery. Expectation just does not capture the way most people think about uncertain propositions. Our task is therefore to replace expectation with a more useful concept.

CERTAIN EQUIVALENT

Lotteries can be meaningfully compared by using the idea of a certain equivalent. The certain equivalent of a lottery is the selling price for the lottery, the amount you would have to be paid to give up the lottery if you already owned it. Of course, you would really like to be paid much more than the lottery is worth to you so we must define the certain

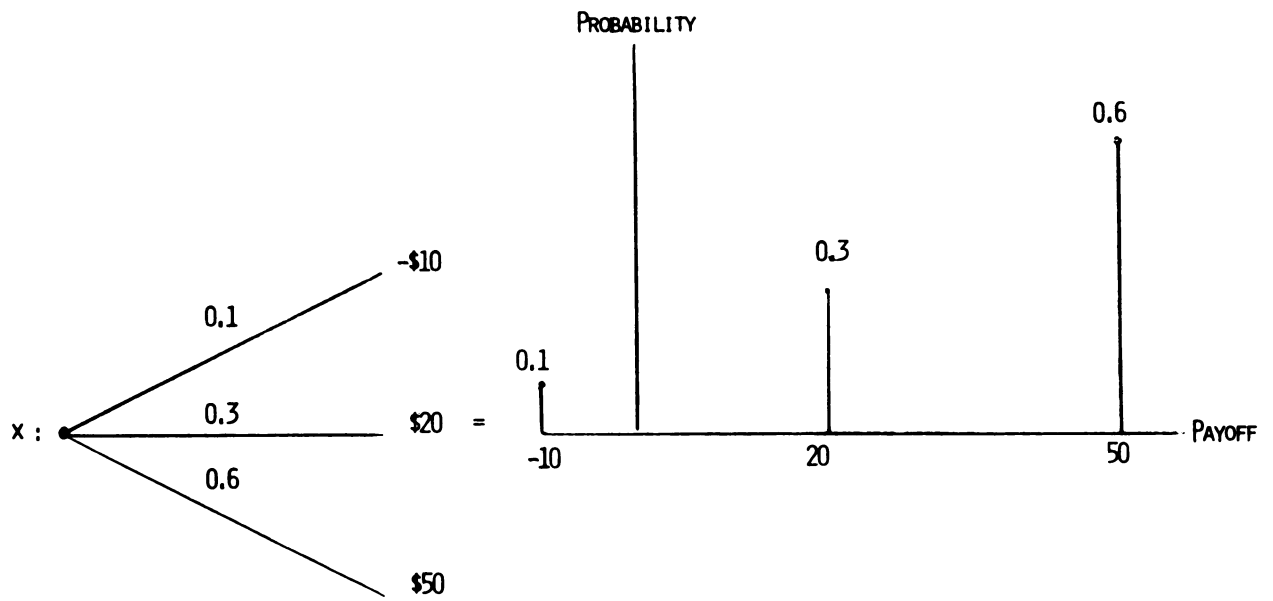


FIGURE 3 A DISCRETE-VARIABLE LOTTERY.




FIGURE 4 A CONTINUOUS-VARIABLE LOTTERY.

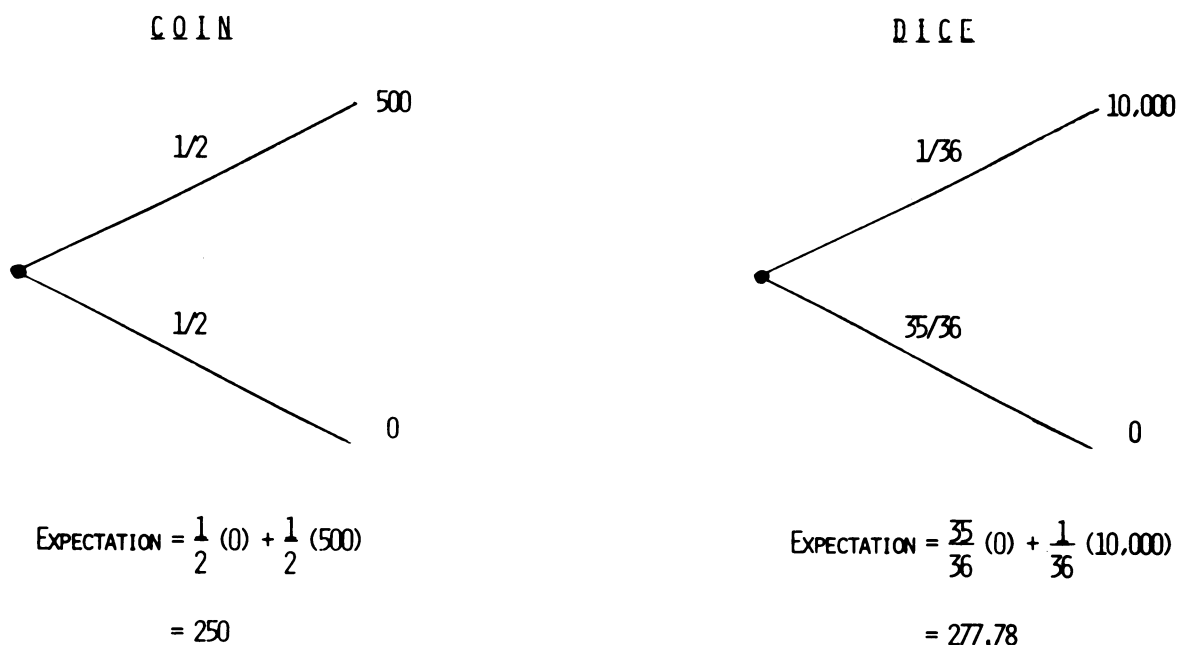


FIGURE 5 COMPUTATION OF EXPECTATION.

equivalent more carefully as the least you would accept in selling the lottery. If someone offered you slightly more than the certain equivalent, you would take it; if he offered slightly less you would refuse his offer and keep the lottery.

Now let us use this idea. Suppose you owned the right to toss the coin for \$500. Someone has written you a contract that you can take to any bank, whereupon the teller will flip a handy coin and pay you \$500 if it falls Heads. That would be a very nice lottery to own--what would you sell it for? While the answers of individuals will differ because of their different financial circumstances and attitudes, let us suppose that you say \$150. That means that you would certainly sell for \$151 and not for \$149. In this case we say that your certain equivalent for the coin lottery is \$150.

Figure 6 illustrates the meaning of certain equivalent graphically. What you have said is that you are indifferent between playing the lottery and receiving \$150 for sure. The diagram makes this statement by showing that you are indifferent (symbol \sim) between the coin lottery and a lottery that pays \$150 with probability one.




FIGURE 6 CERTAIN EQUIVALENT OF THE COIN LOTTERY.

Figure 7 shows similarly that you have stated a \$100 certain equivalent for the dice lottery. If you owned this lottery you would sell it for any offer over \$100 and keep it for any offer under \$100. The reason many people would have a lower certain equivalent for the dice lottery than they would for the coin lottery is that the dice lottery seems like a real "long-shot." However, the point is not what certain equivalents you state, but rather that they constitute a basis for choosing between the lotteries. Clearly, someone who expressed the certain equivalents




FIGURE 7 CERTAIN EQUIVALENT OF THE DICE LOTTERY.

we have used in our illustration would prefer the coin lottery to the dice lottery if he were offered a choice between the two simply because the coin lottery has a higher certain equivalent.

If a lottery is described by a random variable x , then we use \tilde{x} as the notation for its certain equivalent. Since certain equivalents are typically less than expectations, we define the difference between the expectation \bar{x} and the certain equivalent \tilde{x} as the risk premium x_p ,

$$x_p = \bar{x} - \tilde{x} \quad .$$

The risk premium is the amount of expectation that the individual is willing to forego in order to avoid risk. Thus in our example, the risk premium for the coin lottery is $250 - 150 = \$100$, whereas the risk premium for the dice lottery is $277.78 - 100 = \$177.78$.

If $x_p = 0$ then $\bar{x} = \tilde{x}$ and we say that the individual is risk indifferent; if $x_p \neq 0$, we say he is risk-sensitive. If the risk premium x_p is positive, we call the individual "risk-averse"; if negative, "risk-preferring." The case of risk aversion is the case that will occupy our primary interest.

In principle, then, the question of choosing among uncertain propositions becomes one of assessing the certain equivalent of each proposition and then selecting the proposition with the highest certain equivalent. But we still face two problems. The first is how to be consistent in assessing certain equivalents. Is it not possible that we might assess different certain equivalents for the same lottery on successive days even when we have no wish to do so? We would certainly like to make sure that we have a consistent way to make these assessments.

However, the second problem is more serious. Direct assessment of certain equivalent means that every uncertain proposition must be brought directly to us for assessment of its certain equivalent before any decision can be made. We have no way of delegating a policy toward these matters so that our agents can make the choice between uncertain propositions for us in the sure knowledge that they are doing what we would do if we were there. The need for consistency and a procedure for delegation leads directly to the creation of a theory and methodology for the establishment of a risk policy.

THE THEORY OF RISK PREFERENCE

The theory of risk preference rests on a foundation of five axioms. These axioms state a set of beliefs that most people will wish to accept as a basis for decision-making. It is important to understand that they do not describe how people actually make decisions, but rather how they wish they made them.

THE AXIOMS

Axiom 1: Orderability. The decision maker must be able to state his preferences among the prizes of any lottery. That is, if a lottery has prizes A, B, and C he must be able to say which he likes first best, second best, etc. We use the symbol $>$ to mean "is preferred to." Then given any two prizes A and B, he must state either $A > B$, $A \sim B$, or $A < B$.

We further require that his preference be transitive. If $A > B$ and $B > C$ it is necessary that $A > C$. To violate transitivity means that the decision maker could be made into a "money pump." For example, suppose the decision maker should express the intransitive preferences $A > B$, $B > C$, $C > A$. Assume that his statement that he prefers one prize to another means that he would pay a small amount for substituting the preferred prize for the other one and that he is currently going to win prize C. Then we would pay us a small amount to substitute B for C, another small amount to substitute A for B, and finally a third small amount to substitute C for A. Thus we have extracted three payments from him and still left him with the same prize.

Individuals tend to violate transitivity when they consider only one aspect of a prize rather than the whole prize. For example, if A, B, and C are automobiles, A might be preferred to B for its performance, B to C for its convenience, and C to A for its durability. However, the real question is, if you could have only one, which would you pick and if you couldn't have that one, then what would be your second choice.

Axiom 2: Continuity. If the decision maker has expressed the transitive preference $A > B > C$, then we must be able to construct a lottery with prizes A and C and determine a probability p of winning A such that he is indifferent between receiving B for certain and participating in the lottery. This axiom is stated graphically in Figure 8. We observe that for the appropriate p , B is the certain equivalent of the lottery.

For example, suppose the prizes A, B, and C represent winning \$100, \$10, and \$0 respectively. The decision maker might say that if he faced

IF $A \succ B \succ C$, THEN




FIGURE 8 THE CONTINUITY AXIOM.

a lottery that would pay the \$100 or nothing, the probability of winning the \$100 would have to be 0.25 before he would be indifferent between the lottery and receiving a sure \$10.

Axiom 3: Substitutability. If a decision maker has stated his certain equivalent of a lottery, then he must be truly indifferent between the lottery and the certain equivalent: The lottery and its certain equivalent must be interchangeable without affecting preferences. For example, in our discussion of Axiom 2 we postulated that an individual had a \$10 certain equivalent for the lottery (0.25, 100; 0.8, 0). Then if this individual owed us \$10, he would have to be satisfied if we offered to pay him \$100 if a coin fell Heads twice in two tosses and otherwise nothing. This is a "Do you really mean it?" axiom.

Axiom 4: Monotonicity. If a decision maker has a preference between two prizes and if he faces two lotteries which each have these two prizes as their only prizes, then he must prefer the lottery which produces the preferred prize with the higher probability. Figure 9 states the axiom graphically. If $A \succ B$, then the decision maker must prefer the lottery $(p, A; 1 - p, B)$ to the lottery $(p', A; 1 - p', B)$ if and only if $p > p'$. This axiom is so reasonable it needs no further discussion.

IF $A \succ B$, THEN




FIGURE 9 THE MONOTONICITY AXIOM.

Axiom 5: Decomposability. Sometimes an uncertain proposition has a complicated form. A lottery may have other lotteries as prizes. We call these situations "compound lotteries." For example, the Irish Sweepstakes is a compound lottery because first you must draw a ticket of a horse in a horse race and then your horse must win the race.

This axiom states that when faced with such a proposition, the decision maker will consider only the final prize he might win and then compute the probability of winning each prize using the laws of probability. Thus he will reduce a compound lottery to the simple type of lottery we have previously discussed.

Figure 10 illustrates the procedure. The first lottery pays B with probability $1 - p$ and with probability p permits participating in a lottery with probability q of winning A and $1 - q$ of winning B. The decision maker must say that his probability of winning A is just p times q and hence his probability of winning B is $1 - pq$ as shown in the simple lottery.

Sometimes we call this the "no fun in gambling" axiom because the momentary suspense in successive resolutions of the compound lottery is




FIGURE 10 THE DECOMPOSABILITY AXIOM.

considered to have no value or cost. A gambler who satisfied this axiom could be very efficient in his gambling. When he thought of spending an evening in Las Vegas, he could write a book stating just what games he would play and what he would do depending on the amount he won or lost. Then he could take this book to the doorman of a casino who would thank him and ask him to wait a moment. The book would be submitted to the casino's computer which would then simulate the possible evening's plays using the probabilities associated with the various games of chance. Then the doorman would reappear and announce, "You have lost \$200--a check will be acceptable." It is doubtful that this method of wagering will ever catch on in the gambling fraternity, but it would be much cheaper to operate.

Incidentally, if a gambler does assign value to the suspense associated with the gambling process, he can include this value with the prizes in each lottery and still use the theory of risk preference that we are developing.

These are the five axioms of risk preference. If you accept them, they will have far-reaching consequences on the way you choose among uncertain propositions.

THE UTILITY FUNCTION

We shall soon show that an individual whose preferences satisfy the utility axioms may encode these preferences in a "utility function" that assigns a utility number to every prize. This utility function has two important properties:

- 1) The utility of any lottery is the expected utility of its prizes.
- 2) If the decision maker prefers one lottery to another, then it must have the higher utility.

We can think of the utility function as a "preference thermometer." The utility numbers have no meaning in themselves; they serve only to compare the desirability of lotteries. Because of the linear properties of expectation, we can multiply the utility function by any positive number and add any constant to all utilities without changing the preferences they express.

If all prizes are measured in terms of a commodity, then the utility function can be expressed by a curve that assigns a utility number to every value of the commodity. If, furthermore, this commodity is such that more is always better, for example money (you can always give it away if you don't like it) then the utility curve will be monotonically increasing.

USING THE UTILITY CURVE

We shall demonstrate how to use a utility curve before proving that the axioms imply the existence of the utility function. Later we shall show how an individual can establish a utility curve for himself or for his organization.

Consider the utility curve shown in Figure 11. It assigns a utility number $u(x)$ to various dollar amounts x ranging from 0 through \$10,000; however, because of the wide range, only a portion of the curve is shown in detail. The curve is normalized arbitrarily so that the utility of zero is zero, $u(0) = 0$, and the utility of 10,000 is 1, $u(10,000) = 1$. This normalization is logically the same as fixing two points on a temperature scale as the freezing and boiling points of water. We observe that the curve is concave downwards; this type of curve is characteristic of a risk averter.




FIGURE 11 A UTILITY CURVE.

Certain observations on the use of utility curves are important. First, we note that no difficulty would have been introduced if the lottery had more than two prizes--the expectation would simply be computed over all prizes. Second, we see that the scaling of the utility curve has no effect on the choice. For example, if all utilities under discussion were multiplied by 10 and then increased by 100, the utilities of the coin and dice lotteries would become 100.4 and 100.278, but the coin lottery would still be preferred.

The lack of inherent meaning in the utility numbers demonstrates the need for the idea of certain equivalent. Since the individual is indifferent between the certain equivalent and the lottery, the certain equivalent must have the same utility as the lottery. In other words, we find the certain equivalent of a lottery by seeing what dollar value the utility curve shows corresponding to the utility of the lottery. For example, the coin lottery had a utility of 0.04. The utility curve of Figure 11 shows that a dollar value of 150 corresponds to this utility; therefore \$150 is the certain equivalent of the coin lottery. Similarly, the 0.0278 utility of the dice lottery implies the \$100 certain equivalent of this lottery. These certain equivalents are recorded in Figure 12; they agree with those expressed earlier for these lotteries.

Since monotonic utility curves mean that a higher utility will imply a higher certain equivalent, the certain equivalent can serve just as well as the utility number as a basis for comparing lotteries. You simply choose the lottery with the highest certain equivalent. However, the certain equivalent has the further advantage of indicating the rough strength of the preference. Thus the fact that the coin lottery has a certain equivalent \$50 higher than that of the dice lottery means, speaking loosely, that is about \$50 more valuable. We shall soon see under what conditions this statement is precisely correct.

PROOF OF THE PROPERTIES OF THE UTILITY FUNCTION

We shall now see that only a simple proof is required to establish the properties of the utility function that we find so valuable in practice. We shall consider the case where only a finite number N of possible prizes can exist in composing any lottery. Since this will include the case of every conceivable dollar amount, this is not a practical limitation. We shall let R_i designate the prize (reward) i and use the orderability axiom to label the prizes so that

$$R_1 > R_2 > R_3 \dots > R_N \quad .$$

That is, R_1 is the most preferred prize; R_N is the least preferred prize. Such labelling can be performed without loss of generality. The case where two prizes are equally preferred is simply expressed by \sim with no difficulty.

The continuity axiom provides that since R_1 is preferred to any other prizes R_i and since R_i is preferred to R_N , $i = 2, 3, \dots, N - 1$, we can compose a lottery with R_1 and R_N as prizes and find a probability u_i of winning R_1 such that the individual is indifferent between receiving R_i for sure and participating in the lottery. Since u_i is a probability, $0 \leq u_i \leq 1$, $i = 1, 2, \dots, N$. Figure 13 illustrates the typical equivalent lottery.

Now suppose that we consider any lottery A with prizes composed from the prize set. Even if A is a compound lottery, we can use the decomposability axiom to represent it in the form shown in Figure 14, namely, as a simple lottery with some possibly zero probability p_i^A of winning each prize R_i .

Next we use the substitutability axiom to replace each prize in this representation by its equivalent lottery developed as in Figure 13. That




FIGURE 13 APPLICATION OF THE CONTINUITY AXIOM.




FIGURE 14 APPLICATION OF THE DECOMPOSABILITY AXIOM.

is, we are going to substitute the lottery for the equivalent certain prize. The result appears in Figure 15. We see that there are only two ultimate prizes in this compound lottery, the best of all prizes R_1 and the worst of all prizes R_N .

We now employ the decomposability axiom to determine the equivalent simple lottery. The probability of winning R_1 is just

$$\sum_{i=1}^N p_i^A u_i \quad .$$

Therefore the lottery A can be represented in the equivalent form shown at the left of Figure 16.

Suppose that we had performed the same operations on another lottery, lottery B. We would have obtained the same type of result as shown in the right of Figure 16. Both A and B would be represented as lotteries on the best and the worst prize. The probability of winning the best prize would be in each case just the sum over all possible prizes of the




FIGURE 15 APPLICATION OF THE SUBSTITUTABILITY AXIOM.




FIGURE 16 TWO LOTTERIES REDUCED TO A STANDARD FORM.

probability of winning each prize times the number u_i for that prize. By the monotonicity axiom, since $R_1 > R_N$, lottery A must be preferred to lottery B if and only if it has a higher probability of winning R_1 , that is, $A > B$ if and only if

$$\sum_{i=1}^N p_i^A u_i > \sum_{i=1}^N p_i^B u_i \quad .$$

We have established the result we seek. We interpret u_i as the utility of the i^{th} prize. We see that we have proved that one lottery will be preferred to another only if it has a higher utility, where the utility of the lottery is computed as the expected utility of its prizes. We also observe that adding a constant to all u_i 's and multiplying them by any positive constant will leave the expressed preference unchanged. It is interesting to note in passing that we can always normalize the utility function so that it may be interpreted as specifying for any given prize the probability of winning the best prize as opposed to the worst prize in a two-prize lottery whose certain equivalent is the given prize.

ASSESSMENT OF RISK PREFERENCE

A utility function would not be of much use if it were difficult to determine its form for each individual or organization. We shall now discuss just how a utility function can be assessed.

THE ASSESSMENT PROCEDURE

We shall concentrate on establishing an assessment procedure for a utility curve on a commodity, in particular, money. It is easy to deal with a two-prize lottery like that in Figure 17. The lottery pays b with probability p , and a with probability $1 - p$; the certain equivalent of the lottery is \tilde{x} . The four numbers \tilde{x} , p , a , and b express a preference and, consequently, a utility equivalence that we can use in developing the utility curve, namely,

$$u(\tilde{x}) = p u(b) + (1 - p)u(a) \quad . \tag{1}$$

By choosing which three numbers will be supplied initially and therefore which remaining number will be specified by the individual, we have a




FIGURE 17 A STANDARD TWO-PRIZE LOTTERY.

choice of method for carrying out the assessment. Of course, we can always fix two points on the utility curve arbitrarily to begin the process.

Equiprobable Lotteries

In the first method we choose $p = 1/2$, and state two out of three of a , b , and \tilde{x} . Figure 18 illustrates typical questions. In the first case, prizes \$0 and \$100 are postulated and the individual is asked the certain equivalent. This permits evaluating the utility curve at the certain equivalent if the utilities of the prizes a and b are known. We therefore call this interpolation.

In the second case, \tilde{x} is given as 0, b as \$100, and the individual is asked to specify the lower prize that will make him indifferent. The wording would be "If I offered you a lottery that would cause you to win \$100 or x with equal probability, how small would x have to be before you would be indifferent between playing and not playing?". Presumably x would have to be a negative number. When it was specified, we could find the utility curve for an amount below \$0; hence the comment that it will allow us to extrapolate the utility curve downward.





FIGURE 18 EQUIPROBABLE LOTTERY UTILITY ASSESSMENT.

The third case is similar. The certain equivalent is given as \$100 and the lower prize as zero. The individual must specify the upper prize that would make him indifferent between playing and receiving the sure \$100. His answer will allow us to evaluate the utility curve for amounts over \$100 and therefore to extrapolate upward.

Assignment of Probability

The other method always requires that the individual specify the probability of winning that will make him indifferent; it is illustrated in Figure 19. The first question to be asked is "What probability of winning \$100 as opposed to nothing would you have to have before you would be indifferent to playing or receiving a sure \$50?". It would serve to determine the utility of \$50 in terms of utilities of \$0 and \$100; hence it is an interpolation. The second question would have the form "What probability of winning as opposed to losing \$100 would you have to have before you were indifferent to playing?". The utility of -\$100 could then be derived. The third question would allow finding the utility of \$200.



THE UTILITY OF THE CIRCLED NUMBER IS EVALUATED BY THE QUESTION.

FIGURE 19 ASSIGNMENT OF PROBABILITY UTILITY ASSESSMENT.

The assignment of probability method provides a valuable alternative to the equiprobable lottery method. However, they can both be better appreciated by using an example.

A UTILITY ASSESSMENT EXAMPLE

Let us suppose that we wish to assess some individual's utility curve for amounts of the order of hundreds of dollars. We shall be able to do this if the individual satisfies the five axioms of risk preference we have discussed. We might begin by assigning the utility 0 to the amount zero and the utility 1 to the amount \$100,

$$u(0) = 0 \quad , \quad u(100) = 1 \quad .$$

Of course we could have made many other choices for these quantities without changing what follows in any important way, but these choices will serve.

We shall use the equiprobable lottery method and begin by investigating the shape of the curve within the (0,100) region. We can ask him "What is your certain equivalent for an equiprobable lottery on zero and \$100?". He answers, "\$25." We have recorded this answer and subsequent ones in Figure 20. Next we ask him his certain equivalent for an equiprobable lottery on 0 and his answer to the last question, \$25. He replies, "\$10," as recorded in Figure 20. Then we ask him his certain equivalent for an equiprobable lottery on \$25 (his answer to the first question) and \$100. He sets this certain equivalent at \$40. These answers will allow us to determine the rough path of the curve.

To extrapolate below the (0,100) region, we ask him as question four what the losing prize of an equiprobable lottery with one prize of \$100 would have to be before he would be indifferent about playing. His answer is -\$30. To extrapolate above the (0,100) region, we ask what the winning prize of an equiprobable lottery with losing prize \$0 would have to be before he would be indifferent between playing and receiving a sure \$100. His judgment is \$400.

While these questions would suffice for our purposes, let us ask one more question to demonstrate the assignment of probability method. We shall ask what probability of winning \$100 as opposed to winning nothing would you have to have before you would be indifferent between playing and receiving a sure \$40. He states that a probability of 0.8 would cause him to have no preference. (Incidentally, the \$40 was selected because it was his answer to the third question.)

Now we shall process the information he has given us to determine the form of his utility curve for money in this range. The computations are performed under the corresponding question in Figure 20. Thus from his answer to the first question, the normalization of the utility curve, and Equation 1 we find immediately that the utility of \$25 is 0.5. By using this result and his answer to the second question, we determine that the utility of \$10 is 0.25. His answer to the third question allows us to find that the utility of \$40 is 0.75. These results are plotted in Figure 21. We see that the utility curve is generally concave downward in the (\$0,100) region, indicating that the individual is risk averse in this region.

To extrapolate below the (\$0,100) range we use the answer to the fourth question and find $u(-30) = -1$. The answer to the fifth question shows that $u(400) = 2$ and allows us to extrapolate above the (\$0,100) region. These results are also shown in Figure 21. We can readily appreciate that continuing the interrogation process using the same type of questions could allow us to determine the utility curve as closely as we like.




FIGURE 20 RESPONSES TO UTILITY ASSESSMENT QUESTIONS.




FIGURE 21 THE UTILITY CURVE.


Let us turn now to the answer to the sixth question. As we see, his assignment of 0.8 as the probability that causes indifference means that $u(40) = 0.8$. However, we have already determined from the third question that $u(40) = 0.75$. We have discovered an inconsistency: We have two different values for the utility of the same dollar amount. As a practical matter, this inconsistency is small. Indeed, it would be a miracle if anyone could answer a long series of questions of this type without a similar result.

Nevertheless, as a theoretical matter, the inconsistency is important, for we have proved that anyone who satisfies the utility axioms has a unique utility curve. The fact that the curve we have derived is not unique means that an axiom must have been violated. We have now to show that this is the case.



RESOLUTION OF THE INCONSISTENCY

We begin by recalling in Figure 22 his answers to questions 3 and 6. Notice that the lotteries both have certain equivalents of \$40. We

QUESTION 3



QUESTION 6



DECOMPOSABILITY





FIGURE 22 A BROKEN AXIOM.

shall use the individual's agreement to satisfy the axioms to place his answer to question 3 in an alternative form. Because the individual agrees to the substitutability axiom, we can replace the \$25 payment by the lottery known to be equivalent according to his answer to question 1. This replacement is illustrated in Figure 22. The resulting compound lottery has only two prizes, 0 and \$100. We use his agreement to satisfy the decomposability axiom to replace the compound lottery by the equivalent simple lottery on 0 and \$100 with probability 0.75 of winning the \$100. Thus the individual has said in his answer to question 3 that he is indifferent between a sure \$40 and a 0.75 probability of winning \$100 as opposed to nothing. However, his answer to question 3 shows that he is indifferent between a sure \$40 and a 0.80 chance of \$100 as opposed to nothing. Therefore he must be indifferent between a 0.75 probability of winning \$100 and the 0.80 probability of winning \$100 when the alternative in both cases is to win nothing. Since he prefers \$100 to nothing, he has said that he is indifferent between two lotteries when one has a higher probability of winning the better prize. This is a direct contradiction of the monotonicity axiom which requires that he prefer the lottery with the higher probability of winning the better prize. Therefore it is not surprising that we could not develop a unique utility curve from all six answers.

CONSTANT RISK PREFERENCE


The five axioms of risk preference are all that one must accept to have a unique utility curve. However, by accepting additional axioms, the form of the curve may be further restricted, with attendant advantages.

THE DELTA PROPERTY

Consider, for example, the following statement as a possible sixth axiom: An increase of all prizes in a lottery by an amount Δ increases the certain equivalent by Δ . We call this statement the "delta property." A graphical representation of the property appears in Figure 23.

The argument for the acceptance of the delta property can be cogent. Suppose, for example, that you have said that your certain equivalent for an equiprobable lottery on \$0 and \$100 is \$25. Then the person offering you the lottery agrees to pay you an additional \$100 regardless of outcome, so that your final payoffs will be \$100 and \$200 with equal probability. If you feel that your certain equivalent would now be \$125 and reason consistently in all such situations, then you satisfy the delta property.

IF



THEN




FIGURE 23 THE DELTA PROPERTY.

However, acceptance of the delta property has strong consequences. First, as we shall see, the utility curve is restricted to be either a straight line or an exponential. That is, $u(x)$ must have either the form

$$u(x) = a + bx$$

or the form


$$u(x) = a + be^{-\gamma x}$$

where a , b , and γ are constants. Second, the buying and selling prices of a lottery will be the same.



Buying and Selling Price Equivalence

We shall now demonstrate that the acceptance of the delta property requires equivalence of buying and selling prices by using Figure 24. Suppose that the individual is offered the opportunity to buy an equiprobable lottery on \$0 and \$100. If he pays b for the lottery, his new


BUYING PRICE LOTTERY



BY DELTA PROPERTY



SELLING PRICE LOTTERY



THEREFORE,
 $B = S$

FIGURE 24 DEMONSTRATION THAT DELTA PROPERTY IMPLIES EQUIVALENCE OF BUYING AND SELLING PRICES.

lottery will be an equiprobable lottery on 0 - b and 100 - b. His buying price is the value of b such that he is indifferent between buying and not buying, whereupon he will receive nothing. Therefore, his buying price is the value of b that establishes a zero certain equivalent for the lottery that remains after purchase. This is shown in the upper left of Figure 24.

If he already has the right to the equiprobable lottery on \$0 and \$100, his selling price s is the certain equivalent of the lottery, for he will prefer any offer greater than s to the lottery and prefer the lottery to any amount less than s. The upper right portion of Figure 24 is a graphical statement of this result.

Because the individual satisfies the delta property, we can maintain equivalence of preference by adding the same constant to the prizes of a lottery and to its certain equivalent. If we add the constant b in this manner to the buying price of a lottery, we obtain the form shown in the lower left portion of the figure. This shows that the individual's certain equivalent of an equiprobable lottery on \$0 and \$100 is the buying price b. However, the upper right portion of the figure demonstrates that the certain equivalent of this very same lottery is the selling price s. Since any lottery must have a unique certain equivalent, we have b = s, the selling and buying price must be the same. The method of proof clearly carries over to general lotteries.

Proof of the Consequences of the Delta Property

We shall now show that the delta property implies a linear or exponential utility curve. If an individual satisfies the delta property, the utility of a lottery with prizes augmented by Δ must be just the certain equivalent of the original lottery augmented by Δ for any Δ . If $f_x(\cdot)$ is the density function of the variable x describing the lottery, then we have

$$\int dx_0 f_x(x_0) u(x_0 + \Delta) = u(\tilde{x} + \Delta) \quad , \quad \text{for any } \Delta \quad . \quad (2)$$

If we differentiate both sides of this equation twice we obtain, successively,

$$\int dx_0 f_x(x_0) u'(x_0 + \Delta) = u'(\tilde{x} + \Delta) \quad (3)$$

and

$$\int dx_0 f_x(x_0) u''(x_0 + \Delta) = u''(\tilde{x} + \Delta) \quad . \quad (4)$$

Dividing the second of these equations by the first and setting $\Delta = 0$ produces

$$\frac{\int dx_0 f_x(x_0) u''(x_0)}{\int dx_0 f_x(x_0) u'(x_0)} = \frac{u''(x)}{u'(x)} .$$

Many different density functions $f_x(\bullet)$ will produce the same certain equivalent \tilde{x} . We may therefore consider the right-hand side of this equation to be fixed at some constant and ask what relation the ratio of derivatives of the utility curve on the left side must have to preserve the equation if $f_x(\bullet)$ is allowed to vary over all density functions with the same certain equivalent. The answer is that the ratio must be a constant,

$$\frac{u''(x)}{u'(x)} = -\gamma \quad -\infty < x < \infty \quad , \quad (5)$$

where we have chosen the constant to be the negative of the quantity designated by γ .

Now we are almost finished. Integration of Equation 5 produces

$$\begin{aligned} \ln u'(x) &= -\gamma x + k_0 \quad , \\ u'(x) &= k_1 e^{-\gamma x} \quad , \end{aligned} \quad (6)$$

where quantities of the form k_i are constants. If $\gamma = 0$, we have

$$\begin{aligned} u'(x) &= k_1 \quad , \\ u(x) &= k_1 x + k_2 \quad . \end{aligned}$$

The utility curve is linear. If $\gamma \neq 0$, then we integrate Equation 6 to obtain

$$u(x) = k_2 e^{-\gamma x} + k_3$$

The utility curve is exponential. Therefore we have shown that acceptance of the delta property implies restriction of the utility curve to the linear or exponential form.

THE EXPONENTIAL UTILITY CURVE

Satisfying the delta property means that the certain equivalent of any proposed lottery is independent of the wealth already owned. This wealth is just a " Δ " that does not affect the preference. Therefore the two utility curves that satisfy the delta property--linear and exponential--could be called wealth-independent utility curves.

A convenient form for parameterizing the exponential utility curve is

$$u(x) = \frac{1 - e^{-\gamma x}}{1 - e^{-\gamma}} \quad . \quad (7)$$

This form provides the normalization $u(0) = 0$, $u(1) = 1$. As γ approaches zero, we find

$$\lim_{\gamma \rightarrow 0} u(x) = \lim_{\gamma \rightarrow 0} \frac{1 - e^{-\gamma x}}{1 - e^{-\gamma}} \stackrel{\text{L'Hôpital}}{=} x \quad .$$

Therefore the linear case is contained as a specialization of this form and we can consider Equation 8 to be the standard equation for wealth-independent utility curves.

The quantity γ is called the "risk aversion coefficient." When $\gamma = 0$, the individual is risk indifferent. When γ is positive, the individual is risk averse; when negative, risk preferring. Thus acceptance of the delta property ultimately leads to the characterization of risk preference by a single number, the risk aversion coefficient. We sometimes call such an individual a "constant risk averter." We can think of the risk aversion coefficient as playing the same role in risk preference that the single discount factor does in time preference.

A Closed-Form for the Certain Equivalent

Using the exponential utility curve allows us to develop a closed-form expression for the certain equivalent of any lottery. Suppose we

consider a lottery on x described by the density function $f_x(\bullet)$. The certain equivalent of this lottery, \tilde{x} , must satisfy the equation

$$u(\tilde{x}) = \int dx_0 f_x(x_0) u(x_0) \quad .$$

If we substitute the form for $u(\bullet)$ shown in Equation 7, we have

$$\frac{1 - e^{-\gamma\tilde{x}}}{1 - e^{-\gamma}} = \frac{\int dx_0 f_x(x_0) [1 - e^{-\gamma x_0}]}{1 - e^{-\gamma}}$$

or

$$e^{-\gamma\tilde{x}} = \int dx_0 f_x(x) e^{-\gamma x_0} = \overline{e^{-\gamma x}} \quad .$$

Then,

$$\tilde{x} = -\frac{1}{\gamma} \ln \overline{e^{-\gamma x}} = -\frac{1}{\gamma} \ln f_x^e(\gamma) \quad , \quad (8)$$

where we have used $f_x^e(\bullet)$ to represent the exponential transform of the density function $f_x(\bullet)$. The certain equivalent of any lottery is therefore the negative reciprocal of the risk aversion coefficient times the natural logarithm of the exponential transform of the variable evaluated at the risk aversion coefficient.

As γ approaches zero, this expression becomes

$$\begin{aligned} \tilde{x} &= \lim_{\gamma \rightarrow 0} \frac{-\ln f_x^e(\gamma)}{\gamma} = \lim_{\gamma \rightarrow 0} \frac{-f_x^e'(\gamma)}{f_x^e(\gamma)} \\ &= \bar{x} \quad . \end{aligned}$$

The certain equivalent of any lottery to a risk indifferent individual is the mean.

To illustrate the computation of a certain equivalent, consider the lottery $f_x(\bullet)$ described by the normal distribution with mean m and standard deviation σ ,

$$f_x(x_0) = f_n(x_0 | m, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x - x_0)^2}{2\sigma^2}} \quad -\infty < x_0 < \infty ,$$

with exponential transform

$$f_x^e(\gamma) = \int_{-\infty}^{\infty} dx_0 f_x(x_0) e^{-\gamma x_0} = e^{-\gamma m + \frac{\gamma^2 \sigma^2}{2}} .$$

According to Equation 8, the certain equivalent of this lottery is

$$\begin{aligned} \tilde{x} &= -\frac{1}{\gamma} \ln f_x^e(\gamma) = -\frac{1}{\gamma} \ln \left[e^{-\gamma m + \frac{\gamma^2 \sigma^2}{2}} \right] \\ &= m - \frac{1}{2} \gamma \sigma^2 . \end{aligned}$$

The certain equivalent of a normal lottery to a constant risk averter with a risk aversion coefficient γ is just the mean minus one-half γ times the variance.

We can use Equation 8 to determine the effect of playing two lotteries in succession. Suppose that the first lottery pays x_1 , the second x_2 , and that we let x be the total winnings $x = x_1 + x_2$. The certain equivalent of x is given by Equation 8 as

$$\begin{aligned} \tilde{x} &= -\frac{1}{\gamma} \ln \overline{e^{-\gamma x}} = -\frac{1}{\gamma} \ln \overline{e^{-\gamma(x_1 + x_2)}} \\ &= -\frac{1}{\gamma} \ln \overline{e^{-\gamma x_1} e^{-\gamma x_2}} . \end{aligned}$$

If the two variables x_1 and x_2 are independent, then the expectation of the product of their exponentials will be the product of their expectations. We shall make this assumption:

$$\begin{aligned}
\tilde{x} &= -\frac{1}{\gamma} \ln \overline{e^{-\gamma x_1} e^{-\gamma x_2}} \\
&= -\frac{1}{\gamma} \ln \overline{e^{-\gamma x_1}} - \frac{1}{\gamma} \ln \overline{e^{-\gamma x_2}} \\
&= \tilde{x}_1 + \tilde{x}_2 \quad .
\end{aligned}$$

If a constant risk averter engages in several independent lotteries, his certain equivalent for the whole prospect is the sum of the certain equivalents for each individual lottery.

SUMMARY

This discussion of risk preference has been all too brief. We have been able to demonstrate the need for a theory of risk preference and to show how it is developed. We have also discussed how to assess an individual's attitude toward risk and how to use his assessment in selecting among lotteries. Finally, we have introduced an important special type of risk attitude--constant risk aversion--and studied some of its properties.

Perhaps the most important conclusion we could reach is that while the phenomenon of risk preference is very novel, it can be treated with the same ease and precision we apply to other aspects of decision-making processes. It is only a matter of time until all organizations have an expressed risk preference policy.

THE DEVELOPMENT OF A CORPORATE RISK POLICY FOR CAPITAL INVESTMENT DECISIONS

Carl S. Spetzler
Strategic Decisions Group

Reprinted with permission from the IEEE Transactions on Systems Science and Cybernetics,
Vol. SSC-4, No. 3, September 1968. pp. 279-300

Copyright © 1968 by Institute of Electronics and Electrical Engineers.

The Development of a Corporate Risk Policy for Capital Investment Decisions

CARL S. SPETZLER

Abstract—A corporate utility function plays a key role in the application of decision theory. This paper describes how such a utility function was evolved as a risk policy for capital investment decisions.

First, 36 corporate executives were interviewed and their risk attitudes were quantified. From the responses of the interviewees, a mathematical function was developed that could reflect each interviewee's attitude. The fit of the function was tested by checking the reaction of the interviewees to adjusted responses. The functional form that led the interviewees to prefer the adjusted responses to their initial responses was finally accepted. The mathematical form of the function was considered a flexible pattern for a risk policy. The assumption was made that the corporate risk policy would be of this pattern.

With the pattern for a risk policy set, it was possible to simplify the method of deriving a particular individual's risk attitude. Using the simplified method, the corporate policy makers were interviewed once more. The results from these interviews were then used as a starting point in two negotiation sessions. As a result of these negotiation sessions, the policy makers agreed on a risk policy for trial purposes. They also agreed to develop a number of major projects using the concepts of risk analysis and the certainty equivalent.

I. INTRODUCTION

CAPITAL investment analysis is the analytic process of reaching a decision between alternative investment projects. The decision process consists of three basic steps:

- 1) the gathering of information about each project's net effects on corporate profits through estimation and forecasting of the relevant factors;
- 2) the combining of the quantitative information into a decision criterion; and
- 3) the reaching of the decision on the basis of the calculated decision criterion combined with judgment on the nonquantified information.

These steps will simply be referred to as the gathering of information, the calculations, and the decision.

Traditionally, the information about a project's effects on profits has been lumped into two categories: the quantitative or tangible information and the nonquantified or intangible information. Criteria based on tangible information have been developed and are widely accepted. The most notable of these criteria are the present value of a

project's net effects on profits and the rate of return on the investment. However, the intangible information poses major problems in the decision process.

1) Intangible information is difficult to communicate. Therefore, since the decision makers are seldom also the estimators, much of the intangible information can be lost or distorted in communication.

2) It is virtually impossible to formulate a consistent decision strategy toward intangibles. Thus the intangibles are weighted differently from one project to the next.

In other words, the decision maker is commonly confronted with incomplete (or distorted) intangible information and, in addition, has no consistent way of incorporating the intangible information in his decision.

It follows that, to reduce these problems in project analysis, as much information as possible must be quantified. Most analysts are aware of this and try to quantify as many of the relevant variables as possible. However, it is still common practice to estimate only single value forecasts for these variables, and to consider the information about the uncertainty in the forecasts as intangible information. That is, risk or uncertainty is generally not quantitatively considered in the capital investment decision process.

This paper discusses some aspects of a decision tool in which uncertainty is explicitly considered. The approach is largely based on the normative decision model proposed by Savage [1]. The model combines subjective probability with the von Neumann and Morgenstern cardinal utility theory to arrive at a "rational" norm for individual behavior.

While this general approach to decision making is widely hailed in academic circles, it has previously found little application as a corporate decision-making tool.

II. THE CERTAINTY EQUIVALENT (CE) METHOD

Figs. 1, 2, and 3 compare the most commonly accepted capital investment decision process with the proposed procedure. In Fig. 1 the common method of project analysis is represented in the three steps mentioned previously—the information gathering, the calculations, and the decision. In Fig. 2 the three steps are revised to represent project analysis by the technique called risk analysis. Fig. 2 extends the second and third steps in the process to include the expected utility model for decision making.

In the present procedure, as shown in Fig. 1, the three steps take the following form.

Manuscript received March 1, 1968. This paper is based on a Ph.D. dissertation, "The Explicit Consideration of Uncertainty in Capital Investment Analysis," submitted to the Illinois Institute of Technology, Chicago, Ill., 1968.

The author was with the Systems Sciences Group, Stanford Research Institute, Menlo Park, Calif. 94025, and currently is Managing Director of Strategic Decisions Group.

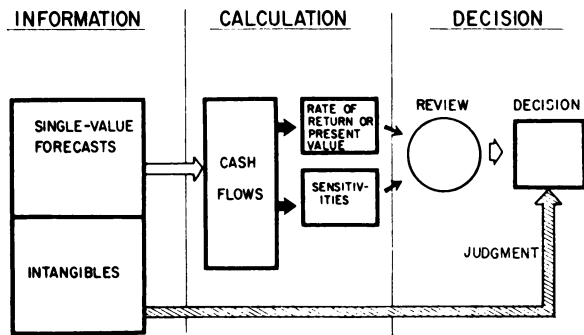


Fig. 1. Typical project analysis.

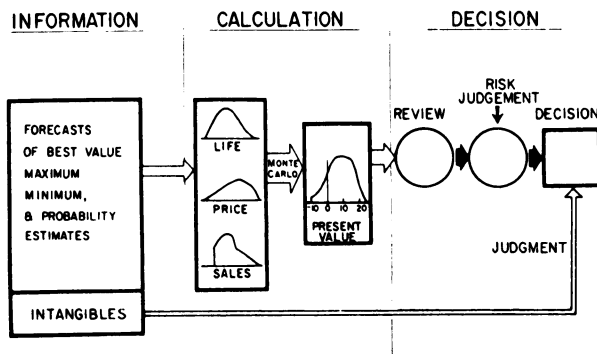


Fig. 2. Risk analysis.

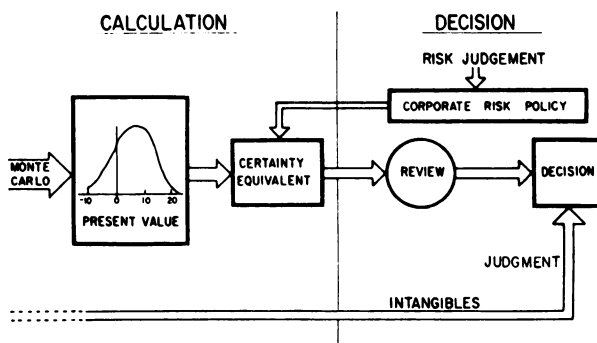


Fig. 3. Certainty equivalent (CE) method.

value with judgment on the intangible information. Note that the sensitivities only call the decision maker's attention to those variables which could strongly affect the criterion. The sensitivities do not quantify the uncertainty in the variables.

Risk analysis as represented in Fig. 2 differs from the preceding procedure in the following way.

1) Gathering information. Instead of quantifying only single best guess values, the uncertainty in the variables is also quantified. This is accomplished by forecasting for each variable the possible states and the probabilities associated with these states. By doing this, a considerable part of previously intangible information is now quantified. Ideally, only the information which is too costly to quantify (or which is overlooked) would now remain in the intangible category.

2) Calculations. In the calculations the uncertainty information is combined into a probability distribution of the decision criterion. This can be accomplished by a Monte Carlo simulation procedure, by means of decision trees, or in simple cases by analytically combining the input distributions.

3) The decision. In risk analysis, the decision maker is confronted with probability distributions of the profitability criterion instead of single values. The additional information that a probability distribution presents increases the difficulty of reaching a decision, since the decision maker must judge the acceptability of the risk in each alternative. In risk analysis, this judgment is left to the individual's intuitive risk judgment.

The certainty equivalent method, as shown in Fig. 3, is an extension of risk analysis. The first and second steps of the decision process are identical up to the development of a probability distribution of the decision criterion. However, the intuitive risk judgment, which is applied in risk analysis, is quantified by means of a corporate utility function. Note that the utility function does not replace judgment, but simply formalizes the judgment so it can be consistently applied. With the utility function, a certainty equivalent is calculated and used as a single-valued decision criterion. The final decision is thereby again reduced to a ranking process. However, in reviewing the inputs and assumptions, the decision maker also needs to review the acceptability of the utility function.

Risk analysis as shown in Fig. 2 has already become operational. A growing number of companies are using such procedures today. However, due to the difficulties associated with developing a meaningful corporate utility function, the CE method is not yet being applied.

The CE Model Specified

The CE method is applicable to many profitability measures. One useful measure is the discounted present value which will be used throughout this paper. If a utility function can be found that represents the preferences of the decision maker over the range of present values, the expected utility of each act can be calculated

1) Gathering information. Only best guess (single-valued) estimates of the relevant variables are quantified. All other information is considered intangible.

2) Calculations. The annual cash flows are developed from the forecast values of the variables. The rate of return, present value, or some other decision criterion is then calculated. In addition, the sensitivities of the decision criterion to changes in the forecasts are often calculated.

3) The decision. First the assumption and estimates are reviewed and then the decision is reached by considering the value of the quantified criterion while tempering this

and the optimal act is simply the act which maximizes the decision maker's expected utility.

Whose Utility Function: In deciding on a source of the utility function, it must be remembered that the expected utility model is for individual rather than for group decisions. However, as Luce and Raiffa [2] state:

The distinction between an individual and a group is not a biological-social one but simply a functional one. Any decision maker—a single human being or an organization—which can be thought of as having a unitary interest motivating its decisions can be treated as an individual in the theory. Any collection of such individuals having conflicting interests which must be resolved, either in open conflict or by compromise will be considered a group.

Thus for investment decisions, a corporation should act as an individual. But the fact remains that widely different attitudes toward risk exist in a company. Even if top management can be convinced that it should adopt a single utility function, the problem of whose utility function remains.

This problem can be resolved in the following manner. Setting policy is top management responsibility. Since a normative utility function is a policy statement, the utility function should be a product of the top policy makers. This statement may seem vague, but in actuality the top policy-making unit of a specific organization can be easily identified. Such a unit may consist of only the president, a board of directors, a management committee, the chairman of the board, etc.

Only in the case of small partnerships or family-held corporations should the actual owners play a direct role. However, in those case the owners are members of the policy-making unit. The author sees the operational utility function as a policy tool, and thus not a direct product of the stockholders of a large corporation. Obviously, any collectively strong feelings on the part of the stockholders will be reflected in the long run.

Whose Probability Estimates: According to decision theory, the utility function of a decision maker should be applied only to his own degrees of belief. This means that in the corporate setting, the policy makers' utility function can only be employed on decisions where the probability information represents the degrees of belief of the policy makers.

This presents a dilemma. Can the probability information on typical projects reflect the policy makers' degrees of belief, when they are not the source of such information? Since members of top management are generally at the mercy of experts in so far as decision information is concerned, they may often be willing to accept the forecasts and probability estimates of the experts as their own.

In some cases, if the experts are not able to agree on particular estimates or if those responsible for the decision disagree with the experts, the decision maker's beliefs should take precedence. It is he who will carry the re-

sponsibility for the decision. When the utility function of the decision maker is matched with probability estimates with which he disagrees, he may find that he intuitively disagrees with the CE value.

Of course, the responsibility of the estimation of probability data can be delegated. However, such delegation should be made only with the understanding that it is a total acceptance of the developed information. Typically, such delegation would be common on routine investments.

What Interest Rate: One more concept in the model needs to be specified, an interest rate for present value calculations. While the interest rate in this model should not differ theoretically from the much discussed cost of capital, it seems worthwhile to review the meaning of the interest rate in the model.

The simplifying assumption that present value is a satisfactory decision criterion requires that the projects can be ordered by preference or indifference using present values, and that this preference or indifference relationship is transitive. Since uncertainty is considered explicitly in the model, this requirement has to hold only in the state of certainty. In other words, it is assumed that there exists an interest rate that can express the true time preference of the decision maker so that he can completely ignore the cash flow pattern and rely solely on the magnitude of present value for reaching decisions. To repeat, however, this must hold only under certainty conditions, i.e., when all sums of money are guaranteed. Thus the interest rate should not include any margin for uncertainty.

If an interest rate actually fulfills this requirement, a unique monotonically increasing utility function for present value should exist for a decision maker. Thus if a decision maker wishes to maximize expected utility, the proposed model should be applicable.

III. METHODOLOGY FOR EVALUATING UTILITY DATA

Introduction

An experimentally determined utility function can be used for one of three purposes. It can be used to describe a subject's risk attitude; it can be used to predict a subject's future behavior in risk situations; or it can be used as a tool for improving future decisions involving risk. Thus the objective may be a descriptive, predictive, or normative model of man. Certainly the approach to an experiment will differ in accordance with the objective. A descriptive or predictive model of a subject should reflect his overall risk attitude and also represent his deviation from rational behavior, while a normative model should lead to consistent and rational decisions that reflect a subject's overall risk attitude.

The Development of a Utility Plot: A utility function for an individual can be developed by recording his preference (or indifference) among alternative investments in risk situations. One efficient method is based on finding alternatives between which an individual is indifferent.

For example, let it be assumed that an individual stated that he is indifferent between the alternatives A and B in the following table.

	Probability	Alternative A Outcome	Alternative B Outcome
1	0.5	\$30 million	A sure \$4 million
	0.5	\$ 0	
2	0.8	\$30 million	A sure \$10 million
	0.2	\$ 0	
3	0.7	\$30 million	A sure \$4 million
	0.3	\$-2 million	

Using the expected utility hypothesis, the following mathematical statements can be made:

$$\begin{aligned}
 0.5 U(\$30 \text{ million}) + 0.5 U(\$0) &= U(\$4 \text{ million}) \\
 0.8 U(\$30 \text{ million}) + 0.2 U(\$0) &= U(\$10 \text{ million}) \\
 0.7 U(\$30 \text{ million}) + 0.3 U(\$-2 \text{ million}) &= U(\$4 \text{ million}).
 \end{aligned}$$

Since the utility scale is only unique up to a linear transformation, two values on the scale can be arbitrarily chosen. Thus let $U(\$30 \text{ million}) = 100$ and $U(\$0) = 0$. Then simple algebraic manipulation and substitution leads to the following:

$$\begin{aligned}
 U(\$4 \text{ million}) &= 50 \\
 U(\$10 \text{ million}) &= 70 \\
 U(\$-2 \text{ million}) &= -66 \frac{2}{3} \sim -67.
 \end{aligned}$$

This information is plotted in Fig. 4.

The results from actual interviews when plotted lead to considerable scatter. Such actual responses from one individual are given in Fig. 5. The specific method of how this plot was developed is discussed later. For the moment it is worthwhile to consider this plot simply a product of statements of indifference between investment alternatives.

The Utility Plot as a Descriptive or Predictive Model of Behavior: Plots such as Fig. 5 are certainly useful for descriptive purposes. The pattern of the points, as represented by the dotted line, describes the level of overall risk aversion of the individual. With a particular choice of scale, the more curved this pattern, the more risk-averse the individual. Of course, a straight line would indicate a risk-neutral individual. However, another aspect of behavior is described by Fig. 5, the degree of inconsistency with utility theory. The scatter in the plot gives an indication of how an individual's responses deviated from the expected utility model. Obviously, the individual's responses given in Fig. 5 deviated considerably from the responses suggested by the expected utility model. The findings show that this plot is not untypical. Therefore, when describing behavior, the scatter in the utility plot should be included in the description. A single best fitting line can only describe a general pattern of a risk attitude.

In using a utility plot to predict behavior, similar problems occur. The assumption must be made that present behavior will be representative of future behavior.




Fig. 4. Utility plot.




Fig. 5. Experimental utility data.

Then, since individuals are presently inconsistent with utility theory, there is every reason to believe that they will continue to be inconsistent in the future. Therefore, at best, the utility plot can be used to predict a general pattern of behavior and the expected deviation from this pattern.

So far, in considering the utility plot for either description or prediction, the assumption has been made that each question to the individual was phrased to elicit a descriptive response. Difficulties arise in asking hypothetical questions under risk. Generally, the greatest problem is to create a real decision situation. Ideally, each question should result in a response that would be identical to an actual decision.

The Utility Plot as a Normative Tool: If the utility plot is to serve as a basis for a normative decision tool, the

questions should elicit responses from the decision maker that are not necessarily descriptive of his behavior. Rather, the responses should represent how he feels that he should behave in light of the corporate goals. This implies a different phrasing of questions.

Once a utility plot has been developed based on the correct type of questions, scatter such as in Fig. 5 is still bound to exist. But if the individual has accepted the axioms of utility theory as a guide to rational behavior, he should be willing to accept some particular curve that fits to his general pattern of responses. Some well-fitting curve, such as the dotted line in Fig. 5, can be used to represent a first approximation of the normative function. With astute questioning, this function can be further adjusted until it becomes an acceptable normative tool.

Definitions: For the rest of the discussion, it will be helpful to use a number of technical terms. The definitions of these terms follow.

1) The certainty equivalent. The certainty equivalent for a risky investment situation is that single value, which under certainty results in the same utility as the expected utility of the risky investment. Thus

$$U(CE) = \sum p(x_i)U(x_i)$$

where

- $p(x_i)$ = probability of the outcome x_i
- $U(x_i)$ = the utility of the outcome x_i
- CE = the certainty equivalent.

Therefore, in accordance with utility theory, a rational individual must be indifferent between a project and his certainty equivalent for the project.

2) The risk premium. The risk premium of a project is the difference between the expected outcome of a project and the certainty equivalent. Thus

$$\text{the risk premium} = EV - CE$$

where

- $EV = \sum p(x_i)x_i$
- x_i = the i th outcome.

Since the certainty equivalent depends on the subject's risk attitude, the risk premium is also dependent on the subject's risk attitude.

3) The probability premium. Given a risky project that has only two possible outcomes x and y , and an alternative project with a certain outcome z , where $x < z < y$, then only one set of probability exists such that the expected value of the risky project is equal to z :

$$p(x)x + p(y)y = z$$

where

$$p(x) = 1 - p(y).$$

A risk-averse individual would have an indifference probability for y , $p^*(y)$, which is greater than $p(y)$. The





Fig. 6. Circular reference chart.

difference between $p^*(y)$ and $p(y)$ is called the probability premium. Thus the probability premium = $p^*(y) - p(y)$.

Alternative Interview Procedures

All interview procedures are based on eliciting preference or indifference responses between alternative investment situations, where at least one of the alternatives is a risky investment, i.e., it has two or more possible outcomes. In the simplest case, the two alternatives are as follows:



Here alternative A has the probability $p(x)$ of the outcome x and the probability $p(y)$ of the outcome y , while alternative B has a certain outcome of z . Since all studies including this one have involved this simple type of decision situation, the discussion will be limited to this particular situation.

Given this simple decision situation, two basic methods are available to elicit indifference responses. The first method is to keep the outcomes x , y , and z of both alternatives constant and vary the probabilities until the interviewee responds with indifference. The second method is to keep the probabilities and two of the three quantities x , y , and z constant and to vary the third. In either case, the interviewee is required to evaluate his certainty equivalent z for a simple risky investment situation and the same type of judgment is required of the individual.

Either of the alternative methods leads to difficulties in practice. The major problem is the personal interpretation of probability weights. A 0.2 probability of an outcome has a different meaning to different individuals. The meaning of probability weights becomes particularly vague to most individuals in the ranges close to 1 or 0. For example, many individuals cannot differentiate between a 0.02 and a 0.01 probability, and any absolute error in the personal interpretation of probabilities causes a greater effect outside the middle range. In addition, many individuals round probability responses to the nearest multiple of 0.05. Responses outside the middle range should, therefore, be avoided.

The problem of subjective probability interpretation can be considerably reduced by the use of a reference process. An example of such a reference process is a circular reference chart, as shown in Fig. 6. This chart is so

designed that a simple twist increases the red area while reducing the green area. It can be used in the following manner. The interviewee is asked to visualize the chart spinning rapidly, with the throw of a single dart determining the outcome. If the dart hits red the interviewee is assured of outcome y , and if the dart hits green he is assured of outcome x .

By showing a 20-percent setting of the chart rather than saying a 0.20 probability, an individual is given the opportunity to evaluate the meaning of an objective 0.20 probability. Careful use of the chart can also eliminate the rounding by the interviewee. The use of a reference process does not improve the judgment of individuals, but it does give them a common fixed scale in evaluating their judgment.

Alternative Methods of Evaluating Responses

At present, there seems to exist three basic methods of converting the responses into a utility plot. Method 1 is based on getting certainty equivalent responses for a risky alternative where the two possible outcomes are always kept the same. Method 2 is based on getting segments of the utility plot and then overlapping these segments. In Method 3 a mathematical form for the utility function is assumed and then the best fit of this function is found. The details of these methods are discussed later.

Method 1: This method, which is often used in textbooks, requires a certainty equivalent response to risky investment situations with two outcomes, x and y . The values of x and y are chosen at the extremes of the range of utility plot. Thus, depending on the various choices of the probabilities of the outcomes, each certainty equivalent response z falls between x and y . The choice of scale is then defined by the values of x and y ; e.g., $x = 0$, $y = 100$. This allows the calculations of the utility of z immediately from the relationship

$$p(x)U(x) + p(y)U(y) = U(z)$$

$$0 + p(y)100 = U(z)$$

or

$$U(z) = 100 \cdot p(y).$$

This approach certainly is the simplest method of evaluating a utility function. However, if the range $y-x$ is to be significant, extreme probabilities are necessary to lead to certainty equivalent responses near x or y .

Method 2: The second method is found as two variants. In the first variant, overlapping segments of a utility function are determined much in the manner of the preceding method. These segments are then plotted and visually adjusted for a good fit.

The second variant of this method starts with a definition of the scale such as in the first method. Some of the utility values that are derived are then assumed to contain no error and are used as new definitions of the scale. With this method, a scale can be extended or broken up, much

the same as the overlapping of segments. Judgment is used in the visual method of the overlapping of line segments where certain points are simply assumed to be exact in the second variant.

Method 3: In this method, a mathematical form of the utility function is assumed prior to plotting of the data. The fit of this functional form is then tested. For descriptive purposes, this method has the disadvantage that the individual's pattern of behavior is somewhat prejudged by the choice of the function form. By varying the mathematical form and by including enough degrees of freedom in the function form, this disadvantage can be reduced or eliminated. For normative purposes, this approach gives the ability of incorporating a "reasonable" pattern of behavior, where the definition of "reasonable" depends on judgment about overall policies in the face of risk.

This method of evolving a normative utility function has been previously suggested by Pratt, Raiffa, and Schlaifer [3]:

... In the construction of a preference function for money it is often preferable to use an analytical curve having certain specified qualitative features but with some adjustable parameters and to determine values of the parameters from questions about particular gambles. As many questions are needed as there are parameters. If extra questions are asked, then the parameters may be determined from some and the rest used to check the consequences, or the parameters may be chosen according to some criterion of 'best fit' to all.

An advantage of this method over the second method is in the treatment of the scatter. In the overlapping of segments, errors in the responses may compound and lead to unintentioned reversals of slope.

For normative purposes, a utility function that has been fitted to data should always be considered a first approximation. The interviewee should have a chance to evaluate the function after it has been smoothed. All adjustments that are made in smoothing the function must be accepted by an individual before it is used as a normative tool.

Results of Previous Studies

The first application-oriented study was by Grayson in his study on oil and gas drilling decision. In this study he phrased the questions ([4], p. 295) in the following way:

I will give you a series of hypothetical drilling deals, and I want you to listen to each deal as I present it and give me an answer as to whether you would accept that deal today or reject it.

This phrasing of the question would tend to lead to an answer which is basically descriptive of the interviewee's risk attitude. It would seem that, before the results could be used in a normative fashion, adjustments should be made. In Grayson's study, each individual was confronted with a gambling situation in the form of the

drilling decision. The interviewee was always asked to respond with an indifference probability. The same approach was used by Cramer and Smith [5], the only other published study that has a normative intent.

In both of these studies, no reference process for probability judgments was used and considerable rounding in the responses is apparent. Virtually all probability responses in both studies end in multiples of 0.05.

The functions were fitted by the overlapping of segments of the function. The alternative of noninvestment, i.e., a certainty equivalent of zero, was used throughout.

Grayson recognized the need for the reevaluation of inconsistencies; however, except in one case, he limited himself to the use of the initial response. He states ([4], p. 303):

Clearly these curves are only approximations, for the operators answered fairly rapidly and an 'error' of a few points in probabilities, particularly in the extremes, throws points far apart. As suggested earlier to remove these errors the operator could be shown the points of inconsistency and asked which one came closest to reflecting his true preferences. There was no time in the research to plot curves and then present them to operators for modification, except with one operator, Bill Beard, who had been visited during the pilot study. He looked at his curve, checked the points of inconsistency, and reduced the variances to a nominal amount.

No attempt to modify the inconsistencies seems to have been made by Cramer and Smith. They took all their subjects from the same company, then assumed that the mean of the responses would represent a corporate response. This assumption is particularly open to criticism. Not only does it mean that executives compromise, but that the degree of each compromise is given by the average. Cramer and Smith did not try to see if their composite was an acceptable compromise to the group. They included individuals of different authority levels without considering that factor. The use of such a composite for actual decisions would most probably run into much opposition. However, even with all of these deficiencies, the composite could provide a reasonable starting point for discussion and review that would result in an acceptable function.

Two descriptive studies that used managers as subjects have been published. The first by Green [6] involved 16 businessmen in a large chemical company. The second by Swalm [7] involved about 100 executives. Both studies used the second variant of the overlapping of line segments to evaluate responses. No reference processes for probability judgments were used, and no attempts to adjust inconsistencies are apparent.

All of the above studies have used methods of evaluation that compound the errors. This may be the cause of the fascinating shapes of utility functions documented by these studies. In the course of this study, intentional behavior among business executives that leads to reversals of slope has not been found.

Experimental Approach

In the experiments for this study, all responses are indifference probability responses. A reference process was used to help the interviewee in visualizing probabilities. The responses were evaluated with a prior assumption of a function as described in Method 3.

The objective of the effort was to evolve a corporate risk policy statement in the form of a utility function. The first step toward this goal involved interviews of individual managers. The purpose of the individual interviews was threefold.

1) To accustom management to the idea of quantifying risk attitudes.

2) To demonstrate the need for an overall risk policy.

3) To see if some analytic function could be found which might form the basis for an overall strategy.

In the second step of the research, the purpose was to evolve a corporate policy from individual attitudes through group meetings and arbitration among the policy makers.

IV. THE INDIVIDUAL MANAGERS INTERVIEWED

Description of the Subject Company

All of the research as reported below was performed in one organization. By request of this company, the name and some of the information will be disguised. However, every effort will be made to make all of the information comparable to the original situation.

Gamma Industrial, Inc., is a major corporation in its industry. Its total annual sales range around \$2 billion with annual after-tax earnings of \$100 million. The average capital investment budget for the last three years has been about \$90 million per year.

The company is organized along functional lines. The chief executive has the overall responsibility; however, the executive committee plays a major role in the management of the company. This committee has weekly meetings where major decisions are discussed and reviewed.

Interview Technique

In total, 36 executives of Gamma Industrial, Inc., were individually interviewed in the initial phase of the research project. This group included all members of the executive committee, i.e., the president and functional vice president, and 15 line managers reporting to the vice presidents. The other individuals who were interviewed were scattered throughout the organization. Most of them were staff members with special responsibilities in economic evaluation. All interviewees were quite familiar with major capital expenditures.

During the introduction to the interview, the author reminded the participants of the fact that Gamma Industrial, Inc., is starting to use risk analysis for the evaluation of projects under uncertainty. They were then confronted with three probability distributions and asked to make a decision between the alternatives. The purpose

TABLE I

Investment (\$MM)	Equivalent Uniform Annual Cash Flow* (\$MM)	Rate of Return†	Present Value at 10% (\$MM)	Chance to Break Even‡	Indifference Chance of Gain
3	0.89	30%	5	50%	—
	-0.22	Negative	-5		
3	4.55	160%	38	25.5%	—
	-1.10	Negative	-13		
50	8.90	15%	30	62.5%	—
	0	Negative	-50		
50	41.00	87%	320	29.7%	—
	-95.00	Negative	-135		

* This annual cash flow figure, if received (for 20 years) after investing \$3 million, will result in the present value.

† The rate of return is a discounted rate of return based on a 20-year project life.

‡ This percent chance of success leads to an expected present value of \$0.

of the three distributions was to underscore the difficulty of making such decisions and to show the need for a subjective "risk judgment" before reaching a decision. Then the purpose of the interviews was explained as being the derivation of a strategy tool to help in reaching decisions between alternative distributions. The interviewees were asked to give responses that would reflect their best judgment for a strategy tool for the company. They were not asked to describe their present decisions.

In addition to the preceding, the following points were emphasized.

- 1) There are no right or wrong answers to the questions.
- 2) All answers will be completely confidential.
- 3) Each response should be made in light of the company's present financial position. (Some highlights of Gamma's capital expenditure program were reviewed here.)

- 4) The emphasis is not on consistency from answer to answer; rather the overall pattern of answers is important.

- 5) This is not a quiz to see what the interviewee knows, but an attempt to quantify his personal feelings about a risk policy for the company.

- 6) Pencil and paper may be used.

Each interviewee was then asked to make a decision in each of 40 hypothetical investment situations. The 40 situations included 20 questions at each of 2 investment levels, \$3 million and \$50 million. The individual questions were phrased in the following manner.

Given a \$3 million initial investment that will result in one of two possible outcomes—either a positive present value of \$5 million (success) or a negative present value of \$3 million (failure)—would you recommend investing in the project if you feel the chances of success are three out of four (i.e., 75 percent)?

Thus the required response was only to accept or not accept. When the interviewee responded affirmatively, the same proposal was repeated with the probability of success reduced to 60 percent. If he still recommended acceptance, the probability of success was again lowered. The indifference probability between acceptance of the project and noninvestment was determined as closely as possible by this search. Most interviewees seemed to be able to respond with probability numbers that bracketed

their indifference point. Typically, in the response range of 30 to 70 percent, a change of 5 percentage points would change their response from acceptance to rejection of the project. For example, at a 50-percent probability of success the interviewee may definitely accept a project, while at a 45-percent probability he would definitely reject the project.

After a number of indifference probabilities were found by using this search technique, most interviewees learned to respond directly with their indifference probability. A direct indifference response was not forced by the interviewer, and at any time when the interviewee seemed to experience difficulties the initial method was again used.

To help the subjects in understanding the alternative outcomes of the hypothetical investment projects, additional measures of profitability were provided to them. An example of four questions from the interview schedule is given in Table I.

To help the interviewees in understanding probability statements, the previously described reference chart was used. When an individual's response indicated that he was indifferent at some probability of success, the chart was set to reflect that portion in green. He was then asked whether he would actually be indifferent between not investing and investing in this project if the throw of a single dart would determine the outcome. The use of this reference process added realism for most individuals. Some individuals adjusted their initial responses considerably after seeing the chart. Others would not have been able to respond at all without reference to the chart, particularly in the extreme probability ranges. There were a number of individuals who personally adjusted the chart to their indifference point. Their indifference probability was then read from the chart.

Results from Interviews

Accuracy of the Results: When evaluating the results from the interviews, the reader should keep in mind that this portion of the study was primarily of an exploratory purpose. Not too much weight should be placed on the details of separate responses, since each interviewee was asked not to try to be consistent, but rather to

express his overall attitude. As previously mentioned, the interviewees generally were unable to express their preferences in probabilities closer than ± 2.5 percentage points. Some other effects which detract from the accuracy of the results were noticed during the interviews. These effects are the following.

1) The learning effect. It was common for individuals to be quite unsure of their answers at the beginning of the interview. However, as the interview progressed, the interviewees learned to make the type of decisions that were required. Because of this, many of the individuals felt better about their later answers. Some even suggested that their first few answers were worthless.

2) The effect of making many decisions. In some interviews it became apparent that the interviewees were affected by the number of questions they were asked. They started feeling that they had many risky investment opportunities available and could thus play closer to the averages. It was possible to keep this effect to a minimum by reminding the interviewee of the number of actual investments that are made by the company at the particular investment levels.

3) The effect of the order of the questions. The order in which the questions were asked seemed to be of particular importance. After a decision about an investment which involved outcomes of large gains and losses, many individuals became more neutral to smaller risks. Had the larger investments been taken first in the interview, presumably the results for the smaller investments could be quite different. This effect was somewhat reduced by using a random order of the questions within each level of investment.

4) Effect of decision rules. In a number of cases, the interviewees adopted a decision rule early during the interview. This resulted in answers that were internally consistent, but that did not express their feelings as well. The author discouraged the use of such rules by explaining the purpose of the interview and stressing that there was no right or wrong answer. Some decision rules that recurred were the following:

- a) use of a probability premium which was directly related to the magnitudes of the outcomes;
- b) use of a probability premium related only to the possible loss;
- c) use of a ratio of expected gain to expected loss in relation to magnitudes of outcome; and
- d) use of a risk premium in present value related to magnitudes of outcome.

However, by far, most individuals relied on their intuitive judgment. Particularly, the top management group was willing to make these decisions on an abstract level.

When asked about the quality of their responses, most individuals answered that they were willing to accept their answers as their recommendation for company policy. A few felt no confidence in their answers because they were unfamiliar with the magnitudes of the outcomes.

The top managers felt their answers were meaningful, but wanted to discuss the policy effects further.

Presentation of Results: Instead of individual utility functions, which would be meaningless to most readers, the raw data are summarized in histogram form in Figs. 7 and 8. The indifference probabilities for subjects are tabulated for each investment situation as a separate histogram.

The horizontal axis of each histogram is shifted so that the left extreme point of each axis is equivalent to the probability of gain which would lead to an expected present value of \$0, i.e., a risk-neutral response. Thus the horizontal distance represents the probability premiums. Since the discounting rate for the present value was the company's minimum acceptable return, a response below this point would, on the average, lead to an opportunity loss. No manager responded below this "break-even" value.

The histograms are ordered according to the magnitudes of the present values of the outcomes. The largest magnitudes are at the top of the figures. The present values for success and failure are given at the left of each histogram. The averages of the responses are shown by the heavy lines.

This type of presentation served very well for the communication of risk attitudes. Different patterns are quickly identified and compared. An individual's attitude can be represented by a discontinuous line such as that shown for the average responses.

Interpretation of Results: From Figs. 7 and 8 the following conclusions can be drawn.

1) The responses vary widely for each investment situation. For example, consider the \$50 million investment that will result in either a \$30 million present value gain or a \$25 million present value loss. (The lowest histogram in Fig. 8.) This can be considered a low-risk project, since the worst that can happen is a \$25 million negative present value, which in this case corresponds to approximately a 1-percent rate of return. On this investment, the responses varied from a 0.50 to 0.95 probability of success. This is a surprisingly wide range, considering that these individuals have been interacting and making decisions within the same organization for years.

2) As the magnitudes of the outcomes increase, the probability premium that is required by the individuals increases. This general tendency is easily recognized in the trend of the averages. This is also quite as expected. The trend does not become strong, however, until the \$50 million investment level is reached.

3) The spread in the responses also increases as the magnitudes of the outcomes increase. The average absolute deviation from the mean of the probability premiums at the \$3 million level increases from a low of 0.08 to a high of 0.12. At the \$50 million level, the average deviation goes from a low of 0.12 to a high of 0.17. The average deviation can be interpreted as a measure of the consensus of feelings about a risk strategy. Thus there exists more of an agreement on the level of acceptable risks at the lower




Fig. 7. Probability premiums for all responses at the \$3 million investment level.




Fig. 8. Probability premiums for all responses at the \$50 million investment level.

level. However, as the magnitudes of the outcomes increase, this agreement dissolves.

4) The probability premium seems to be much more influenced by the magnitudes of the outcomes than the investment level. This was to be expected, since the outcomes were given in net present value and the outcomes are frequently much larger than the investment. In comparing the top of Fig. 7 with the bottom of Fig. 8, it becomes apparent that the influence of the investment level is minor if the outcomes are stated as net effect. This leads to the conclusion that a utility function on net present value should be applicable over wide ranges of initial investment level.

Each interviewee was sent a report on the results of the interviews that included similar presentations of the data histograms. Their personal responses were entered in the histograms with a red line. This allowed each individual to evaluate his responses in respect to all other participants. The same analysis as the preceding was made, but no explanations for the differences were given. The reports led to widespread discussions of the results of the interviews. A number of interviewees openly discussed their results with others.

All participants were highly interested in the study and the results. There seemed to be a general agreement among the participants that the findings represented a reasonably accurate description of their individual risk attitudes. However, some were surprised at their relative standing to others. A few that generally thought of themselves as "risk takers" found that they were not as willing as most of their colleagues to take risks.

V. FITTING A UTILITY FUNCTION TO THE DATA

Initial Choice of a Function

One part of the objectives of the first phase of the study was to find a mathematical form for utility functions that would be able to represent adequately the attitudes of the individuals over the range of the outcomes. All such functions must be monotonically increasing in utility, since individuals prefer the larger of two present values.

A number of such functions have previously appeared in the literature. A common assumption has been a section of a quadratic, i.e., the increasing section of a parabola [8]. Dolbear [9] used straight line segments to express the measured preferences of his subjects. Kaufman [10] found that the function of $U(x) = A + B[\ln(x + c)]$ fits very closely to one of the subjects from Grayson's study.

Since for normative purposes the function has to have a good fit only over the range under consideration, segments of a multitude of functions could be used. However, theoretical considerations do give some indications as to the choice of function.

Let $U(x)$ be the utility function of the present value x over the range of x under consideration. The function $r(x) = -U''(x)/U'(x)$ can then be interpreted as a measure of local risk aversion as shown by Pratt [11]. Thus a utility function that would lead to a present value risk

premium ≥ 0 for all possible investments with probabilistic outcomes must have $r(x) \geq 0$ over the whole range of x . Such risk-averse or risk-neutral attitudes were exhibited by all subjects in the study. Since a businessman would seldom be willing to pay a premium simply for the sake of taking a risk, this should be part of the overall pattern of a normative utility function.

Another overall pattern of attitude that should be incorporated in a normative utility function is that $r(x)$ should be monotonically decreasing or constant over the range of x . This means that for identically shaped distributions with different expected values, the risk premium should either decrease or stay constant as the expected value increases. In other words, as a decision maker expects more of a net gain, he becomes willing to take more of a chance. Certainly in most business situations, other things being equal, a decision maker would not knowingly want to become more risk averting as the expected outcome increases.

Summarizing from the preceding, a prescriptive utility function should most commonly be

- 1) continuous and twice differentiable,
- 2) lead to a function $r(x) = -U''(x)/U'(x)$ which is ≥ 0 over the range of x , and
- 3) $r(x)$ should be constant or monotonically decreasing, i.e., $r'(x) \leq 0$ over the range of x .

It is of interest that the widely used quadratic utility function does not satisfy the last condition over any range ([11], p. 122).

One function that does satisfy the stated conditions is the logarithmic function

$$U(x) = A + B \log(x + c) \quad (1)$$

where $x + c > 0$. Here

$$U'(x) = B(x + c)^{-1}$$

$$U''(x) = -B(x + c)^{-2}$$

$$r(x) = (x + c)^{-1} > 0,$$

since $x + c > 0$, and

$$r'(x) = -(x + c)^{-2} < 0.$$

As mentioned previously, this function was used by Kaufman to fit the data to one of Grayson's subjects. The same function was first investigated in this research.

The Method of Fitting the Function

The method of fitting the function to the 20 responses at each investment level was built on a least squares approach. The expected utility $EU(x)$ of each hypothetical investment is

$$EU(x) = p_s U(x_s) + (1 - p_s) U(x_f) \quad (2)$$

where $U(x)$ = the utility of \$ x present value, s stands for success, and f stands for failure. Thus p_s = probability of success, and $(1 - p_s)$ = probability of failure.

Since each response indicated indifference between investing and not investing, the expected utility of the hypothetical investments should be equal to $U(\$0)$ if the decision maker had accepted utility as his decision criterion and were infinitely sensitive.

For the experimental data, obviously some deviation from this criterion can be expected. However, parameters for a given function can be determined by minimizing the sum of the squares of the deviations. Thus

$$\sum [U(\$0) - p_i U(x_i) + (1 - p_i) U(x_f)]^2 = \text{minimum.} \quad (3)$$

Since the choice of scale for a utility function is arbitrary, let $U(\$0) = 0$ and $U(\$K \text{ million}) = K$. Then for the function $U(x) = A + B \log(x + C)$,

$$U(0) = A + B \log C = 0$$

and

$$A = -B \log C$$

$$U(K) = A + B \log(K + C) = K$$

$$= -B \log C + B \log(K + C) = K.$$

Thus $B = K / \log[(K + C)/C]$ and substituting,

$$U(x) = \frac{K \log \frac{x + C}{C}}{\log \frac{K + C}{C}}. \quad (4)$$

Substituting into (3) leads to the following criterion:

$$\sum \left(\frac{K p_i \log(x_i + C)/C - K(1 - p_i) \log(x_f + C)/C}{\log(K + C)/C} \right)^2 = \text{minimum.} \quad (5)$$

To find the value of C which minimized the above, a simple computer search routine was written. This routine calculated the sum of the squares of the deviations for various values of C , and then searched for that C which led to the minimum.

Note that the sum of the squares of the deviations is not a very meaningful measure of the tightness of the fit, since it depends on the arbitrary choice of the scale K . Using the same K , however, the sums of the squares of the deviations do give a relative measure of the fit of the function to the different subjects. For most of the studies, K was chosen to be 50. This choice was based purely on considerations for the plotting of the functions.

Utility Plots

To get an idea of the absolute fit of the function to various subjects, it was decided to plot the utility functions. However, here a problem appears. Assuming that any deviation from the function is due to error in judgment, the questions remain as to how the error should be apportioned between the utility of the gain and loss. It




Fig. 9. Utility plot of responses for all errors in positive quadrant.

could be assumed that all the errors occurred in the utility of the gain. Then the utility of the gain could be calculated from the following equation, which is a direct result of (2):

$$p_i U(x_i) = -(1 - p_i) U(x_f)$$

and substitution of (4) leads to

$$U^*(x_i) = - \frac{(1 - p_i) K \log [(x_f + C)/C]}{p_i \log [(K + C)/C]} \quad (6)$$

where the quantities marked with an asterisk are experimentally implied values.

Fig. 9 shows a utility function of Subject 2 using this assumption.

Similarly, all errors could be assumed to lie in the negative quadrant. However, for lack of any better apportionment and since it is more pleasing to the eye, it seems reasonable that each quadrant should bear half the burden of the error. Thus from (4) and (6),

$$U^*(x_i) = \frac{1}{2} \left(\frac{K \log [(x_i + C)/C]}{\log [(K + C)/C]} - \frac{(1 - p_i) K \log [(x_f + C)/C]}{p_i \log [(K + C)/C]} \right) \quad (7)$$

and similarly,

$$U^*(x_f) = \frac{1}{2} \left(\frac{K \log [(x_f + C)/C]}{\log [(K + C)/C]} - \frac{p_i K \log [(x_i + C)/C]}{(1 - p_i) \log [(K + C)/C]} \right) \quad (8)$$

where the quantities marked with an asterisk are again the experimentally implied values.




Fig. 10. Utility plot with errors proportioned




Fig. 11. Utility plot—Subject 2.

In Fig. 10, the same data as in Fig. 9 is plotted using the preceding method of calculating the experimentally implied utility values.

Fig. 11 shows the plot for Subject 2 for the responses at the \$50 million investment level. Similar plots of the subjects which had a relatively poor fit using the function showed that many of the utility values fell below the function. This tendency was particularly strong close to the origin. The conclusion was reached that the function is not able to convey fully the feelings of all subjects.

A Subjective Test of the Acceptability of the Function

To confirm that these systematic deviations of the data points were of significance, adjusted indifference probabilities for each investment were calculated using the best fitting function for each subject. Substituting

TABLE II
DATA FOR BEST FIT OF FUNCTION AT \$50 MILLION
INVESTMENT LEVEL
SUBJECT 2 C = 260

NO.	XL	XG	PG	U(XL, T)	U(XG, T)
1	-50.00	30.00	0.850	-60.71	31.04
2	-135.00	320.00	0.600	-208.19	228.08
3	-185.00	950.00	0.500	-353.40	437.12
4	-90.00	60.00	0.850	-120.78	59.03
5	-115.00	320.00	0.500	-166.00	228.08
6	-150.00	230.00	0.750	-244.53	180.15
7	-25.00	60.00	0.450	-28.74	59.03
8	-200.00	600.00	0.500	-416.83	340.06
9	-70.00	160.00	0.650	-89.16	136.33
10	-135.00	230.00	0.600	-208.19	180.15
11	-25.00	30.00	0.850	-28.74	31.04
12	-170.00	450.00	0.500	-301.57	285.57
13	-200.00	750.00	0.450	-416.83	385.76
14	-115.00	100.00	0.750	-166.00	92.51
15	-135.00	160.00	0.650	-208.19	136.33
16	-200.00	950.00	0.400	-416.83	437.12
17	-90.00	100.00	0.700	-120.78	92.51
18	-150.00	600.00	0.600	-244.53	340.06
19	-70.00	30.00	0.900	-89.16	31.04
20	-185.00	450.00	0.750	-353.40	285.57

NO.	U(XL, E)	U(XG, E)	PGT	DEV (K)	DEV of PG
1	-118.31	20.88	0.662	17.28	-0.1883
2	-275.15	183.44	0.477	53.57	-0.1228
3	-395.26	395.26	0.447	41.86	-0.0530
4	-227.63	40.17	0.672	32.05	-0.1783
5	-197.04	197.04	0.421	31.04	-0.0788
6	-392.48	130.83	0.576	73.98	-0.1742
7	-38.52	47.07	0.327	10.76	-0.1225
8	-378.44	378.44	0.551	-38.39	0.0507
9	-171.17	92.17	0.395	57.41	-0.2546
10	-239.20	159.47	0.536	24.81	-0.0639
11	-102.32	18.06	0.481	22.07	-0.3693
12	-293.57	293.57	0.514	-8.00	0.0136
13	-366.23	447.61	0.519	-55.67	0.0694
14	-221.76	73.92	0.642	27.88	-0.1079
15	-230.68	124.21	0.604	15.75	-0.0457
16	-354.12	531.18	0.488	-75.25	0.0881
17	-168.32	72.14	0.566	28.52	-0.1337
18	-377.31	251.54	0.418	106.22	-0.1817
19	-184.27	20.47	0.742	19.02	-0.1582
20	-605.06	201.69	0.553	125.83	-0.1969

(4) into (2) and solving for the indifference probability of success p_s leads to the following equation:

$$p_s = \frac{-\log [(x_f + C)/C]}{\log [(x_s + C)/C] - \log [(x_f + C)/C]} \quad (9)$$

Here p_s represents the probability of gain consistent with the best fitting function. As is to be expected, p_s depends solely on the magnitude of gain and loss and the value of C .

The adjusted values of p_s were calculated by (9) for all hypothetical investments for each subject. Example output from the computer program which searched for the best fitting function by (5) and then calculated the various quantities of interest for Subject 2 is given in Table II. In this table the first two columns, labeled XL and XG, give the present value magnitudes of loss and gain in millions of dollars. The PG column gives the indifference probabilities with which the subject responded,

i.e., the raw data from the interview. The columns $U(XL, T)$ and $U(XG, T)$ give the utility values of loss and gain as calculated from the best fitting C value using (4). The next two columns, $U(XL, E)$ and $U(XG, E)$, represent the experimentally implied utilities with the errors proportioned to gain and loss in accordance with (7) and (8). These two columns were used to develop the utility plots as shown in Figs. 10 and 11. The column labeled PGT consists of the "theoretically adjusted" probability values as calculated from (9). The DEV (K) column gives the values of the deviations as calculated by (5), and the DEV of PG column simply gives the difference between the PG and PGT columns.

After the values of the adjusted probabilities in the PGT column had been calculated, selected subjects were confronted with their original responses and these adjusted responses. The subjects were then advised that the adjustments represent only one type of adjustment which does not need to represent their feelings. A number of subjects considered the adjustments to be an improvement without distorting their feelings. Some that deviated considerably were not satisfied with these adjustments, particularly around the origin. It was therefore decided that the function was not flexible enough to represent the feelings of all interviewees.

Improvement in the Initial Function

Fig. 11 suggests a further improvement in the shape could be made by allowing a stronger change in the slope around the origin. This can be achieved by adding another parameter D in the following manner:

$$U(x) = A + B \log(x + C - D|x|) \quad (10)$$

where $x + C - D|x| > 0$ and $0 \leq D < 1$. This adjustment in the function diminishes the positive values and emphasizes the negative values, thereby leading to a break at the origin. The effect of the parameter D is shown in Figs. 12 and 13. Fig. 12 shows the utility data for Subject 3 with $D = 0.05$ and the best fit of C , while Fig. 13 shows the same data plotted using $D = 0.50$. The fit is considerably better with the higher D . Except at the origin, this utility function still meets all the conditions mentioned previously.

Using the same procedure as previously (namely, the calculation of adjusted responses) to check the acceptability of this revised function, more acceptance was found. However, in the immediate range around the origin some individuals were still dissatisfied with the adjustments. This was somewhat expected because of the break at the origin. The function was, therefore, again revised to include one more parameter E , which has the effect of smoothing the function through the origin. The function now reads as follows:

$$U(x) = A + B \log\{x + C - D[(x^2 + E^2)^{1/2} - E]\} \quad (11)$$




Fig. 12. Utility plot—Subject 3.




Fig. 13. Utility plot with break at the origin—Subject 3.

where $0 \leq D < 1$, $E \geq 0$, and $\{x + C - D[(x^2 + E^2)^{1/2} - E]\} > 0$ for the range of x .

This function includes as special cases the previous two functions. Thus when $D = 0$, $U(x) = A + B \log(x + C)$ and when $E = 0$, $U(x) = A + B \log(x + C - D|x|)$.

It can be shown that the function meets all of the previously stated conditions for certain combinations of the parameters C , D , and E . However, the condition that the risk aversion be monotonically decreasing as x is increasing is violated by some combinations of C , D , and E that were the best fitting parameters for a number of individuals. The violation of that condition occurs close to the origin and is due to the so-called zero illusion [12]. The zero illusion is the sudden change in curvature in the vicinity of zero and in Fig. 13. While this effect

may not be acceptable in a well worked out normative utility function, it can be expected to occur (as it did) in descriptions of attitudes.

The same procedure as described above was used in testing the acceptability of the final function (11). The best fitting parameters *C*, *D*, and *E* were found by a somewhat more complicated three-way search routine. These parameters were then used to calculate adjusted responses, and the interviewees were confronted with the adjustments. At this point, no individual was found that was dissatisfied with the proposed adjustments. In fact, everyone preferred the adjusted responses to their original responses, since the adjustments did not distort the overall attitude and clearly led to greater internal consistency. At this point, the functional form was judged acceptable as a reasonable starting point for an overall risk policy.

VI. THE DEVELOPMENT OF A CORPORATE RISK POLICY

First Presentation to Top Management

The purpose of this first presentation to the top management group was to explain the use of a corporate utility function in decision making, to review the results of the previous interviews, and finally to gain approval to proceed in trying to evolve a group strategy.

The presentation was given at a regular weekly meeting of the top management. As usual, the attendance at such a meeting included the president and the various vice presidents. Before the meeting each person had received a report covering the results of the 36 interviews. In addition, each one had been previously exposed to the material through the personal interview which had lasted an average of one and one-half hours.

The presentation was started with a comparison between the present method of project evaluation and risk analysis. A chart was used containing the equivalent of Figs. 1 and 2. The difficulty of reaching decisions from probability data was demonstrated with three alternative probability distributions. The certainty equivalent method of reaching such decisions was proposed as a tool for reaching these decisions, and the need for a risk policy (a utility function) was explained. How a risk attitude can be quantified was shown with a personal example.

The use of a risk policy in the calculation of a certainty equivalent was also demonstrated. Next the dilemma of having different attitudes among the corporate decision makers was pointed out and the inconsistency among the decision makers was emphasized by examples from the recent interviews. The need for a consistent corporate strategy was discussed at length. Finally it was proposed 1) that all top management members be reinterviewed, using a simplified technique (this technique will be discussed below), and 2) that the results of these interviews be used to check the reproducibility of risk attitudes and as a starting point in evolving a corporate risk policy through group discussion.

The presentation lasted about 45 minutes. After discussion of another 45 minutes, the proposal was accepted.

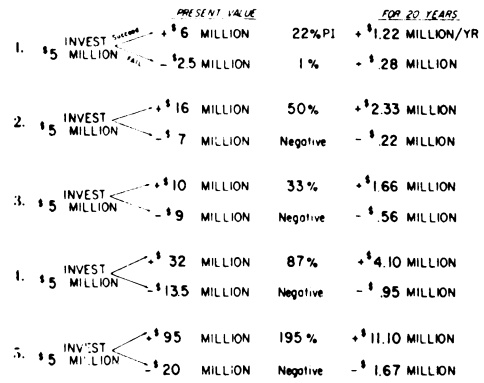


Fig. 14. Simplified interview schedule for \$5 million investment level.

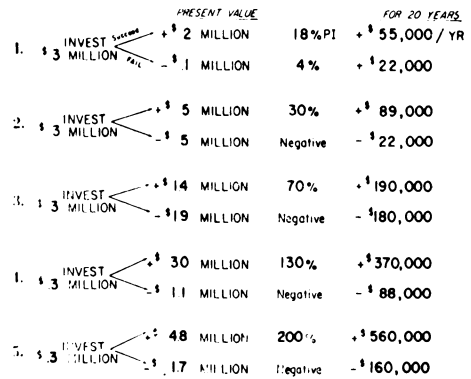


Fig. 15. Simplified interview schedule for \$0.3 million investment level.

In the discussion, the executives requested that the levels of investment be reduced by a factor of 10.

The acceptance of the proposal was by no means assured before the meeting, since the study demanded considerable time of top management. The initial interviews had been one and one-half hours, the presentation one and one-half hours, and the proposal meant at least another two hours of each top executive's time. The top managers' willingness to spend this time was a sign of their belief in the possible usefulness of the certainty equivalent as a decision-making tool.

The Revised Interview Technique

With the assumption of (11) as a reasonable function for the expression of individual risk attitudes, it is possible to simplify the interview technique considerably. Since the function contains only three parameters, a minimum of three investment decisions suffices to specify the complete function.

In the new interview technique, five responses at each investment level are required. These five responses provide the parameters of the function and two checkpoints which assure an acceptable fit of the function over the range of the outcomes.

TABLE III
 PROBABILITIES OF GAIN CONSISTENT WITH THE FINAL FUNCTION
 $U = A + B \log(X + C - D \cdot (\text{sort}(X^{**2} + E^{**2}) - E))$,
 where $D = 0.40$, $E = 2.00$ (IN MILLIONS OF DOLLARS)

	Gain =	6.00	16.00	10.00	32.00	95.00
	Loss =	-2.50	-7.00	-9.00	-13.50	-20.00
C =	28	0.442	0.556	0.720	0.660	0.763
	30	0.440	0.550	0.714	0.646	0.690
	32	0.438	0.545	0.709	0.633	0.649
	34	0.437	0.540	0.705	0.623	0.619
	36	0.435	0.536	0.701	0.614	0.596
	38	0.434	0.532	0.697	0.606	0.577
	40	0.433	0.529	0.694	0.599	0.561
	45	0.431	0.522	0.688	0.584	0.530
	50	0.429	0.517	0.683	0.573	0.506
	55	0.427	0.512	0.679	0.564	0.488
	60	0.426	0.509	0.675	0.556	0.473
	65	0.425	0.505	0.672	0.550	0.461
	70	0.424	0.503	0.670	0.545	0.450
	75	0.423	0.500	0.668	0.540	0.441
	80	0.422	0.498	0.666	0.536	0.434
	85	0.421	0.497	0.664	0.532	0.427
	90	0.421	0.495	0.663	0.529	0.421
	100	0.420	0.492	0.661	0.524	0.411
	110	0.419	0.490	0.659	0.520	0.403
	120	0.418	0.488	0.657	0.516	0.396
	130	0.418	0.487	0.655	0.513	0.390
	140	0.417	0.485	0.654	0.510	0.385
	150	0.417	0.484	0.653	0.508	0.381
	160	0.417	0.483	0.652	0.506	0.377
	170	0.416	0.482	0.652	0.504	0.374
	180	0.416	0.482	0.651	0.503	0.371
	200	0.415	0.480	0.650	0.500	0.366
	250	0.415	0.478	0.647	0.495	0.357
	300	0.414	0.476	0.646	0.492	0.351
	350	0.414	0.475	0.645	0.490	0.346
	400	0.413	0.474	0.644	0.488	0.343
	450	0.413	0.474	0.644	0.487	0.341
	500	0.413	0.473	0.643	0.486	0.339
	550	0.413	0.473	0.643	0.485	0.337
	600	0.413	0.472	0.642	0.484	0.336
	650	0.412	0.472	0.642	0.484	0.334
	700	0.412	0.472	0.642	0.483	0.333
	800	0.412	0.471	0.642	0.482	0.332
	900	0.412	0.471	0.641	0.481	0.331
	1000	0.412	0.471	0.641	0.481	0.330
	1100	0.412	0.470	0.641	0.481	0.329
	1200	0.412	0.470	0.641	0.480	0.328
	1300	0.412	0.470	0.640	0.480	0.327
	1400	0.412	0.470	0.640	0.480	0.327
	2000	0.412	0.469	0.640	0.479	0.325
	10 000	0.411	0.468	0.639	0.477	0.321
	100 000	0.411	0.468	0.639	0.476	0.320

Figs. 14 and 15 are examples of the new interview schedule. For each investment situation, the information which had been found most useful was provided. This information includes the net present value, the profitability index (continuously discounted rate of return), and the uniform annual cash flow assuming a 20-year duration. At the request of the interviewee, the probability of success, which leads to an expected present value equal to zero, was given orally. The investment situations were ordered by the magnitude of the outcomes and, as requested by the executives, the questions were based on \$5 million and \$0.3 million investments.

For each investment situation the probability response which is consistent with the function and a combination of the three parameters C , D , and E can be calculated. This calculation is a direct application of (9). Using this equation, a short computer program was written which developed an extensive book of tables giving

indifference probabilities of success, which correspond to various combinations of the parameters. One page from this book of tables is shown as Table III. The complete book of tables consists of pages similar to Table III for all combinations of D and E , where D ranges from 0 to 0.9 in intervals of 0.1, and E from 0 to 5.0 in intervals of 0.5. Thus there are 100 pages for each investment level.

In a brief introduction, each executive was reminded of the purpose of the interviews. The interview schedules for both investment levels were then introduced simultaneously, and the executive was familiarized with the 10 investment situations. Copies of the schedules were given to the interviewee as worksheets. To familiarize him completely with each investment situation, the interviewee was first asked qualitative questions which led to a ranking of the investment situations by the required chance of success. Only after he had carefully ranked the outcomes at each level of investment was he asked to respond with an indifference probability of success for the investment situations. The circular reference chart was used again to help in the understanding of the probability statements. The quantitative responses were reviewed in light of the previous ranking and any discrepancies were immediately removed by the subject. After this point, the interviewee was asked to give his preference, if any, between various pairs of alternative investment situations. If any preferences remained, the responses were readjusted to remove these, since the interviewee should show indifference between all investment situations and non-investment.

During this period, while the interviewee adjusted his responses, a combination of probabilities was sought from the book of tables which closely resembled the interviewee's responses and maintained his ranking. Generally, only one or two of the probabilities out of the five differed somewhat from the interviewee's. The interviewee was then offered these responses as an alternative to his responses. He was given to understand that he need not adjust his responses, but that the adjusted responses were internally consistent with a particular utility function. In every case, responses were found in the table that were fully acceptable to the interviewee.

After a set of acceptable responses was found, the interview was terminated with a short discussion of the planned session for reconciling the attitudes among the policy makers. A copy of his answers was retained by the interviewee.

Results of the Interviews

The new simplified technique of interviewing was found to be very successful. The executives termed it "more meaningful." The use of one of five investment situations at each level of investment allowed them to spend more time on each one and to understand fully the differences among the various choices. The patterns of attitudes offered by the three-parameter function seemed sufficient, since no individual insisted on deviating from

TABLE IV

Subject	Investment Level \$300 000			\$5 Million		
1	$C = 1\ 000\ 000$	$D = 0$	$E = 0$	$C = 1\ 000\ 000$	$D = 0$	$E = 0$
2	$C = 5$	$D = 0.2$	$E = 2$	$C = 38$	$D = 0.6$	$E = 1$
3	$C = 65$	$D = 0.1$	$E = 2$	$C = 65$	$D = 0.1$	$E = 2$
4	$C = 11$	$D = 0.4$	$E = 2$	$C = 110$	$D = 0.4$	$E = 2$
5	$C = 70$	$D = 0.4$	$E = 2$	$C = 70$	$D = 0.4$	$E = 2$
6	$C = 35$	$D = 0.7$	$E = 1$	$C = 50$	$D = 0.7$	$E = 2$
7	$C = 100$	$D = 0.7$	$E = 0.5$	$C = 65$	$D = 0.7$	$E = 0$

the patterns even though this alternative was repeatedly offered.

Most of the problems that had led to inaccuracy in the previous responses had been eliminated. By allowing the individuals to go back and forth between the alternatives, the effect of the order of the questions was almost completely eliminated. The learning effect was also eliminated, since the subject was encouraged to go back and change answers if he so desired. By using only five questions, the illusion of having the opportunity to invest in many investments of this type was also reduced.

The improved technique in conjunction with a more enlightened audience on the topic of decision theory led to meaningful results.

The parameters of the utility function

$$U(x) = A + B \log\{x + C - D[(x^2 + E^2)^{1/2} - E]\}$$

for the seven top executives are given in Table IV.

The relative level of conservatism can be seen better in Fig. 16. This figure presents the indifference probabilities in the same manner as in Section IV, except the responses of each executive are given by a discontinuous line and the investment situations are arranged in decreasing order. Again, the left limit of the scales represents a risk-neutral response; therefore, the horizontal distance stands for the probability premium. One executive, Subject 1, insisted on risk neutrality. A \$300 000 investment is a rather common investment for this company, while a \$5 million investment occurs somewhere between four and eight times each year.

Even though these results are somewhat different, comparison with the previous responses of each executive showed considerable consistency. About six months had elapsed since the first interview, and each manager had received a report detailing the results of these interviews. Thus any consistency over time is surprising. When some middle management was tested with the new interview technique, very little consistency with their previous responses was found. Instead a rather interesting phenomenon had occurred in the middle management. Members of middle management seemed to have revised their attitudes toward the average of the responses as given in the report. This difference between middle and top management leads the author to believe that middle management would welcome a corporate risk strategy statement by top management.

Evolving an Acceptable Corporate Risk Policy

The First Session: About two months after the initial presentation to top management, a meeting was scheduled with the express purpose of finding some acceptable corporate risk policy.

In the introduction to the meeting, the author emphasized the need for an overall risk policy and the broad implications of such a risk policy. The main points of this introduction were as follows.

1) The choice between a highly risk-averting policy and a virtually risk-neutral policy is essentially a choice between a safer low return on investment and a higher average return on investment. However, the higher average return carries with it a greater chance of fluctuations over time.

2) That the present risk attitudes vary greatly within the company was demonstrated by both groups of interviews. It is, therefore, reasonable to conclude that, in the absence of a stated risk policy, different standards are applied in various departments of the company. Also, many risky investments are probably suppressed before ever reaching top management.

3) The results from the interviews of the top executives show that everyone is willing to become virtually risk-neutral in projects that have outcomes within the range of $\pm \$300\ 000$. However, actual decisions indicate more conservative decisions on these projects.

4) Going to risk analysis and communicating the risk attitudes of top management are beginning steps to improve the situation. Agreement by the top executives is necessary to achieve a uniform risk policy for the company.

The results of the recently completed interviews, as shown in Fig. 16, were then presented. Each individual knew which line represented his responses, since he had in front of himself a copy of the interview schedule with his answers. The differences among the responses were pointed out. The executives were again reminded that no theoretically best policy exists. Thus the choice among the strategies must be a subjective one. This choice, however, was to be considered as an important policy decision.

At this point, the executives were invited to discuss the choice of a particular policy. To the great surprise of the author, Subject 1, whose responses had been risk-neutral, volunteered himself as the target of discussion. First he identified himself as the individual that prefers to play




Fig. 16. Probability premiums of the executive committee.

the averages up to \$5 million. He then stated that the following assumptions were the basis for his reasoning.

- 1) We have virtually unlimited funds available at 10 percent per annum.
- 2) An average of 10-percent return seems to be a satisfactory minimum to us in line with our present policies.
- 3) We have enough investments of these types to play the averages.
- 4) We must assume that the possible outcomes and probabilities can be accurately calculated.

He continued, that if these assumptions did not hold, we would need to change present company policy.

The discussion now started centering around the four points. Points 1) and 2) were quickly accepted. Point 3), covering the frequency of investments, was discussed at length. At the lower level of investment there resulted an agreement that not much danger lies with playing the averages. However, at the higher level of investment, where about five projects are approved annually, no agreement could be reached at this time. The fourth point, that probability calculations are feasible, was tentatively accepted but was of major concern.

Agreement appeared that a risk-neutral attitude seems to make sense at the lower investment level. On examining the present practices regarding such investments, the group concluded that considerable risk aversion is apparent. Thus a policy of risk neutrality would allow much more capital to be invested in these types of projects. In fact, one executive seemed to feel that all of the present capital investment budget could be absorbed in these projects using a risk-neutral policy. However, no executive felt that the overall capital budget should be increased, thus the company seemed to operate under a budget constraint in the short run. This is somewhat in contradiction to point 1) made by Subject 1. As a solution to this budget limitation, one of the participants suggested an increase in the discount rate. Thus, possibly, only projects with an average return of 12 percent per annum should be accepted.

The discussion returned one more time to whether there really existed any everyday risky projects, which could be added into the investment mix. In answer to this, some actual examples were given by one of the executives. He demonstrated that a conservative risk attitude is really applied and that many riskier projects could be considered. He also stated that the number of projects which can be considered is in direct relation to the manpower applied to the search and that they are presently constrained by the availability of manpower.

When the discussion slowed, the question was asked if a decision toward a specific risk policy could be made in light of the discussion. In answer, the top executive stated that this is "too big of a policy matter to be decided on the spur of the moment." The group then asked for another meeting to consider this matter further. A number of examples showing the effect of various risk policies on projects were also requested.

The meeting closed after two hours without a specific choice of a risk policy. However, at the lower investment level a consensus close to risk-neutral had appeared, while considerable disagreement still existed at the higher investment level.

The Second Session: Before the second meeting, two example projects were prepared. The first was a hypothetical project which was based on the output of a previous risk analysis. Using the same risk aversion, it was demonstrated that the risk premium increases as the magnitude of the outcomes increases.

The second example project consisted of a complete certainty equivalent analysis of a recently made decision. The probabilistic inputs were developed using the circular reference chart. The problem structure was then developed, using a stochastic decision-tree framework. The probability distribution of the present value was developed with a computer run of 750 samples. For the present value distribution, certainty equivalents using various utility functions were calculated.

Due to some computer programming problems, the material was presented to top management five months after the previous meeting. The presentation was again made and the attendance was virtually the same.

Because of the long lapse of time between the two meetings, the author reviewed the results of the first session in considerable detail. The executives were reminded that they seemed to be willing to play the averages on small investments following the discussion of the four points made by Subject 1. The four points were then reviewed. The decision to play the averages had led the executives to the question of capital budgeting and the conclusion had been reached that the discount rate should be raised rather than being risk-averse at small investments. The first session had ended without a choice of a risk policy, and the present meeting had been requested.

The purpose of this second meeting was then stated as the following:

- 1) to demonstrate how the certainty equivalent of a project is affected by changes in a risk policy and changes in magnitudes of outcomes;
- 2) to show in detail how a real project analysis can be made using the certainty equivalent method; and
- 3) to get closer to adopting an operationally acceptable risk policy.

Before discussion of the two example projects, the concepts of risk analysis and the certainty equivalent were reviewed. The previous exposure of management to these concepts had by now been considerable and their grasp of the subject was good.

After presenting the example projects and discussion of the usefulness of the certainty equivalent as a decision tool, the following proposals were introduced:

- 1) that specific risk policies be adopted for trial purposes at both levels of investment;
- 2) that the risk policy for investments of less than \$300 000 be communicated to middle management;

- 3) that the analysis of a specific very large investment decision, presently under consideration, be made using the certainty equivalent method;
- 4) that a risk analysis be made on a group of typical investment projects at the lower investment levels to demonstrate the effect of probability analysis on common investments.

Considerable discussion about the necessity of a specific risk policy ensued. The conclusion of this discussion resulted in a choice of a risk-neutral policy at the lower level of investment, and the policy represented by Subject 3 at the higher level of investment. The executives made it clear that their choice was not final, but that it represented their best present choice for a policy. Since the choice was tentative pending further examples of the effect of the policy, they did not accept the proposal to communicate the policy. Both proposals for analysis of presently important investment projects were accepted. After the results of these studies, the communication of the risk policy was to be reconsidered.

In total, the meeting resulted in a tentative choice of a risk policy and a strong affirmation of the effort toward using the concepts of risk analysis and certainty equivalents. In essence, many man-months of work on the projects were approved. This is of particular importance since at this time the required skills for such analysis are in very short supply.

Since the tentative choice of a corporate risk policy, progress has been made on the approved projects, but neither project has been completed.

Once fully established, the risk policy is expected to remain a pliable policy tool through periodic reviews. For that matter, each decision using the CE method will be an automatic review of the effect of the risk policy.

VII. SUMMARY AND CONCLUSION

In this paper, the feasibility of developing a corporate utility function was investigated. Such a utility function provides an important link in the application of decision theoretic models.

The corporate utility function was viewed as a policy statement by top management. Such a policy statement was evolved in a unique experiment.

First, 36 corporate executives were interviewed and their risk attitudes were quantified. From the responses of the interviewees, a mathematical function was developed that was able to reflect each interviewee's attitude. The fit of the function was tested by checking the reaction of the interviewees to adjusted responses. The functional form that led the interviewees to prefer the adjusted responses to their initial responses was finally accepted. The mathematical form of the function was considered a flexible pattern for a risk policy. The assumption was made that the corporate risk policy would be of this pattern.

With the pattern for a risk policy set, it was possible to simplify the method of deriving a particular individual's

risk attitude. Using the simplified method, the corporate policy makers were interviewed once more. The results from these interviews were then used as a starting point in two negotiation sessions. As a result of these negotiation sessions, the policy makers agreed on a risk policy for trial purposes. They also agreed to develop a number of major projects using the concepts of risk analysis and the certainty equivalent.

It is still too early to report on the company's experience of using a risk policy. However, the company's willingness to commit the required resources first to carry out this project and now to continue with application indicates a belief in the approach.

The specific choice of the trial policy is of particular significance. The choice was drastically different from the average of the group members. At the lower investment level, the choice of a risk-neutral policy by the group was surprising since only one member of the group had exhibited this attitude. At the higher level of investment, the group consensus was also for more risk taking than the average of the individual attitudes. The findings that group consensus tends to be more risk taking than the average of the individual group members is not unique. Social psychologists have presented similar findings in a number of articles.¹ These articles generally used examples of subjective real-life situations to measure risk attitudes. While the results are not directly comparable, the shift toward more risk taking by a group as a general phenomenon has been well documented. This increased willingness of a group to take risk is, therefore, not an isolated phenomenon and can be expected to occur in further group risk policy statements.

Some of the major effects of the study on the company have been educational. Because of the exposure, many managers have become more aware of the decision process under uncertainty. They are now able to discern between the risk in a project and their risk attitude which they bring to bear on the project. Thus they know the difference between the conservative estimator and the conservative decision maker. The managers are also more fully aware of the shortcomings of the single-value analysis. Therefore, the way is paved for further steps toward probabilistic decision techniques.

The success of this study greatly depended on the receptiveness of management to quantitative decision tools. In this respect, it must be mentioned that the managers of the subject company were quite open-minded. While most of the top management did view the sophistication of the decision model with suspicion, they were willing to be convinced. The job of selling the managers took considerable skill and effort. Clearly this task was made easier by management's generally technical backgrounds and their relative familiarity with capital investment analysis. The company presently ranks among

¹ This finding has been called to the author's attention by Prof. G. C. Hoyt of Iowa University, Ames, and is discussed in Hoyt and Stone [13].

the leaders in the use of sophisticated capital investment analysis techniques. Thus the introduction of the certainty equivalent method is just one more step, albeit a major one, in the path of improving their decision tools. Discounted cash flow methods had been applied in the company for better than a decade.

There are two major improvements in the corporate decision process that can be derived from developing a risk policy. Firstly, the communication of a risk policy throughout an organization could help to avoid opportunity costs that are incurred by premature rejection of risky projects. Secondly, when incorporated into capital investment analysis, the risk policy permits the interpretation of risk information by means of a certainty equivalent, which is consistent with the goal of maximization of expected utility. The communication of the policy may lead to a greater improvement in the decision process than the improvement in the analytical technique. However, to be effective, the risk policy statement must be supported by corresponding control procedures that do not penalize managers for taking risks. In the control procedure, a good decision must be carefully distinguished from a good outcome of a decision.

REFERENCES

- [1] L. J. Savage, *The Foundation of Statistics*. New York: Wiley 1954.
- [2] R. D. Luce and H. Raiffa, *Games and Decisions*. New York: Wiley, 1957, p. 13.
- [3] J. W. Pratt, H. Raiffa, and R. Schlaifer, *Introduction to Statistical Decision Theory*. New York: McGraw-Hill, 1965 (preliminary edition), pp. 4-15, 16.
- [4] C. J. Grayson, *Decisions Under Uncertainty: Drilling Decisions by Oil and Gas Operators*. Boston, Mass.: Graduate School of Business, Harvard University, 1960.
- [5] R. H. Cramer and B. E. Smith, "Decision models for the selection of research projects," *Engrg. Economist*, vol. 9, no. 2.
- [6] P. E. Green, "Risk attitudes and chemical investment decision," *Chem. Engrg. Progr.*, January 1963.
- [7] R. O. Swalm, "Utility theory—insights into risk taking," *Harvard Business Rev.*, November-December 1966.
- [8] H. M. Markowitz, *Portfolio Selection, Efficient Diversification of Investments*. New York: Wiley, 1959. See also D. E. Farrar, *The Investment Decision Under Uncertainty*. Englewood Cliffs, N. J.: Prentice-Hall, 1962.
- [9] F. T. Dolbear, "Individual choice under uncertainty," *Yale Econ. Essays*, Fall 1963.
- [10] G. M. Kaufman, *Statistical Decision and Related Techniques in Oil and Gas Exploration*. Englewood Cliffs, N. J.: Prentice-Hall, 1963.
- [11] J. W. Pratt, "Risk aversion in the small and in the large," *Econometrica*, vol. 32, January-April 1964.
- [12] J. S. Hammond, III, "Better decisions with preference theory," *Harvard Business Rev.*, pp. 123-141, November-December 1967.
- [13] G. C. Hoyt and J. A. F. Stone, "Leadership and group decision involving risk" (to be published).

Reprinted from IEEE TRANSACTIONS ON SYSTEMS SCIENCE AND CYBERNETICS
Vol. SSC-4, No. 3, September 1968
pp. 279-300

Copyright 1968, and reprinted by permission of the copyright owner
PRINTED IN THE U.S.A.

THE USED CAR BUYER

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Man is called upon to make decisions about his home, his business, and his pleasure. These decisions vary in importance, but they have one property in common: most people do not have an orderly procedure for thinking about them. Of course, it is not practical to spend much time and effort thinking about the minor decisions in our lives--yet how can we judge what is practical until we develop a logical framework for decision problems? Our present task is the construction of such a decision procedure.

There are three main points we shall attempt to make about the science of decision making.

1. Probabilistic considerations are essential in the decision-making process;
2. The lessons of the past must be included;
3. The implications of the present decision for the future must be considered.

Let us discuss each of these points. The importance of probability is revealed when we realize that decisions in situations where there is no random element can usually be made with little difficulty. It is only when we are uncertain about which of a number of possible outcomes will occur that we find ourselves with a real decision problem. Consequently, much of our discussion of decision-making will be concerned with the question of how best to incorporate probabilistic notions in our decision procedure.

The question of using previous information in making decisions seems to incite some statisticians to riot, but most of the rest of us think it would be unwise to make a decision without using all our knowledge. If we were offered an opportunity to participate in a game of chance by our best friend, a tramp, and a business associate, we would generally have different feelings about the fairness of the game in each case. Although we might agree on the necessity of considering prior information, it is not clear just how we shall accomplish this objective. The problem is intensified because the prior information available to us may range in form from a strong belief that results from many years of experience to a vague feeling that arises from a few haphazard observations. The decision formalism to be described will allow us to include prior information of any form.

The influence of present decisions upon the future is a point often disregarded by decision-makers. Unfortunately, a decision that seems appropriate in the short run may, in fact, place the decision-maker in a

very unfavorable position with respect to the future. For example, a naive taxi driver might be persuaded to take a customer on a long trip to the suburbs by the prospect of the higher fare for such a trip. He might not realize, however, that he will have to return in all likelihood without a paying passenger, and that when all alternatives are considered, it could be more profitable for him to refuse the long trip in favor of a number of shorter trips that could be made within the city during the same time period. The solution of such problems requires slightly more sophisticated reasoning than the first two points we have discussed, but it is just as amenable to an analytic approach.

Let us now begin our analysis of decision problems with an example that is so commonplace that there will be every possibility of understanding the environment of the problem, and yet is sufficiently detailed that it is not obvious at first glance just how the decision should be approached. A fellow named Joe, of our acquaintance, is in the market for a new car. He has decided to buy a three-year-old Spartan Six sedan, and has surveyed the used-car dealers for such a car. After searching for a while, he has found a car like the one he wants on one dealer's lot. The going rate for a three-year-old Spartan is \$1100, but the price asked by the dealer is only \$1000. Consequently, Joe figures that he will make \$100 profit by buying this particular car.

Unfortunately, just as Joe is about to close the deal, he overhears the salesman who has been serving him talking with another salesman. His salesman says, "This used-car business is a tough racket. I have a customer interested in the Spartan on our lot, but the practices of our business prevent me from warning him that he may get stuck if he buys it." The other salesman asks, "What do you mean?" Joe's salesman replies, "I worked at a Spartan dealership when that car first came on the market. Spartan made 20% of its cars in a new plant where they were still having production line troubles; those cars were lemons. The other 80% of total production were pretty good cars." The other salesman asks, "What is the difference between a 'lemon' and a 'peach'?" "Well," says Joe's salesman, "every car has 10 major mechanical systems--steering, brakes, transmission, differential, fuel, electric, etc. The peaches all had a serious defect in only one of these 10 systems, but the lemons had serious defects in 6 of the 10 systems." The other salesman replies, "Well, don't feel so bad; maybe some cars didn't have any defects, or maybe the defects in this car have already been fixed."


"No, that's just it," says Joe's salesman. "Every car produced had either 1 or 6 defects in the ratio I mentioned; and I happen to know, because the previous owner was a friend of mine, that this particular car has never been repaired." "If it is bothering you so much, why don't you

tell the guy it's a lemon and forget about it?" says the other salesman. "Ah," answers Joe's man, "that's the trouble. I personally don't know whether or not it is a lemon, and I'm certainly not going to take the chance of losing a sale by worrying a customer unnecessarily." To which the other salesman replies, "It's time for coffee."

We can now imagine the state of our friend Joe. What seemed like a real bargain has turned into a potential nightmare; he can no longer make the \$100 profit he had hoped for. Joe's first reaction is to turn and flee, but he has the icy nerves of a decision-maker and so soon regains his composure. Joe realizes that he would be foolish to forego the chance to buy the car he thought he wanted, at this price, without good reason. He decides to call an acquaintance who is a mechanic and get his estimate of what the possible repairs might cost. The mechanic reports that it costs about \$40 to repair a single serious defect in one of a car's major systems, but that if 6 defects were to be repaired, the price for all 6 would be only \$200.

Now Joe considers the possibilities open to him. He can either buy the car or refuse it. If he decides to buy the car, then his outcome is uncertain. If the car turns out to be a peach, then only one defect will develop and Joe will have made a profit of \$60: \$100 from buying the car at a low price, less \$40 for repairing the one defect. However, if the car should be a lemon, then Joe will lose \$100 because it will cost him \$200 to repair the 6 defects to be found in a lemon. If, on the other hand, he refuses to buy, then he gains and loses nothing.

We can represent the decision structure of Joe's problem by drawing a decision tree like that shown in Figure 1. The direction of the arrows refers to the time flow of the decision process. In this figure, each directed line segment represents some event in the decision problem. We have used B to indicate the event of Joe's buying the car, and R to indicate his refusing it. P is the event of the car's ultimately turning out to be a peach, while L is the event of the car's being a lemon. The tree as drawn in Figure 1 shows that the car may turn out to be a peach or a lemon regardless of Joe's action. Note that different symbols are used for the node joining the B-R branches and the nodes joining the P-L branches. The X is used to indicate points in the decision tree where the decision-maker must decide on some act; the • is used for nodes where the branch to be taken is subject to chance rather than decision. We shall call these two types of nodes "decision" nodes and "chance" nodes, respectively. In this example, Joe's only decision is whether to buy or refuse to buy; consequently, only the node joining the B-R branches requires an X. The ultimate outcome as to whether the car is a peach or a lemon is governed only by chance and so the P-L branches are joined by a •.



- B: Joe buys the car
- R: Joe refuses to buy
- P: Car is a peach
- L: Car is a lemon

Joe's Original Decision Tree

FIGURE 1


Generally, traversing each branch on the decision tree will bring some reward, positive or negative, to the decision-maker. We shall choose as a convention to write this reward under each branch. In Figure 1 we have written 100 under the branch labeled B to represent the immediate profit to Joe in buying the car; 0 is written under R branch, because Joe will neither gain nor lose by refusing to buy. The numbers under the P and L branches refer to the cost of repairing a peach and a lemon, respectively. If the decision-maker follows a tree from its unique starting node to all of its tips, then he will experience some pattern of gains and losses according to the branches he actually traverses. The net profit of all such traversals is written at each tip of the tree. Each tip may be designated by the sequence of branches that lead to it. Thus in this case the tip BP is given the value \$60 as the net profit in buying the car and then finding that it is a peach. The tip BL corresponds to a loss of \$100 from buying a lemon, while the tip R is evaluated at zero because the car is refused. These three tips of the tree represent the three possible outcomes of this decision problem. The outcome BP is favorable to Joe, the outcome BL is unfavorable, and the outcome R is indifferent.

Naturally, Joe would like the outcome to be BP with a profit of \$60, but after hearing the salesman's conversation he realizes that the likelihood of this outcome will be controlled by Nature rather than by himself. We can think of Nature as playing a game with Joe, as follows. When she

placed the car on the used-car lot, she made it a lemon with probability 0.2 and a peach with probability 0.8. She performed her selection by tossing a coin with probability of "heads" equal to 0.8 and made the car a lemon if the coin came up "tails." Thus the nodes that were chance nodes in Joe's decision tree we can imagine to have been performed by an opponent called Nature who is not malevolent and who selects actions using chance mechanisms.


We can draw a tree to show Nature's options, as is done in Figure 2. In Nature's tree, all nodes are chance nodes. We shall write above the beginning of each branch the probability that Nature will follow that branch. In the present example, we know that the probability of a peach is 0.8, the probability of a lemon is 0.2. We also write at each tip of Nature's tree the probability that Nature will produce an outcome corresponding to that tip. In general, these probabilities are calculated by multiplying together the probabilities on all the branches that lead from the initial node on Nature's tree to each tip. In this simple case, all we must do is write 0.8 and 0.2 at the end of both the P and L branches.

The importance of Nature's tree, as we shall see, is that it provides all the probabilistic information that is necessary for the decision tree. To illustrate this point, we recall that we have yet to write probabilities on each chance node of the decision tree. The results of the calculations in Nature's tree allow us to draw Figure 1 in the form of Figure 3. The various features of Figure 3 will be explained gradually. At the moment, our example has such a simple form that it is not at all clear why it is necessary to consider a separate tree for Nature. As our example becomes more complex, the need for Nature's tree will be evident. The numbers in the square boxes at each node in Figure 3 represent the



Nature's Tree

FIGURE 2



Joe's Decision Tree with Probabilities from Nature's Tree

FIGURE 3

net profit to Joe from future activities if he should arrive at such a node. Thus, if Joe is at node B (we label nodes by the branches that must be traversed to reach them), then he expects to earn \$60 with the probability 0.8, and lose \$100 with probability 0.2. His expected earnings are $0.8(60) + 0.2(-100) = \$28$. Of course, if Joe decides not to buy the car, then he will earn nothing, and so 0 appears in the square box appended to node R.

As a result of evaluating each possible action that Joe might take in terms of its expected value equivalent, we are in a position to help Joe with his decision. If Joe buys the car, then he expects to earn \$28; if he refuses to buy, he will earn nothing. If Joe is an expected-value decision-maker, he should decide to buy the car. His recommended action is shown by drawing a solid arrowhead on the B branch leading from the decision node. We then write his expected profit from taking that action, \$28, in the square box over the decision node.

As a result of this analysis of the problem, Joe feels a little better than he did before. He has forsaken all hope of a \$100 profit and is coming around to the idea that it might be wise to settle for an expected profit of \$28. However, while he is becoming reconciled to the forces of fate, a stranger approaches him and says, "I couldn't help overhearing you talking to yourself about your problems. Perhaps I can help you. You see, I worked in the factory where the substandard Spartans, or lemons as you called them, were made. I can tell you whether the car sitting on this lot is a lemon simply by looking at the serial number." Joe can hardly believe his ears. At last, a possibility of finding out whether the car is a lemon before buying it.

Joe looks at the man, decides he has an honest appearance, and says, "You are just the kind of help I need. Let's go over to the car and take a look at it. I am eager to find out whether or not it is a good deal." The stranger smiles and replies, "I am sure you are, but you can hardly expect me to go to all the trouble of examining the car and getting myself dirty without some financial consideration." At first Joe is angry about the stranger's mercenary attitude, but then he remembers he is not in a position to throw away potentially useful information if it can be obtained at a reasonable price. He asks for and is granted a few moments to think over the stranger's offer.


The problem is this; how much is Joe willing to pay the stranger for his information? He reasons as follows. On the basis of the stranger's appearance and manner, Joe decides that he can be trusted in his claim of being able to distinguish peaches from lemons. If the stranger reports that the car is a peach, then Joe will buy it and make an expected profit of \$60. If the stranger says it is a lemon, then Joe will refuse to buy it and make nothing. The probability that the stranger will find a peach is 0.8; the probability of finding a lemon is 0.2. Consequently, Joe's expected profit after receiving the information is $0.8(60) + 0.2(0) = \$48$. Therefore, is the information worth \$48? No, because even without it Joe expects to make \$28, according to our original analysis. Hence, the net value of the stranger's information to Joe is \$20. That is, Joe as an expected-value decision-maker should be willing to pay any amount up to \$20 for the stranger's advice.

This figure of \$20 seems high to Joe, so he decides to check it in the following way. Joe thinks, without this new information I would buy the car and make an expected profit of \$28. If I buy the information, then with probability 0.8 the stranger will report that the car is a peach and his information will be worthless because I am going to buy the car anyway. On the other hand, with probability 0.2 the stranger will find that the car is a lemon, and in this case the information is worth \$100 since that is the amount that I would lose if I bought the car and it turned out to be a lemon. Consequently, the expected value of the information to me is $0.8(0) + 0.2(100) = \$20$, the same as before. Now Joe is convinced that he should pay as much as \$20.

We shall call this quantity the expected value of perfect information, or the EVPI. It represents the maximum price that should be paid for any experimental results in a statistical decision situation. This follows since no partial knowledge could ever be worth more than a report of the actual outcome of nature's process. We shall have much more to say of this quantity in our later discussion.

Joe now decides to offer the stranger \$15 in hopes of getting the information at a bargain price. However, when he confronts the stranger with this offer, the stranger replies that he couldn't consider the job for less than \$25 and suggests that Joe think it over for a while. Joe is upset by this turn of events, but quickly regains his composure. He thinks to himself that the real reason for his difficulties is that he doesn't have a wide enough range of alternatives from which to select an appropriate action. Suddenly he has a brainstorm--maybe he can get the dealer to give him the guarantee on the car! He inquires of the dealer whether a guarantee is available. The dealer says, "Yes, there is a guarantee plan; it costs \$60 and covers 50% of repair cost." Joe thinks fast and replies, "You certainly don't have much confidence in your cars. If I bought a car and it turned out to be a lemon, I could go broke even on my 50%." The dealer says, "All right. Just for you I will include an anti-lemon feature in the guarantee. If total repairs on the car cost you \$100 or more, I will make no charge for any of the repairs. How's that for meeting a customer half-way?" Joe says that's fine and now he would like to think it over again.

At this point Joe realizes that he has a new decision tree. It is shown in Figure 4. This tree differs from the preceding one because there are now three possible actions at the decision node. The new alternative is to buy the car with the guarantee; that is, to hedge against the possibility of getting a lemon by spending \$60. This alternative is given the symbol G. We see that although the car might still turn out to be a lemon if this alternative is followed, the costs associated with the two outcomes



Joe's Decision Tree Including Guarantee Possibility

FIGURE 4

are strikingly different from what they are in the case where the car is bought without such a guarantee.

Let us examine Figure 4 in some detail. The figures written below each branch are again the expected profit from traversing that branch. The numbers on the tips are the total expected profit of the chain of branches leading to that tip. Now, as before, we shall choose to calculate the expected value of each node by using the number on the tips rather than on the branches. However, this choice is arbitrary and will be reversed when a reversal is convenient.

The expected value of the nodes B and R are calculated as before. The value of \$40 written under the G branch refers to the fact that our initial profit from buying the car with the guarantee is only \$40 because the guarantee itself costs \$60. The value of -\$20 over the P branch following the G action arises because even a peach will require one repair at a cost of \$40, but half of this \$40 will be paid by the guarantee. The 0 under the corresponding L branch is a result of the anti-lemon feature of the guarantee. Since the cost of repairs on a lemon will exceed \$100, there will be no charge for repairs. Thus the net profit of buying the car with a guarantee and having it turn out to be a peach is \$20, while the profit if it turns out to be a lemon is \$40. Since Nature's tree of Figure 2 still applies to this case, the probabilities of these two events have values of 0.8 and 0.2, respectively. Hence, the expected earnings from buying the car with the guarantee is $0.8(20) + 0.2(40) = \$24$. Since this is less than the \$28 profit to be expected if the car is bought without the guarantee, the guarantee does not look like a good idea. The choice should once more be to buy the car without any protection, as is indicated by the heavy arrowhead on the B branch.

At this point our knowledgeable stranger returns and once more offers his advice--for a price. Has the advent of the guarantee changed what Joe should pay? Let's find out. If the information is bought, the stranger will find that the car is a peach with probability 0.8. If a peach is reported, then Joe will buy it without a guarantee and make an expected profit of \$60. With probability 0.2 the stranger will discover a lemon. In this case, however, Joe is best advised not to refuse the car and make nothing as he did before, but rather to buy it with the guarantee. As the number on the tip of the branch GL in Figure 4 indicates, by taking this action he will earn an expected profit of \$40. Thus, the amount that Joe expects to earn by buying the car is $0.8(60) + 0.2(40) = \$56$. Since Joe expects to earn \$28 anyway by buying the car without this information, the value of the additional information to him is \$28.

It may at first seem strange that the expected value of perfect information, or EVPI, should increase simply because an alternative has been added to those already available to Joe. However, such an increase has taken place as a result of the fact that Joe is in a better position to make use of information that the car is a lemon than he was previously. We can verify the figure of \$28 using the same method employed before. If the stranger reports a peach, then Joe's decision to buy the car will be unchanged; but if a lemon is reported, then Joe will buy the car with, rather than without, the guarantee and so will turn a loss of \$100 into a profit of \$40. Consequently, his expected profit will increase by \$140 with probability 0.2. Thus, the information is worth $0.2(140) = \$28$ to Joe.

Now, of course, the stranger's asking price of \$25 for the perfect information seems quite reasonable. Joe is about to purchase the information when he has another brainstorm. He knows that perfect information is worth \$28 to him, and so he reasons that if he can get partial information at a price sufficiently lower than \$28 he may be able to increase his profits. He first asks the dealer if he can take the car to his mechanic friend for a checkup. The dealer is willing to allow this, but places a time limit of one hour on the car's absence from the lot. Somewhat elated, Joe calls his friend to ask what kind of tests could be performed in an hour and how much they would cost. The mechanic says that he can only do at the most one or two tests on the car in the time available. He then supplies Joe with the following test alternatives:

1. He can test the steering system alone, at a cost of \$9;
2. He can test two systems--the fuel and electrical systems--for a total cost of \$13;
3. He can perform a two-test sequence, in which Joe will be able to authorize the second test after the result of the first test is known. Thus, under this alternative, the mechanic will test the transmission, at a cost of \$10, report the outcome of the test to Joe, and then proceed to check the differential, at an additional cost of \$4, if he is requested to do so.

All the tests will find a defect in each system tested, if a defect exists. The test alternatives are summarized in Table 1.

Including the possibility of no testing, Joe now looks over these test alternatives and decides that it is worthwhile at least to consider testing because the cost of each of these tests is significantly less than the \$28 value of perfect information. If all tests had cost over

Table 1

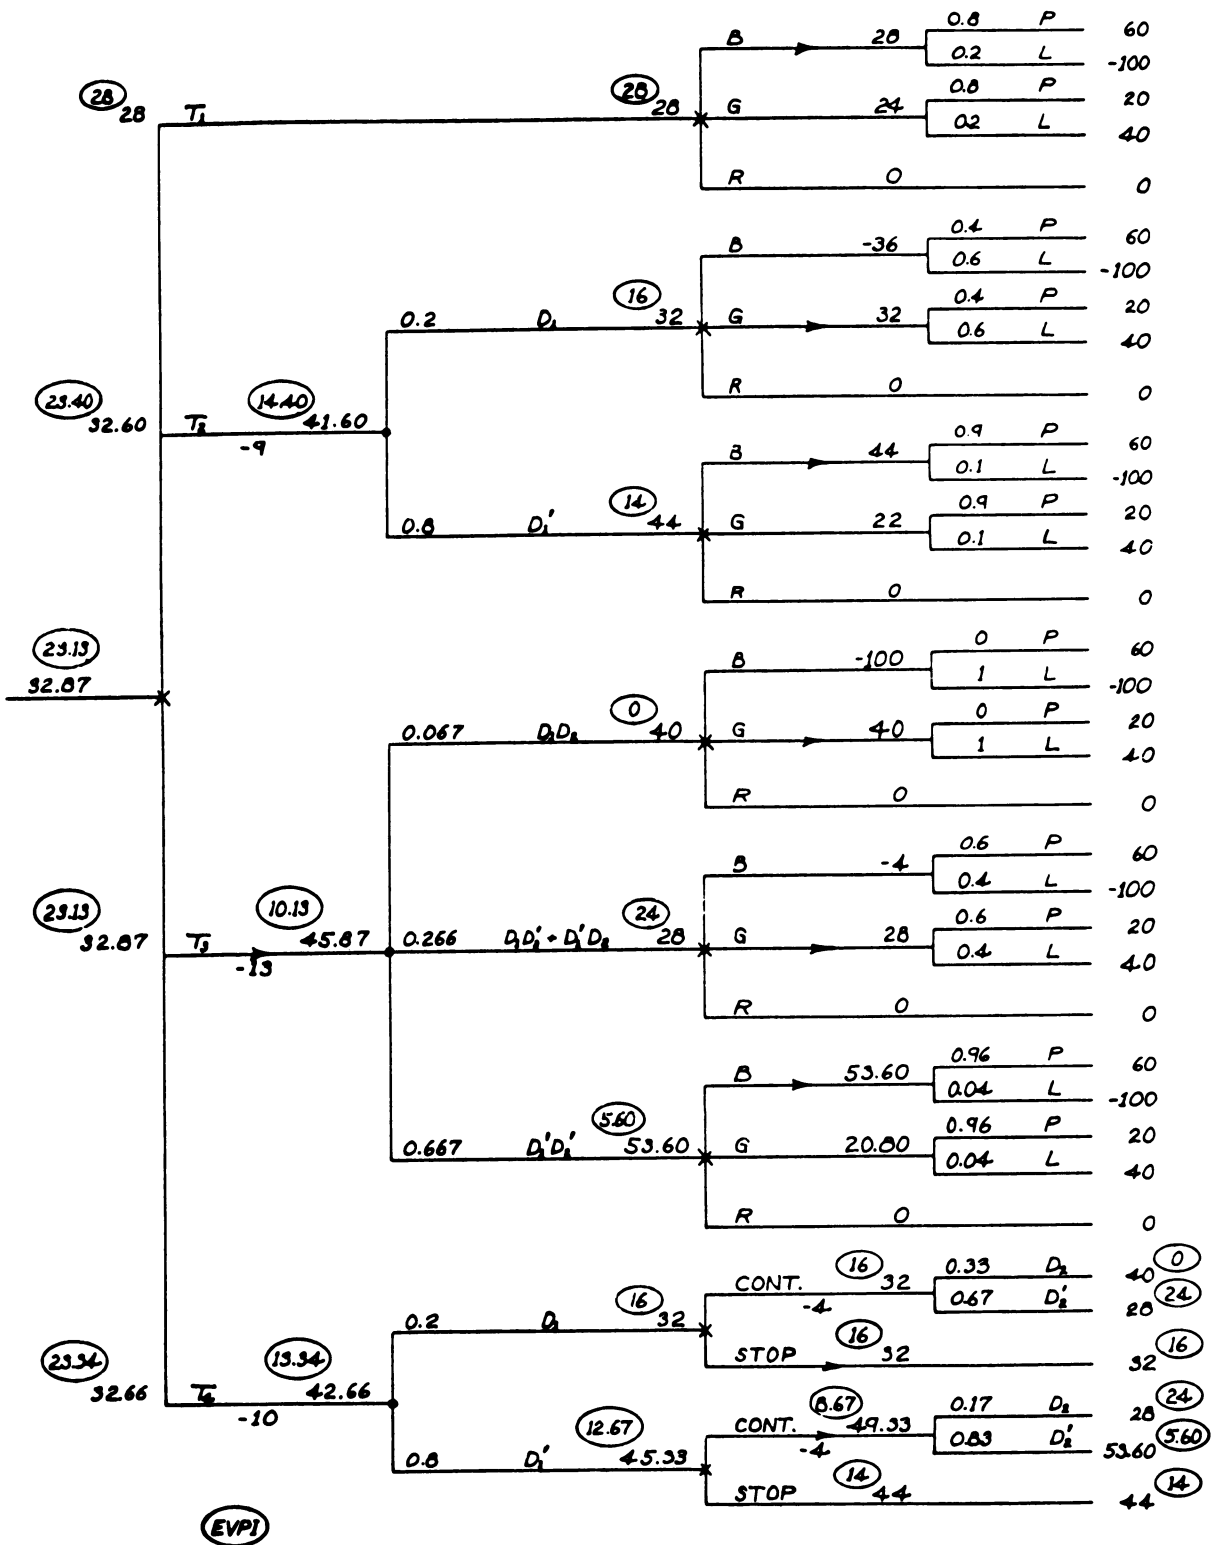
THE TEST ALTERNATIVES

<u>Test</u>	<u>Description</u>	<u>Cost</u>
T ₁	Perform no tests	\$ 0
T ₂	Test steering system	9
T ₃	Test fuel and electrical systems (2 systems)	13
T ₄	Test transmission with option on testing differential for	10 4

\$28, then there would be no point in considering a testing program because each test will generally provide only partial information, and even perfect information is worth a maximum of \$28. However, it is still not clear which test, if any, should be performed. Furthermore, Joe would like to know the value of the stranger's information under these new circumstances. These problems will be approached by drawing a new decision tree for Joe and a new tree for Nature. The general structure of the decision tree is shown in Figure 5.

This tree is quite complicated, so we shall explain it in gradual steps. Notice that the first decision to be made is which of the four test options--T₁, T₂, T₃, T₄--to follow. If some tests are made, the mechanic will report the results, and then a decision about buying the car must be made. If the test T₄ is used, of course, then there will also be a step in which the mechanic is advised whether or not to continue the test procedure. Let us now examine the situation resulting from each test in more detail.

If test T₁ is selected, then no physical test is made and Joe is required to make a decision about buying the car immediately. The decision tree from this point on looks just like that of Figure 4. In fact, the numbers that appear in Figure 4 have been reproduced exactly in Figure 5, with the exception that only the numbers on the tips of the branches have been copied because they are sufficient for our purposes. Indeed, a little reflection will reveal that regardless of the test program we follow, we must end up with a decision tree like that of Figure 4. However, although the numbers on the tips of the branches will be the same in all cases, the



Joe's Complete Expected Value Decision Tree


FIGURE 5

probabilities to be written on the branches will differ in each case. The probability of the final outcome of a peach or a lemon will generally depend on the findings of the experimental program until the time the decision on buying the car must be made. For example, if two defects have been found, then the car is a lemon with probability one.

We see that what is now required is a mechanism that will give for each possible result of the experimental program the appropriate probabilities for the ultimate outcome of a peach and a lemon. Nature's tree is just such a mechanism. It is drawn for this problem in Figure 6. In this figure we have used D_1 to represent the event that a defect is discovered in the first test on the car, if such a test is performed, and D_2 is used similarly to indicate the finding of a defect on a second test, if any. The numbers on each branch represent the conditional probabilities of going to each following node, given that the present node has been reached. The numbers on the nodes represent the unconditional probability of occupying that node. The tree can then be explained as follows. Nature first decides whether the car is to be a peach or a lemon with probabilities 0.8 and 0.2, respectively, using some random process like the biased coin-flipping described earlier; thus, $p(P) = 0.8$, $p(L) = 0.2$.

Suppose that the car has turned out to be a peach. Then, using our convention that a node is labeled by the letters on the branches that must be traversed to reach it, we are at node P. Now suppose that one major system of the car is tested. Since the car is a peach, there is probability 1 in 10, or 0.1, that the one defective system will be checked and found defective; thus $p(D_1|P) = 0.1$. If this happens, we proceed to the node PD_1 ; then, $p(PD_1) = p(P) p(D_1|P) = 0.08$. On the other hand, with probability 0.9 no defect is discovered and we reach node PD_1' . Suppose, further, that a second test on another system is now performed. If we are at node PD_1 , then the only defective system in the car has already been discovered and there is probability 0 of finding another defect and reaching node PD_1D_2 . Under these circumstances, we shall be certain to proceed to node PD_1D_2' . The overall probability of such event as PD_1D_2 is determined by multiplying together the probabilities on all the branches that lead to that tip of the tree. Thus, $p(PD_1D_2) = 0$ and $p(PD_1D_2') = 0.08$.

If the car were a peach, but no defect had been found on the first test, then we would be at node PD_1' . If, now, a second test is performed, it will yield a defect with the probability that the system tested is the one defective system in the remaining nine, or 1/9. Of course, the probability of finding no defect in this situation is then 8/9. The overall probabilities $p(PD_1'D_2) = 0.08$ and $p(PD_1'D_2') = 0.64$ can then be calculated.



Nature's Tree for Complete Decision Problem

FIGURE 6

If Nature selects a lemon initially, then the same sort of reasoning applies. The probability of finding a defect in the first test on a lemon is equal to the chance of testing one of the 6 defective systems out of the 10 systems on the car, or 0.6. If one defect has been found in a lemon, then the probability of finding another is the chance that one of the 5 defective systems among the remaining 9 systems will be inspected, or 5/9. If, on the other hand, no defect is found in the first test on a lemon, then the probability of finding one during the second test is the chance of testing one of the 6 defective systems among the 9 systems remaining, or 2/3. The probabilities of all final outcomes pertaining to the lemon branch of the tree are then computed and written on the tips of the branches.

Figure 6 contains all the information necessary to answer any question about the probabilistic structure of the decision process. We can best see this by returning at this point to our discussion of the test alternatives in Figure 5.

If the alternative T_2 , test one system is followed, then the first requirement is that Joe pay \$9 for the services of the mechanic. This payment is indicated by the -9 on the T_2 branch. The next event to take place is the report of the mechanic on whether or not he found a defect. His report is a chance event, so indicated by the solid dot that follows branch T_2 . The mechanic reports either that he found a defect, D_1 , or did not find a defect, D_1' . However, the probability that each of the branches D_1 or D_1' will occur must yet be determined. But $p(D_1) = p(PD_1) + p(LD_1)$ since P and L are mutually exclusive and collectively exhaustive events. By using the results of Nature's tree in Figure 4, we have $p(PD_1) = 0.08$, $p(LD_1) = 0.12$ and so $p(D_1) = 0.2$; of course, $p(D_1') = 0.8$. These two probabilities are recorded on the branches D_1 and D_1' that follow branch T_2 to indicate the nature of the chance point. Once D_1 or D_1' has occurred, Joe faces a decision tree like that of Figure 4, but with different probabilities that will be calculated from Nature's tree in Figure 6. In particular, we require the probabilities $p(P|D_1)$, $p(P|D_1')$ and their complements. These probabilities are easy to obtain because $p(P|D_1) = p(PD_1)/p(D_1)$ by definition, and we have just calculated both probabilities involved in this expression. Thus, $p(P|D_1) = 0.08/0.2 = 0.4$, and $p(L|D_1) = 0.6$. These numbers are entered as the probabilities of a peach and a lemon, respectively, on the branches that follow node T_2D_1 in Figure 5. Similarly, $p(P|D_1') = p(PD_1')/p(D_1') = 0.72/0.80 = 0.9$, again using the results of Figure 6, and $p(L|D_1') = 0.1$. The branches for peach and lemon that follow node T_2D_1' in Figure 5 are labeled with these probabilities.

We have now obtained the complete probabilistic structure of the test T_2 . The branches emanating from every chance point have been assigned the appropriate probabilities. It is now possible to determine the expected profit to be obtained by following test T_2 . First, we shall compute the decision to be made if a defect is reported. If, in this case, Joe decides to buy the car without a guarantee, he will earn \$60 with probability 0.4 and lose \$100 with probability 0.6. His expected profit is then $-\$36$. If he hedges by buying with the guarantee, his expected profit is $0.4(20) + 0.6(40) = \$32$. If he refuses to buy, he earns nothing. Since \$32 is a better result than no earnings or a loss, Joe should decide to buy the car with a guarantee if he finds himself at this situation. His expected return will be \$32, as indicated in the square boxes following node T_2D_1 .

On the other hand, if the mechanic finds no defect in the steering, then Joe will be at node $T_2D'_1$ and will again be faced by a decision. If he buys without a guarantee, his expected profit is $0.9(60) + 0.1(-100) = \$44$. If he buys with a guarantee, his expected profit is $0.9(20) + 0.1(40) = \$22$. Again, he makes nothing if he refuses to buy. Since \$44 is the maximum return, he should decide to buy the car without the guarantee. The expected earnings of \$44 are written at the end of branch $T_2D'_1$.

There is but one step remaining in the analysis of test option T_2 . If the mechanic reports a defect, Joe expects to earn \$32. If he reports no defect, then Joe expects to earn \$44. These two events happen with probability 0.2 and 0.8, respectively, according to the earlier calculations using Nature's tree. Hence, the expected profit before the results of the test are known, but after the test has been paid for, is $0.2(32) + 0.8(44) = \$41.60$. Since Joe must pay the mechanic \$9 to reach this position, his expected earnings from test T_2 , including the payment to the mechanic, are $\$41.60 - \$9 = 32.60$. This number is entered at the left of branch T_2 to indicate the expected profit from following this test program. Since we have already calculated the expected profit of program T_1 to be \$28, it is clear that Joe is better advised to proceed with the test on the steering rather than to make the decision in the absence of this information. By so doing he will increase his expected earnings by \$4.60. Of course, it is still not proved that T_2 is the best test alternative to follow--we have only shown that it is better than T_1 . It remains to investigate T_3 and T_4 .

Before we do so, however, let us return once more to the concept of the value of perfect information. We have already shown that the partial information supplied by option T_2 is more valuable than its cost. How has this revelation affected our evaluation of the stranger's information? Before the test alternatives were introduced, Joe had calculated that the expected value of perfect information was \$28. As you recall, this figure

was determined by calculating first the amount of money Joe could make if the perfect information was available to him (\$56) and then subtracting from this quantity the amount he could expect to earn in the absence of this information (\$28); thus, EVPI equalled $\$56 - \28 . Now what has changed in these calculations? The \$56 profit to be expected by using perfect information has remained unchanged since the introduction of the guarantee plan. However, Joe's expectation without the stranger's information has been increased from \$28 to \$32.60. Hence, the expected value of perfect information has been lowered to $\$56 - \$32.60 = \$23.40$.

It is interesting to note how we have vacillated about the value of the stranger's information. Before the advent of the guarantee plan, it was \$20 and the stranger's price of \$25 seemed too high. Then the guarantee possibility was introduced and the value of perfect information rose to \$28. At that point the stranger's \$25 price seemed like a bargain. Finally, however, Joe calculated the results to be expected using the test alternative T_2 and saw that the value of perfect information had decreased to \$23.40, a figure below the stranger's price. Consequently, Joe is not in a mood to buy at the moment. Although he has not yet evaluated the value of perfect information under test plans T_3 and T_4 , at this point he is sure that it cannot possibly be greater than \$23.40.

The value of perfect information at each point in the tree will be shown in Figure 5 in the ovals at pertinent nodes. In every case the EVPI is calculated simply by subtracting the expected earnings at each node from the profit to be expected if the perfect information were available. At the two nodes that begin and end branch T_2 , the result of the test is not known and so the expected profit using perfect information is still \$56. Thus the node to the right of branch T_2 bears the EVPI \$14.40 since $\$56 - \$41.60 = \$14.40$. Perfect information is worth \$9 less than it was to the right of branch T_2 because of the payment to the mechanic.

The calculation of the value of perfect information is performed in the same way when the test results are known, but in this case, the expected profit from using the perfect information is different. Consider the situation where a defect has been reported. Joe knows that if the car is a peach he should buy it without the guarantee and make \$60, and that if it is a lemon he should buy it with the guarantee and make \$40. In the absence of any test result, the stranger would report a peach with probability 0.8 and a lemon with probability 0.2, so that Joe's expected profit would be $0.8(60) + 0.2(40) = \$56$. However, now that a defect has been reported, the probabilities of a peach and a lemon have changed to 0.4 and 0.6, respectively. Thus, the expected profit using perfect information is now $0.4(60) + 0.6(40) = \$48$. It is from this quantity that

the expected value of state T_2D_1 , \$32, must be subtracted in order to obtain the EVPI of \$16 entered in the oval above node T_2D_1 .

Similarly, we see that if no defect had been reported, the probabilities of peach and lemon would be 0.9 and 0.1, and the expected profit of using perfect information would be $0.9(60) + 0.1(40) = \$58$. When we subtract the \$44 value of node T_2D_1' , we obtain the \$14 figure for the EVPI that is pertinent to that node.

There is one other observation we should make. The values of perfect information at nodes T_2D_1 and T_2D_1' are \$16 and \$14. The probabilities of arriving in each of these states is 0.2 and 0.8, respectively. Consequently, the expected value of what the expected value of perfect information will be after the mechanic report is $0.2(16) + 0.8(14) = \$14.40$, in agreement with our previous value for this quantity entered in the oval at node T_2 . Thus, it is possible to compute the expected value of perfect information at each point in the tree by using only the values of perfect information pertinent to the final decision on buying the car and the probabilistic structure of the tree. We shall have more to say of these quantities later.

Let us now move forward to an analysis of test option T_3 . In this case, as you recall, two systems on the car--the fuel and electrical systems--are subjected to test and then the results of both tests are reported to Joe. The possible reports are that 2, 1, or 0 defects were found. These three events are represented by the three branches, D_1D_2 , $D_1D_2' + D_1'D_2$, and $D_1'D_2'$ that are drawn to the right of node T_3 in the tree of Figure 5. Note that once more we have written under branch T_3 the amount to be paid to the mechanic for performing the tests. When the mechanic's report is known, Joe must make a decision on buying the car, using a decision tree similar to that shown in Figure 4. The expected earnings at the tips of the tree remain the same, but once more we require a new assignment of the ultimate probabilities of a peach and a lemon as a result of the mechanic's report. These probabilities may be found from Nature's tree in Figure 6. The probabilities necessary are: $p(D_1D_2)$, $p(D_1D_2' + D_1'D_2)$, $p(D_1'D_2')$, $p(P|D_1D_2)$, $p(P|D_1D_2' + D_1'D_2)$ and $p(P|D_1'D_2')$. By using the numbers on the nodes of Nature's tree and the basic relations of probability theory, we obtain the following results:

$$p(D_1 D_2) = p(PD_1 D_2) + p(LD_1 D_2) = 0 + 1/15 = 1/15 = 0.067$$

$$\begin{aligned} p(D_1 D'_2 + D'_1 D_2) &= p(PD_1 D'_2) + p(LD_1 D'_2) + p(PD'_1 D_2) + p(LD'_1 D_2) \\ &= 6/75 + 4/75 + 6/75 + 4/75 = 4/15 = 0.266 \end{aligned}$$

$$p(D'_1 D'_2) = p(PD'_1 D'_2) + p(LD'_1 D'_2) = 48/75 + 2/75 = 2/3 = 0.667$$

$$p(P|D_1 D_2) = p(PD_1 D_2)/p(D_1 D_2) = \frac{0}{1/15} = 0$$

$$\begin{aligned} p(P|D_1 D'_2 + D'_1 D_2) &= [p(PD_1 D'_2) + p(PD'_1 D_2)]/p(D_1 D'_2 + D'_1 D_2) \\ &= \left[\frac{6/75 + 6/75}{4/15} \right] = 3/5 = 0.6 \end{aligned}$$

$$p(P|D'_1 D'_2) = p(PD'_1 D'_2)/p(D'_1 D'_2) = \frac{48/75}{2/3} = 24/25 = 0.96$$

Thus, we see that after Joe has committed himself to the test, there are probabilities of 0.067, 0.266, and 0.667 that the mechanic will report 2, 1, or 0 defects. These numbers are entered in Figure 5 on the three branches leaving the chance node T_3 . If two defects are reported, $p(P|D_1 D_2)$ shows that Joe will make his decision with the satisfying, but disappointing, knowledge that the car is certain to be a lemon. This information is indicated on the tree by the 0 and 1 entered on the branches P and L that originate in chance nodes $T_3 D_1 D_2 B$, $T_3 D_1 D_2 G$, and $T_3 D_1 D_2 R$. The expected earnings from making each of the decisions B, G, and R are -100, 40, and 0. Consequently, the most profitable act for Joe is to buy the car with the guarantee, even though it is a lemon, and thus earn the \$40 profit. This preferred decision is shown by the solid arrowhead on the branch G following node $T_3 D_1 D_2$; the profit of \$40 is recorded in the square box above that node.

The situation when only one defect is reported is very similar. In this case, we observe from $p(P|D_1 D'_2 + D'_1 D_2)$ that the probabilities of a peach and a lemon are 0.6 and 0.4. These probabilities appear on the P and L branches at the ends of the sub-tree that follows node $T_3(D_1 D'_2 + D'_1 D_2)$. The expected earnings of the three acts B, G, and R are $0.6(60) + 0.4(-100) = -\$4$; $0.6(20) + 0.4(40) = \$28$; and \$0. Once more, the highest expected profit will result if Joe buys the car with the guarantee. Note that he does this even though the car is still more likely to be a peach

than a lemon. Again we record the expected profit of \$28 in the square boxes over the decision node and indicate the preferred decision with a solid arrowhead.

If no defects are reported, the car is almost certain to be a peach; there is only a 4 percent chance of its being a lemon. When we compute the expected profit of the three decisions following node $T_3 D_1' D_2'$, using the probability 0.96 for a peach and 0.04 for a lemon, we find that buying the car without a guarantee pays \$53.60, buying it with a guarantee pays \$20.00, and not buying it at all pays nothing. Thus, Joe is best advised to buy the car without the guarantee, as represented by the solid arrowhead on the B branch following node $T_2 D_1' D_2'$ and by the \$53.60 entered in the square box over that node.

We have now calculated the optimum decision and maximum expected earnings for each possible mechanic's report under test plan T_3 . As we know, chance determines the actual reporting, but we also have learned the probabilities of the mechanic's reporting 2, 1, or 0 defects, and have entered them in the decision tree. The expected profit to Joe when he is waiting to learn the test results is thus $0.067(40) + 0.266(28) + 0.667(53.60)$, or \$45.87. Of course, in order to reach a situation with this expected value, Joe had to pay out \$13. Hence, his expected earnings from test T_3 are \$32.87. Since this number is higher than the expected profit under either the policy of no testing or of testing only one system, the option of testing two systems for \$13 is the most favorable yet evaluated. However, its margin over test plan T_2 is only \$0.28.

We might, at this point, examine once again the value of the perfect information offered by the stranger. As we found earlier, this quantity can be calculated at each node of the decision tree simply by subtracting from the expected earnings with perfect information the expected earnings at that node as given in the pertinent square boxes. Accordingly, since the expected profit using perfect information is still \$56 before the test results are known, the value of perfect information when Joe has decided to use test T_3 is \$23.13 (i.e., $\$56 - \32.87) before he has paid the mechanic, and \$10.13 (i.e., $\$56 - \45.87) after the mechanic has received his \$13.

However, after the test results have been reported, the expected profit using perfect information is different from \$56. Remember that Joe can make a profit of \$60 if he knows the car is a peach, and of \$40 if he knows it is a lemon. From our tree we see that the pair $[p(P), p(L)]$ takes on the values (0,1), (0.6,0.4), and (0.96,0.04) according to whether 2, 1, or 0 defects were discovered. Joe's expected profit using perfect information is thus \$40, \$52, or \$59.20, depending on the defect situation.

Since we have already calculated the expected values of these states to be \$40, \$28, and \$53.60 without perfect information, the EVPI's for them must be \$0, \$24, and \$5.60, respectively. As before, if we weigh these three numbers with the respective probabilities of 2, 1, or 0 defects being reported, namely, 0.067, 0.266, and 0.667, we obtain the figure of \$10.13, formerly computed as the value of perfect information at node T_3 .

An observation of particular importance may be based on these numbers: Although we would expect the amount Joe would be willing to pay the stranger for his perfect information to decrease after he is committed to a test plan, it is not necessary for this situation to obtain for any experimental outcome, but only on the average. Thus, after Joe has decided to follow test plan T_3 , he establishes that the value of perfect information to him is only \$23.13. However, if the mechanic should report that he had found exactly one defect in the car, Joe now notices that the value of perfect information has increased to \$24, a net gain of \$0.87. This means that if Joe had decided on T_3 , and the stranger's price for his information was \$23.50, Joe would refuse the information and go ahead with the test, but then willingly pay \$24 for the same information if the mechanic reports only one defect.

This result is really not too surprising when we realize that Joe had already considered the change of being placed in a situation where the expected value of perfect information is \$24 when he made his optimum decision at node T_3 . When Joe contracted for test plan T_3 he had to consider how every possible outcome of the test--2, 1, or 0 defects--would affect his state of knowledge about the type of car on the lot. If no defects were found, then Joe would be very confident that the car is a peach and would be willing to pay only \$5.60 to remove his remaining uncertainty. If two defects were found, then the car is surely a lemon and the stranger cannot tell Joe anything of value. However, if the mechanic reports one defect, then Joe does not expect to make any more money from this point into the future than he would have made if no tests whatever had been performed; \$28. It is important to note that the value of perfect information is \$24 in this situation rather than the \$28 figure applicable in the absence of tests. This difference is, of course, due to the fact that the probability that the stranger will discover that the car is good has fallen from 0.8 to 0.6. Thus, we see that although the expected value of perfect information cannot increase on an average value basis in such trees, it is possible for it to increase for some of the chance outcomes.

Now let us turn to the evaluation of test plan T_4 . Under this option the transmission is tested for \$10; when the outcome of this test is reported, it is possible to have the mechanic test the differential

for an additional cost of \$4. Such a test procedure is representative of a large class of experimental plans which we may call sequential tests. Such processes are characterized by the option to decide whether or not to continue testing after the results of the initial tests are known.

The decision tree pertinent to T_4 is shown in Figure 5. The development of this tree is once more most easily understood by considering the chronological sequence of the decisions that must be made and their outcomes. The payment of \$10 to initiate this test plan is indicated by a -10 under the branch T_4 . The next event that will occur is the report of the mechanic about whether he found a defect in the transmission. Thus, we establish a chance point that generates branches D_1 and D_1' . Regardless of whether or not a defect has been found, Joe must make a decision on the continuation of the test. His two possible actions, continue on to test the differential, and stop testing, are shown by the two branches named CONTINUE and STOP that leave decision nodes T_4D_1 and T_4D_1' . Both of the CONTINUE branches are labeled -4 to indicate the cost of requesting the testing of the differential.

If Joe decides to stop the testing program after hearing the report on the transmission, he will have to make his final decision on buying the car having only the information that either a defect was or was not found. But these two situations were also encountered under test plan T_2 after the mechanic had made his report. Since Joe finds himself in the same positions they must have the same value to him. (Remember that the money paid out for the performance of the test is a fixed cost at this point and so does not affect the future expected earnings.) Consequently, we should enter in the tree at the tips of the T_4D_1 STOP and T_4D_1' STOP branches the same values to be found at nodes T_2D_1 and T_2D_1' , respectively. We shall denote these values by $v(T_2D_1)$ and $v(T_2D_1')$; we see that $v(T_2D_1) = \$32$, $v(T_2D_1') = \$44$.

The situation if Joe decides to continue testing after hearing the mechanic's report on his first test is analogous but not identical. If the CONTINUE option is followed, the next event to take place is the report by the mechanic on whether he found a defect on his second test. Thus, we create chance points at the T_4D_1 CONTINUE and T_4D_1' CONTINUE nodes and D_2 and D_2' branches emanating from them. However, when we receive the second report from the mechanic, our total information is that in two tests 2, 1, or 0 defects have been found in the car. Thus, we are in the same positions as we were under test option T_3 after the mechanic's report was known. The appropriate value for T_4D_1 CONTINUE D_2 is, therefore, $v(T_3D_1D_2) = 40$; for T_4D_1 CONTINUE D_2' and T_4D_1' CONTINUE D_2 it is $v(T_3D_1D_2' + D_1'D_2) = 28$; and for T_4D_1' CONTINUE D_2' it is $v(T_3D_1'D_2') = 53.60$. These numbers have been placed at the pertinent tips of the T_4 test plan tree.

We have now been able to evaluate the terminal points of the T_4 tree by identifying them with nodes that had been considered earlier. It remains to place the relevant probabilities on the chance nodes in this tree so that we can proceed to make a judgment about the utility of this option. Once more we find that Nature's tree of Figure 6 supplies the probabilistic information we require. The probabilities of the branches D_1 and D_1' that leave node T_4 have already been computed in the tree for test plan T_2 ; they are 0.2 and 0.8. The only remaining probabilities are $p(D_2|D_1)$ and $p(D_2'|D_1)$ to go to the right of node T_4 D_1 CONTINUE and the probabilities $p(D_2|D_1')$ and $p(D_2'|D_1')$ to go in the analogous place on the D_1' fork. Our task is again simplified by the fact that the sum of all probabilities emerging from a chance node must be 1. From the definition of conditional probability we can write:

$$p(D_2|D_1) = p(D_1 D_2)/p(D_1)$$

and

$$p(D_2'|D_1') = p(D_1' D_2')/p(D_1')$$

From Figure 6 we find

$$\begin{aligned} p(D_2|D_1) &= \frac{p(D_1 D_2)}{p(D_1)} = \frac{p(PD_1 D_2) + p(LD_1 D_2)}{p(PD_1 D_2) + p(LD_1 D_2) + p(PD_1 D_1') + p(LD_1 D_1')} \\ &= \frac{1/15}{1/5} = 1/3 \end{aligned}$$

and

$$\begin{aligned} p(D_2'|D_1') &= \frac{p(D_1' D_2')}{p(D_1')} = \frac{p(PD_1' D_2') + p(LD_1' D_2')}{p(PD_1' D_2') + p(LD_1' D_2') + p(PD_1' D_1') + p(LD_1' D_1')} \\ &= \frac{2/15}{4/5} = 1/6 \end{aligned}$$

Of course, most of the probabilities in this calculation were computed earlier in the evaluation of test options T_2 and T_3 . However, their repetition at this time serves to emphasize the basic role of Nature's tree. Finally we have

$$p(D_2'|D_1) = 1 - p(D_2|D_1) = 2/3$$

and

$$p(D_2|D_1') = 1 - p(D_2'|D_1') = 5/6$$

When the four conditional probabilities we have just found are entered in their appropriate places in the tree for test option T_4 , we are ready to proceed with the expected value computation.

At node T_4D_1' CONTINUE there is a $1/3$ probability of the value 40 and a $2/3$ probability of the value 28. The expected value of this node is thus $1/3(40) + 2/3(28) = \$32$, as indicated in the square box. The node T_4D_1 STOP also has a value of $\$32$; however, in order to reach node T_4D_1 CONTINUE, $\$4$ must be paid and so when viewed from the left end of the T_4D_1 CONTINUE branch, this action is worth only $\$28$. Consequently, Joe is best advised to take the stop branch at this juncture and thereby make the value of decision node T_4D_1 equal to $\$32$. Such a decision has been indicated on the tree.

At node T_4D_1' CONTINUE we see a $1/6$ probability of the value 28 and a $5/6$ probability of the value 53.60. The expected value of node T_4D_1' CONTINUE is $1/6(28) + 5/6(53.60) = \$49.33$. Even after the $\$4$ expense for continuing the test has been included, this act still has an expected value of $\$45.33$, an amount slightly in excess of the $\$44$ value to be expected if branch T_4D_1' STOP is followed. The solid arrowhead and the number in the square box at node T_4D_1' correspond to this decision.

At chance node T_4 there is an 0.2 probability of the mechanic's reporting that he found a defect on the first test and thus causing us to expect a profit of $\$32$. With probability 0.8 we shall expect earnings of $\$45.33$ because he has reported no defect. Therefore, the expected value of being at decision node T_4 is $0.2(32) + 0.8(45.33) = \$42.66$. Since it is necessary to pay $\$10$ for the first test, the expected value of test plan T_4 is $\$32.66$, as shown in the square box to the left of branch T_4 .

The expected value of perfect information can be easily calculated for this test plan. All that is necessary is to copy the EVPI numbers corresponding to the value expressions at the tips of the T_3 tree. For example, the EVPI in the oval at node $T_3D_1D_2$ is 0; this figure is placed in the oval at the node T_4D_1 CONTINUE D_2 where $v(T_3D_1D_2)$ has already been copied. When this has been done for all six terminating nodes of the T_4 tree, the EVPI of all other nodes in the tree can be obtained by taking expected values of these quantities at chance nodes and taking the route indicated by the solid arrowhead at decision nodes. The solid arrowhead will always correspond to the act that minimizes the expected value of perfect information. To illustrate, at node T_4D_1 CONTINUE, the value of perfect information will be 0 if a second defect is reported and 24 if not. Weighting with the $(1/3, 2/3)$ probabilities of these events, we obtain $\$16$ at this node, or $\$20$ before the $\$4$ cost of the second test is paid. At node T_4D_1 STOP the expected value of perfect information is

\$16--therefore, the STOP alternative should be selected and the EVPI at node T_4D_1 is \$16. The reader should finish the calculation of the EVPI's in the T_4 tree to satisfy himself that the entries in Figure 5 are correct.

We have now evaluated all four test plans. From Figure 5 we can see that the expected profits from options T_1 , T_2 , T_3 , and T_4 are, respectively, \$28, \$32.60, \$32.87, and \$32.66. Since plan T_3 , that of testing two systems, has the highest expected profit, it is the one indicated by a solid arrowhead after the initial decision node. However, the evidence of the tree should be interpreted not to mean that T_3 is the best test plan, but rather that any of the plans T_2 , T_3 , T_4 will be slightly less than \$5 better than the option of no testing, on the average. The big payoff is not in the selection of a particular test plan, but rather in the decision to do some testing.

Let us review these test plans to show their operational character. If Joe does no testing, he will buy the car without a guarantee. If he follows plan T_2 , he will buy the car with the guarantee if a defect is found in the system tested and he will buy it without the guarantee if no defect is discovered. Our evaluation of plan T_3 shows that Joe should buy the car without a guarantee only if no defects are found in the two systems tested, and buy it with the guarantee otherwise.

Finally, if T_4 is chosen, Joe should stop further testing if a defect is discovered on the first test and continue testing otherwise. If a defect is found in the first test on the transmission, then Joe should buy the car with a guarantee, as we see from the decision at node T_4D_1 . However, if the transmission is not defective, then depending on whether the further test of the differential does or does not reveal a defect, Joe will either buy the car with or without the guarantee, in that order. This is determined by locating the ultimate outcomes of the T_4D_1 CONTINUE D_2 and T_4D_1 CONTINUE D_2' branches in the T_3 tree. It is interesting to note that the reason the nodes T_4D_1 CONTINUE and T_4D_1 STOP have the same values is that even if the tests were continued at this point, Joe's decision would be to buy the car with a guarantee regardless of how the second test came out. Since the test cannot affect the decision, it is not worthwhile to pay anything for the privilege of making it. The tree implies just this result.

We have now seen that after all the calculations have been performed, the final decision offers no real problem. Since test plan T_3 is most favorable by a small amount, Joe will probably decide to follow it. The expected value of perfect information is \$23.13 when plan T_3 is used; therefore, the stranger's \$25 price for this information once more looks too high. Unless the price is lowered below \$23.13, Joe should proceed

with having the fuel and electrical systems tested at a cost of \$13. He will buy the car without the guarantee only if no defects are found and with it otherwise. Joe's expected profit from this plan of action is \$32.87, an increase of \$4.87 over what he expected to make without considering testing. Of course, by this time Joe may have decided that he would rather walk than do all this calculation!

The stranger with the perfect information has witnessed a good deal of vacillation in what Joe is willing to pay him. The EVPI was \$20 initially, \$28 after the guarantee was introduced, and \$23.13 under test plan T_3 . From the stranger's point of view, the guarantee was good news, but the test options were bad news. However, even if Joe decides to follow T_3 , the stranger can still sell his knowledge to Joe by reducing its price below \$23.13. Joe will realize an increase in profit equal to the difference between \$23.13 and what he pays the stranger.

Let's suppose, however, that the stranger had stepped away by the time Joe had completed his deliberations and that when he had reappeared, Joe had already paid the mechanic the \$13 necessary to carry out test plan T_3 . Even at this point, the stranger can make some money if he considers this situation carefully. His immediate problem is that: should he offer his perfect information to Joe at a reduced price before or after Joe has received the test results from the mechanic, and what should his price be? Since Joe already has paid the mechanic, the EVPI to Joe is now \$10.13 according to the figure in the rounded box above node T_3 ; Joe will presumably pay any amount less than \$10.13 to get perfect information. Now the probabilities that the mechanic will report 2, 1, or 0 defects are $1/15$, $4/15$, and $2/3$. In fact, it is on the basis of these probabilities and the EVPI of 0, 24, and 5.60 recorded at nodes $T_3D_1D_2$, $T_3(D_1D_2' + D_1'D_2)$ and $T_3D_1'D_2'$ that Joe established the EVPI at node T_3 to be \$10.13. However, let us suppose that the stranger had determined the one piece of information that Joe does not have; namely, he has found out whether or not the car is a lemon simply by observing the serial number. Using this information, the stranger can calculate new probabilities of the various reports of the mechanic according to whether the car is a peach or a lemon. He will thus obtain an expected value of EVPI after the report is known that will be different from Joe's estimate of \$10.13. If the stranger's estimate is higher than Joe's, he will do better in his expected value by not offering his perfect information until the outcome of the test is known. On the other hand, if the stranger's estimate is lower than Joe's, he should offer his information immediately.

The calculations involved are quite straightforward. If the stranger determines that the car is a peach, then the three probabilities that should be used to weigh the numbers 0, 24, and 5.60 should be $p(D_1D_2 | P)$,

$p(D_1 D_2' + D_1' D_2 | P)$, and $p(D_1' D_2' | P)$. If the car is found to be a lemon, then the appropriate probabilities are $p(D_1 D_2 | L)$, $p(D_1 D_2' + D_1' D_2 | L)$, and $p(D_1' D_2' | L)$. These probabilities are computed from Nature's tree of Figure 6 as follows:

$$p(D_1 D_2 | P) = \frac{p(PD_1 D_2)}{p(P)} = \frac{0}{0.8} = 0$$

$$p(D_1 D_2' + D_1' D_2 | P) = \frac{p(PD_1 D_2') + p(PD_1' D_2)}{p(P)} = \frac{0.08 + 0.08}{0.8} = 1/5$$

$$p(D_1' D_2' | P) = \frac{p(PD_1' D_2')}{p(P)} = \frac{0.64}{0.8} = 4/5$$

$$p(D_1 D_2 | L) = \frac{p(LD_1 D_2)}{p(L)} = \frac{5/75}{0.2} = 1/3$$

$$p(D_1 D_2' + D_1' D_2 | L) = \frac{p(LD_1 D_2') + p(LD_1' D_2)}{p(L)} = \frac{8/75}{0.2} = 8/15$$

$$p(D_1' D_2' | L) = \frac{p(LD_1' D_2')}{p(L)} = \frac{2/75}{0.2} = 2/15$$

The expected value of the expected value of perfect information that will exist after the results of the test are known is computed for the states of knowledge of Joe, of the stranger when the car is a peach, and of the stranger when the car is a lemon in Table II.

The important thing to note is that the stranger expects the EVPI to be only \$9.28 after the results of the experiment are known if the car is a peach, but \$13.55 if the car is a lemon. In other words, considering that the EVPI of perfect information is \$10.13 in Joe's eyes, the stranger expects the EVPI to be lower than Joe's when the results are reported if the car is a peach and higher if it is a lemon. It is, therefore, prudent for the stranger to sell the perfect information to Joe before the mechanic calls for, say, \$10 if the car is a peach, but to wait until after the mechanic's report before offering it if the car is a lemon. Since $p(D_1 D_2) = p(D_1 D_2 | P) \times p(P) + p(D_1 D_2 | L) \times p(L)$, etc., and since $p(P) = 0.8$, $p(L) = 0.2$, $10.13 = 0.8(9.28) + 0.2(13.55)$. That is, the expectation of the EVPI from Joe's point of view is the expected value of the EVPI from the stranger's

Table II

The expected value of what the expected value of perfect information will be when the results of test T_3 are known.

EVPI of Report	Probabilities of the Report as Seen by			
	Joe	The stranger when the car is a peach	The stranger when the car is a lemon	
Two defects, $D_1 D_2$	0	$p(D_1 D_2) = 1/15$	$p(D_1 D_2 P) = 0$	$p(D_1 D_2 L) = 1/3$
One defect, $D_1 D'_2 + D'_1 D_2$	24	$p(D_1 D'_2 + D'_1 D_2) = 4/15$	$p(D_1 D'_2 + D'_1 D_2 P) = 1/5$	$p(D_1 D'_2 + D'_1 D_2 L) = 8/15$
No defects, $D'_1 D'_2$	5.60	$p(D'_1 D'_2) = 2/3$	$p(D'_1 D'_2 P) = 4/15$	$p(D'_1 D'_2 L) = 2/15$
EVPIs weighed with probabilities		10.13	9.28	13.55

point of view. This computation provides the essential reconciliation between the viewpoints of the buyer and seller of perfect information.

However, to show that our problem still has hidden facets, we suddenly realize that a competitive-game aspect has appeared. If the stranger offers his information before the test results are known, and if Joe knows that the stranger has reasoned according to the previous paragraph, then Joe is certain that the car is a peach. Similarly, if the offer is made after the test, Joe is certain that the car is a lemon. In either case, Joe will have received perfect information without paying for it. This forces the stranger to randomize his strategy, and so on and on and on. We shall give up trying to help Joe at this point.

Well, at last Joe is driving away in his Spartan, having used test plan T_3 and abided by the results. A most human question is: Did he make a good decision or didn't he? The answer to this question does not depend at all on whether his new car is actually a peach or a lemon. We must make a distinction between a good decision and a good outcome. Joe made a good decision because he based it on logic and his available knowledge. Whether or not the outcome is good depends on the vagaries of chance.

INFLUENCE DIAGRAMS

Ronald A. Howard
James E. Matheson
Strategic Decisions Group

1. Introduction

The rapid growth of electronic computation continues to challenge our ability to conceptualize and describe the world around us. Mathematical tools and formal descriptions serve poorly as a communication device with the majority of people not trained in nor used to mathematical means of expression. Yet virtually everyone has information useful in the solution of his own problems or the problems of others if only it could be tapped.

The subject of this paper is a new form of description, the influence diagram, that is at once both a formal description of the problem that can be treated by computers and a representation easily understood by people in all walks of life and degrees of technical proficiency. It thus forms a bridge between qualitative description and quantitative specification.

The reason for the power of this representation is that it can serve at the three levels of specification of relation, function, and number, and in both deterministic and probabilistic cases. In the deterministic case, relation means that one variable can depend in a general way on several others; for example, profit is a function of revenue and cost. At the level of function we specify the relationship; namely, that profit equals revenue minus cost. Finally, at the level of number, we specify the numerical values of revenue and cost and hence determine the numerical value of profit.

In the probabilistic case, at the level of relation we mean that given the information available, one variable is probabilistically dependent on certain other variables and probabilistically independent of still other variables. At the level of function, the probability distribution of each variable is assigned conditioned on values of the variables on which it depends. Finally, at the level of number, unconditional distributions are assigned on all variables that do not depend on any other variable and hence determine all joint and marginal probability distributions.

As an example of the probabilistic case, we might assert at the level of relation that income depends on age and education and that education depends on age. Next, at the level of function we would assign the conditional distribution of income given age and education and the distribution of education given age. Finally, at the level of number, we would assign the unconditional distribution on age.

The successive degrees of specification can be made by different individuals. Thus, an executive may know that sales depend in some way on price, but he may leave to others the probabilistic description of the relationship.

Because of its generality, the influence diagram is an important tool not only for decision analysis, but for any formal description of relationship and thus for all modeling work.

In the present paper, we shall focus on the probabilistic use of influence diagrams since the deterministic use is a special, but important, case of the probabilistic. We now proceed to development of the influence diagram concept, to examination of its properties, and to illustration of its use.

2. Probabilistic Dependence*

One of the most perplexing aspects of making decisions under uncertainty is the problem of representing and encoding probabilistic dependencies. A probabilistic dependency is one that arises as a result of uncertainty. For example, if a and b are known variables and $c = a + b$, then it is clear that c depends on both a and b , both in a vernacular sense and in a mathematical sense. However, suppose a is known and b is uncertain. Then c is probabilistically dependent on b but not on a . The reason is that knowing the specific value of b tells us something new about c , but there is no such possibility with respect to a .

3. Probabilistic Independence

Probabilistic independence, like the assigning of probability itself, depends on the state of information possessed by the assessor. Let x , y , and z be aleatory state variables of interest, which can be either continuous or discrete. Then $\{x|S\}$ is the probability distribution assigned to x given the state of information S . Two variables x and y are probabilistically independent given the state of information S if

$$\{x,y|S\} = \{x|S\} \{y|S\}$$

or equivalently, if $\{x|y,S\} = \{x|S\}$.

4. Expansion

Regardless of whether x and y are probabilistically independent, we can write

$$\begin{aligned} \{x,y|S\} &= \{x|y,S\} \{y|S\} \\ &= \{y|x,S\} \{x|S\} \end{aligned}$$

*This entire discussion applies as well to events as to variables.

We call this the "chain rule of probabilities". Note that for three events there are six possible representations:

$$\begin{aligned}
 \{x,y,z|S\} &= \{x|y,z,S\} \{y|z,S\} \{z|S\} \\
 &= \{x|y,z,S\} \{z|y,S\} \{y|S\} \\
 &= \{y|x,z,S\} \{x|z,S\} \{z|S\} \\
 &= \{y|x,z,S\} \{z|x,S\} \{x|S\} \\
 &= \{z|x,y,S\} \{x|y,S\} \{y|S\} \\
 &= \{z|x,y,S\} \{y|x,S\} \{x|S\}
 \end{aligned}$$

For n variables there are $n!$ possible expansions, each requiring the assignment of a different set of probabilities and each logically equivalent to the rest. However, while the assessments are logically equivalent there may be considerable differences in the ease with which the decision maker can provide them. Thus the question of which expansion to use in a problem is far from trivial.

5. Probability Trees

Associated with each expansion is a probability tree. The expansion

$$\{x,y,z|S\} = \{x|y,z,S\} \{y|z,S\} \{z|S\}$$

implies the tree shown in Figure 5.1. The tree is a succession of nodes with branches emanating from each node to represent different possible values of a variable. The first assignment made is the probability of various values of z . The probability of each value of y is assigned conditioned on a particular value of z , and placed on the portion of the tree indicated by that value. Finally, the probabilities of various levels of x are assessed given particular values of z and y and placed on the portion of the tree specified by those values. When this has been done for all possible values of x , y , and z the tree is complete. The probability of any particular path through the tree is obtained by multiplying the values along the branches and is $\{x,y,z|S\}$. Notice that the tree convention uses small circles to represent chance nodes. If we wish to focus on the succession in the tree rather than the detailed connections, we can draw the tree in the generic form shown in Figure 5.2.

6. Decision Trees

If a variable is controlled by a decision maker, it is represented in a tree by a decision node. Thus if y were a decision variable, Figure 5.2 could be redrawn as Figure 6.1. This tree states that the decision maker is initially uncertain about z and has assigned a probability distribution




FIGURE 5.1 A PROBABILITY TREE




FIGURE 5.2 A GENERIC PROBABILITY TREE

$\{z|S\}$ to it. However, he will know z at the time he must set y , the decision variable. This node is represented, like all decision nodes, by a small square box. Once z and y are given, the decision maker will still be uncertain about x ; he has represented this uncertainty by $\{x|y,z,S\}$. Notice that a decision tree implies both a particular expansion of the probability assessments and a statement of the information available when a decision is made.




FIGURE 6.1 A GENERIC DECISION TREE

7. Probability Assignment for Decision Trees

The major problem with decision trees arises from the first of these characteristics. The order of expansion required by the decision tree is rarely the natural order in which to assess the decision maker's information. The decision tree order is the simplest form for assessment only when each variable is probabilistically dependent on all preceding aleatory and decision variables. If, as is usually the case, many independence assertions can be made, assessments are best done in a different order from that used in the decision tree. This means that we first draw a probability tree in an expansion form convenient to the decision maker and have him use this tree for assignment; it is called a probability assignment tree. Later the information is processed into the form required by the decision tree by representing it in one of the alternative expansion orders. This is often called "using Bayes's Rule" or "flipping the tree." It is a fundamental operation permitted by the arbitrariness in the expansion order.

Consider, for example, the decision tree of Figure 6.1 with one additional aleatory variable v added, as shown in Figure 7.1. We interpret z as a test result that will become known, y as our decision, x as the outcome variable to which the test is relevant, and v as the value we shall




FIGURE 7.1 A FOUR NODE DECISION TREE

receive if the test indicates z , we decide y , and x is the value of the outcome variable. Often y will not affect x in any way, even though y affects v ; we write

$$\{x|y, z, S\} = \{x|z, S\}$$

to represent this assertion.

With this independence assertion we have the tree shown in Figure 7.2. This tree requires the specification of $\{z|S\}$ and $\{x|z, S\}$: the probability of various test results and the probability of various outcomes given test results. But typically in situations of this kind, the decision maker would prefer to assign directly the probabilities of different outcomes $\{x|S\}$ and then the probabilities of various test results given the outcome, $\{z|x, S\}$. In other words, he would prefer to make his assessments in the probability tree of Figure 7.3 and then have them processed to fit the decision tree of Figure 7.2. Since

$$\{x|S\} \{z|x, S\} = \{z|S\} \{x|z, S\} = \{x, z|S\}$$

this is no more than choosing one expansion over the other. The exact processing required for the decision tree is then summation,

$$\{z|S\} = \int_x \{z|x, S\} \{x|S\}$$

and division,

$$\{x|z, S\} = \frac{\{z|x, S\} \{x|S\}}{\{z|S\}}$$




FIGURE 7.2 A FOUR NODE DECISION TREE GIVEN THE ASSERTION THAT y WILL NOT AFFECT x




FIGURE 7.3 THE PROBABILITY ASSIGNMENT TREE

Recall, however, that this whole procedure was possible only because variable x did not depend on the decision variable y .

8. Influence Diagrams

An influence diagram is a way of describing the dependencies among aleatory variables and decisions. An influence diagram can be used to visualize the probabilistic dependencies in a decision analysis and to specify the states of information for which independencies can be assumed to exist.

Figure 8.1 shows how influence diagrams represent the dependencies among aleatory variables and decisions. An aleatory variable is represented by a circle containing its name or number. An arrow pointing from aleatory variable A to aleatory variable B means that the outcome of A can influence the probabilities associated with B . An arrow pointing to a decision from either another decision or an aleatory variable means that the decision is made with the knowledge of the outcome of the other decision or aleatory variable. A connected set of squares and circles is called an influence diagram because it shows how aleatory variables and decisions influence each other.

The influence diagram in Figure 8.2 (a) states that the probability distribution assigned to x may depend on the value of y , whereas the influence diagram in Figure 8.2 (b) asserts that x and y are probabilistically independent for the state of information with which the diagram was drawn. Note that the diagram of Figure 8.2 (a) really makes no assertion about the probabilistic relationship of x and y since, as we know, any joint probability $\{x,y|S\}$ can be represented in the form

$$\{x,y|S\} = \{x|y,S\} \{y|S\}.$$

However, because

$$\{x,y|S\} = \{y|x,S\} \{x|S\},$$

the influence diagram of Figure 8.2 (a) can be redrawn as shown in Figure 8.2 (c); both are completely general representations requiring no independence assertions. While the direction of the arrow is irrelevant for this simple example, it is used in more complicated problems to specify the states of information upon which independence assertions are made.

Similarly, with three variables x , y , z there are six possible influence diagrams of complete generality, one corresponding to each of the possible expansions we developed earlier. They are shown in Figure 8.3. While all of these representations are logically equivalent, they again




THE PROBABILITIES ASSOCIATED WITH ALEATORY VARIABLE B DEPEND ON THE OUTCOME OF ALEATORY VARIABLE A



THE PROBABILITY OF ALEATORY VARIABLE D DEPENDS ON DECISION C




THE DECISION MAKER KNOWS THE OUTCOME OF ALEATORY VARIABLE E WHEN DECISION F IS MADE



THE DECISION MAKER KNOWS DECISION G WHEN DECISION H IS MADE


FIGURE 8.1 DEFINITIONS USED IN INFLUENCE DIAGRAMS



(a) A SIMPLE INFLUENCE DIAGRAM



(b) AN EVEN SIMPLER INFLUENCE DIAGRAM



(c) AN ALTERNATE REPRESENTATION

FIGURE 8.2 TWO-NODE INFLUENCE DIAGRAMS




FIGURE 8.3 ALTERNATE INFLUENCE DIAGRAMS FOR $\{x, y, z\}$

differ in their suitability for assessment purposes. In large decision problems, the influence diagrams can display the needed assessments in a very useful way.

9. Graphical Manipulation


Since there are many alternative representations of an influence diagram, we might ask what manipulations can be performed on an influence diagram to change it into another form that is logically equivalent.

The first observation we should make is that an arrow can always be added between two nodes without making an additional assertion about the independence of the two corresponding variables (as long as no loops are created). That is, saying that x may depend on variable y is not equivalent to saying that x must depend on y. Thus the diagram of Figure 8.2 (b) can be changed into either of the diagrams shown in Figures 8.2 (a) and 8.2 (c) without making an erroneous assertion. However, the reverse procedure could lead to an erroneous assertion. Creating additional influence arrows will not change any probability assessment, but may destroy explicit recognition of independencies in the influence diagram.


Thus, Figures 8.2 (a) and 8.2 (c) are two equivalent influence diagrams. They are equivalent in that they imply the same possibility of dependencies between x and y given the state of information on which the diagram was based.

An arrow joining two nodes in an influence diagram may be reversed provided that all probability assignments are based on the same set of information. For example, consider the influence diagram of Figure 9.1 (a). Since the probability assignment to both x and y are made given knowledge of z the arrow joining them can be reversed as shown in Figure 9.1 (b) without making any incorrect or additional assertions about the possible independence of x and y. Figure 9.1 (c) shows another example where the assignment of probability to x does not depend on the value of z, and so it might appear that no reversal of the arrow from x to y is possible. However, recall that we can always add an arrow to a diagram without making an incorrect assertion. Thus we can change the diagram of Figure 9.1 (c) to that of Figure 9.1 (a), and then that of Figure 9.1 (a) to that of Figure 9.1 (b). The influence arrow between x and y can be reversed after an influence arrow is inserted between z and x.


The graphical manipulation procedure may yield more than one result. For example, consider the reversal of the three-node influence diagram shown in Step 1 of Figure 9.2 (a). Suppose we first attempt to reverse the y to x arrow. For x and y to have only common influences, we must provide x with an influence from z (Step 2), before performing the reversal (Step 3). Since both x and z now are based on the same state of information (there are no impinging influences, i.e. arrows into x or z from any other node), the influence joining them may be reversed (Step 4). Finally, since both z and y are assigned probabilities after x is known, the influence joining them can be reversed (Step 5).



(a) AN INFLUENCE DIAGRAM

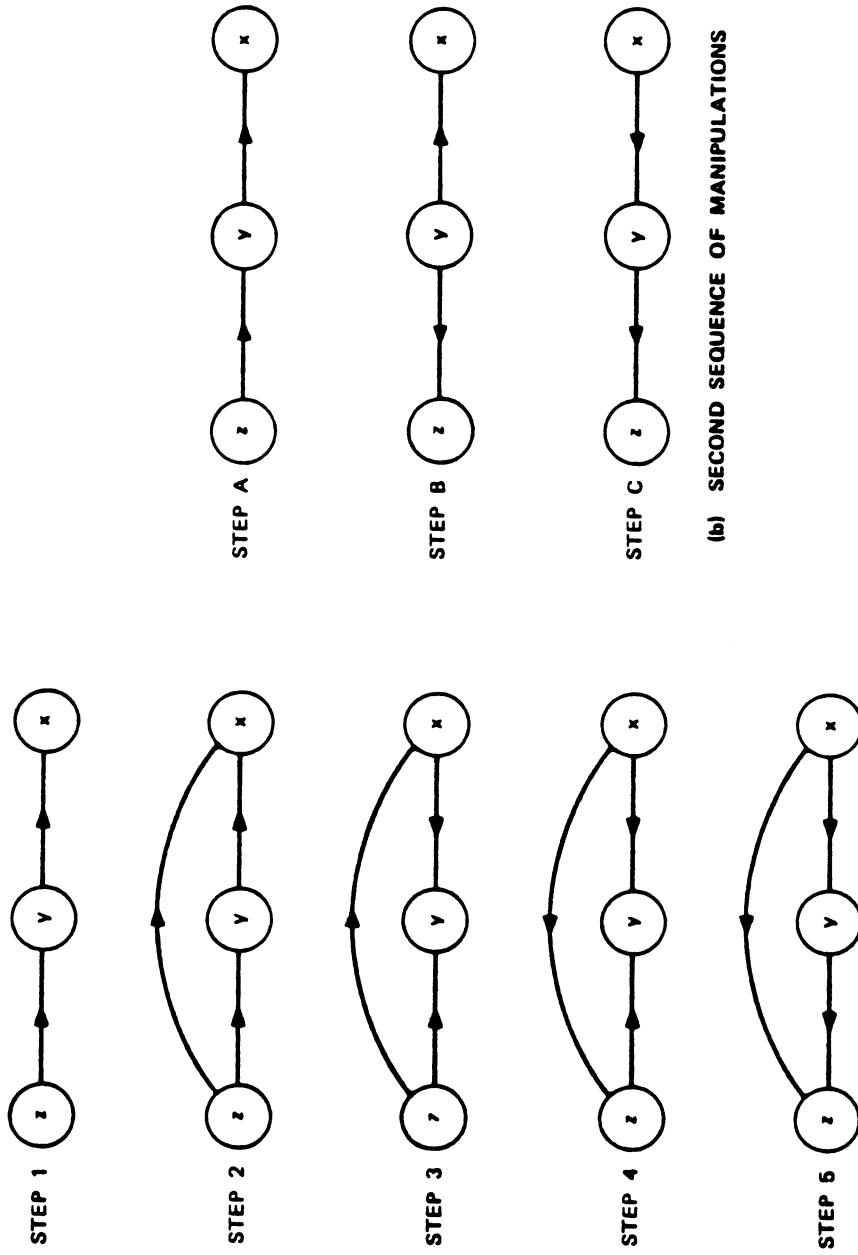


(b) ARROW BETWEEN x AND y REVERSED



(c) ANOTHER INFLUENCE DIAGRAM

FIGURE 9.1 GRAPHICAL MANIPULATION OF INFLUENCE DIAGRAMS



(a) FIRST SEQUENCE OF MANIPULATIONS

(b) SECOND SEQUENCE OF MANIPULATIONS

FIGURE 9.2 GRAPHICAL MANIPULATIONS PRODUCING NON-UNIQUE RESULTS

Suppose, however, that the same diagram (Step A, Figure 9.2 (b)) were transformed by first reversing the arrow joining z and y (Step B), which is possible since y and z are based on the same state of information (i.e. there are no impinging influences). Then the arrow joining x and y can be reversed (Step C) because neither x nor y now have impinging influences. Both this transformation and the one in Figure 9.2 (a) are correct. However, Step C of Figure 9.2 (b) shows that there is no need to indicate conditioning of z on x. Step 5 of Figure 9.2 (a) contains this unnecessary but not incorrect influence.

10. Influence Diagrams with Decision Variables

We shall now extend the concept of influence diagrams to include decision variables. We begin with a formal definition of influence diagrams.

An influence diagram is a directed graph having no loops. It contains two types of nodes:

- Decision nodes represented by boxes (\square)
- Chance nodes represented by circles (\circ)

Arrows between node pairs indicate influences of two types:

- Informational influences represented by arrows leading into a decision node. These show exactly which variables will be known by the decision maker at the time that the decision is made.
- Conditioning influences, represented by arrows leading into a chance node. These show the variables on which the probability assignment to the chance node variable will be conditioned.

The informational influences on a decision node represent a basic cause/effect ordering whereas the conditional influences into a chance node represent, as we have seen, a somewhat arbitrary order of conditioning that may not correspond to any cause/effect notion and that may be changed by application of the laws of probability (e.g. Bayes's Rule).

Figure 10.1 is an example of an influence diagram. Chance node variables a, b, c, e, f, g, h, i, j, k, l, m, and o all indicate aleatory variables whose probabilities must be assigned given their respective conditioning influences. Decision node variables d and n represent decision variables that must be set as a function of their respective informational influences. For example, the probability assignment to variable i is conditioned upon variables f, g, and l, and only these variables. In inferential notation, this assignment is

$$\{i | f, g, l, E\},$$




FIGURE 10.1 AN INFLUENCE DIAGRAM WITH DECISION NODES

where E represents a special S, the initial state of information upon which the construction of the entire diagram is based. As another example, the decision variable d is set with knowledge of variables a and c, and only these variables. Thus, d is a function of a and c.

11. Node Terminology

One of the most important, but most subtle, aspects of an influence diagram is the set of possible additional influences that are not shown on the diagram. An influence diagram asserts that these missing influences do not exist.

To illustrate this characteristic of influence diagrams more clearly we must make a few more definitions.

- A path from one node to another node is a set of influence arrows connected head to tail that forms a directed line from one node to another.

With respect to any given node we make the following definitions:

- The predecessor set of a node is the set of all nodes having a path leading to the given node.
- The direct predecessor set of a node is the set of nodes having an influence arrow connected directly to the given node.
- The indirect predecessor set of a node is the set formed by removing from its predecessor set all elements of its direct predecessor set.
- The successor set of a node is the set of all nodes having a path leading from the given node.
- The direct successor set of a node is the set of nodes having an influence arrow connected directly from the given node.
- The indirect successor set of a node is the set formed by removing from its successor set all elements of its direct successor set.

We refer to members of these sets as predecessors, direct predecessors, indirect predecessors, successors, direct successors, and indirect successors. Figure 11.1 shows the composition of each of these sets in relation to node g.

12. Missing Influences

We now are prepared to investigate the implications of influences not shown in a diagram. A given node could not have any arrows coming into it

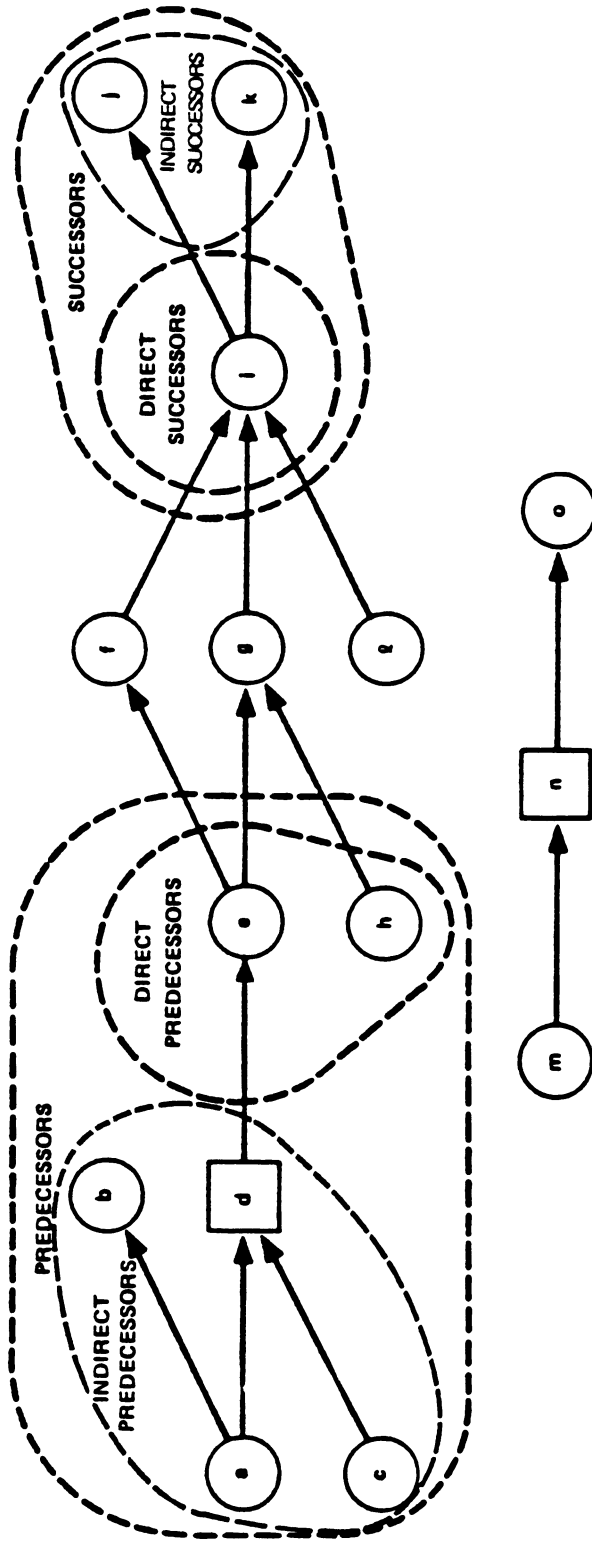


FIGURE 11.1 SOME SETS DEFINED BY THE NODE g

from successor nodes because this addition would form a loop in the diagram. A loop is prohibited since it could not represent any possible expansion order. However, the given node could conceivably have an additional arrow coming from any predecessor node.

The situation for decision nodes is relatively simple. The diagram asserts that the only information available when any decision is made is that represented by the direct predecessors of the decision. The addition of a new arrow, or informational influence, would usually add to the information available for decision making, and destroy the original logic of the diagram. The influence diagram asserts that this information is not directly available; however, all or part of it might be inferred indirectly from the direct predecessor set.

The situation for chance nodes is more complex. The diagram partially constrains the probabilistic conditioning (expansion) order for chance nodes. In general, the probability assignment for a given chance node, x , might be conditioned on all non-successors (except for x itself). Let us call this set N_x , and let D_x be the set of direct predecessors of x . The set D_x is, of course, contained in N_x . The diagram asserts that the probability assignment to x given N_x is the same as to x given D_x ; that is,

$$\{x|N_x, E\} = \{x|D_x, E\}$$

The addition of a new arrow or conditioning influence from an element of N_x to x would increase the set of direct predecessors and seem to increase the dimensionality of the conditional probability assignment. While this addition would not violate the logic of the diagram, it would cause a loss of information regarding independence of the added conditioning influence. The original diagram asserts that all information in the set N_x that is relevant to the probability assignment to x is summarized by the direct predecessors D_x . In classical terms, with respect to x , D_x is a sufficient statistic for N_x .

Returning to Figure 10.1 as an example, the probability assignment to variable g is in principle conditioned on all variables except g , i , j , and k . However, the diagram asserts that the variables on which g depends are sufficiently summarized by only e and h . This means

$$\{g|a,b,c,d,e,f,h,l,m,n,o,E\} = \{g|e,h,E\}$$

This strong and useful assertion is based as much on the lack of arrows as on the ones that are present.

We have seen that an influence diagram indicates a specific, but possible non-unique order for conditioning probability assignments as well as the information available as the basis for each decision. When decision rules are specified for each decision node and probability assignments are made for each chance node, the influence diagram relationships can be used to develop the joint probability distribution for all variables.

13. Relationship of Influence Diagrams to Decision Trees

Some influence diagrams do not have corresponding decision trees. As in a decision tree, all probability assignments in an influence diagram--including the assignment limitations represented by its structure--must be founded on a base state of information, E. Unlike the nodes in a decision tree, the nodes in an influence diagram do not have to be totally ordered nor do they have to depend directly on all predecessors. The freedom from total ordering allows convenient probabilistic assessment and computation. The freedom from dependence on all predecessors allows the possibility of decisions in the diagram being made by decision makers who agree on the common base state of information E, but who differ in their ability to observe certain variables in the diagram. If the diagram represents a single decision maker who does not forget information, then the direct predecessor set of one decision must be a subset of the direct predecessor set of any subsequent decision. In the influence diagram of Figure 11.1, decisions d and n have mutually exclusive direct predecessor sets, (a,c) and (m). This situation could not be represented by a conventional decision tree.

If the informational arrows shown as dashed lines in Figure 13.1 are added to Figure 11.1, then the influence diagram can be represented by a decision tree. Many different valid decision trees can be constructed from this new influence diagram. The only conditions are that they must (1) preserve the ordering of the influence diagram and (2) not allow a chance node to be a predecessor of a decision node for which it is not a direct predecessor. For example, the chance node m must not appear ahead of decision node d in a decision tree because this would imply that the decision rule for d could depend on m, which is not the case.

The situation becomes more complex when we add a node such as p in Figure 13.2. If we were to construct a decision tree beginning with chance node p it would imply that the decision rules at nodes d and n could depend on p, which is not the actual case. Node p represents a variable that is used in the probability assignment model but that is not observable by the decision maker at the time that he makes his decisions. In this situation, we would normally use the laws of probability (e.g. Bayes's Rule) to eliminate the conditioning of c on p. This process would lead to a new influence diagram reflecting a change in the sequence of conditioning, and could result in the inclusion of additional influences.

In Figure 13.3, the dashed arrow represents the influence as "turned around" by Bayes's Rule. The resulting diagram can be developed into a decision tree without further processing of probabilities. Also note that the change in the influence diagram required only information already specified by the original influence diagram (Figure 13.2) and its associated numerical probability assignments. Thus it can be carried out by a routine procedure.

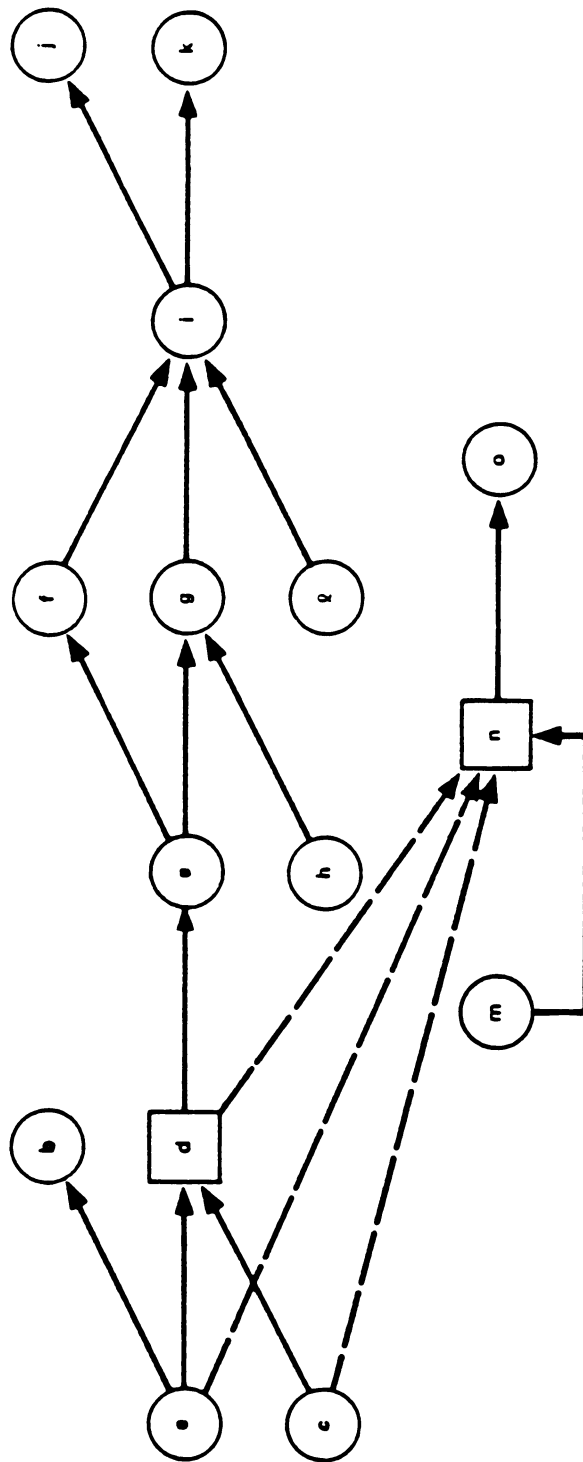


FIGURE 13.1 AN INFLUENCE DIAGRAM REPRESENTABLE BY A DECISION TREE




FIGURE 13.2 INFLUENCE DIAGRAM REQUIRING PROBABILISTIC MANIPULATION BEFORE DECISION TREE CONSTRUCTION




FIGURE 13.3 INFLUENCE DIAGRAM READY FOR DEVELOPMENT INTO A DECISION TREE

The forgoing considerations motivate two new definitions.

- A decision network is an influence diagram:
 - (i) that implies a total ordering among decision nodes,
 - (ii) where each decision node and its direct predecessors directly influence all successor decision nodes.
- A decision tree network is a decision network:
 - (iii) where all predecessors of each decision node are direct predecessors.


Requirement (i) is the "single decision maker" condition and requirement (ii) is the "no forgetting" condition. These two conditions guarantee that a decision tree can be constructed, possibly after some probabilistic processing. Requirement (iii) assures that no probabilistic processing is needed so that a decision tree can be constructed in direct correspondence with the influence diagram.

As an example consider the standard inferential decision problem represented by the decision network of Figure 13.4 (a). This influence diagram cannot be used to generate a decision tree directly because the decision node c has a non-direct predecessor that represents an unobservable chance variable. To convert this decision network to a suitable decision tree network we simply reverse the arrow from a to b, which is permissible because they have only common predecessors, namely none. We thus achieve the decision tree network of Figure 13.4 (b), and with redrawing we arrive at Figure 13.4 (c).


Specifying the limitations on possible conditioning by drawing the influence diagram may be the most significant step in probability assignment. The remaining task is to specify the numerical probability of each chance node variable conditioned on its direct predecessor variable by a probability assessment procedure.

14. Example: The Used Car Buyer


As an illustration of the use of influence diagrams, we consider a problem known as "The Used Car Buyer"² presented elsewhere in detail. For our purposes, we need only specify that the buyer of a used car can select among various tests T at different costs, observe their results R, choose a purchase alternative A, and then receive some value V that depends on the state of the car he bought, the outcome O. Figure 14 (a) shows the influence diagram. The arrows show that the test results R depend on the test selected T and the state of the car O. The buying alternative A is chosen knowing the test selected T and its results R. The value V depends on the buying alternative chosen A, on the test selected T (as a result of the cost of the test), on the outcome O, and on the test results R. This last influence allows for the possibility that the value may depend



(a)



(b)



(c)

FIGURE 13.4 THE PROCESS OF CONVERTING A DECISION NETWORK TO A DECISION TREE NETWORK







FIGURE 14 THE INFLUENCE DIAGRAM FOR THE USED CAR BUYER CONVERTED TO A GENERIC DECISION TREE

directly on the results of the test; for example, if the testing is destructive. The outcome O does not depend on any other variable, and in particular, not on the test T,

$$\{O|E\} = \{O|T,E\}$$

This assumption is based on the belief that the seller of the car will not switch the car to be tested as a result of the test selected.

This influence diagram is a decision network, but not a decision tree network because node O is a predecessor of node A, but not a direct predecessor. To create a decision tree network, we must reverse the arrow connecting node O to node R. The first step in this reversal is to assure that these nodes have a common information state. We accomplish this by adding an influence from node T to node O as shown in Figure 14 (b). Then we reverse the arrow from node O to node R and redraw the diagram as a decision tree network in Figure 14 (c).

This reversal means, of course, that the original probability assessments $\{R|T,O,E\}$ and $\{O|E\} = \{O|T,E\}$ must be changed to the probability distributions $\{R|T,E\}$ and $\{O|T,R,E\}$ according to the equation

$$\{R|T,E\} = \int_0^1 \{R|T,O,E\} \{O|T,E\}$$

and Bayes's equation

$$\{O|T,R,E\} = \frac{\{R|T,O,E\} \{O|T,E\}}{\{R|T,E\}}$$

The resulting generic decision tree appears in Figure 14 (d), where the value assigned to each path through the tree, $\langle V|T,R,A,O,E \rangle$, is recorded at the endpoint of the path. The detailed calculations are shown in Reference 1.

15. Toxic Chemical Testing Example

To illustrate the power of influence diagrams to solve complex problems of decision-making and information acquisition, we shall apply this method to a problem of toxic chemical testing. We shall carry out the analysis under the assumption that an automated influence diagram system is available to provide the flavor of its use.

Let us suppose that a chemical having some benefits also is suspected of possible carcinogenicity. We wish to determine whether to ban, restrict, or permit its use, and also whether to undertake any information gathering regarding cancer-producing activity of the chemical or its degree of exposure to humans.

The primary decision problem can be formulated by drawing the influence diagram on the input screen of the system as in Figure 15.1. This

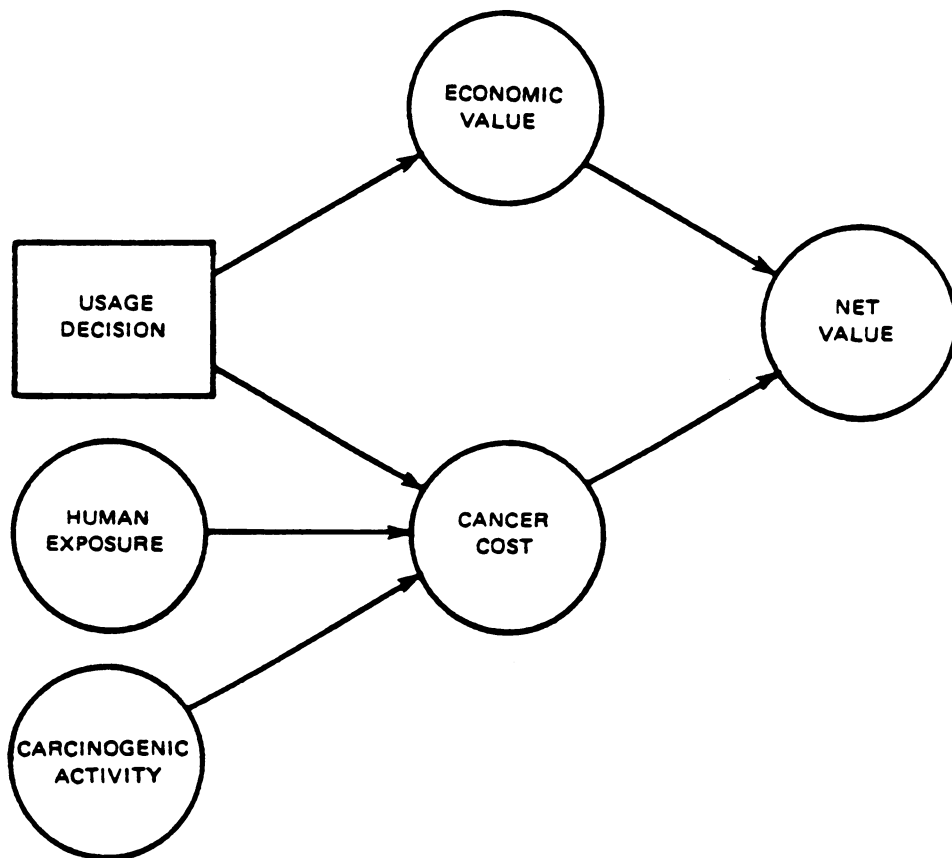


FIGURE 15.1 INFLUENCE DIAGRAM FOR PRIMARY DECISION

figure shows that the system has been told that the economic value of the product and the cancer cost attributed to it both depend on the decision regarding usage of the chemical. The (probability assignment on) economic value given the usage decision is independent of the human exposure, carcinogenic activity and the cancer cost. However, the cancer cost is dependent upon the usage decision as well as on both the carcinogenic activity and human exposure levels of the chemical. The net value of the chemical given the economic value and the cancer cost is independent of the other variables. Also, human exposure and carcinogenic activity are independent.

These relationships are not necessarily obvious ones; they depend on knowledge of the problem at hand. For example, the economic value of a particular chemical might well depend on its chemical activity which in turn might be closely related to its carcinogenic activity. In such a case an arrow might have to be added from "carcinogenic activity" to "economic value".

The next step is to obtain probability and value assessments corresponding to the influence diagram. The automated influence diagram system asks for a list of usage decision alternatives. In this case they are BAN, RESTRICT, and PERMIT. Next it asks for the economic value given each of these alternatives. In this case the permit alternative is considered to have a reference value of zero, the restrict alternative a substitute process cost of \$1 million, and the ban alternative a substitute process cost of \$5 million.

The next request is to assess possible outcomes for human exposure and carcinogenic activity along with their corresponding (unconditional) probabilities. The probability trees of Figure 15.2 illustrate these assignments. Then we are asked for the cancer cost given human exposure and carcinogenic activity levels as well as the usage decision. We assess the expected values of this cost as given in Table 15.1. Finally we state that the net value is simply the sum of the economic value and cancer cost.

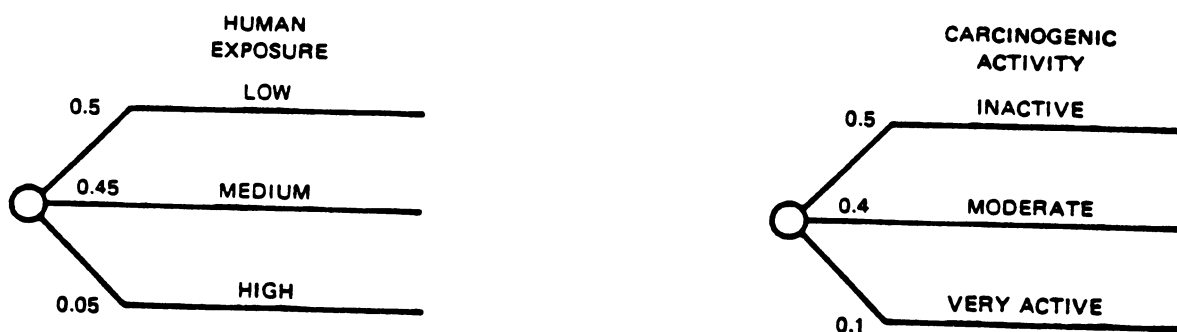


FIGURE 15.2 INITIAL PROBABILITY ASSIGNMENTS

Table 15.1

CANCER COST (\$MILLIONS)

<u>Activity</u>	<u>PERMIT ALTERNATIVE</u>			<u>RESTRICT ALTERNATIVE</u>			<u>BAN ALTERNATIVE</u>		
	<u>Exposure</u>			<u>Exposure</u>			<u>Exposure</u>		
	<u>Low</u>	<u>Med</u>	<u>High</u>	<u>Low</u>	<u>Med</u>	<u>High</u>	<u>Low</u>	<u>Med</u>	<u>High</u>
Inactive	0	0	0	0	0	0	0	0	0
Moderate	0.5	5	50	0.05	0.5	5	0	0	0
Very Active	10	100	1000	1	10	100	0	0	0

All of this information is based on detailed modeling and expert judgment regarding the decision situation. Once it has been captured with the influence diagram, analysis can proceed. The automated influence diagram procedure generates the appropriate decision tree, displays it if desired by the user, and determines that the best decision is to restrict usage. The expected value given this decision is a cost of \$2.2 million. An example display containing this information is shown in Figure 15.3. In the example we consider only the expected value or risk-neutral case although the case of risk aversion can be treated without difficulty.

15.1 The Value of Clairvoyance (Perfect Information)

Before investigating actual information gathering alternatives, the usual decision analysis practice is to determine the value of clairvoyance (perfect information) on the uncertain variables. The value of clairvoyance furnishes an upper limit on the value of real information gathering.

With the automatic influence diagram procedure these calculations are trivial. For example, to calculate the value of the problem with clairvoyance on carcinogenic activity we need only add the influence arrow indicated by a dotted line in Figure 15.1.1. This modification states that the decision maker knows the degree of carcinogenic activity when he makes the usage decision. The result is an expected cost of \$1.1 million and a decision rule to permit if inactive, restrict if moderate, and ban if very active. This means that the expected value of clairvoyance is the original \$2.2 million minus this \$1.1 million which is \$1.1 million. Figure 15.1.2 shows a more complete display of the decision tree for this case that would be automatically generated upon request of the user.

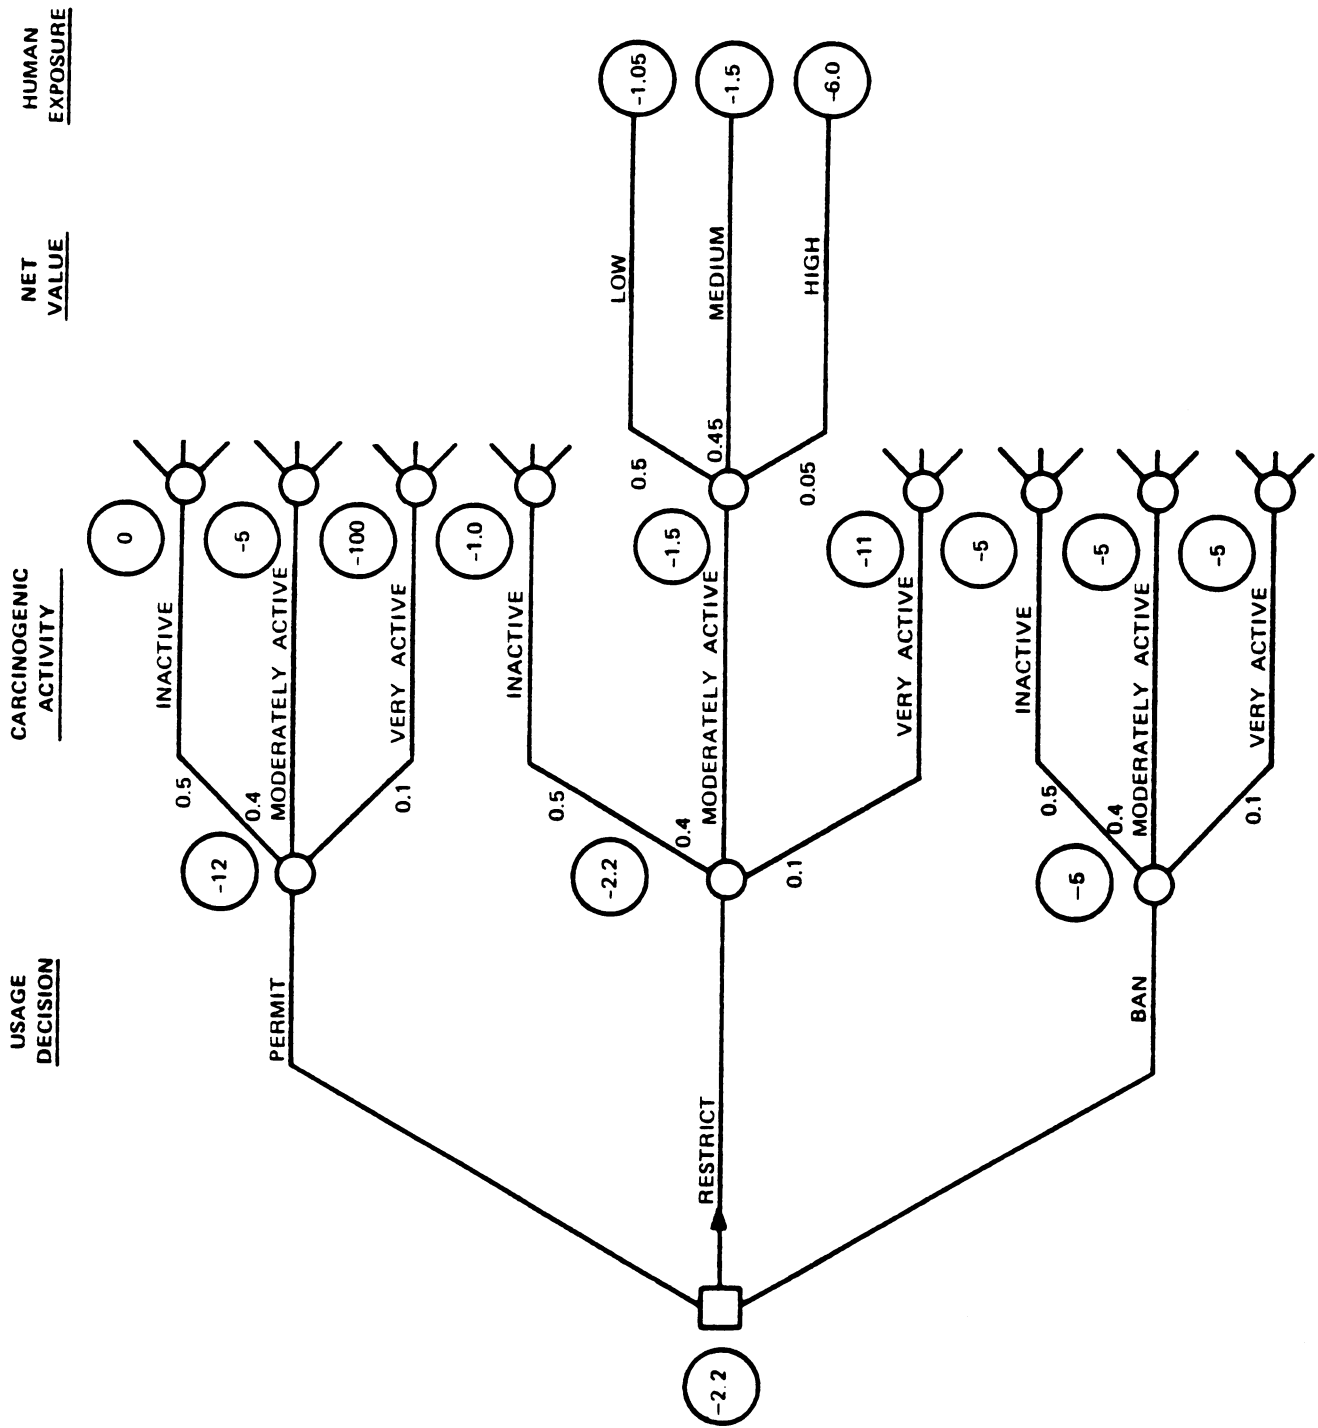


FIGURE 15.3 PARTIAL DECISION TREE FOR PRIMARY DECISION

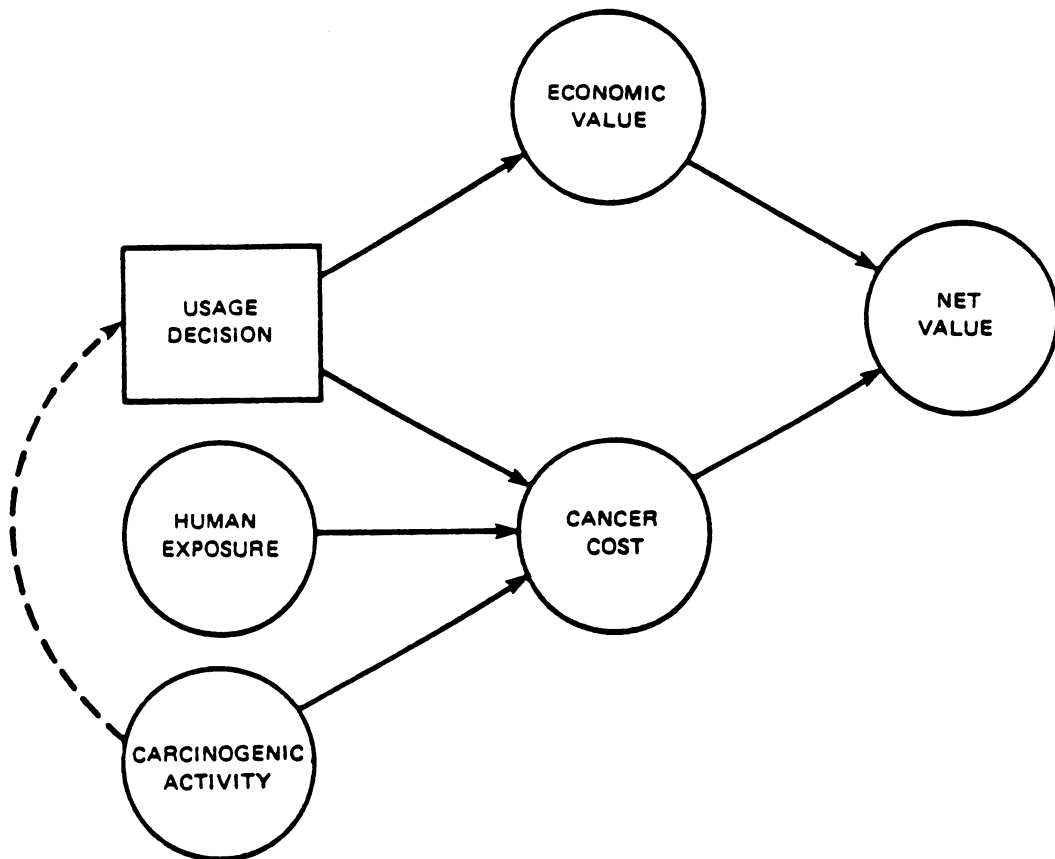


FIGURE 15.1.1 INFLUENCE DIAGRAM MODIFICATION TO DETERMINE THE VALUE WITH PERFECT INFORMATION ON CARCINOGENIC ACTIVITY

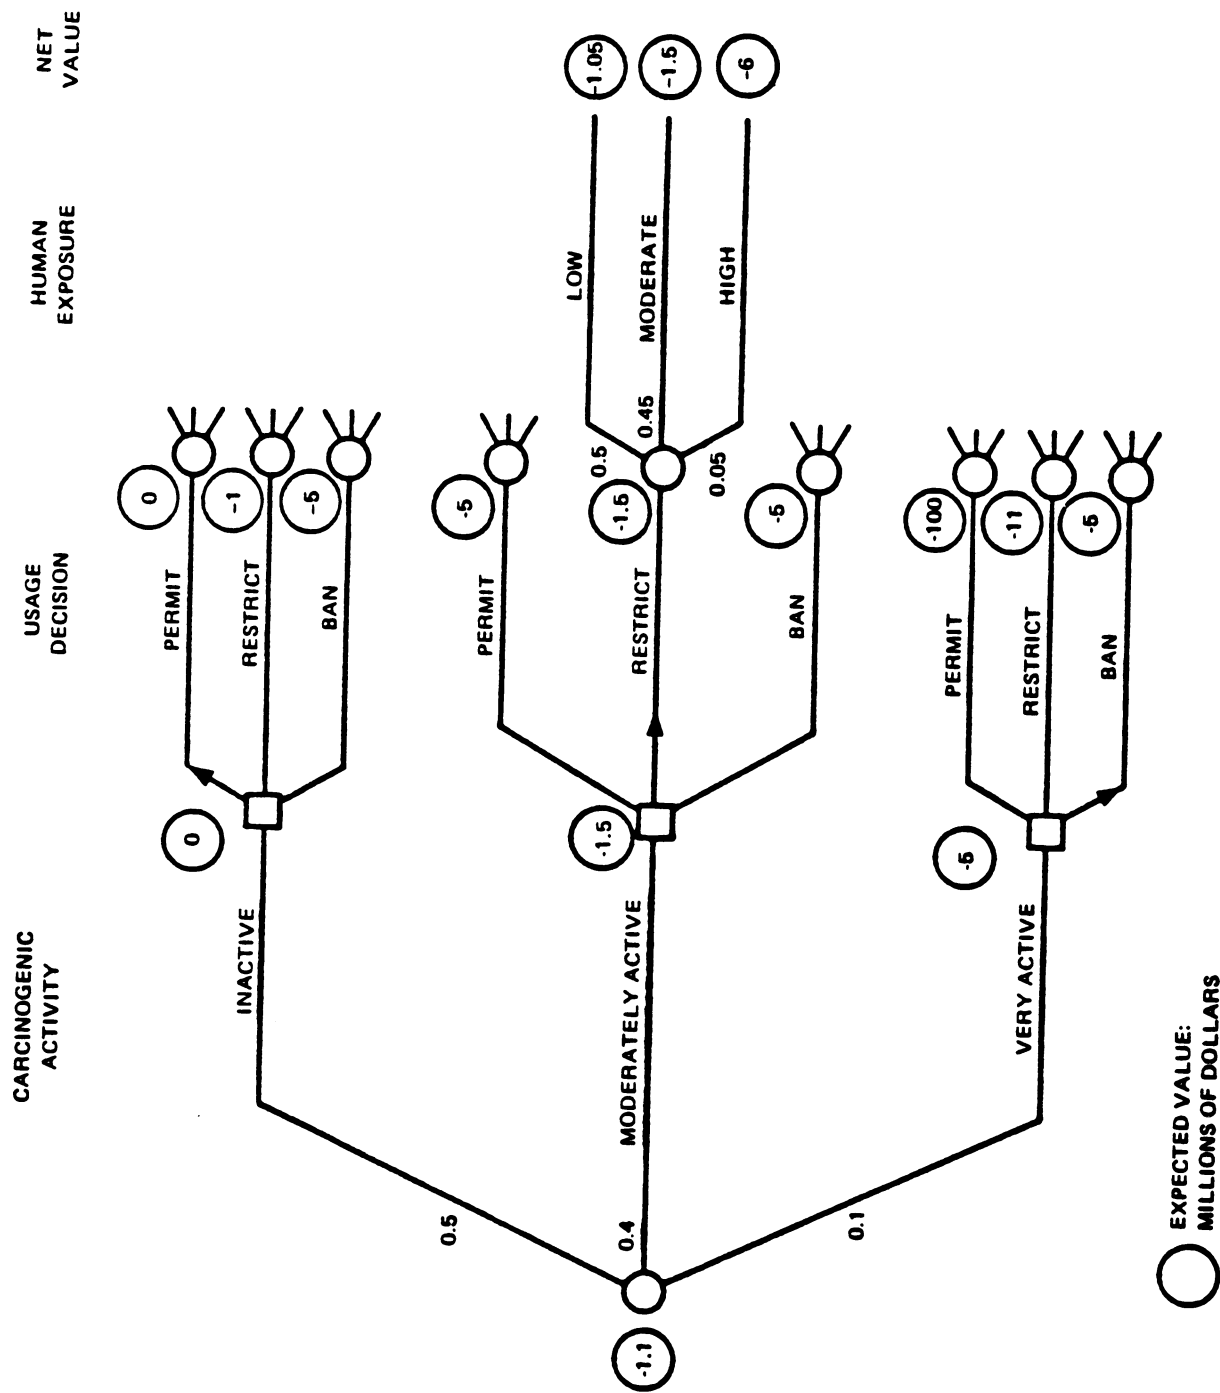


FIGURE 15.1.2 PARTIAL DECISION TREE FOR PERFECT ACTIVITY INFORMATION

The value of clairvoyance on exposure can be calculated to be \$0.4 million by adding on influence arrows from the human exposure made to the usage decision node in Figure 15.1. The associated decision rule is to restrict if exposure is low or medium, and to ban if exposure is high.

Finally by adding influence arrows from both the carcinogenic activity node and the human exposure node to the usage decision node, we find the value of clairvoyance on both activity and exposure to be \$1.38 million, which is less than the sum of the values of clairvoyance on each quantity separately. The decision rule is shown later in Table 15.2.1.

15.2 Value of Imperfect Information

To place a value on imperfect information we must model the information source. To be useful the informational report must depend probabilistically on one or more of the uncertain variables in the problem. To incorporate this dependence we augment the influence diagram with a model of the information gathering activity.

In the example at hand, it might be possible to carry out a laboratory test of the carcinogenic activity of the chemical. In this case we begin by adding a chance node to represent the report from the activity test. In Figure 15.2.1 we have added an activity test node, we have drawn an arrow to it from the carcinogenic activity node showing that the test result depends on the actual carcinogenic activity of the chemical, and we have drawn an arrow from the activity test to the usage decision showing that the decision-maker will know the test result when he makes the usage decision. We must also check the logic of each probabilistic statement represented in the diagram because additional knowledge, in principle, could change the probabilistic dependence elsewhere in the diagram.

The automated system would now ask us to define the test results. We reply that there are three test results called "INACTIVE", "MODERATELY ACTIVE", and "VERY ACTIVE" corresponding to the possibilities for the actual activity. However, unlike the case of perfect information, these test indications may be misleading. The system now asks us to supply the probabilities of these test results for each state of carcinogenic activity (i.e., to supply the likelihood function). Figure 15.2.2 shows a possible display with the assigned probabilities.

All of the information needed to determine the value of the carcinogenic activity test has now been supplied. However, the influence diagram of Figure 15.2.1 is a decision network, rather than a decision tree network, so it must be manipulated into decision tree network form before a decision tree can be generated and evaluated. The problem is that the carcinogenic activity node precedes the usage decision node, but activity is unknown to the decision maker when he makes the usage decision. A decision tree beginning with resolution of carcinogenic activity would incorrectly give this information to the decision maker. The problem is resolved by turning around the influence arrow between carcinogenic activity and the

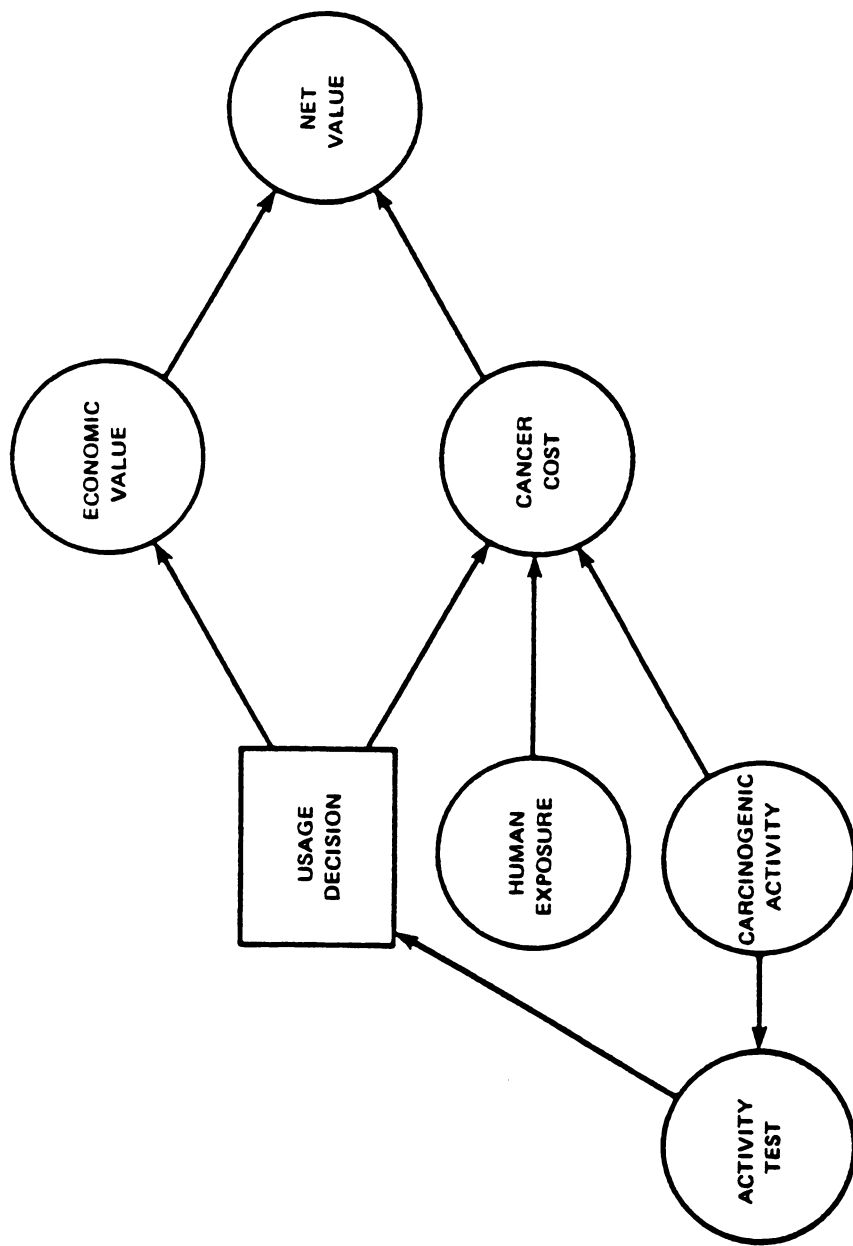


FIGURE 15.2.1 INFLUENCE DIAGRAM TO DETERMINE THE VALUE WITH IMPERFECT INFORMATION ON CARCINOGENIC ACTIVITY

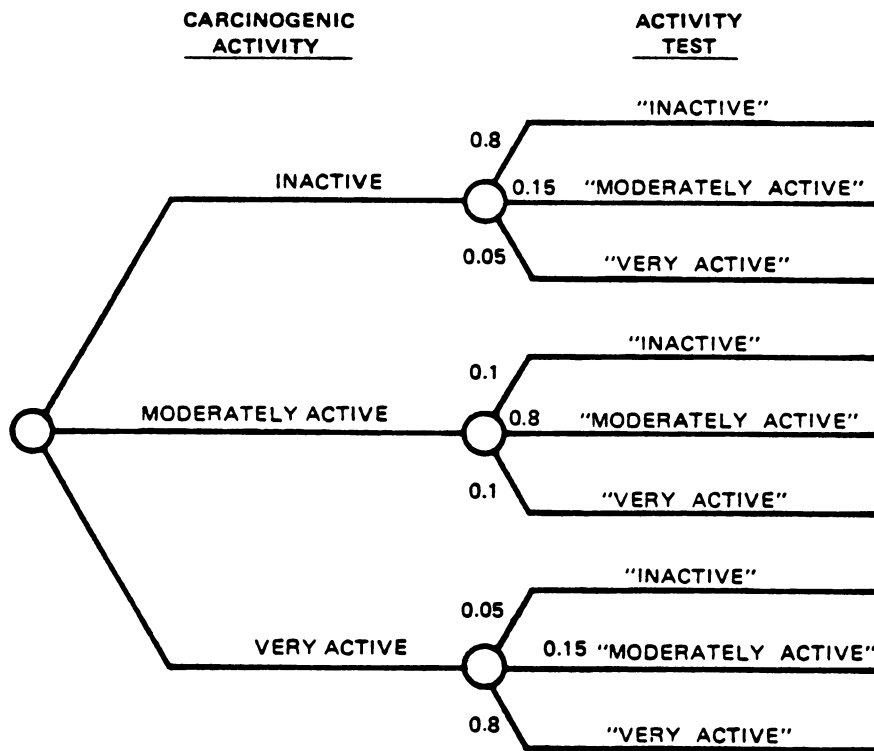


FIGURE 15.2.2 ACTIVITY TEST PROBABILITY ASSIGNMENTS

activity test; the reversal is possible because both nodes have no impinging influences. This manipulation requires the application of Bayes's rule to determine from the original probability assignments new assignments conditional in the opposite order. The procedure is straightforward for an automated system and results in the desired decision tree network. In fact, a sophisticated system could determine that this manipulation was required and carry it out without being asked by the user.

Evaluation of this network yields an expected cost, given the activity test option, of \$1.96 million. Subtracting this cost from the original cost of \$2.20 million yields an expected value of \$0.24 million from a free activity test. This is the upper limit on the price the decision maker should pay for the actual test.

A test of the degree of human exposure also could be treated by adding an exposure test node to the influence diagram. The necessary probability assignments are shown in Figure 15.2.3. Finally, the value of testing both carcinogenic activity and human exposure could be determined by making both modifications as illustrated in Figure 15.2.4. This influence diagram indicates that given human exposure and carcinogenic activity, exposure test and activity test results are probabilistically independent.

In this example, we have shown how influence diagrams can be used to designate the initial structure of the problem. The automated system can then interact with the user to request and develop values for the probability assignments that are implicitly specified in the influence diagram. The automated system can then process the information to solve the decision problem. The automated system, not the user, develops the decision tree from the influence diagram specifications. This method allows the user to ask value of perfect information questions through simple modifications of the initial influence diagram, and to ask value of imperfect information questions by augmenting the influence diagram to model the information gathering activities.

For this example, the decision rules and values for all information gathering possibilities are displayed in Tables 15.2.1 and 15.2.2. The value of the situation with the specified information is given as well as this value less the value of the primary decision (-2.2 in this case). This difference is the value of the specified information. The decision rules for joint information are given in matrix form and the ones for individual information are given along the edges of the matrix. These summaries allow the user to see easily which information is most useful. For example, Table 15.2.2 shows that imperfect exposure information is useless because the decision rule is to restrict usage regardless of the

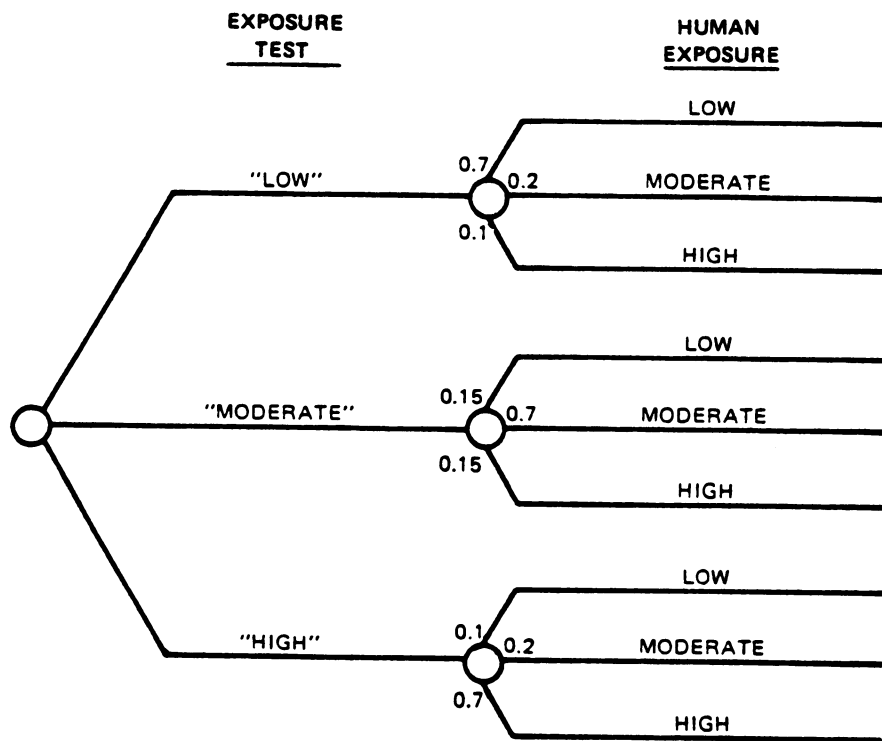


FIGURE 15.2.3 EXPOSURE TEST PROBABILITY ASSIGNMENTS

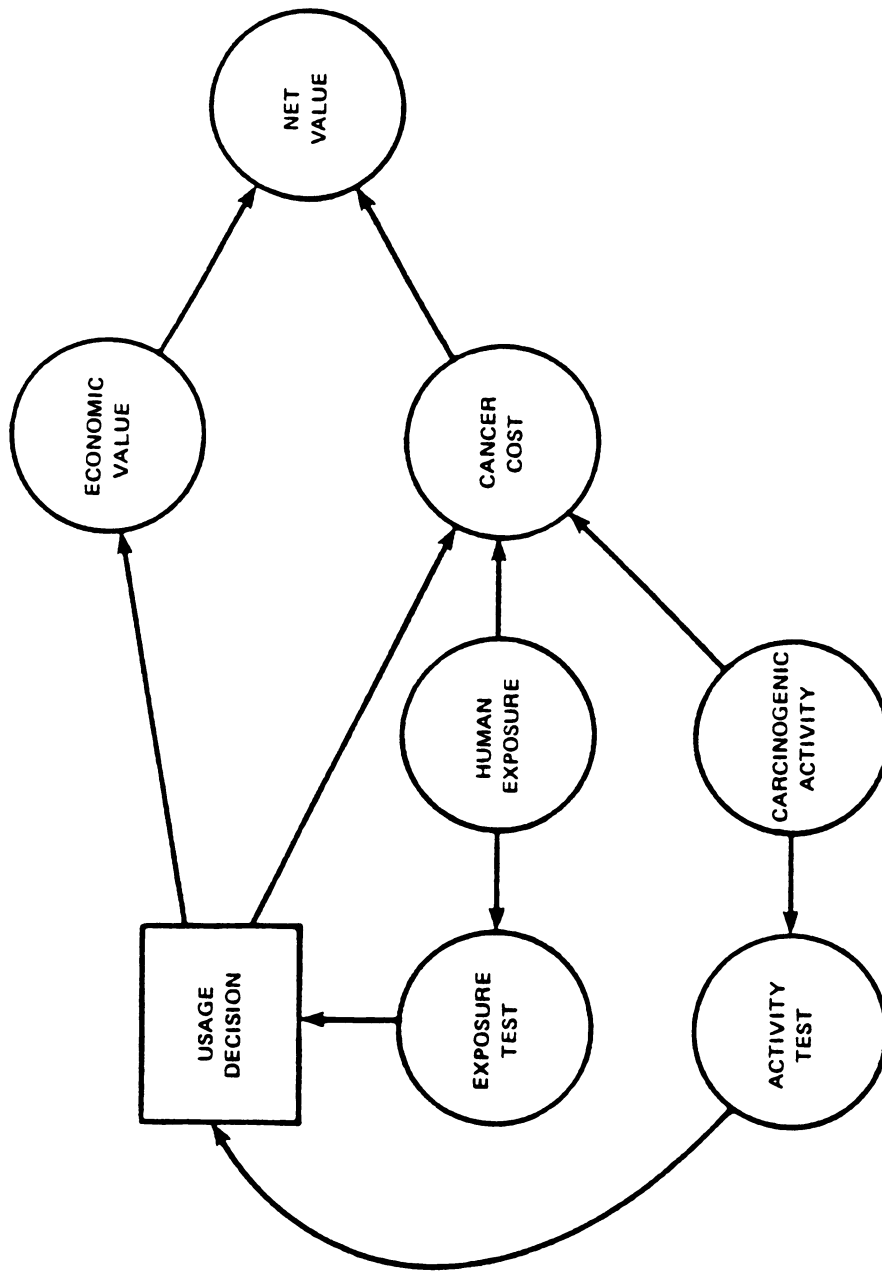


FIGURE 15.2.4 INFLUENCE DIAGRAM TO DETERMINE THE VALUE OF IMPERFECT INFORMATION ON BOTH CARCINOGENIC ACTIVITY AND HUMAN EXPOSURE

EXPOSURE INFORMATION

VALUE WITH = 1.8 VALUE OF = 0.40

JOINT INFORMATION

VALUE WITH - 0.82

VALUE OF 1.38

DECISION RULE

		EXPOSURE INFORMATION		
		LOW	MEDIUM	HIGH
DECISION RULE		RESTRICT	RESTRICT	BAN
INACTIVE	PERMIT	PERMIT	PERMIT	PERMIT
MODERATELY ACTIVE	RESTRICT	PERMIT	RESTRICT	BAN
VERY ACTIVE	BAN	RESTRICT	BAN	BAN

ACTIVITY INFORMATION
 VALUE WITH = -1.10
 VALUE OF + 1.10
 DECISION RULE

TABLE 15.2.1 PERFECT INFORMATION SUMMARY

ACTIVITY INFORMATION
 VALUE WITH = -1.96
 VALUE OF = 0.25
 DECISION RULE

JOINT INFORMATION
 VALUE WITH = -1.74
 VALUE OF = 0.46

EXPOSURE INFORMATION
 VALUE WITH = -2.2 VALUE OF = 0.00

		DECISION RULE		
		"LOW"	"MEDIUM"	"HIGH"
		RESTRICT	RESTRICT	RESTRICT
"INACTIVE"	RESTRICT	PERMIT	RESTRICT	RESTRICT
"MODERATELY ACTIVE"	RESTRICT	RESTRICT	RESTRICT	RESTRICT
"VERY ACTIVE"	BAN	RESTRICT	BAN	BAN

TABLE 15.2.2 IMPERFECT INFORMATION SUMMARY

outcome of the test, even though as shown in Table 15.2.1, perfect information would be valuable. Examination of the two decision-rule matrices for the joint information cases shows four differences in choice of alternatives between the perfect and imperfect information cases. Perfect joint information is three times more valuable than imperfect information.

We have shown in this example how influence diagrams can be used to model the primary decision problem, to determine the value of perfect information on the uncertain variables, and finally to determine the value of actual, but imperfect, information. The latter calculation usually requires the application of Bayes's law. Decision tree methods require the user to apply Bayes's law and supply the answers, or at least the formulas, for the appropriate probabilities on the decision tree. Because the influence diagram captures the logic of the problem in a more fundamental way, the user need only supply the initial probabilities that represent his model of the information gathering activity, and an automated system can carry out the rest of the analysis. This example shows how influence diagrams can greatly simplify the probabilistic modeling and decision making process.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the contributions of Dr. Miley W. Merkhofer and Dr. Allen C. Miller, III to many of the early discussions in which the concept of influence diagrams was developed.

REFERENCE

1. Howard, R. A., "Bayesian Decision Models for System Engineering," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-1, Vol. 1, November, 1965; reprinted in Readings in Decision Analysis, 2nd Edition, Stanford Research Institute, Menlo Park, CA 1977.
2. Howard, R. A., "The Used Car Buyer," Readings in Decision Analysis, 2nd Edition, Stanford Research Institute, Menlo Park, CA 1977.
3. Miller, III, A. C., Merkhofer, M. M., Howard, R. A., "Development of Automated Aids for Decision Analysis," prepared for the Defense Advanced Research Projects Agency. Menlo Park, California: Stanford Research Institute, May 1976.

THE USE OF INFLUENCE DIAGRAMS IN STRUCTURING COMPLEX DECISION PROBLEMS

Daniel L. Owen
Strategic Decisions Group

THE USE OF INFLUENCE DIAGRAMS

IN STRUCTURING COMPLEX DECISION PROBLEMS

Daniel Owen
Decision Analysis Group
SRI International
Menlo Park, California

Abstract

An influence diagram is a theoretically-based aid for obtaining the decision maker's structure for a complex decision problem under uncertainty. Advantages of using an influence diagram are rapid identification of important state and decision variables, a more balanced decision model, and the direct construction of the decision tree. A new mathematical characterization of "influence" is also presented.

1. INTRODUCTION

The idea of an influence diagram was originally the result of a need to communicate with computers about the structure of decision problems. Under contract to the Defense Advanced Research Projects Agency, researchers in the Decision Analysis Group of SRI International were working to develop automated aids for decision analysis.(1) They hoped that the decision problem structure could be described to a computer that could partially automate the solution of the decision problem.

My experience with influence diagrams has been in two other areas. First, at SRI International our current use of influence diagrams is not for interacting with computers, but rather for

communicating among people. One might describe our use of influence diagrams as participative modeling (2) rather than interactive modeling. Second, in my research at Stanford University, I have attempted to extend the notion of an influence diagram so that it can be used by the decision analyst to conceptualize the relationship between the probability distributions on different variables in a decision model.


After a brief review of influence diagram fundamentals in the next section, the remainder of this paper deals with the use of influence diagrams in participative modeling and the extension of the concept of influence.

2. INFLUENCE DIAGRAM FUNDAMENTALS

An influence between two random variables, x and y , is said to exist when the variables are not probabilistically independent:

Definition. An influence exists between two random variables x and y if and only if $\{y|x, S\} \neq \{y|S\}$, where $\{y|S\}$ denotes the probability distribution for y , conditioned on the state of information S .

The existence of this influence can be shown diagrammatically by placing the names of the variables within circular nodes and connecting those two nodes with an arrow, such as



which may be read "x influences y."

Using this definition, some rules for the manipulation of influence diagrams can be derived and are discussed thoroughly in Reference (1). For example, the influence diagram of Figure 1a represents the expansion $\{x,y,z|S\} = \{z|x,y,S\} \{y|x,S\} \{x|S\}$. An alternative expansion, represented by Figure 1b, is $\{x,y,z|S\} = \{x|y,z,S\} \{z|y,S\} \{y|S\}$ and therefore, Figure 1b is an allowable rearrangement of the influences of Figure 1a. Comparing Figures 1b and 1c shows that the influence between y and z has been removed in the latter case, and $\{x,y,z|S\} = \{x|z,y,S\} \{z|S\} \{y|S\}$. Each influence diagram is an assertion of probabilistic dependence.

Since an influence diagram implies the existence of a joint probability distribution over the variables in the diagram, those variables must be precisely defined. Variables such as "quality," "market conditions," and "attractiveness" do not take on values that are identifiable events. Consequently, probability distributions cannot be

defined over these variables, and they should not be included in an influence diagram.

Decision nodes are represented as squares on influence diagrams. An arrow from a decision node to a state variable node indicates that the probability distribution on the state variable depends on the setting of the decision variable. An arrow from a state variable to a decision variable means that the value of the uncertain state variable will be known at the time the decision is made.

3. USING INFLUENCE DIAGRAMS FOR PARTICIPATIVE MODELING

Three important difficulties arise in structuring a decision problem: unfamiliarity, complexity, and numerous participants. First, unfamiliarity means the analyst's initial unfamiliarity with the political, technical, and economic factors in a particular decision problem. The generality of the decision analysis methodology permits its application to decision problems regardless of the particular discipline or setting in which the problem occurs. Consequently, the decision analyst may be unfamiliar with the relationships of the problem variables. Secondly, complexity refers to the fact that decision models often involve a large number of random variables. Finally, if the problem is complex, there are likely to be many participants in structuring the problem. Even when there is a single decision maker, many experts may be consulted regarding the relationships of the variables. Moreover, several decision analysts may be involved.

An important property of the definition of the existence of influence given above is that it appears to coincide with the decision maker's intuitive use of the word influence. In our several experiences with influence diagrams at SRI International, we found that when a decision maker identified the existence of an influence, the variables later turned out to be probabilistically

dependent. Furthermore, influences that were identified as being "strong" represented, roughly speaking, more probabilistic dependence than influences that were identified as "weak." Thus, influence diagrams provide a language through which those untrained in the modeling of complex probabilistic systems can describe their perception of the problem.

An influence diagram is constructed jointly with the decision maker by beginning with the value attributes and working backwards from there. As an example, suppose we have a client who must decide how much to expand his production facilities. If we seek the amount of additional capacity that maximizes the present value of the business, then the value attributes are the cash flows for each future year n . Since those future cash flows are uncertain, the cash flow for each year n is drawn on an influence diagram as a chance node (circle). We ask the decision maker what variable he would most like to know the value of in order to reduce his uncertainty about the cash flow in year n . If he answers "revenue in year n ," he is asserting that $\{ \text{cash flow year } n \mid \text{revenue year } n, S \} \neq \{ \text{cash flow year } n \mid \text{revenue year } n, S \}$.

We then ask him if, given the value of revenue for year n , there is another variable, the knowledge of which would further reduce his uncertainty about cash flow. His answer is that he would like to know his total costs for year n . In that case, he is asserting that $\{ \text{cash flow year } n \mid \text{revenue year } n, S \} \neq \{ \text{cash flow year } n \mid \text{revenue year } n, \text{costs year } n, S \}$. These influences are shown in Figure 2a.

Next, we select one of the two influencing variables, say revenue in year n , and ask what information the decision maker would like about revenue in year n in order to reduce his uncertainty about the revenue in year n . He may answer that he would like to know the number of units sold in year n . The price of those units is a decision variable and is displayed in a square

node. The influence diagram now appears as in Figure 2b.

Notice that in Figure 2a revenues in year n are undoubtedly influenced by costs in year n . Knowing that costs are low would suggest that revenues are also low. Hence, an arrow could have been drawn from costs to revenue. However, in Figure 2b no arrow needs to be drawn from costs to revenue because, given units sold and price, the revenues are independent of costs, i.e., $\{ \text{revenue} \mid \text{units sold, unit price, } S \} = \{ \text{revenue} \mid \text{units sold, unit price, costs, } S \}$.

If this procedure is repeated, an influence diagram such as that in Figure 2c may result. This diagram displays several noteworthy features. First, it is evident from the diagram that the decisions about production capacity and price are closely related. Price must be optimized for each setting of the additional capacity decision variable.

A second feature is the absence of a node called "market size in year n ," which might influence units sold. If the decision maker is certain about the market size in year n because of its predictability, then the revelation of its actual value would do nothing to reduce his uncertainty about units sold. In this case, no node for market size appears on the diagram, as it would in a flow or block diagram. Similarly, unit cost for existing production capacity is well known and not shown as influencing cost in year n .

A final noteworthy feature of the diagram is the noncausal influence of the competitor's unit cost on the decision maker's unit cost for the new plant. The decision maker may be uncertain about the performance of the additional production facilities and therefore about the unit cost. If a competitor has a plant of a design similar to the decision maker's new capacity, knowing the competitor's unit cost might considerably reduce the decision maker's uncertainty about his own

unit cost. Notice that the relationship between competitor's unit cost and the unit cost for the new plant is stochastic rather than deterministic. Such stochastic and noncausal relationships are difficult to represent with conventional flow or block diagrams.

Since the influence diagram contains the important probabilistic relationships and decisions as perceived by the decision maker, a consistent decision tree can be constructed directly from this diagram and the necessary conditional probability distributions can be assessed. Because the diagram is generated backwards from the value attributes, the identified variables tend to fan away from the value attributes uniformly, resulting in a more balanced model. Excessive modeling of detail in one area at the expense of other areas is avoided. Furthermore, the backwards procedure helps the decision maker to identify stochastic and noncausal relationships as well as deterministic and casual relationships.

4. EXTENDING THE CONCEPT OF INFLUENCE

There are some deficiencies with the influence method of structuring decision problems presented above. In my opinion, the most serious is the lack of a mathematical characterization of the strength of an influence. The definition presented earlier gives only the condition for the existence of an influence, but does not characterize the influence itself. In generating the influence diagram, we ask the decision maker to identify the strongest influences. However, we do not have a mathematical expression corresponding to the intuitive notion of the strength of an influence, and in some instances the relative strengths of two influences may be ambiguous.

My research to characterize the influences has involved distinguishing possible types of influence and then determining how to quantify the

strength of each type. (3) The diagram of Figure 3 displays the influence as transforming a marginal distribution $\{x | S\}$ into another marginal distribution $\{y | S\}$, through the equation

$$\{y | S\} = \int_{-\infty}^{\infty} \{y | x, S\} \{x | S\} dx.$$

Similarly, one can think of an influence as transforming the moments of $\{x | S\}$ into the moments of $\{y | S\}$, as in Figure 3b. For these two models of influence to be equivalent, the distributions $\{x | S\}$ and $\{y | S\}$ must be uniquely determined by their moments, which requires analyticity of the exponential transforms of $\{x | S\}$ and $\{y | S\}$. In this case, specifying how strongly each of the moments of $\{y | S\}$ depends on the moments of $\{x | S\}$ completely characterizes the influence between the two variables.

The strength of each type of influence can be measured by a derivative denoted as

$$(1) \frac{d \langle (y - \bar{y})^m | S \rangle}{d \langle (x - \bar{x})^n | S \rangle} \quad n > 1, m > 1$$


where $\langle . \rangle$ denotes expectation. (The extension of this expression when $n = 1$ or $m = 1$ should be obvious.) This derivative is the change in the m th central moment of $\{y | S\}$ due to a differential change in the n th central moment of $\{x | S\}$. It can be computed from knowledge of the derivatives of the conditional moments, e.g.,

$$\left. \frac{d^n \langle y^m | S \rangle}{dx^n} \right|_{\langle x | S \rangle}.$$

Therefore, the influence of x on y is characterized by an infinite matrix that has its n, m element given by expression (1) and is denoted by

$$\frac{d \{y | S\}}{d \{x | S\}}$$

This influence matrix has several useful properties. First, the matrix is null if and only if the moments of $\{y|S\}$ are unaffected by the moments of $\{x|S\}$ and y is probabilistically independent of x . Hence, the matrix is null if and only if no influence exists. Second, when the matrix has nonzero elements, those elements characterize the type and strength of the influence. Finally, because the elements of the matrix follow the rules of differential calculus, a calculus of influences appears possible. For example, in the diagram:



the influence matrix of x on z may be computed as

$$\frac{d\{z|S\}}{d\{x|S\}} = \frac{d\{y|S\}}{d\{x|S\}} \frac{d\{z|S\}}{d\{y|S\}}$$

5. CONCLUSION

There is a theoretical basis for influence diagrams: each one corresponds to a mathematical statement of probabilistic dependence. As demonstrated in this paper, these diagrams can be used to obtain the decision maker's structure for a complex decision problem and to communicate that structure among those trying to solve the problem. Further research to characterize an influence may lead to a new way to conceptualize the relationship between probabilistically related variables.


References

1. Allen C. Miller, et al., Development of Automated Aids for Decision Analysis, research report to Defense Advanced Research Projects Agency by SRI International, Menlo Park, California (May 1976).


2. Raymond Fritz, "System Methods in the Public Sector: A Dialogue," Systems Man. and Cybernetic Review, Vol. 7, No. 1 (February 1978).

3. Daniel Owen, "The Concept of Influence and Its Use in Structuring Complex Decision Problems," Ph.D. dissertation, Engineering-Economic Systems Department, Stanford University, October 1978.

a. $\{x,y,z|S\} = \{z|x,y,S\}\{y|x,S\}\{x|S\}$



b. $\{x,y,z|S\} = \{x|y,z,S\}\{z|y,S\}\{y|S\}$



c. $\{x,y,z|S\} = \{x|z,y,S\}\{z|S\}\{y|S\}$




Figure 1. THE CORRESPONDENCE BETWEEN INFLUENCE DIAGRAMS AND ASSERTIONS OF PROBABILISTIC INDEPENDENCE.

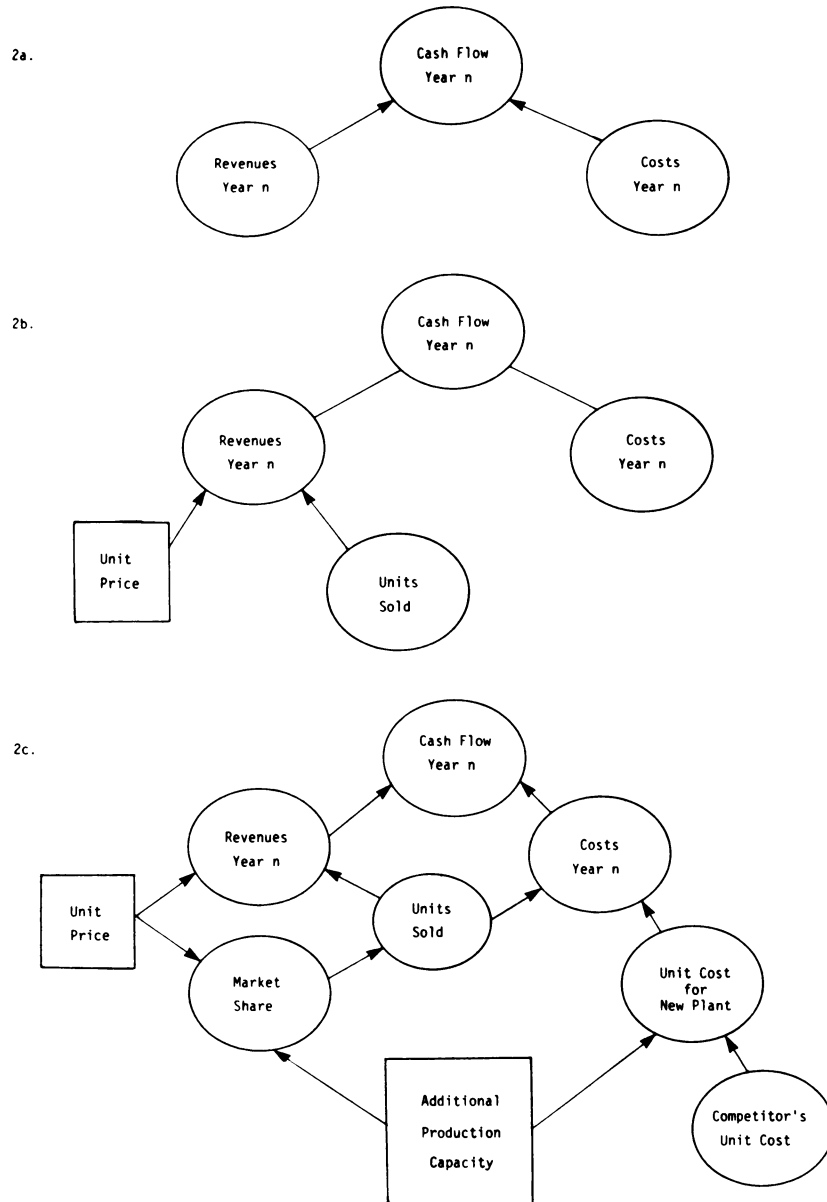


Figure 2. SEQUENTIAL GENERATION OF AN INFLUENCE DIAGRAM.

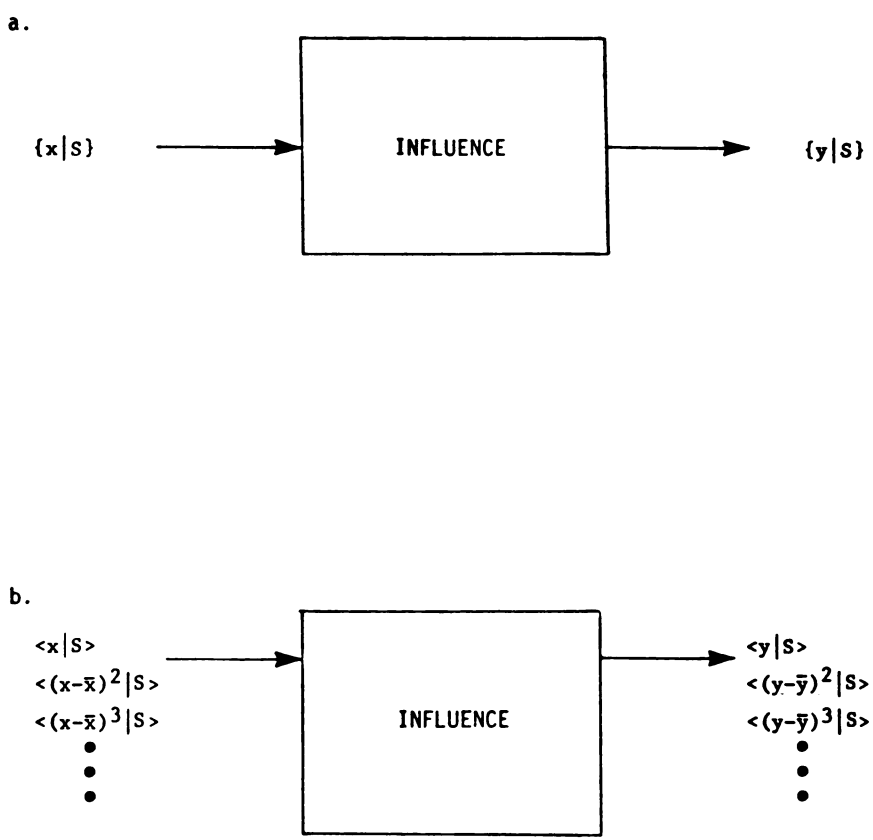


Figure 3. TWO ALTERNATE CHARACTERIZATIONS OF INFLUENCE

TECHNICAL CONTRIBUTIONS

Preface

These technical papers are written at the engineering level of mathematical training. However, the main ideas may be of interest to the general reader.

"Information Value Theory" shows how to place a monetary value on information in a decision problem. It uses a bidding example to demonstrate that the values of information on a pair of independent variables may exceed the sum of the values of information on each separately.

"The Value of Information Lotteries" shows how the profit lottery for the bidding example changes when information is available. The results are strikingly non-intuitive.

"The Economic Value of Analysis and Computation" extends the value of information idea to place a monetary value on further analysis or computation. It uses the bidding problem of the previous two papers as a computational example.

"Competitive Bidding in High-Risk Situations" shows that in oil-lease bidding the winner tends to be one who most overestimates reserves potential. This paper presents and discusses this phenomenon, which is sometimes called the winner's curse, and develops methods for bidding that take this phenomenon into account.

"Decision Analysis: Perspectives on Inference, Decision, and Experimentation" uses the formalism of decision analysis to provide a complete conceptual and methodological framework for designing experiments and interpreting their results. An example of tossing a biased coin illustrates the framework.

"Bayesian Decision Models in Systems Engineering" also treats experimentation, dealing with experimental determination of reliability. It describes the use of decision trees and the use of conjugate distributions for conceptual and computational convenience.

"Proximal Decision Analysis" develops a methodology for using deterministic sensitivity results to perform approximate probabilistic and inferential calculations. Although the methods are especially useful in complex situations, they can also provide insight and guidance in simple situations.

"Risk-Sensitive Markov Decision Processes" expands the risk probabilistic modeling structure of Markov decision processes to allow expression of risk attitude. A numerical example shows how the optimum decision policies change with risk attitude.

INFORMATION VALUE THEORY

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from the IEEE Transactions on Systems Science and Cybernetics,
Vol. SSC-2, No. 1, August 1966. p. 22-26

Copyright © 1966 in Institute of Electronics and Electrical Engineers.

Information Value Theory

RONALD A. HOWARD, SENIOR MEMBER, IEEE

Abstract—The information theory developed by Shannon was designed to place a quantitative measure on the amount of information involved in any communication. The early developers stressed that the information measure was dependent only on the probabilistic structure of the communication process. For example, if losing all your assets in the stock market and having whale steak for supper have the same probability, then the information associated with the occurrence of either event is the same. Attempts to apply Shannon's information theory to problems beyond communications have, in the large, come to grief. The failure of these attempts could have been predicted because no theory that involves just the probabilities of outcomes without considering their consequences could possibly be adequate in describing the importance of uncertainty to a decision maker. It is necessary to be concerned not only with the probabilistic nature of the uncertainties that surround us, but also with the economic impact that these uncertainties will have on us.

In this paper the theory of the value of information that arises from considering jointly the probabilistic and economic factors that affect decisions is discussed and illustrated. It is found that numerical values can be assigned to the elimination or reduction of any uncertainty. Furthermore, it is seen that the joint elimination of the uncertainty about a number of even independent factors in a problem can have a value that differs from the sum of the values of eliminating the uncertainty in each factor separately.

Manuscript received March 1, 1965.
 The author is Professor of Engineering-Economic Systems,
 Stanford University, Stanford, Calif.

NOTATION

A SPECIAL NOTATION will be used to make as explicit as possible the conditions underlying the assignment of any probability. Thus,

- x = a random variable
- A = an event
- S = the state of information on which probability assignments will be made
- $\{x|S\}$ = the density function of the random variable x given the state of information
- $\{A|S\}$ = the probability of the event A given the state of information
- $\langle x|S \rangle$ = the expectation of the random variable x , which equals $\int_x x\{x|S\}$
- \mathcal{E} = the experience brought to the problem, the special state of information represented by total a priori knowledge
- $\{x|\mathcal{E}\}$ = the density function of a random variable x assigned on the basis of only a priori knowledge \mathcal{E} ; designated the prior on x .

Our notation does not emphasize the difference in probability assignment to a random variable and to an event because the context always makes clear the appro-

appropriate interpretation. We should interpret the generalized summation symbol \int used in the definition of $\langle x|S \rangle$ as an integral if the random variable is continuous and as a summation if it is discrete.

AN INFERENCE CONCEPT

A most important inferential concept is expansion; it allows us to encode our knowledge in a problem in the most convenient form. The concept of expansion permits us to introduce a new consideration into the problem. Suppose, for example, that we must assign a probability distribution to a random variable u . We may find it much easier to assign a probability distribution to u if we had previously specified the value of another random variable v and we may also find it easy to assign a probability distribution directly to v . In this case the expansion equation,

$$\langle u|S \rangle = \int_v \langle u|vS \rangle \langle v|S \rangle \quad (1)$$

shows that all we have to do to find $\langle u|S \rangle$ is multiply $\langle u|vS \rangle$ by $\langle v|S \rangle$ and sum over all possible values of the random variable v . Mathematically this equation is no more than a consequence of the definition of conditional probability, but in the operational solution of inference problems it is extremely valuable. If we want to find only the expectation of u , $\langle u|S \rangle$, rather than the entire density function $\langle u|S \rangle$ we apply expansion in the form,

$$\langle u|S \rangle = \int_u u \langle u|S \rangle = \int_v \int_u u \langle u|vS \rangle \langle v|S \rangle = \int_v \langle u|vS \rangle \langle v|S \rangle. \quad (2)$$

This equation shows that in this case we must only provide the expectation of u conditional on v and the probability distribution of v in order to find the expectation of u .

The inferential concept of expansion allows us to go directly from the statement of an inference problem to its solution in simple logical steps. It provides a link between the formalism of probability theory and the path of human reasoning; it is a tool of thought.

A BIDDING PROBLEM

To illustrate our approach we shall consider a specific problem. Suppose that our company is bidding on a contract against a number of competitors. We shall let p be our company's cost of performing on the contract; unfortunately, we are uncertain of this cost. We let ℓ be the lowest bid of our competitors, and as you might expect, this too is uncertain. Our problem is to determine b , our company's bid on the contract. Our objective in this determination is to maximize the expected value of v , our company's profit or value from the contract.

Naturally, our company will not win the contract if its bid b is higher than the lowest bid of our competitors ℓ , thus, in this case our company's profit is zero. However, if our bid b is less than the lowest competitive bid ℓ our company will obtain the contract and will make a profit

equal to the difference between our bid b and our cost p . Thus, profit v to our company is defined by

$$v = \begin{cases} b - p & \text{if } b < \ell \\ 0 & \text{if } b > \ell \end{cases} \quad (3)$$

This equation indicates that the probability distribution of profit given our bid $\{v|b\mathcal{E}\}$ would be trivial to determine if we only knew our company's cost p and the lowest bid of our competitors ℓ , since by the expansion concept

$$\{v|b\mathcal{E}\} = \int_{p,\ell} \langle v|bp\ell\mathcal{E} \rangle \{p,\ell|b\mathcal{E}\}. \quad (4)$$

In this equation $\{p,\ell|b\mathcal{E}\}$ represents the joint distribution of our cost and our competitor's lowest bid given our bid and the state of knowledge \mathcal{E} that we brought to the problem. However, since we are interested only in the expected profit by assumption we can write (4) in the expectation form,

$$\langle v|b\mathcal{E} \rangle = \int_{p,\ell} \langle v|bp\ell\mathcal{E} \rangle \langle p,\ell|b\mathcal{E} \rangle. \quad (5)$$

At this point we shall make two assumptions. The first is that our cost p and the lowest competitor's bid ℓ do not depend upon our bid b ; that is,

$$\langle p,\ell|b\mathcal{E} \rangle = \langle p,\ell|\mathcal{E} \rangle. \quad (6)$$

The second assumption is that our cost p is independent of the lowest competitive bid ℓ , or,

$$\langle p,\ell|\mathcal{E} \rangle = \langle p|\mathcal{E} \rangle \langle \ell|\mathcal{E} \rangle. \quad (7)$$

With these assumptions (5) becomes

$$\langle v|b\mathcal{E} \rangle = \int_{p,\ell} \langle v|bp\ell\mathcal{E} \rangle \langle p|\mathcal{E} \rangle \langle \ell|\mathcal{E} \rangle. \quad (8)$$

In view of (3) we have immediately that the expectation of profit conditional on our bid, our cost, and the lowest competitive bid is simply

$$\langle v|bp\ell\mathcal{E} \rangle = \begin{cases} b - p & \text{if } b < \ell \\ 0 & \text{if } b > \ell \end{cases} \quad (9)$$

The next step is to assign prior probability distributions $\langle p|\mathcal{E} \rangle$, $\langle \ell|\mathcal{E} \rangle$ to our cost and to the lowest competitive bid. The probability distribution on our cost $\langle p|\mathcal{E} \rangle$ would be based upon the information that our company had gathered in performing similar contracts as amended by the technical considerations involved in the new contract. The prior distribution on the lowest competitive bid $\langle \ell|\mathcal{E} \rangle$ would be based on our experiences in bidding against our competitors on previous occasions and on other information such as reports in the trade press. Although we shall not have the opportunity to digress on the subject of how these assignments are made, suffice it to say that effective procedures for this purpose are available.

To gain ease of computation we shall state simply that in this problem the prior distribution on our company's cost p is a uniform distribution between zero and one and that our prior distribution on our competitor's lowest bid is a uniform distribution between zero and two. These prior




Fig. 1. Priors on cost of performance and lowest competitive bid.

$$\langle v|bE \rangle = \frac{1}{2}(2-b)(b-\frac{1}{2}) = -\frac{1}{2} + \frac{5}{4}b - \frac{1}{2}b^2 \quad b \leq 2$$




Fig. 2. Expected profit as function of our bid.

distributions are shown in Fig. 1 along with the cumulative and complementary cumulative distributions that they imply. From the symmetry of the density functions we see immediately that our expected cost $\langle p|E \rangle = \bar{p}$, is just 1/2 and that the expected lowest bid of our competitors $\langle l|E \rangle$ is 1. However, it would be folly to assume that these expected values will occur.

We can now return to (8) and substitute these results,

$$\begin{aligned} \langle v|bE \rangle &= \int_p \int_{\ell} v|bp\ell E \rangle \{p|E\} \{l|E\} \\ &= \int_p \int_{\ell=b}^{\infty} (b-p) \{p|E\} \{l|E\} \\ &= \int_p (b-p) \{p|E\} \int_{\ell=b}^{\infty} \{l|E\} \\ &= \{l > b|E\} \int_p (b-p) \{p|E\} \\ &= \{l > b|E\} (b - \langle p|E \rangle) \\ &= \{l > b|E\} (b - \bar{p}). \end{aligned} \quad (10)$$

Our expected profit, given that we bid b , is therefore the product of the probability that our competitor's lowest bid will exceed ours and the difference between our bid and our expected cost. From Fig. 1 we have that our expected profit, given that we bid b , is simply

$$\langle v|bE \rangle = \frac{1}{2}(2-b)\left(b - \frac{1}{2}\right) \quad 0 \leq b \leq 2. \quad (11)$$

Figure 2 shows this expected profit as a function of our bid. Note that a maximum profit of 9/32 is achieved by making the bid equal to 5/4. Of course, the same result is obtained by setting the derivative of (11), with respect to b , equal to zero and solving for b . Note that the expected profit is positive for bids in the range between 1/2 and 2,

and negative for bids between 0 and 1/2. Formally,

$$\langle v|E \rangle = \text{Max}_b \langle v|bE \rangle = \left\langle v|b = \frac{5}{4}, E \right\rangle = \frac{9}{32} = \frac{27}{96}. \quad (12)$$

We have, therefore, found that the best bidding strategy is to bid 5/4 and that this strategy will have an expected profit of 9/32.

CLAIRVOYANCE

Even though we have devised a bidding strategy that is optimum in the face of the uncertainties involved we could still face a bad outcome. We may not get the contract, and if we do get it, we may lose money. It is reasonable, therefore, that if a perfect clairvoyant appeared and offered to eliminate one or both of the uncertainties in the problem, we would be willing to offer him a financial consideration. The question is how large should this financial consideration be.

We shall let C represent clairvoyance and C_x represent clairvoyance about a random variable x . Thus any probability assignment conditional on C_x is conditional on the fact that the value of x will be revealed to us. We define v_{C_x} as the increase in profit that arises from obtaining clairvoyance about x . It is clear that the expected increase in profit owing to clairvoyance about x , $\langle v_{C_x}|E \rangle$ is just the difference between the expected profit that we shall obtain as a result of the clairvoyance $\langle v|C_xE \rangle$ and the expected profit that we would obtain without clairvoyance $\langle v|E \rangle$; thus,

$$\langle v_{C_x}|E \rangle = \langle v|C_xE \rangle - \langle v|E \rangle. \quad (13)$$

We compute the expected profit given clairvoyance about x , $\langle v|C_xE \rangle$, by evaluating the expected profit given that we know x , $\langle v|xE \rangle$, for each possible value of x that the clairvoyant might reveal and then summing this expectation with respect to the probability assignment on x , $\{x|E\}$,

$$\langle v|C_xE \rangle = \int_x \langle v|xE \rangle \{x|E\}. \quad (14)$$

We use the prior probability distribution $\{x|E\}$ for x in this calculation because up to the moment that the clairvoyant actually reveals the value of x the probability that we must assign to his statement about x is based only on our prior knowledge E .

By using (13) and (14) we can assign the expected value in monetary units of eliminating any uncertainty in the problem. We shall now apply these results to computing the value of eliminating the uncertainty in our cost and in the lowest competitive bid in the problem.

ANALYSIS OF CLAIRVOYANCE ABOUT OUR COST

We now determine $\langle v_{C_p}|E \rangle$, the expected value to us of eliminating uncertainty about our cost p . From (13) and (14) we write

$$\langle v_{C_p}|E \rangle = \langle v|C_pE \rangle - \langle v|E \rangle \quad (15)$$

$$\langle v|C_pE \rangle = \int_p \langle v|pE \rangle \{p|E\}. \quad (16)$$

The expected profit $\langle v|p\varepsilon \rangle$ that we would make if we knew our cost is just the expected profit that we would make if we bid b and if our cost was p maximized with respect to our bid b ; that is,

$$\langle v|p\varepsilon \rangle = \underset{b}{\text{Max}} \langle v|bp\varepsilon \rangle. \tag{17}$$

Using the assumptions of (6) and (7) we invoke the expansion concept to write

$$\begin{aligned} \langle v|bp\varepsilon \rangle &= \int_{\ell} v|bp\varepsilon \ell \varepsilon \rangle \{ \ell | \varepsilon \} \\ &= \int_{\ell > b} (b - p) \{ \ell | \varepsilon \} = (b - p) \{ \ell > b | \varepsilon \} \\ &= (b - p) \frac{1}{2} (2 - b) \end{aligned} \tag{18}$$

Note that if we take the expectation of this equation with respect to the prior on p we immediately obtain (11).

We determine our bid by maximizing $\langle v|bp\varepsilon \rangle$ with respect to b . By setting its derivative with respect to $b = 0$ we find immediately that

$$\frac{d}{db} \langle v|bp\varepsilon \rangle = 0 \rightarrow 2 - b - (b - p) = 0, 2b = 2 + p.$$

$$b = 1 + \frac{p}{2} \tag{19}$$

The bid should be equal to 1 plus 1/2 the value of our cost. When we insert this result in (17) we find that the expected profit given that our cost is p is

$$\langle v|p\varepsilon \rangle = \left\langle v|b = \left(1 + \frac{p}{2} \right), p\varepsilon \right\rangle = \frac{1}{2} \left(1 - \frac{p}{2} \right)^2. \tag{20}$$

Now we compute the expected profit given clairvoyance about p from (16),

$$\langle v|C_p\varepsilon \rangle = \int_p v|p\varepsilon \rangle \{ p | \varepsilon \} = \int_0^1 dp \cdot \frac{1}{2} \left(1 - \frac{p}{2} \right)^2 = \frac{7}{24} = \frac{28}{96} \tag{21}$$

This expected profit is 28/96, 1/96 more than we obtained in (12) in the case of no clairvoyance. Therefore, following (15) we have found

$$\langle v_{c_p} | \varepsilon \rangle = \langle v|C_p\varepsilon \rangle - \langle v | \varepsilon \rangle = \frac{28}{96} - \frac{27}{96} = \frac{1}{96}. \tag{22}$$

The expected increase in our profit that will result from having our cost revealed to us by a clairvoyant is consequently 1/96.

If the value of this analysis depended on the actual existence of clairvoyance then it would have only theoretical interest. However, since clairvoyance represents complete elimination of the uncertainty about a random variable it follows that what we would pay for clairvoyance should be an upper bound on any experimental program that purports to aid us in eliminating uncertainty about that variable. Thus, in the present problem the company

should not hire any cost accounting or production experts to aid in eliminating uncertainty about p unless the cost of these services is considerably below $\langle v_{c_p} | \varepsilon \rangle$. The concept of clairvoyance plays the same role in analyzing decision problems that the concept of a Carnot engine plays in analyzing thermodynamic problems. These theoretical constructs provide bench marks against which practical realizations can be tested.

ANALYSIS OF CLAIRVOYANCE ABOUT THE LOWEST COMPETITIVE BID

Now we determine what it is worth to know the lowest competitive bid. We write immediately,

$$\langle v_{c_\ell} | \varepsilon \rangle = \langle v|C_\ell\varepsilon \rangle - \langle v | \varepsilon \rangle \tag{23}$$

$$\langle v|C_\ell\varepsilon \rangle = \int_{\ell} v | \ell \varepsilon \rangle \{ \ell | \varepsilon \} \tag{24}$$

and

$$\langle v | \ell \varepsilon \rangle = \underset{b}{\text{Max}} \langle v|b\ell\varepsilon \rangle. \tag{25}$$

We can perform expansion in terms of our cost p just as we did in (18) in terms of ℓ . We write,

$$\langle v|b\ell\varepsilon \rangle = \int_p v|bp\ell\varepsilon \rangle \{ p | \varepsilon \}. \tag{26}$$

From (9) we have immediately that

$$\langle v|b\ell\varepsilon \rangle = \begin{cases} \int_p (b - p) \{ p | \varepsilon \} = b - \bar{p} & \text{if } b < \ell \\ 0 & \text{if } b > \ell. \end{cases} \tag{27}$$

Therefore if $\ell < \bar{p}$ do not bid; if $\ell > \bar{p}$ bid ℓ^- , just below ℓ .

It is easy to see why this is the best bid. Since we know the lowest competitive bid ℓ we can get the contract by bidding just under ℓ . However, if ℓ is already less than our expected cost then we would expect to lose money by such a strategy for it would be better not to get the contract at all. Consequently, if ℓ is less than our expected cost \bar{p} we should not bid; while if ℓ is greater than \bar{p} we should bid just below ℓ . Therefore we have

$$\langle v | \ell \varepsilon \rangle = \underset{b}{\text{Max}} \langle v|b\ell\varepsilon \rangle = \begin{cases} \ell - \bar{p} & \text{if } \ell > \bar{p} \\ 0 & \text{if } \ell < \bar{p}. \end{cases} \tag{28}$$

The expected profit given the lowest competitive bid ℓ , $\langle v | \ell \varepsilon \rangle$, is plotted as a function of ℓ in Fig. 3. To find the expected profit given clairvoyance about ℓ , $\langle v|C_\ell\varepsilon \rangle$, we integrate this function with respect to the ℓ density function of Fig. 1 and obtain

$$\langle v|C_\ell\varepsilon \rangle = \int_{\ell} v | \ell \varepsilon \rangle \{ \ell | \varepsilon \} = \int_{1/2}^1 dl \left(\ell - \frac{1}{2} \right) \cdot \frac{1}{2} = \frac{9}{16} = \frac{54}{96}. \tag{29}$$

The expected profit given clairvoyance about ℓ is 54/96, twice as large as the expected profit of 27/96 obtained in (12) where no clairvoyance was available. The expected




Fig. 3. Our expected profit as function of lowest competitive bid.

increase in profit with clairvoyance about ℓ is, therefore,

$$\langle v_{c\ell} | \mathcal{E} \rangle = \langle v | C_{\ell} \mathcal{E} \rangle - \langle v | \mathcal{E} \rangle = \frac{54}{96} - \frac{27}{96} = \frac{27}{96} \quad (30)$$

By comparing (30) and (22) we see that the expected increase in profit because of clairvoyance about the lowest competitive bid ℓ is 27 times as great as the expected increase in profit because of clairvoyance about our cost p . Yet the density functions for p and ℓ shown in Fig. 1 reveal that the range of uncertainty in ℓ is only twice the range of uncertainty in p . Why, therefore, is information about the lowest competitive bid so much more valuable than information about our cost? The explanation is that we can use information about ℓ far more effectively in controlling the profitability of the situation than we can use information about our costs. If we know ℓ then we can control immediately whether or not we get the contract, although we may still lose money even if we get it. Knowing p prevents us from bidding on any contract that would be unprofitable for us if we got it but is no help in determining whether or not we will get it.

ANALYSIS OF CLAIRVOYANCE ABOUT OUR COSTS AND THE LOWEST COMPETITIVE BID

Suppose that we are now offered clairvoyance about our cost p and the lowest competitive bid ℓ . What would this joint information be worth? We let $v_{c_{p\ell}}$ represent the increase in profit because of clairvoyance about both p and ℓ . Then in direct analogy with our earlier results for clairvoyance about a single random variable we can write

$$\langle v_{c_{p\ell}} | \mathcal{E} \rangle = \langle v | C_{p\ell} \mathcal{E} \rangle - \langle v | \mathcal{E} \rangle \quad (31)$$

$$\langle v | C_{p\ell} \mathcal{E} \rangle = \int_{p,\ell} v(p,\ell) \{p,\ell | \mathcal{E}\} \quad (32)$$

and

$$\begin{aligned} \langle v | p\ell \mathcal{E} \rangle &= \text{Max}_b \langle v | b p \ell \mathcal{E} \rangle \\ \langle v | b p \ell \mathcal{E} \rangle &= \begin{cases} b - p & \text{if } b < \ell \\ 0 & \text{if } b > \ell. \end{cases} \end{aligned} \quad (33)$$

Therefore, if $\ell < p$, do not bid; if $\ell > p$, bid ℓ . Since we know both our cost p and the lowest competitive bid ℓ we can get the contract by bidding slightly less than ℓ and we want it if this bid exceeds our cost p . Consequently, our expected profit given p and ℓ will be $\ell - p$ if ℓ exceeds p because we shall bid just less than ℓ on the contract and get it. The expected profit will be zero if ℓ is less than p because we shall not bid, e.g.,

$$\langle v | p\ell \mathcal{E} \rangle = \begin{cases} \ell - p & \text{if } \ell > p \\ 0 & \text{if } \ell < p. \end{cases} \quad (34)$$

Now we substitute the results of (34) and (7) to obtain

$$\begin{aligned} \langle v | C_{p\ell} \mathcal{E} \rangle &= \int_p \int_{\ell} v(p,\ell) \{p | \mathcal{E}\} \{ \ell | \mathcal{E} \} \\ &= \int_0^1 dp \{p | \mathcal{E}\} \int_{p-\frac{1}{2}}^2 d\ell \{ \ell | \mathcal{E} \} (\ell - p) \\ &= \frac{1}{2} \int_0^1 dp \int_p^2 d\ell (\ell - p) = \frac{7}{12} = \frac{56}{96}. \end{aligned} \quad (35)$$

The expected profit given clairvoyance about both p and ℓ is therefore 56/96, the highest expected profit we have seen thus far. The expected increase in profit resulting from clairvoyance about p and ℓ is, therefore

$$\langle v_{c_{p\ell}} | \mathcal{E} \rangle = \langle v | C_{p\ell} \mathcal{E} \rangle - \langle v | \mathcal{E} \rangle = \frac{56}{96} - \frac{27}{96} = \frac{29}{96} \quad (36)$$

The expected value 29/96 of knowing p and ℓ jointly is, therefore, greater than the sum of the 1/96 value of knowing p alone and the 27/96 value of knowing ℓ alone. The further advantage provided by the joint knowledge of p and ℓ is illustrated by the fact that now the company can never lose money whereas in each of the two earlier cases it could.

However, the most striking result of our analysis is the way in which the uncertainty about the lowest competitive bid towers over the uncertainty in our company's cost as a concern of management. In this problem it is worth far more to know what the competition is doing than it is the internal performance of our own company. You might say facetiously that we have demonstrated the dollars and cents payoff in industrial espionage.

CONCLUSION

The observations made in this paper have wide application. We can treat the case where our clairvoyance is imperfect rather than perfect. The imperfection can be because of a statistical effect arising in nature or the result of incompetence or mendacity or both in the source of the information.

Placing a value on the reduction of uncertainty is the first step in experimental design, for only when we know what it is worth to reduce uncertainty do we have a basis for allocating our resources in experimentation designed to reduce the uncertainty. These remarks are as applicable to the establishment of a research laboratory as they are to the testing of light bulbs.

If information value theory and associated decision theoretic structures do not in the future occupy a large part of the education of engineers, then the engineering profession will find that its traditional role of managing scientific and economic resources for the benefit of man has been forfeited to another profession.

REFERENCES

- [1] R. A. Howard, "Bayesian decision models for system engineering," *IEEE Trans. on Systems Science and Cybernetics*, vol. SSC-1, pp. 36-40, November 1965.
- [2] ———, "Dynamic inference," *J. Operations Research*, vol. 13, pp. 712-733, September-October 1965.

VALUE OF INFORMATION LOTTERIES

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from IEEE Transactions on Systems Science and Cybernetics,
Vol. SSC-3, No. 1, June 1967. pp. 54-60.

Copyright © 1967 by Institute of Electronics and Electrical Engineers.

Value of Information Lotteries

RONALD A. HOWARD, SENIOR MEMBER, IEEE

Abstract—The essence of decision-making is understanding the economic impact of uncertainty. In this paper a previous discussion of information value theory is extended to illustrate how the availability of information on the uncertain factors of a problem affects the probability density function of profit, the profit lottery. A bidding problem serves to demonstrate the type of calculations required and their implications.

INTRODUCTION

A THEORY for the value of information, with particular emphasis on the expected value of information of various kinds, was discussed in a previous paper.⁽¹⁾ Here the discussion will be extended to illustrate the uncertainty in the value of information by developing

Manuscript received March 20, 1967. The research reported in this paper was partially supported by the Office of Naval Research under Contract N00014-67-A-0112-0010. Reproduction in whole or in part is permitted for any purpose of the United States Government.

The author is with the Division of Engineering-Economic Systems, Stanford University, Stanford, Calif.

probability distributions for the value. The mechanism for the discussion is the bidding example of Howard;⁽¹⁾ however, the statement and results of that example are repeated here to preserve the integrity of the present discussion.

NOMENCLATURE

The same special notation introduced in Howard⁽¹⁾ will be used here. Briefly, the definitions are as follows:

x	a random variable
A	an event
\mathcal{S}	the state of information on which probability assignments will be made
$\{x \mathcal{S}\}$	the density function of the random variable x given the state of information \mathcal{S}
$\{A \mathcal{S}\}$	the probability of the event A given the state of information \mathcal{S}
$\langle x \mathcal{S}\rangle$	the expectation of the random variable x , which equals $\int_2 x \{x \mathcal{S}\}$

- ξ the experience brought to the problem, the special state of information represented by total a priori knowledge
- $\{x|\xi\}$ the density function of a random variable x assigned on the basis of only a priori knowledge ξ ; designated the *prior* on x .

The notation does not emphasize the difference in probability assignment to a random variable and to an event because the context always makes clear the appropriate interpretation. The generalized summation symbol \int used in the definition of $\langle x|S \rangle$ should be interpreted as an integral if the random variable is continuous and as a summation if it is discrete.

THE BIDDING PROBLEM

The bidding problem⁽¹⁾ may be stated as follows: Our company is bidding on a contract against a number of competitors. We let p be our company's cost of performing on the contract; unfortunately, we are uncertain of this cost. We let ℓ be the lowest bid of the competitors; as one might expect, this too is uncertain. The problem is to determine b , our company's bid on the contract.

Naturally, the company will not win the contract if its bid b is higher than the lowest bid of the competitors ℓ ; in this case the company's profit is zero. However, if b is less than the lowest competitive bid ℓ , the company will obtain the contract and will make a profit equal to the difference between its bid b and its cost p . Thus profit v to the company as a function of p , ℓ , and b is defined by

$$v = \begin{cases} b - p & \text{if } b < \ell \\ 0 & \text{if } b > \ell \end{cases} \quad (1)$$

If p and ℓ were known, then all one would have to do would be to maximize v by adjusting b . However, because p and ℓ are random variables, v will be a random variable whose distribution will depend on b , the bid. This probability distribution of profit is called the profit lottery. The problem of selecting a bid is, then, the problem of selecting from among the several profit lotteries that can be produced by varying b , the one profit lottery that is regarded most desirable. Of course, the corresponding value of b would be the bid the company should make.

One way to attack the problem would be to assign a utility function that would produce a scalar index of desirability for any profit lottery. The value of b would then be determined by maximizing this index. However, rather than considering a general utility function, Howard⁽¹⁾ used a very special one: the straight line. This

utility function is appropriate to the risk-indifferent individual. The straight-line utility function results in measuring the desirability of a profit lottery by its expected profit. This criterion of desirability will be used throughout the present paper with the understanding that the results could be developed for any other utility function that




Fig. 1. Priors on cost of performance and lowest competitive bid.

TABLE I
OPTIMUM BIDS, EXPECTED PROFITS, AND EXPECTED VALUES OF CLAIRVOYANCE BASED ON EXPECTED VALUE MAXIMIZATION

State of Information	Optimum Bid	Expected Profit	Expected Increase in Value of this State of Information over ξ
ξ	$b = 5/4$	$\langle v \xi \rangle = 27/96$	—
C_p	$b = 1 + (p/2)$	$\langle v C_p\xi \rangle = 28/96$	$\langle v_{C_p} \xi \rangle = 1/96$
C_ℓ	$b = \begin{cases} \ell & \text{if } \ell > p \\ p = 1/2 & \text{if } \ell = 1/2 \\ \infty & \text{if } \ell < 1/2 \end{cases}$	$\langle v C_\ell\xi \rangle = 54/96$	$\langle v_{C_\ell} \xi \rangle = 27/96$
$C_{p\ell}$	$b = \begin{cases} \ell & \text{if } \ell > p \\ \infty & \text{if } \ell < p \end{cases}$	$\langle v C_{p\ell}\xi \rangle = 56/96$	$\langle v_{C_{p\ell}} \xi \rangle = 29/96$

might be desired. Thus, our problem was, and is, to find the value of bid b that maximizes expected profit.

To derive specific results, it is of course necessary to assign probability distributions to the cost of performance p and the lowest competitive bid ℓ . These probability distributions are assigned on the basis of prior information ξ . Assume that p and ℓ are independent and assign them the uniform probability distributions shown in Fig. 1. Note that the expected cost of performance p is $\bar{p} = 1/2$ and that the expected lowest competitive bid ℓ is $\bar{\ell} = 1$.

Previously⁽¹⁾ the expected profit was computed as a function of b based on the probability assignment, and then the value of b was found that would maximize the expected profit. The result is shown in Table I in the row corresponding to the state of information ξ . The optimum bid is $b = 5/4$; this optimum bid produces an expected profit $\langle v|\xi \rangle = 27/96$.

If no source of additional information were available, the problem would end at this point. However, even though a bidding strategy was devised that is optimum in the face of the uncertainties involved, the outcome could still be bad. The company may not get the contract; and if it does, it may lose money. It is reasonable, therefore, that if a perfect clairvoyant appeared and offered to eliminate one or both of the uncertainties in the problem, the company would be willing to offer him a financial

consideration. The question is how large should this nancial consideration be.

Let C represent clairvoyance and C_x represent clairvoyance about a random variable x . Thus any probability assignment conditional on C_x is conditional on the fact that the value of x is revealed. Define v_{C_x} as the increase in profit that arises from obtaining clairvoyance about x . It is clear that the expected increase in profit owing to clairvoyance about x , $\langle v_{C_x} | \xi \rangle$, is merely the difference between the expected profit that will be obtained as a result of the clairvoyance $\langle v | C_x \xi \rangle$ and the expected profit that we would obtain without clairvoyance $\langle v | \xi \rangle$; thus

$$\langle v_{C_x} | \xi \rangle = \langle v | C_x \xi \rangle - \langle v | \xi \rangle. \tag{2}$$

The expected profit given clairvoyance about x , $\langle v | C_x \xi \rangle$, is computed by evaluating the expected profit given that x , is known $\langle v | x \xi \rangle$, for each possible value of x that the clairvoyant might reveal and then summing this expectation with respect to the probability assignment on x , $\{x | \xi\}$,

$$\langle v | C_x \xi \rangle = \int_x \langle v | x \xi \rangle \{x | \xi\}. \tag{3}$$

The prior probability distribution $\{x | \xi\}$ for x is used in this calculation because, up to the moment that the clairvoyant actually reveals the value of x , the probability assigned to his statement about x is based only on prior knowledge ξ .

By using (2) and (3) the expected value in monetary units of eliminating any uncertainty in the problem can be assigned.

In Howard⁽¹⁾ three possible types of clairvoyance were considered: C_p , clairvoyance about cost of performance p ; C_ℓ , clairvoyance about the competitor's lowest bid ℓ ; and $C_{p\ell}$, clairvoyance about both of these quantities. The results of these considerations are summarized in Table I. It was found that if we had clairvoyance about p , the bid would be $b = 1 + p/2$, or 1 plus 1/2 the cost of performance predicted by the clairvoyant. The expected profit if one had this clairvoyance would be $\langle v | C_p \xi \rangle = 28/96$. Therefore, the increase in expected profit due to this clairvoyance instead of merely prior information is

$$\langle v_{C_p} | \xi \rangle = \langle v | C_p \xi \rangle - \langle v | \xi \rangle = 28/96 - 27/96 = 1/96.$$

When considering clairvoyance about the lowest competitive bid ℓ , it was found that one should not bid (or equivalently, bid a very large amount) if the lowest competitive bid ℓ were less than the expected cost of performance $\bar{p} = 1/2$, and that one should bid ℓ^- , slightly less than ℓ , if ℓ were greater than $\bar{p} = 1/2$. The expected profit with this clairvoyance is $\langle v | C_\ell \xi \rangle = 54/96$, an increase of $\langle v_{C_\ell} | \xi \rangle = 27/96$ over the profit to be expected based on only prior information. Note that the expected increase in profit due to clairvoyance about ℓ is 27 times that due to clairvoyance about p .

Finally, consider clairvoyance about p and ℓ jointly. In this case, it is clear that one should not bid if ℓ is less

than p and that one should bid ℓ^- if ℓ exceeds p . The expected profit with this clairvoyance is $\langle v | C_{p\ell} \xi \rangle = 56/96$; the expected increase in profit due to the clairvoyance rather than prior information alone is $\langle v_{C_{p\ell}} | \xi \rangle = 29/96$. Clairvoyance about both p and ℓ , then, is not worth much more than clairvoyance about ℓ alone. However, the forthcoming consideration of the probability distributions of profit and of the value of clairvoyance illustrates an important feature that separates the two types of clairvoyance.

THE UNCERTAINTY OF PROFIT

All considerations up to this time have been based on the expectation of profit and, in particular, on its maximization. However, only a rare decision-maker would be content to measure his ventures in terms of their expected returns alone; most would require in addition the probability distribution of profit to show the nature of the risk that must be borne. Therefore, the previous analysis will be supplemented by developing the probability distributions for profit for the various cases discussed.

To begin, a general expression for the probability distribution of profit as it depends on the state of information is written. Thus, if the state of information is S , the probability distribution of profit is given by expansion in the form

$$\{v | S\} = \int_p \int_\ell \{v | p, \ell, b(S), S\} \{p, \ell | S\}. \tag{4}$$

Here $b(S)$ indicates the optimal bid for the state of information S . The probability distribution for profit when p , ℓ , and the optimal bid are known, $\{v | p, \ell, b(S), S\}$, is an impulse distribution since the profit is uniquely determined from p, ℓ , and the bid. This is indicated by writing

$${}^D \langle v | p, \ell, b(S), S \rangle = \begin{cases} b - p & \text{if } b < \ell \\ 0 & \text{if } b > \ell \end{cases} \tag{5}$$

where the superscript D connotes that v is given deterministically by this expression.

The states of information S in (4) may be just prior experience ξ or prior experience coupled with one of the clairvoyant events. The effect of including a clairvoyant event is to change the form of the function $b(S)$ in (4). For the states of information S that are being considered, the joint distribution of p and ℓ , $\{p, \ell | S\}$ in (4) will always be the joint distribution based on only prior information $\{p, \ell | \xi\}$. For convenience this joint distribution is shown in Fig. 2; it is just the product of the two independent marginal distributions from Fig. 1. Thus, by using (4) and (5), the probability distribution of profit, the profit lottery, can be computed for any of the states of information previously considered.

We start by calculating the profit lottery based only on prior information ξ . Table I showed that the optimal bidding strategy for this case is to make the bid equal to




Fig. 2. Prior joint density function on cost of performance and lowest competitive bid.




Fig. 3. Profit lottery given prior information. (a) Sample space. (b) Profit function. (c) Profit density function. (d) Profit complementary cumulative distribution.

5/4. Then, from (5),

$$\begin{aligned}
 {}^D\langle v|p, \ell \xi \rangle &= {}^D\langle v|p, \ell, b(\xi), \xi \rangle = {}^D\langle v|p, \ell, b = 5/4, \xi \rangle = \\
 &\begin{cases} 5/4 - p & \text{if } 5/4 < \ell \\ 0 & \text{if } 5/4 > \ell. \end{cases} \quad (6)
 \end{aligned}$$

The calculation of the probability distribution for profit in this case $\{v|\xi\}$ is shown in Fig. 3. Figure 3(a) shows the functional form of the profit in the p, ℓ sample space; Fig. 3(b) shows its actual height point by point over the sample space. Since we thus have the profit as a function of p and ℓ and since we know the joint probability distribution of p and ℓ from Fig. 2, it is an elementary, if not trivial, application of probability mechanics to derive the density function of profit as indicated by (4). The resulting density function is given in Fig. 3(c). Note that the density function has an impulse of area 5/8 at zero corresponding to the 5/8 probability of zero profit. A zero profit is, of course, primarily a result of not winning the contract. The expected profit $\langle v|\xi \rangle$ is again 27/96 in agreement with Table I; however, we can now see that the expectation is far from an adequate description of the profit lottery.

Figure 3(d) shows the complementary cumulative probability distribution of profit, the probability that the profit exceeds any given value. The height of the com-

plementary cumulative distribution at any point is just the area under the density function to the right of that point. For positive random variables, the area under the complementary cumulative distribution is the mean of the random variable. The complementary cumulative distribution for profit is often a very convenient form in which to represent the profit lottery.

Thus we see that the strategy of bidding 5/4 involves the decision-maker in a highly risky venture. With probability 5/8, the profit is zero; however, it could be as high as 5/4. Further, it is impossible for the company to lose money.

UNCERTAINTY ASSOCIATED WITH CLAIRVOYANCE ABOUT COST

How would the profit lottery change if clairvoyance about p could be bought? Note from Table I that the optimal bid is $b = 1 + (p/2)$. Then (5) becomes

$$\begin{aligned}
 {}^D\langle v|p, \ell, C_p \xi \rangle &= {}^D\langle v|p, \ell, b(C_p \xi), C_p \xi \rangle \\
 &= {}^D\langle v|p, \ell, b = 1 + \frac{p}{2}, C_p \xi \rangle \quad (7) \\
 &= \begin{cases} 1 - \frac{p}{2} & \text{if } 1 + \frac{p}{2} < \ell \\ 0 & \text{if } 1 + \frac{p}{2} > \ell. \end{cases}
 \end{aligned}$$

Figure 4 illustrates the calculation of $\{v|C_p \xi\}$ following the pattern of Fig. 3. Thus the joint density function for p and ℓ from Fig. 2 is again used to find the profit lottery implied by the profit function from (7). There is still a 5/8 probability of a zero profit and a zero probability of a loss. The expected profit given clairvoyance about p , $\langle v|C_p \xi \rangle$, is 28/96, in agreement with Table I. Note that it is in fact impossible for the profit to equal 28/96 under these circumstances, because this point is not in the domain of the random variable. The actual profit must be either zero or a number between 1/2 and one. Observe that it is not possible to earn a profit larger than one when the bid is based on clairvoyance about p , although the firm could possibly earn a profit as high as 5/4 when acting on prior experience alone.

To examine further the effect of clairvoyance about p , consider the question of how much profit is increased as a result of having clairvoyant information. Since this increase is itself a random variable, it has a probability distribution that can be called the lottery for clairvoyance. Thus the lottery for clairvoyance about some variable or set of variables x given the prior information ξ is $\{v_{C_x}|\xi\}$ and is defined by

$$\{v_{C_x}|\xi\} = \iint \{v_{C_x}|p, \ell, \xi\} \{p, \ell|\xi\}. \quad (8)$$

Since v_{C_x} is known when p and ℓ are known, one can write

$${}^D\langle v_{C_x}|p, \ell, \xi \rangle = {}^D\langle v|p, \ell, C_x \xi \rangle - {}^D\langle v|p, \ell, \xi \rangle. \quad (9)$$

The value of clairvoyance at a particular point (p, ℓ) is




Fig. 4. Profit lottery given clairvoyance about the cost p . (a) Sample space. (b) Profit function. (c) Profit density function. (d) Profit complementary cumulative distribution.




Fig. 5. Lottery for clairvoyance about the cost p . (a) Sample space. (b) Clairvoyance value function. (c) Clairvoyance value density function. (d) Clairvoyance value complementary cumulative distribution.

just the increase in profit it produces at that point. Since v_{c_x} as a function of p and ℓ and the joint prior on p and ℓ are known, (8) can be used to derive $\{v_{c_x}|\xi\}$, the lottery for clairvoyance about x .

To apply this result to clairvoyance about cost p , write (9) as

$${}^D\langle v_{c_p}|p, \ell, \xi \rangle = {}^D\langle v|p, \ell, C_p \xi \rangle - {}^D\langle v|p, \ell, \xi \rangle. \quad (10)$$

Note that the value of clairvoyance about p at a point (p, ℓ) is just the difference in the profit functions of Figs. 4 and 3 at that point. This difference is shown in Fig. 5(a) and (b). Observe that the value of clairvoyance may be negative for some values of p and ℓ . When the joint prior on p and ℓ is used, the density function and complementary cumulative distribution is obtained for the value of clairvoyance about p shown in Fig. 5(c) and (d). The expected value of clairvoyance about p , $\langle v_{c_p}|\xi \rangle$ is $1/96$ as found in Table I. Note that there is a $19/32$ probability that the information received on p is worthless and a $7/32$ probability that it has a negative value. As the density function for v_{c_p} is examined, we see that the mean value of $1/96$ provides only the crudest indication of the value of clairvoyance about the cost.

UNCERTAINTY ASSOCIATED WITH CLAIRVOYANCE ABOUT THE LOWEST COMPETITIVE BID

To develop the profit lottery given clairvoyance about ℓ , recall that the optimal bidding strategy was to refrain from bidding if ℓ was less than the mean cost \bar{p} , and to bid ℓ^- otherwise. Since refraining from bidding is equivalent to bidding a very large number, (5) can be written as

$${}^D\langle v|p, \ell, C_\ell \xi \rangle = {}^D\langle v|p, \ell, b(C_\ell \xi), C_\ell \xi \rangle = {}^D\langle v|p, \ell, b = \ell^-, \ell > \bar{p} = 1/2, C_\ell \xi \rangle = \begin{cases} \ell - p & \text{if } 1/2 < \ell \\ 0 & \text{if } 1/2 > \ell. \end{cases} \quad (11)$$

Figure 6, then, shows the calculation of the profit lottery given clairvoyance about ℓ . The expected profit $\langle v|C_\ell \xi \rangle$ is $54/96$, as was first found in Table I. Observe that there is $1/4$ probability of a zero profit and a $1/16$ probability of a loss. Contrast this result with the profit lotteries of Figs. 3 and 4, where it was found that acting on prior information or clairvoyance about p could not possibly produce a loss. Thus, although clairvoyance about the lowest competitive bid ℓ promises greater profits than does clairvoyance about the cost of performance p , it also introduces the possibility of a loss.

To establish the lottery for clairvoyance about ℓ , write (9) in the form

$${}^D\langle v_{c_\ell}|p, \ell, \xi \rangle = {}^D\langle v|p, \ell, C_\ell \xi \rangle - {}^D\langle v|p, \ell, \xi \rangle. \quad (12)$$

The increase in profit due to clairvoyance about ℓ at the point (p, ℓ) is merely the difference in the profit functions of Figs. 6 and 3, a difference indicated in Fig. 7(a) and (b). Note that the value of clairvoyance about ℓ can take on negative values. When the joint density function for p and ℓ is introduced from Fig. 2, the density function and complementary cumulative distribution for the value of clairvoyance about ℓ shown in Fig. 7(c) and (d) are developed. The expected value of clairvoyance about ℓ ,




Fig. 6. Profit lottery given clairvoyance about the lowest competitive bid l . (a) Sample space. (b) Profit function. (c) Profit density function. (d) Profit complementary cumulative distribution.




Fig. 8. Profit lottery given clairvoyance about both p and l . (a) Sample space. (b) Profit function. (c) Profit density function. (d) Profit complementary cumulative distribution.




Fig. 7. Lottery for clairvoyance about the lowest competitive bid l . (a) Sample space. (b) Clairvoyance value function. (c) Clairvoyance value density function. (d) Clairvoyance value complementary cumulative distribution.

$\langle v_{cl} | \ell \rangle$, is 27/96, in accordance with Table I. With probability 1/4, the information on l is worthless, and with probability 1/16, of negative value. By comparing Figs. 5(c) and 7(c), we can see that the value of clairvoyance about the lowest competitive bid l is different from the value of clairvoyance about cost p in the shape of its distribution as well as in expected value.

UNCERTAINTY ASSOCIATED WITH CLAIRVOYANCE ABOUT BOTH COST AND THE LOWEST COMPETITIVE BID

Recall that when both the cost p and the lowest competitive bid l are known, the optimal bidding strategy is to refrain from bidding when l is less than p and otherwise to bid l . Thus (5) becomes

$$\begin{aligned}
 {}^D(v|p, l, C_{p\ell}) &= {}^D(v|p, l, b(C_{p\ell}), C_{p, \ell}) = \\
 {}^D(v|p, l, b = l, \ell > p, C_{p\ell}) & \\
 {}^D(v|p, l, b = \infty, \ell < p, C_{p\ell}) & \\
 &= \begin{cases} \ell - p & \text{if } p < \ell \\ 0 & \text{if } p > \ell. \end{cases} \quad (13)
 \end{aligned}$$

Figure 8(a) and (b) shows the nature of the profit function. By using the joint density function on p and l from Fig. 2, the profit density function and comprehensive cumulative distribution shown in Fig. 8(c) and (d) is derived. The probability of not winning the contract and, therefore, having profit zero is 1/4; the expected profit $\langle v | C_{p\ell} \rangle$ is 56/96, as in Table I. Note that clairvoyance about both p and l increases the maximum profit to 2 and eliminates any possibility of a loss.

The value of clairvoyance about both p and l is shown by (9) to be

$${}^D(v_{cpl} | p, \ell) = {}^D(v | p, C_{p\ell}) - {}^D(v | p, \ell). \quad (14)$$

This increase in profit at the point (p, ℓ) is the difference in the profit functions of Figs. 8 and 3 at that point. The resulting function for the value of clairvoyance about both p and l appears in Fig. 9(a) and (b). The density function and complementary cumulative distribution




Fig. 9. Lottery for clairvoyance about both p and l . (a) Sample space. (b) Clairvoyance value function. (c) Clairvoyance value density function. (d) Clairvoyance value complementary cumulative distribution.

for the value of clairvoyance in this case is obtained in the usual manner; they appear in Fig. 9(c) and (d). The expected value of clairvoyance about both p and l $\langle v_{c,p,l}|\xi \rangle$ is 29/96, as in Table I. Note that a negative value for clairvoyance about both p and l cannot exist.

CONCLUSION

The results that have been obtained show the inadequacy of characterizing profit lotteries by their ex-

pected values. Until such time as decision makers routinely use utility curves to express risk preference, one excellent alternative is to display profit lotteries of the type developed in this paper.

An elementary problem was used here to illustrate the concept of placing monetary values on information in a decision process. These results can be readily extended to the case where the information purchased is imperfect, but the essential methodology remains the same. We should point out that placing a value on a reduction of uncertainty is the first step in experimental design, for only when the worth of reducing uncertainty is known does one have a basis for allocating resources in experimentation designed to reduce the uncertainty.

However, in addition to illustrating the concept of the economic value of information, we have shown that even an elementary problem of this type may be far from trivial in the familiarity with probabilistic operations required to derive the results one would like to examine. In view of this observation, it is not surprising that so few managers and engineers make use of formal decision models in decision-making.

Yet it is inevitable that in the future both technical and managerial decision-makers will employ formal logical methods in decision-making. The transition probably will be painful.

REFERENCES

- [1] R. A. Howard, "Information value theory," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 22-26, August 1966.
- [2] —, "Bayesian decision models for system engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-1, pp. 36-40, November 1965.
- [3] —, "Dynamic inference," *Operations Research*, vol. 13, pp. 712-733, September-October 1965.
- [4] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory*. Cambridge, Mass.: Harvard University Press, 1961.

THE ECONOMIC VALUE
OF ANALYSIS AND COMPUTATION

James E. Matheson

Strategic Decisions Group

Reprinted with permission. IEEE Transactions on
Systems Science and Cybernetics, Vol. SSC-4, No. 3,
September 1968, pp. 325-332.

The Economic Value of Analysis and Computation

JAMES E. MATHESON, MEMBER, IEEE

Abstract—This paper shows how the decision analysis approach can be used to determine the most economic method of carrying out computations or analyses. A primary decision problem is first formulated to obtain a structure for the analysis. Then several computational or analytical procedures, which can be used to analyze the primary decision problem in greater detail, are evaluated to select the most economic procedure. The purpose of each of these procedures is to increase the available information about uncertain parameters before making the primary decision, thereby yielding a “better” decision. Each procedure is evaluated by combining the value structure of the primary decision problem with a model of that procedure. The procedures considered in this paper are clairvoyance, complete analysis, Monte Carlo analysis, and numerical analysis. An example of a bidding problem is used to illustrate the results.




Fig. 1. Trees for the primary decision problem.

The symbol \int is used as a generalized summation symbol representing either an integral or a discrete summation, depending on whether the random variable is continuous or discrete.

INTRODUCTION

MANY computations and analyses are carried out in support of decision-making processes. The primary value of these computations is derived from improvement in the decision that follows from the computational results. Usually the extensiveness of such an analysis is determined on the basis of intuitive judgment and budgetary constraints, rather than on an economic analysis of the value of the computational results to the primary decision problem.

The purpose of this paper is to show how computational or analytical procedures can be economically evaluated and to illustrate the results with an example. In order to evaluate the procedures, it will also be necessary to show how the results of each procedure can be taken into account in the decision policy.

NOTATION

The following notation will be used to describe probability assignments:

- x \triangleq a random variable,
- S \triangleq the state of information on which probability assignments will be made,
- $\{x|S\}$ \triangleq the density function of the random variable x given the state of information S ,
- $\{x \leftarrow y|S\}$ \triangleq the above density function evaluated at y , and
- $\langle x|S \rangle$ $\triangleq \int x \{x|S\} \triangleq$ the expectation of the random variable x given the state of information S .

Manuscript received March 25, 1968. This paper was presented at the IEEE Systems Science and Cybernetics Conference, Washington, D. C., October 17, 1966.

The author is with the Stanford Research Institute, Menlo Park, Calif. 94025.

THE PRIMARY DECISION PROBLEM

The additional notation used to describe the primary decision problem is as follows:

- E \triangleq the experience or prior knowledge brought to the problem,
- a \triangleq the action variable, the primary decision is to select the value of a ,
- o \triangleq the outcome random variable,
- $V(a, o)$ \triangleq the value function that assigns a value to each action–outcome pair,
- v \triangleq the value random variable,
- $\hat{a}(v|S)$ \triangleq the value of a that maximizes $\langle v|a, S \rangle$, and
- $a^*(v|S)$ \triangleq the intention to use the decision rule $\hat{a}(v|S)$ at that point in time when the information symbolized by S becomes available.

The structure of a decision problem is most easily represented by drawing both the decision tree and nature’s tree for the problem (Howard [1]). Fig. 1 illustrates these trees for the primary decision problem. The crossed nodes indicate decision points and the uncrossed nodes indicate chance points.

In this paper the term outcome is used to signify the revelation of a “state-of-nature,” and as such the outcome is not influenced by the action. This assumption is represented in nature’s tree shown in Fig. 1. However, the value of the outcome is determined by both the action and the outcome, and thus information about the outcome is useful in determining the best action. Models in which the action influences the outcome may be incorporated, essentially by changing nature’s tree.

Assuming the decision maker would like to maximize his expected value, the best decision is given by action $\hat{a}(v|E)$. This action is the one that maximizes

$$\langle v|a, E \rangle = \int_o \langle v|a, o, E \rangle \{o|a, E\}. \quad (1)$$

By the definition of the value function

$$\langle v|a, o, E \rangle = V(a, o) \tag{2}$$

so that

$$\langle v|a, E \rangle = \int_o V(a, o) \{o|E\}. \tag{3}$$

It is assumed that the outcome is independent of the action so

$$\{o|a, E\} = \{o|E\}. \tag{4}$$

Thus the decision maker obtains the expected value

$$\begin{aligned} \langle v|a^*(v|E), E \rangle &= \langle v|\hat{a}(v|E), E \rangle \\ &= \max_a \int_o V(a, o) \{o|E\}. \end{aligned} \tag{5}$$

In (5) the distinction between $a^*(v|E)$, the intention to use $\hat{a}(v|E)$, and $\hat{a}(v|E)$ itself is not critical. The importance of this distinction will become apparent in the cases to follow.

THE BIDDING PROBLEM

The results will be illustrated by a competitive bidding problem, which was introduced by Howard.¹ We will play the part of a decision analyst for a company that has the opportunity to place a bid for a contract. We then assume the company wishes to maximize its expected value of performance on the contract, less any costs it has to pay for further analyses or computations that we undertake. The variables in this problem are

- $b \triangleq$ our bid (our action),
- $l \triangleq$ the lowest bid of our competitors (a random variable),
- $p \triangleq$ our cost of performance on the contract (a random variable), and
- $v \triangleq$ the value of performance on the contract.

We assume that our cost and the lowest competitive bid are independent random variables that are also independent of our bid. That is,

$$\{p, l|b, E\} = \{p|E\} \{l|E\}. \tag{6}$$

We further assume that according to our prior information (E) the distribution of our cost is uniform between 0 and 1, and that the distribution of the lowest competitive bid is uniform between 0 and 2, as shown in Fig. 2.

If our bid is less than that of our competitors, $b < l$, we receive the contract and make a profit of $b - p$. Otherwise, we lose the contract and make nothing. Thus our value function is

$$V(b, p, l) = \begin{cases} b - p & \text{if } b < l \\ 0 & \text{if } b \geq l. \end{cases} \tag{7}$$

¹ Howard [2] introduced this bidding problem and presented analysis of the value of clairvoyance.

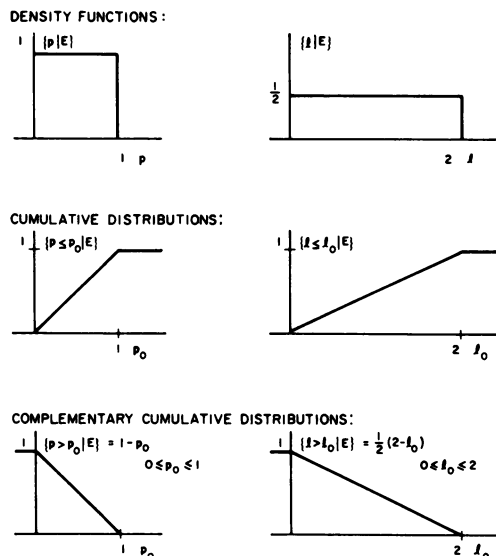


Fig. 2. Priors on cost of performance and lowest competitive bid.




Fig. 3. Decision tree for the primary bidding problem.

Our bid b is our action and the p, l pair is our outcome. Since in this paper only p will be treated in detail, the role of l will be suppressed by combining its effect into the value function, and creating a new value function with parameters b and p . We do this by writing

$$\begin{aligned} \langle v|b, p, S \rangle &= \int_l V(b, p, l) \{l|b, p, S\} \\ &= \int_{l=b}^{\infty} (b - p) \{l|b, p, S\} \\ &= (b - p) \{l > b|b, p, S\} \end{aligned} \tag{8}$$

where S represents E along with any additional conditions that we wish to consider. Since l is independent of b and p and of any additional variables to be used in this paper,

$$\{l > b|b, S\} = \{l > b|b, E\} = \frac{1}{2} (2 - b) \tag{9}$$

so

$$\langle v|b, p, S \rangle = \frac{1}{2} (b - p)(2 - b). \tag{10}$$

For the remainder of the paper, the role of l will be suppressed by using the value function

$$V(b, p) = \langle v|b, p, S \rangle = \frac{1}{2} (b - p)(2 - b) \tag{11}$$

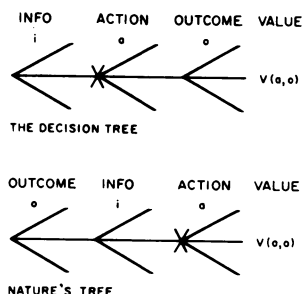


Fig. 4. Trees for additional information.

and regarding b as the action variable and p as the outcome (Fig. 3). The best bid $\hat{b}(v|E)$ is the bid maximizing

$$\begin{aligned} \langle v|b, E \rangle &= \int_p V(b, p)\{p|E\} \\ &= \int_p \frac{1}{2}(b - p)(2 - b)\{p|E\} \\ &= \frac{1}{2}(b - \langle p|E \rangle)(2 - b). \end{aligned} \tag{12}$$

Thus by differentiation we can determine that

$$\hat{b}(v|E) = 1 + \frac{1}{2}\langle p|E \rangle \tag{13}$$

and

$$\langle v|b^*(v|E), E \rangle = \frac{1}{2} \left(1 - \frac{1}{2}\langle p|E \rangle \right)^2. \tag{14}$$

For the prior on cost given in Fig. 2,

$$\langle p|E \rangle = \frac{1}{2} \tag{15}$$

so that

$$\hat{b}(v|E) = 1.25 \tag{16}$$

and

$$\langle v|b^*(v|E), E \rangle = 0.28125. \tag{17}$$

THE VALUE OF INFORMATION

After a primary decision problem has been structured utilizing the available prior knowledge E , it usually is possible to obtain more knowledge or information, which will increase the expected value of the decision at some additional cost. Let

i = additional information (random variable).

In this case, the primary decision is postponed until the information is received, and then the best decision is made utilizing the information. This sequence is illustrated in the decision tree in Fig. 4.

Our prior knowledge about the information is encoded as a probability distribution of the information, given an

outcome and prior knowledge, $\{i|o, E\}$. This corresponds to nature's tree in Fig. 4. Now

$$\langle v|i, a, E \rangle = \int_o \langle v|i, a, o, E \rangle \{o|i, a, E\}. \tag{18}$$

The assumption that

$$\{o|i, a, E\} = \{o|i, E\} \tag{19}$$

the insertion of the value function, and the application of Bayes' formula yield

$$\langle v|i, a, E \rangle = \int_o V(a, o) \frac{\{i|o, E\}\{o|E\}}{\{i|E\}}. \tag{20}$$

The best action $\hat{a}(v|i, E)$ is the one that maximizes the above expression, which yields

$$\langle v|i, a^*(v|i, E), E \rangle = \langle v|i, \hat{a}(v|i, E), E \rangle = \max_a \langle v|i, a, E \rangle. \tag{21}$$

Prior to receiving the information, the expected value given this procedure is

$$\begin{aligned} \langle v|a^*(v|i, E), E \rangle &= \int_i \{i|E\} \langle v|i, a^*(v|i, E), E \rangle \\ &= \int_i \{i|E\} \max_a \int_o V(a, o) \frac{\{i|o, E\}\{o|E\}}{\{i|E\}} \end{aligned} \tag{22}$$

where it is assumed that

$$\{i|a^*(v|i, E), E\} = \{i|E\}. \tag{23}$$

In some cases it is convenient to eliminate the factors $\{i|E\}$, yielding

$$\langle v|a^*(v|i, E), E \rangle = \int_i \max_a \int_o V(a, o) \{i|o, E\}\{o|E\}. \tag{24}$$

Equations of the form (24) can be written by inspection by writing integrations and maximizations in the order of the decision tree, followed by the value function from the tips of the tree, multiplied by the joint density of the random variables given the decision variables. The latter term may be factored using conditional probability expressions in the order of nature's tree and simplified on the basis of independence assumptions. Equations in the form of (22) retain the natural "updating" of the probability distributions, while equations in the joint density form of (24) are sometimes easier to derive. For a given problem, either form may turn out to be more straightforward for analysis.

VALUE OF CLAIRVOYANCE

One particular type of information is clairvoyance about the outcome. In this case i is a discrete random variable with probability distribution defined by

$$\{i = o|o, E\} = 1. \tag{25}$$

That is, i will have the value of o with certainty. Combining (24) with (25) yields

$$\begin{aligned}
 \langle v|a^*(v|i, E), E \rangle &= \int_i \max_a \int_o V(a, o) \{i|o, E\} \{o|E\} \\
 &= \int_i \max_a V(a, i) \{o \leftarrow i|E\} \\
 &= \int_o \max_a V(a, o) \{o|E\}. \quad (26)
 \end{aligned}$$

The optimum action $\hat{a}(v|i, E)$ is therefore the one that maximizes $V(a, i) \{o \leftarrow i|E\}$.

Clairvoyance about p

In the bidding problem, if we consider obtaining clairvoyance about our cost of performance p , we must make the bid that maximizes

$$V(b, i) \{p \leftarrow i|E\} = \frac{1}{2} (b - i)(2 - b). \quad (27)$$

This bid is given by

$$\hat{b}(v|i, E) = 1 + \frac{1}{2} i \quad (28)$$

and for this bid

$$V\left(1 + \frac{1}{2} i, i\right) \{p \leftarrow i|E\} = \frac{1}{2} \left(1 - \frac{1}{2} i\right)^2. \quad (29)$$

Applying (26) yields

$$\begin{aligned}
 \langle v|b^*(v|i, E), E \rangle &= \int_0^1 \frac{1}{2} \left(1 - \frac{1}{2} p\right)^2 dp = \frac{7}{24} \\
 &= 0.291667. \quad (30)
 \end{aligned}$$

THE VALUE OF A COMPLETE ANALYSIS

In the primary decision problem, the distribution of outcome given the prior knowledge $\{o|E\}$ was used in arriving at the best action. Suppose that this analysis was preliminary analysis of the problem using only easily available knowledge and techniques to arrive at $\{o|E\}$. After the primary analysis, it is proposed that a more complete analysis be made; for example, by structuring the process of outcome generation in finer detail and seeking out the best experts to obtain data for each detail. Hopefully this would lead to a narrower distribution of outcome that would in turn lead to better action, so that the decision process would have higher expected value. This section deals with the value of such a complete analysis.

The Model of a Complete Analysis

A complete analysis will, in general, lead to a new distribution of outcome. Let

- d \triangleq an index over possible distributions of outcome resulting from the complete analysis and
- $\{o|d, E\}$ \triangleq the density function of outcome given the value of the index d and our prior knowledge E .




Fig. 5. Densities for a complete analysis.

The preceding densities encode our prior knowledge of the possible results of the complete analysis. The entire complete analysis is viewed as the determination of the value of the index d . In this view, the complete analysis is the identification of one of the densities from the above family as the result of the analysis.

In addition, the prior probability distribution over the preceding densities must be specified. Thus the probability distribution of occurrence of each possible resulting distribution is encoded by

$$\{d|E\} = \text{the density function of index } d \text{ given the experience or prior knowledge } E.$$

For consistency with the primary analysis, it is required that

$$\{o|E\} = \int_d \{o|d, E\} \{d|E\}. \quad (31)$$

Fig. 5 illustrates the modeling of the complete analysis.

The Decision Equations for a Complete Analysis

The decision tree and nature's tree for a complete analysis are given in Fig. 6.

The best action $\hat{a}(v|d, E)$ is the one that maximizes

$$\langle v|d, a, E \rangle = \int_o \langle v|d, a, o, E \rangle \{o|d, a, E\}. \quad (32)$$

It is assumed that

$$\{o|d, a, E\} = \{o|d, E\}. \quad (33)$$

Using (33) and inserting the value function, yields

$$\langle v|d, a, E \rangle = \int_o V(a, o) \{o|d, E\}. \quad (34)$$

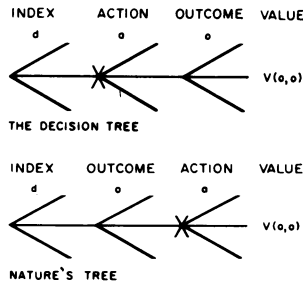


Fig. 6. Trees for a complete analysis.




Fig. 7. Prior densities for a complete analysis of p .




Fig. 8. Decision tree for a complete analysis of p .

So

$$\begin{aligned} \langle v|a^*(v|d, E), d, E \rangle &= \langle v|\hat{a}(v|d, E), d, E \rangle \\ &= \max_a \int_o V(a, o)\{o|d, E\}. \end{aligned} \quad (35)$$

Prior to performing the complete analysis, the expected value is

$$\begin{aligned} \langle v|a^*(v|d, E), E \rangle &= \int_d \{d|E\} \langle v|a^*(v|d, E), d, E \rangle \\ &= \int_d \{d|E\} \max_a \int_o V(a, o)\{o|d, E\} \end{aligned} \quad (36)$$

assuming of course that

$$\{d|a^*(v|d, E), E\} = \{d|E\}. \quad (37)$$

Complete Analysis of p

Suppose that a complete analysis of our cost of performance p will result in one of two possible distributions for p with equal probability, as given in Fig. 7. The decision tree for a complete analysis in the notation of the bidding problem is given in Fig. 8. From (34) and (11), we see that in order to find $\hat{b}(v|d, E)$ we must maximize

$$\langle v|d, b, E \rangle = \int_p \frac{1}{2} (b - p)(2 - b)\{p|d, E\} \quad (38)$$

over our bid b . Noting that this (38) is identical to (12) with E replaced by d, E , we can immediately write

$$\hat{b}(v|d, E) = 1 + \frac{1}{2} \langle p|d, E \rangle \quad (39)$$

and

$$\langle v|b^*(v|d, E), d, E \rangle = \frac{1}{2} \left(1 - \frac{1}{2} \langle p|d, E \rangle \right)^2. \quad (40)$$

Finally the application of (36) yields

$$\langle v|b^*(v|d, E), E \rangle = \int_d \{d|E\} \frac{1}{2} \left(1 - \frac{1}{2} \langle p|d, E \rangle \right)^2. \quad (41)$$

From Fig. 7 we see that

$$\langle p|d, E \rangle = \begin{cases} \frac{2}{3} & \text{if } d = 1 \\ \frac{1}{3} & \text{if } d = 2 \end{cases} \quad (42)$$

and

$$\{d|E\} = \frac{1}{2} \quad d = 1, 2. \quad (43)$$

Hence our bidding policy is

$$\hat{b}(v|d, E) = \begin{cases} \frac{8}{6} = 1.3333 & \text{if } d = 1 \\ \frac{7}{6} = 1.1666 & \text{if } d = 2 \end{cases} \quad (44)$$

and the value of the complete analysis is

$$\begin{aligned} \langle v|b^*(v|d, E), E \rangle &= \frac{1}{4} \left[1 - \frac{1}{2} \left(\frac{2}{3} \right) \right]^2 + \frac{1}{4} \left[1 - \frac{1}{2} \left(\frac{1}{3} \right) \right]^2 \\ &= \frac{41}{144} = 0.284722. \end{aligned} \quad (45)$$

THE VALUE OF A MONTE CARLO ANALYSIS

In many cases, the cost of carrying out a complete analysis is extremely high, due to the complexity and difficulty of performing the analysis. However, it may be much less expensive to put the problem in the form of a complete analysis, gather the relevant data, and then generate a number of sample outcomes using Monte Carlo simulation techniques. This procedure will not be as valuable as a complete analysis, but it may attain most of the value at a much lower cost.

This section deals with the value of a Monte Carlo analysis as a function of the number of samples taken, and the rules by which the samples should influence the choice of action.

The Model of a Monte Carlo Analysis

The prior probability structure for the Monte Carlo analysis is identical to that of the complete analysis, which is illustrated in Fig. 5. It is assumed that nature selects a value of d which would result from a complete analysis if it were carried out, and that each Monte Carlo sample is then drawn independently from $\{o|d, E\}$.

The Decision Equations for a Monte Carlo Analysis

The decision tree and nature's tree for a Monte Carlo analysis are presented in Fig. 9, where

- $N \triangleq$ the number of independent samples taken,
- $o_j \triangleq$ the j th sample value, and
- $\mathbf{o} \triangleq (o_1, o_2, \dots, o_N) =$ the vector of the N sample values.

At the action point in the decision tree, the best action $\hat{a}(v|N, \mathbf{o}, E)$ is the action which maximizes

$$\langle v|N, \mathbf{o}, a, E \rangle = \int_d \int_o V(a, \mathbf{o}) \cdot \{o|N, \mathbf{o}, a, d, E\} \cdot \{d|N, \mathbf{o}, a, E\}. \quad (46)$$

It is assumed that

$$\{o|N, \mathbf{o}, a, d, E\} = \{o|d, E\} \quad (47)$$

that is, o is generated by an independent sample from $\{o|d, E\}$.

The use of Bayes' formula yields

$$\{d|N, \mathbf{o}, a, E\} = \frac{\{o|N, a, d, E\} \{d|N, a, E\}}{\{o|N, a, E\}}. \quad (48)$$

It is also assumed that

$$\{d|N, a, E\} = \{d|E\} \quad (49)$$

$$\{o|N, a, d, E\} = \{o|N, d, E\} \quad (50)$$

$$\{o|N, a, E\} = \{o|N, E\}. \quad (51)$$

Since the samples are independent

$$\{o|N, d, E\} = \prod_{j=1}^N \{o \leftarrow o_j|d, E\} \quad (52)$$

the combination of (46) through (52) yields

$$\langle v|N, \mathbf{o}, a, E \rangle = \int_d \int_o V(a, \mathbf{o}) \cdot \{o|d, E\} \frac{\{d|E\} \prod_{j=1}^N \{o \leftarrow o_j|d, E\}}{\{o|N, E\}}. \quad (53)$$

Also

$$\langle v|a^*(v|N, \mathbf{o}, E), N, \mathbf{o}, E \rangle = \max_a \langle v|N, \mathbf{o}, a, E \rangle. \quad (54)$$

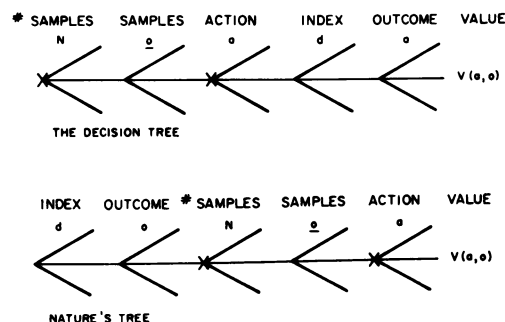


Fig. 9. Trees for a Monte Carlo analysis.




Fig. 10. Decision tree for a Monte Carlo analysis of p .

Thus the equation for the expected value at the sample node in the decision tree is

$$\begin{aligned} \langle v|a^*(v|N, \mathbf{o}, E), N, E \rangle &= \int_o \{o|N, E\} \langle v|a^*(v|N, \mathbf{o}, E), N, \mathbf{o}, E \rangle \\ &= \int_o \{o|N, E\} \max_a \int_d \int_o V(a, \mathbf{o}) \{o|d, E\} \\ &\quad \cdot \frac{\{d|E\} \prod_{j=1}^N \{o \leftarrow o_j|d, E\}}{\{o|N, E\}}. \quad (55) \end{aligned}$$

This equation can either be left in the above form or the factors $\{o|N, E\}$ can be canceled out leaving

$$\langle v|a^*(v|N, \mathbf{o}, E), N, E \rangle = \int_o \max_a \int_d \int_o V(a, \mathbf{o}) \{o|d, E\} \cdot \{d|E\} \prod_{j=1}^N \{o \leftarrow o_j|d, E\}. \quad (56)$$

The calculations may be carried out with either (55) or (56). The best choice of N depends on the cost of samples. If c is the cost of the N samples, then the optimum number of samples is denoted

$$\hat{N}(v - c|a^*(v|N, \mathbf{o}, E), E).$$

In the example the number of samples will be left as a parameter until the summary section. The formal equations for determining the best N are left to the interested reader.

The Value of a Monte Carlo Analysis of p

The decision tree for a Monte Carlo analysis of our cost of performance is given in Fig. 10. Applying (53) and (11), we find that $\hat{b}(v|N, \mathbf{p}, E)$ is the bid which maximizes

$$\begin{aligned} \langle v|N, \mathbf{p}, b, E \rangle &= \int_d \int_p \frac{1}{2} (b - p)(2 - b) \\ &\quad \cdot \{p|d, E\} \frac{\{d|E\} \prod_{j=1}^N \{p \leftarrow p_j|d, E\}}{\{\mathbf{p}|N, E\}} \\ &= \frac{1}{2} (b - \langle p|N, \mathbf{p}, E \rangle)(2 - b) \end{aligned} \quad (57)$$

where

$$\langle p|N, \mathbf{p}, E \rangle = \int_d \langle p|d, E \rangle \frac{\{d|E\} \prod_{j=1}^N \{p \leftarrow p_j|d, E\}}{\{\mathbf{p}|N, E\}}. \quad (58)$$

Evaluating the various quantities, we have

$$\langle p|d, E \rangle = \begin{cases} \frac{2}{3} & \text{if } d = 1 \\ \frac{1}{3} & \text{if } d = 2 \end{cases} \quad (59)$$

$$\{d|E\} = \frac{1}{2}, \quad d = 1, 2 \quad (60)$$

$$\prod_{j=1}^N \{p = p_j|d, E\} = \begin{cases} 2^N \pi_1 & \text{if } d = 1 \\ 2^N \pi_2 & \text{if } d = 2 \end{cases} \quad (61)$$

where

$$\pi_1 = \prod_{j=1}^N p_j, \quad \pi_2 = \prod_{j=1}^N (1 - p_j) \quad (62)$$

$$\begin{aligned} \{\mathbf{p}|N, E\} &= \sum_{d=1}^2 \{d|E\} \prod_{j=1}^N \{p \leftarrow p_j|d, E\} \\ &= 2^{N-1} (\pi_1 + \pi_2). \end{aligned} \quad (63)$$

So

$$\begin{aligned} \langle p|N, \mathbf{p}, E \rangle &= \left(\frac{2}{3}\right) \frac{\left(\frac{1}{2}\right)^{2^N \pi_1}}{2^{N-1}(\pi_1 + \pi_2)} + \left(\frac{1}{3}\right) \frac{\left(\frac{1}{2}\right)^{2^N \pi_2}}{2^{N-1}(\pi_1 + \pi_2)} \\ &= \left(\frac{1}{3}\right) \frac{2\pi_1 + \pi_2}{\pi_1 + \pi_2}. \end{aligned} \quad (64)$$

Recognizing the form of (57) is that of (12), we have immediately

$$\begin{aligned} \hat{b}(v|N, \mathbf{p}, E) &= 1 + \left(\frac{1}{2}\right) \langle p|N, \mathbf{p}, E \rangle \\ &= 1 + \left(\frac{1}{6}\right) \frac{2\pi_1 + \pi_2}{\pi_1 + \pi_2} \end{aligned} \quad (65)$$

and

$$\begin{aligned} \langle v|b^*(v|N, \mathbf{p}, E), N, \mathbf{p}, E \rangle &= \frac{1}{2} \left(1 - \frac{1}{2} \langle p|N, \mathbf{p}, E \rangle\right)^2 \\ &= \left(\frac{1}{2}\right) \left[1 - \left(\frac{1}{6}\right) \frac{2\pi_1 + \pi_2}{\pi_1 + \pi_2}\right]^2. \end{aligned} \quad (66)$$

Inserting this expression in (55) yields

$$\begin{aligned} \langle v|b^*(v|N, \mathbf{p}, E), N, E \rangle &= \int_p \left(\frac{1}{2}\right) \left[1 - \left(\frac{1}{6}\right) \frac{2\pi_1 + \pi_2}{\pi_1 + \pi_2}\right]^2 \{\mathbf{p}|N, E\} \\ &= \int_p \left(\frac{1}{2}\right) \left[1 - \left(\frac{1}{6}\right) \frac{2\pi_2 + \pi_2}{\pi_1 + \pi_2}\right]^2 2^{N-1} (\pi_1 + \pi_2). \end{aligned} \quad (67)$$

The Bidding Policy

The bidding policy to be employed after the samples have been observed was given in (65). For $N = 1$

$$\hat{b}(v|N = 1, \mathbf{p}, E) = 1 + \frac{1}{6} (p_1 + 1) \quad (68)$$

which ranges from 1.1666 to 1.3333, and

$$\begin{aligned} \langle v|b^*(v|N, \mathbf{p}, E), N = 1, E \rangle &= \int_{p_1} \left(\frac{1}{2}\right) \left[1 - \left(\frac{1}{6}\right) (p_1 + 1)\right]^2 \\ &= \frac{61}{216} = 0.282407. \end{aligned} \quad (69)$$

For $N = 2$ and 3, the expected value increases to 0.28308 and 0.28352, respectively. When many samples are taken, if $d = 1$ the ratio π_2/π_1 will approach zero, while if $d = 2$ the ratio π_1/π_2 will approach zero. Thus

$$\hat{b}(v|N \rightarrow \infty, \mathbf{p}, E) \rightarrow \begin{cases} 1 + \left(\frac{1}{6}\right) \cdot \left(\frac{2}{1}\right) = 1.3333 & \text{if } d = 1 \\ 1 + \left(\frac{1}{6}\right) \cdot \left(\frac{1}{1}\right) = 1.1666 & \text{if } d = 2. \end{cases} \quad (70)$$

In this limiting case, we are actually identifying the value of d that nature has selected. Thus the bidding strategy approaches that of the complete analysis and the expected value approaches the expected value of a complete analysis, 0.284722. Fig. 11 shows how the expected value of a Monte Carlo analysis increases as a function of N , expressed as a percentage of the expected value gain from a complete analysis; that is,




Fig. 11. Expected value increase of a Monte Carlo analysis as a percentage of that of a complete analysis.

$$\frac{\langle v|b^*(v|N, p, E), E \rangle - \langle v|b^*(v|E), E \rangle}{\langle v|b^*(v|d, E), E \rangle - \langle v|b^*(v|E), E \rangle} \times 100 \text{ percent.}$$

This curve reaches 99 percent at approximately $N = 15$. Thus for this bidding problem, only a small number of Monte Carlo samples are required to capture most of the value of a complete analysis.

THE VALUE OF NUMERICAL ANALYSIS

A numerical analysis can be evaluated using the same approach as that of the previous examples. If the numerical analysis is a very refined one, then it can be treated as a complete analysis. If the numerical analysis is a less refined one, then a model must be established for the information which might be generated by the numerical analysis. This model will basically determine the probability distribution of the index d , and by inference, the probability distribution of the outcome given the results of the numerical analysis. In the spirit of the Monte Carlo analysis, it is possible to model and determine the expected value of a numerical analysis as a function of the "degree of refinement" of the analysis in order to select the most economical analysis.

SUMMARY OF THE BIDDING PROBLEM

Suppose the values we have been dealing with in the bidding problem are millions of dollars. This makes the primary decision problem have an expected value of \$281 250. On the cost side, suppose that we can buy clairvoyance by building a prototype at a cost of \$20 000, that we can buy a complete analysis for a cost of \$3000, and that we can buy a Monte Carlo analysis for \$400

TABLE I
DETERMINATION OF THE BEST ANALYSIS

Analysis	Expected Value	-	Additional Cost	=	Net Expected Value
Primary					
Decision problem	\$281 250	-	\$ 0000	=	\$281 250
Clairvoyance	\$291 667	-	\$20 000	=	\$271 667
Complete analysis	\$284 722	-	\$ 3000	=	\$281 722
Monte Carlo analysis:					
$N = 1$	\$282 400	-	\$ 750	=	\$281 650
$N = 2$	\$283 080	-	\$ 1100	=	\$281 980
$N = 3$	\$283 520	-	\$ 1450	=	\$282 070
$N = 4$	\$283 830	-	\$ 1800	=	\$282 030
$N = 5$	\$284 060	-	\$ 2150	=	\$282 010

plus \$350 per sample. Table I lists expected value, cost, and expected value minus cost for each possibility considered. The highest net expected value of \$282 070 is achieved by a Monte Carlo analysis with $N = 3$.

CONCLUSIONS

This paper has shown how the decision analysis approach can be used to determine the economic value of computations or analyses of various levels of complexity. The key to the selection of a computational procedure is the recognition that the procedure obtains its value from the influence of its results on the decision problem it is serving. This approach can also be applied to other methods of reducing uncertainty, such as conducting experiments or research programs.

This paper also provides an indirect answer to the question, "What is meant by a good prior probability distribution?" Instead of wondering whether a prior is good or bad, we can determine the economic value of using available procedures to improve our state of knowledge and obtain a "better" prior distribution.

ACKNOWLEDGMENT

The author wishes to express his thanks to Prof. R. A. Howard for several stimulating discussions during the development of this paper.

REFERENCES

- [1] R. A. Howard, "Bayesian decision models for system engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-1, pp. 36-40, November 1965.
- [2] —, "Information value theory," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 22-26, August 1966.
- [3] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory*. Boston, Mass.: Graduate School of Business, Harvard University, 1961.

COMPETITIVE BIDDING IN HIGH-RISK SITUATIONS

E. C. Capen

R. V. Clapp

W. M. Campbell

Atlantic Richfield Co.



Competitive Bidding in High-Risk Situations

E. C. Capen, Atlantic Richfield Co.

R. V. Clapp, Atlantic Richfield Co.

W. M. Campbell, Atlantic Richfield Co.

Introduction

We would like to share with you our thoughts on the theory of competitive bidding. It is a tough business. We are not sure we understand as much as we ought to about the subject. As in most scientific endeavors, we think there is more knowledge to be gained by talking with others than by keeping quiet.

Our first attempt at actually using a probability model approach to bidding was in 1962. We borrowed heavily from Lawrence Friedman's fine paper on the subject.¹ But the further our studies went, the more problems we noticed for our particular application. We decided to strike out on our own. By 1965 we had our model just as it is today. But having a model and completely understanding its workings are not the same thing. We are still learning.

While we refer to the "model" as though it were some inanimate object, it is not. What we want to describe to you is a system for taking the best judgments of people — properly mixed, of course, with historical evidence — and putting those judgments together in a rational way so they may be used to advantage.

Lest the reader be too casual, thinking that since he is not personally involved in lease sales he need not pay the closest attention, we offer this thought. There is a somewhat subtle interaction between competition and property evaluation, and this phenomenon — this culprit — works quietly within and without the specific lease sale environment. We would

venture that many times when one purchases property it is because someone else has already looked at it and said, "Nix." The sober man must consider, "Was he right? Or am I right?" The *method of analysis* we will describe is strictly for sealed bid competitive lease sales, but the *phenomenon* we will be talking about pervades all competitive situations.

Industry's Record in Competitive Bidding

In recent years, several major companies have taken a rather careful look at their records and those of the industry in areas where sealed competitive bidding is the method of acquiring leases. The most notable of these areas, and perhaps the most interesting, is the Gulf of Mexico. Most analysts turn up with the rather shocking result that, while there seems to be a lot of oil and gas in the region, the industry probably is not making as much return on its investment there as it intended.²⁻⁵ In fact, if one ignores the era before 1950, when land was a good deal cheaper, he finds that the Gulf has paid off at something less than the local credit union.

Why? Have we been poor estimators of hydrocarbon potential? Have our original cost estimates been too conservative? Have we not predicted allowables well? Was our timing off? Or have we just been unlucky?

It is our view that none of the factors these questions suggest has been the major cause of the in-

If it is true, as common sense tells us, that a lease winner tends to be the bidder who most overestimates reserves potential, it follows that the "successful" bidders may not have been so successful after all. Studies of the industry's rate of return support that conclusion. By simulating the bidding game we can increase our understanding and thus decrease our chance for investment error.

dustry's performance, though certainly all may have contributed. Poor luck might affect a few offshore participants. But the whole industry? Not likely. Industry has had enough opportunities in the Gulf to invoke the law of averages — if we may be so loose with mathematics.

We believe that in the competitive bidding environment normal good business sense utterly failed to give people the return they expected. Since many industry folk have not understood the rather complex laws of probability at work in competitive bidding, they have been inclined to make serious errors in arriving at their dollar bid for a particular tract. We are not saying that all of the bids turned out poorly. But enough of them have, throughout the industry, leading to lower rates of return than people planned for.

A new wrinkle appeared in the 1970 Offshore Louisiana wildcat sale (an \$850 million sale). Because of a Federal Power Commission order, some of the gas companies assumed they would be able to include their bonus investment in rate base. If they are correct, then their risk in offshore exploration has been effectively removed. They will make their legal return regardless of how much or how little reserves they find. This most recent sale, then, is very different from the others we have mentioned. The bidding model we would like to describe does not apply if lease bonus can be included in rate base.

We want to emphasize that we are not criticizing competitive bonus bidding as a method for acquiring leases from selling authorities. We believe this method is fair for all concerned. If the industry has not performed as well as it hoped, perhaps it is only because the industry has failed to understand the laws of probability that seem to govern the whole estimation-bidding process.

A "Think" Sale

Let us play a little game. Think of yourself as a manager whose task it is to set bids on parcels in an impending sale. On any one of your parcels you have a consensus property value put together by your experts. (We will not worry for the moment about how you handled risk, what your discount rate is, if you have one, or how you arrived at your reserves and costs.) One thing you can be sure of: Your value is either too high or too low; it has no chance of being exactly the true value.

Not to belabor a simple point, there are people in our business who fall in love with a number and fail to recognize the uncertainty associated with it. If a company's estimate happens to be \$5 million, who knows what the actual worth might be? If the tract is dry, the owner will have a loss — bonus plus exploration costs. If the tract produces — how much? There are fields discovered 50 years ago where we still do not know the reserves. And the uncertainty in field size before drilling is fantastic. So we repeat: Reserve estimates are either high or low—and maybe not even close.

We will assume, however, that on the average your value estimates are correct. (This does not contradict what we have already said. Most people are aware that they are high on some and low on others, but

over the long haul, they ought to come out about right on their value estimates.) You realize that other managers are going through the same agony you are. You ask yourself, "What do my competitors think these tracts are worth?" You know that some of your opponents may have better information than you, some worse. There will be, on sale date, quite a divergence of opinion as to value among the bidders. If you doubt this, look at the published bids by serious competitors at any recent sale. Bid ratios between the highest and lowest serious competitors range to as much as 100 and are commonly 5 or 10. (See Table 1.)

Implications of Divergence

What are the implications of this divergence of opinion? We could certainly argue that some people may have overestimated the true value of the parcel, and others may have underestimated it. Consider a piece of land that has exactly 10 million bbl of recoverable oil. If you let five different people in your own company interpret the seismic data, logs on nearby wells, and other sundry information, you will get five different estimates of reserves — even though they all use the same basic information. The problem becomes more confounding if we look at reserve estimates (before drilling now) of five different companies. They may each have different seismic data and different logs. Isn't it likely that some companies will come up with more than 10 million bbl? And some less? We have already admitted that while our estimates of reserves may be all right on the average, on any one tract we are going to be either high or low.

In Table 1 we saw evidence of this wide variation in value estimates by different competitors. Perhaps the several bidders had somewhat different exploration information. We all know the difference one properly placed seismic line can make in our mapping. Whatever the reasons, it is clear that different information leads to different value estimates.

Let us look at what different competitors can do given the same basic information. In the 1969 Alaska North Slope Sale, we find Atlantic Richfield and

TABLE 1—BIDS BY SERIOUS COMPETITORS
IN RECENT SALES
(All bids in millions of dollars)

Offshore Louisiana, 1967 Tract SS 207	Santa Barbara Channel, 1968 Tract 375	Offshore Texas, 1968 Tract 506	Alaska North Slope, 1969 Tract 059
32.5	43.5	43.5	10.5
17.7	32.1	15.5	5.2
11.1	18.1	11.6	2.1
7.1	10.2	8.5	1.4
5.6	6.3	8.1	0.5
4.1		5.6	0.4
3.3		4.7	
		2.8	
		2.6	
		0.7	
		0.7	
		0.4	
Ratio of Highest to Lowest Bid			
10	7	109	26

Humble bidding independently of each other. Since the two companies are equal partners in much exploration and development, both probably had essentially the same information; but each company took that information and developed its own evaluations without consulting the other. Table 6 shows the ratio of the Humble bid to the Atlantic Richfield bid for 55 tracts on which the companies competed against each other. At one extreme we find Humble making bids of about 0.03 of Atlantic Richfield's bid; at the other, Humble's bid is about 17 times higher than Atlantic Richfield's. And between these two extremes, we find a smooth gradation of ratios.

We have portrayed the same information a bit differently in Fig. 1. Here you will see a cross-plot of Humble's bids and Atlantic Richfield's bids for the same 55 tracts. No one has yet been able to identify any pattern or hint of correlation in these numbers. Clearly, the fact that companies have much the same seismic lines and well logs does not mean that those companies will come up with similar bids or property values.

On seeing such an exhibit, some ask if the wide range might not be due to differing discount rates or differing market conditions. But those items offset all of a company's bids in the same direction. A lower discount rate by one company, for instance, would force all of its evaluations up in dollars. There still would be large differences in bids.

Now more often than not, he who "sees" the most barrels will "see" the most dollar value. (Again, we recognize the effect of risk, cost estimates, production rates, pricing, discount rates and all that. But for the moment, let us focus on concepts and not clutter the picture with all these other items.) Can we not then conclude that he who thinks he sees the most reserves, will *tend* to win the parcel in competitive bidding? This conclusion leads straightway to another: *In competitive bidding, the winner tends to be the player who most overestimates true tract value.* And yet

another: He who bids on a parcel what he thinks it is worth will, in the long run, be taken for a cleaning.

A chorus enters sobbing, "But you told us earlier that our evaluations were correct on the average, albeit high sometimes and low sometimes. Doesn't the law of averages save us from ruin?" First, the so-called law of averages never guaranteed salvation for anyone, though it often gives some courage to act. Second, it is true (or we assume it so) that one's evaluations are correct on the average — but it is not true that one's evaluations on *tracts he wins* are correct on the average. There is a difference. Only in a noncompetitive environment, can one counter his overevaluated parcels with his underevaluated parcels and expect to do well on average. In bidding, however, he has a poor chance of winning when he has underestimated value and has a good chance of winning when he has overestimated it. So we say the player *tends* to win a *biased set* of tracts — namely, those on which he has overestimated value or reserves.

Note that we are talking now about *trends* and *tendencies* — not about what will happen every time one purchases a tract. It is possible that everyone will underestimate the value of a particular parcel. The winner will, under those circumstances, have a very attractive investment. But that is like winning the Irish Sweepstakes on your first ticket and then going around claiming that buying sweepstakes tickets is going to be a great investment for the future. As we make our investment decisions we must distinguish among the lucky event, the unlucky one, and the average of what occurs year after year.

Some may argue that the industry is smarter now — has new exploration techniques — and will not make the same kind of mistakes in the future. It is certainly true that we are better able to make exploration judgments these days; but it still does not mean we are very good. Anyway, even when technology was not so advanced, we were probably still "about right on average".

For example, before the "new technology" one might have expected a particular reservoir to contain 10 million bbl. If he had examined his uncertainties, he would have said the reservoir, if it exists, might have any amount between 2 million and 50 million bbl. With better information, he might still say he expects 10 million bbl, but his uncertainty has decreased and now ranges from 3 million to 35 million bbl. We claim that the effect of new technology only narrows our uncertainties — and does not necessarily change our expected values — again on average.

Bid Strategy

So what is the best bid strategy? We cannot tell you and will not even try. The only thing we can do is show you one approach to the mathematical modeling of competitive sales. The theory, as we interpret it, agrees well with what we perceive has happened in the real world.

For some competitive environments, in order to reach some specified return on investment, the model suggests a lower bid than one might come up with otherwise. What are these environments? The following rules are not without exceptions; but for the nor-




Fig. 1—Atlantic Richfield bids vs Humble bids, 1969 Alaska North Slope Lease Sale.

mal level of competition and the large uncertainties underlying our value estimates, the rules seem to apply.

1. The less information one has compared with what his opponents have, the lower he ought to bid.
2. The more uncertain one is about his value estimate, the lower he should bid.
3. The more bidders (above three) that show up on a given parcel, the lower one should bid.

How do we know these rules? Call it simulation. We modeled the competitive bidding process on a computer as closely as we knew how and then sat back to let the machine churn away. We allowed for such things as different numbers of bidders, different value estimates by the opponents, different information positions for the opponents, different bid levels* by the opponents, and the proper ranges of uncertainty about each of these. We let the computer take our estimates of competition (with the associated uncertainties) and play the lease sale game over and over again. After some thousands of runs the computer tells us, for our various bid levels, the probability of our winning the parcel and its value to us. Looking at the results, we simply choose a bid level that assures us (in a probability sense) of not investing incremental dollars at less than some specified rate of return.

We made all kinds of sensitivity tests to see "what if". We examined the effect of low rate of return criteria for opponents and checked on few opponents vs many. We looked into the influence of an opponent's superior information. We varied every significant variable we could identify.

When it was all over, we concluded that the competitive bidding environment is a good place to lose your shirt.

Previously we listed three reasons for lowering one's bid. The first two are easy enough to understand. But the third takes some work. Most people assume that the tougher the competition (i.e., the more serious bidders there are) the more they must bid to stay with the action. What action are they wanting to stay with? If they are trying to maximize the number of acres they buy, they are right. If they would like to maximize the petroleum they find, they are probably right. But if they are trying to invest money at some given rate of return, our model says they are probably wrong.

Although the concept may not be clear to everyone, we are convinced that if one's mistakes tend to be magnified with an increase in number of opponents, then he must bid at lower levels in the face of this stiffer opposition in order to make a given rate of return. Let us reinforce this with an example.

Assume we have a 10-tract sale. Also, for the sake of simplicity, let us assume that all tracts will be productive and that after exploratory drilling costs, each will be worth \$10 million at a 10 percent discount rate. Each competitor in this sale correctly estimates the total value of the sale acreage but on any one tract he may be too high or too low. (This assumption

merely means that one tends to be unbiased in his estimate of value. He may not be correct on any one parcel, but he does all right on the average.)

As in the real world, let us have the competitors disagree as to the value of the individual tracts — and let that divergence of opinion be about the same as we see in major lease sales. But let the average of all the competitors' value estimates be very close to the true value. (Here we are saying that when they estimate value the competitors are not misled in the same direction.)

Finally, assume that to protect himself from the risks and uncertainties of the estimating procedure, each competitor chooses to bid one-half *his* value estimate. What we want to do is check the rate of return of the winners as we increase the number of bidders.

Table 2 reflects the sale as if only Company A bids. Remember, he correctly estimates that the 10 tracts are worth \$100 million to him and he bids one-half of his value estimate on each tract. The sum of his 10 bids is then \$50 million. He wins all tracts since there is no competition. Since he pays \$50 million for what is worth \$100 million (at a 10 percent discount rate) his rate of return for the sale will be about 17 percent** after tax. This is his reward even though he has overestimated value on Tracts 2, 6, and 8.

Table 3 examines the consequences of adding one competitor, Company B. Since both companies are unbiased in their estimates, use the same discount rate for calculating value, and bid the same fraction of their respective values, then we would expect each to win half the time. As it turns out, that is exactly what happens. But see what else happens. In Table 1 we saw that Company A won all 10 tracts — on seven of which he had underestimated value and on three of which he had overestimated. Now along comes Company B and wins five of the seven tracts on which Company A had underestimated value. Remember our contention that one *tends* to lose those tracts on which he has underestimated value? Company A has spent more than 70 percent as much money as he spent when he was the only bidder, but now he gets only half as much acreage. The only thing that saves him is his strategy to bid one-half his value estimates. His rate of return drops to 14 percent. The "industry" consisting of the two companies has about the same return.

Now go to Table 4 and see what happens if we raise the number of bidders to four. More and more of Company A's undervalued tracts have been grabbed off by the competition. Company A is left with only Tract 8, which he evaluated at \$35 million. (It is worth only \$10 million, remember.) The selling authority's take has climbed to about \$92 million — the sum of all the high bids. Company A's return drops to about 5 percent, whereas the industry's return is about 11 percent. Company A turns out to be a little unlucky in that its return is lower than the industry's. Somebody has to be unlucky. That should not detract from our argument. We could pick any

*Bid level is the fraction of his value estimate a player will bid.

**We estimated this return and others in the example from studies of cash flows from typical projects.

of the competitors and see the same trend toward lower returns.

Table 5 shows the results of eight bidders. Company A still retains its Tract 8. Bidders E through H pick up five of the 10 tracts. The seller gets about \$26 million more than he did with the four competitors. Since the tracts did not pick up any more reserves, the additional expenditure must mean a decreased rate of return for the industry. We estimate about 8 percent — even though each bidder is bidding only half his value estimate.

There is no table to show the results for 16 bidders, but the trend continues onward to lower returns. The 16 bidders spent a total of \$162.6 million for a return of about 6 percent.

What if the industry had wanted to make about 10 percent on its investment? What percent of value would each competitor have had to bid to accomplish that goal? Just taking the results of our example, the bid levels would have been something like this:

Number of Bidders	Total Value Estimates for Highest Estimators on Each Tract	Bid Level for 10 Percent Return
1	\$100 million	1.00
2	\$139 million	0.72
4	\$184 million	0.54
8	\$237 million	0.42
16	\$325 million	0.31

(The bid levels that appear in the third column are valid for only the particular example we have just gone through, where everyone uses the same return criterion and everyone uses the same bidding strategy. Companies, in the real world, are not so inclined to play that way. Nevertheless, the phenomenon of decreasing rate of return with increasing numbers of

TABLE 2—CASE 1—ONLY COMPANY A BIDS ON PROPERTY

Tract Number:	1	2	3	4	5	6	7	8	9	10
A's bid*	1.9	5.6	2.6	3.4	3.7	5.2	1.9	17.5	3.9	4.3
						Company A		Industry		
Winning bids**						50		50		
Value of acreage won*						100		100		
Present-worth profit*						50		50		
Investor's rate of return, percent						17		17		

*In millions of dollars.

TABLE 3—CASE 2—ONE COMPETITOR ENTERS SALE WITH COMPANY A

Tract Number:	1	2	3	4	5	6	7	8	9	10
A's bid	1.9	5.6*	2.6	3.4	3.7*	5.2*	1.9	17.5*	3.9*	4.3
B's bid	3.8*	5.1	4.0*	4.9*	0.6	4.2	5.9*	4.5	1.8	15.2*
						Company A		Industry		
Winning bids**						35.9		69.7		
Value of acreage won**						50.0		100.0		
Present-worth profit**						14.1		30.3		
Investor's rate of return, percent						14		14		

*Winning bid.

**Millions of dollars.

bidders appears to us a general rule of sealed bidding.)

It is certainly true that the value of the tracts does not change just because there are more bidders. What does change drastically as the number of bidders increases is the set of tracts one wins. Not only does that set get smaller with increasing competition, but also its quality tends to decrease compared with what the winner thought it would be ahead of time.

The more serious bidders we have, the further from true value we expect the top bidder to be. If one wins a tract against two or three others, he may feel fine about his good fortune. But how should he feel if he won against 50 others? Ill. He would wonder why 50 others thought it was worth less. On the average, one misjudges true value much worse when he comes out high against 50 other bidders than when he beats only two or three. Hence, our bidding model usually tells us to move toward lower bids as competition increases in order to protect ourselves from the winner's curse. True, the probability of purchasing property decreases — but so does the chance of losing that shirt.

Some Mathematics

The theory of competitive bidding obviously involves mathematics. For those so inclined, we will lay out here and in the Appendix analytical procedures for examining the effects we have spoken of. (Then we will say, "But the analytical approach is so difficult from the practical side that we must try a simulation.") What we will try for analytically is the *expected value of the winning bid*. We simply compare that value with

TABLE 4—CASE 3—THREE COMPETITORS ENTER SALE WITH COMPANY A

Tract Number:	1	2	3	4	5	6	7	8	9	10
A's bid	1.9	5.6	2.6	3.4	3.7	5.2	1.9	17.5*	3.9	4.3
B's bid	3.8	5.1	4.0	4.9	0.6	4.2	5.9*	4.5	1.8	15.2*
C's bid	5.7	3.1	2.6	6.5*	9.8*	9.8*	4.0	1.5	3.3	3.7
D's bid	6.5*	8.3*	7.8*	6.4	3.3	2.2	3.3	5.0	4.5*	2.7

	Company A	Industry
Winning bids**	17.5	91.8
Value of acreage won**	10.0	100.0
Present-worth profit**	- 7.5	8.2
Investor's rate of return, percent	5	11

*Winning bid.

**Millions of dollars.

TABLE 5—CASE 4—SEVEN COMPETITORS ENTER SALE WITH COMPANY A

Tract Number:	1	2	3	4	5	6	7	8	9	10
A's bid	—	—	—	—	—	—	—	17.5	—	—
B's bid	—	—	—	—	—	—	5.9	—	—	15.2
C's bid	—	—	—	6.5	—	—	—	—	—	—
D's bid	—	—	7.8	—	—	—	—	—	—	—
E's bid	—	10.3	—	—	—	—	—	—	—	—
F's bid	—	—	—	—	14.3	13.0	—	—	—	—
G's bid	23.3	—	—	—	—	—	—	—	—	4.7
H's bid	—	—	—	—	—	—	—	—	—	—

	Company A	Industry
Winning bids*	17.5	118.5
Value of acreage won*	10.0	100.0
Present-worth profit*	- 7.5	- 18.5
Investor's rate of return, percent	5	8

*Millions of dollars.

true parcel value to see whether a particular bidding strategy can lead to trouble.

Let

$f_i(x)$ = probability density function for i th opponent's bid.

And let

$F_i(x)$ = probability that the i th opponent bids a value less than x .

Therefore,

$\prod_{i=1}^n F_i(x)$ = probability that n independent opponents all bid a value less than x .

Now let

$g(x)$ = probability density function for our bid.

Define

$h(x) = K_n \left[\prod_{i=1}^n F_i(x) \right] g(x)$ = probability density function for our winning bid,

where

K_n = constant to make the integral of that density = 1

$$K_n = 1 \int_{-\infty}^{\infty} \left[\prod_{i=1}^n F_i(x) \right] g(x) dx.$$

Then it is a simple matter to get the expected value of our winning bid, $E(X_w)$

$$\begin{aligned} E(X_w) &= \int_{-\infty}^{\infty} xh(x) dx \\ &= \int_{-\infty}^{\infty} xK_n \left[\prod_{i=1}^n F_i(x) \right] g(x) dx. \end{aligned}$$

Then under some very simple assumptions (too simple for the real world), we can define some $F_i(x)$ and $g(x)$ in such a way that we can evaluate the integral. In fact, we can show that if $f_i(x)$ and $g(x)$ are uniform on the interval of 0 to 2, and all competitors bid their full value estimate, then:

$$K_n = n + 1$$

$$E(X_w) = 2 \left(\frac{n + 1}{n + 2} \right).$$

These uniform distributions imply a true value of 1 (the mean of each is 1). If there are no opponents ($n = 0$), then:

$$E(X_w) = 2 \left(\frac{1}{2} \right) = 1.$$

That is what we hope if we bid our value estimate against no opposition. On the average, we win tracts at our value. But what if there are five opponents?

$$E(X_w) = 2 \left(\frac{5 + 1}{5 + 2} \right) = \frac{12}{7} \approx 1.71.$$

That means that on the average, we would expect to pay 71 percent more than value on the tracts we won. That is not good.

One might think he could take the reciprocal of 1.71 to get his "break-even" bid level. Not so. The subtleties of competition force the "break-even" bid level to be even lower than that reciprocal, although perhaps not too much lower.

We can set up the mathematics, but for the real world, we cannot solve the equations. Instead, we simulate the whole process. And that is all right, for by simulation, we can do many things we would not even try with strict mathematical analysis.

How Can a Bidding Strategist Win Tracts?

Some will claim he cannot — we believe they are wrong.

An analyst comes in claiming a tract is worth X . The bidding strategist then recommends a bid of, say, $X/2$. A voice from the rear cries, "That bid won't be competitive." The voice is usually forgetting about the large divergence in value estimates by competitors. There is a very good chance some other competitor will see a much larger value than X . We could not be competitive with any bid we would reasonably try. So our chance of winning depends more upon our reserves estimate than upon our particular bid level. The bid level adjustment is primarily for the purpose of achieving a certain profitability criterion.

Some interesting evidence to back up these comments comes from the 1969 Alaska North Slope Sale.⁶ Examine the second-high bids for that sale. The sum of those second-high bids was only \$370 million compared with the winning bid sum of \$900 million. Said another way, the fellow who liked the tract second best was willing to bid, on the average, only 41 percent as much as the winner. In this respect, the sale was not atypical.

If that is not shocking enough, try this one. For 26 percent of the tracts, had the second-high bidder increased his bid by a factor of 4, he still would not have won the tract. A 50-percent increase in bid by the second-high man would not have won 77 percent of the tracts. Turn the idea around. If every tract winner had bid only two-thirds as much as he did, the winners still would have retained 79 percent of the tracts they won. (The apparent discrepancy, 77 percent vs 79 percent, comes from the 15 tracts that drew only one serious bidder.) We therefore conclude, based on historical study, that bid manipulation to achieve desired profitability does not drastically impair one's chances of winning acreage.

TABLE 6—ALASKA LEASE SALE, 1969 —
RATIO OF HUMBLE BID TO ARCO BID
For the 55 tracts on which both companies bid

0.03	0.32	0.50	1.11	2.53
0.03	0.32	0.51	1.13	2.56
0.04	0.33	0.51	1.31	3.82
0.06	0.33	0.60	1.39	5.25
0.08	0.36	0.69	1.39	5.36
0.11	0.36	0.76	1.40	6.14
0.12	0.36	0.77	1.79	7.98
0.16	0.39	0.78	2.02	9.19
0.18	0.41	0.79	2.41	13.32
0.22	0.45	0.82	2.41	15.45
0.24	0.45	1.00	2.50	16.80

How Far Off Might the Winner Be?

We have been saying that the winner of a tract tends to be the one who most overestimates value. You may say, "So, if we win, we wish we hadn't. If we lose, we wish we hadn't. You mathematicians are really saying to stay away from lease sales." That is not what we are saying. The bidding model gives us a bid that we can make with confidence, and be happy with when we win. Yes, we may have overestimated value. But we have bid lower than our value estimate — hedging against expected error. In a probability sense, we "guarantee" that we obtain the rate of return we want.

As to how far off the highest estimator might be, we have resorted to simulation of the estimating process. We perhaps could have got the result through use of extreme value theory, but we chose not to. Also, we want to caution the reader that we are examining what we think will happen on the average — not what will happen on a particular tract. If the wildcat fails, obviously everyone was too high in his value estimate. If the well hits, it is entirely possible everyone was too low. That is not the kind of problem we are talking about. The question is more likely: "If I win 10 parcels at a sale, how many barrels will they all contain compared with my pre-sale estimate?"

Fig. 2 shows the results of our simulations (using log-normal distributions) for various numbers of competitors and degrees of uncertainty. We use the variance of a distribution — measure of its spread — to quantify general uncertainty as to value among competitors. One can get a rough idea of the magnitude of variance by measuring the parameter on sets of bids on tracts in past sales. That variance, however, will be too high since the actual bids contain "noise" items apart from property evaluation — for example, various company discount rates and bid levels. Obviously, there is not so much uncertainty in drainage sales as there is in North Slope-type wildcat sales. We use variance to account for these differences.

Intuition would argue that the greatest potential for large errors in estimating reserves exists on the frontier — Alaska. The simulation agrees wholeheartedly. For 12 serious bidders in an environment of uncertainty such as the North Slope, the one estimating the largest amount of expected reserves can




Fig. 2—Relation of mean high estimate to true value under various conditions of uncertainty.

expect to be off by a factor of 4 on average. In the Louisiana Offshore, facing the same kind of competition, he would expect to miss by a factor of only 2.5.

Nature of the Model

We must choose a probability distribution for the value estimates of various companies. The log-normal seems to us the best. Many writers have documented the variables in our business that seem to follow the log-normal. Here is a partial list of them:

1. Reservoir volume
2. Productive area
3. Net pay thickness
4. Recoverable hydrocarbons
5. Bids on a parcel in a lease sale
6. Land value estimates calculated by companies.

The first four items have been ordained by Nature. The last two are man-made. Why should they perform like Nature? There is an amazing theorem in mathematics — the Central Limit Theorem — that says if you take sums of random samples from any distribution with finite mean and variance, the sums will tend toward a normal or Gaussian distribution. The tendency will be stronger the more numbers there are in each sum. If the original numbers come from a normal distribution, the sum is guaranteed to be normal. If we insert the word "product" for "sum" we can then insert the word "log-normal" for "normal." Since we arrive at value through a series of multiplications of uncertain parameters (reservoir length \times reservoir width \times net pay \times recovery \times after-tax value per barrel), it is not surprising that bids and land-value estimates seem to take on this log-normal characteristic.⁷

There are certain problems in applying the theorem. Negative dollars (a loss or lower-than-criterion rate of return) will not fit the log-normal distribution. No one knows how to take the logarithm of a negative number. And we all know that the value calculation involves more than simple multiplication. Even so, the error in our assumption does not appear to be great, and we happily use the log-normal distribution in our computer simulation.

The evaluation of a potential cash flow stream by different investment criteria has been the subject of much study. We believe that methods involving the discounting of the cash flow stream are effective for the decision maker. The criterion we prefer is present worth or present value (PW), using as the discount rate the Internal or Investor's Rate of Return (IRR) expected to be earned by the investor in the future.⁸ The very essence of PW is that it is the value or *worth* we place on an investment opportunity at the *present* time. In a situation where the future cash flow is known with certainty, we can discount this cash flow to the present.

We do not know the future cash flow with certainty, however, and resort to using the *expected value* concept. Expected value can mean different things to different people, but we use it in the accepted probabilistic sense: Expected value is the sum of all possible events multiplied by their chance of occurrence. Arithmetic mean is a common term for expected

value. Expected value is *not* necessarily the mode (most probable value), nor the median (the value that is exceeded half the time). We do not specify all the possible events, since this would be an outrageous number. But we do try to specify enough possible events so that the calculations with these relatively few discrete values will yield a good value. The “good” value should be close to that expected from a consideration of *all* possible events.

The tract value plays a much smaller role in our model than one might think. We essentially normalize everything to value = 1.0. The model tells us what *fraction* of our value (bid level) to bid in order to maximize expected present worth for the competition we put in. The bid level can change only if our idea of the competition somehow changes. If we think the degree of competition is independent of tract value, then value need never be discussed. But sometimes there are tracts that, because of their potential, may cause competitors to deviate from past or expected performance. We allow for this by considering the competition the way we think it will be for a given tract. In that sense, then, value gets into the model.

Our model differs from some other models that have been discussed. An earlier philosophy reasoned thus: “Our value may be incorrect on a given tract, but it is correct on the average. So let our value estimate serve as the mean of the distribution from which our opponents draw.” We think that tack can lead to trouble. It is inconsistent with the idea that when we win, our estimated value was probably higher than true value. Instead, we let the true value of a tract be 1.0 and simply take our value estimate from a distribution with mean = 1, the same as everyone else. We treat all value estimates as independent random variables. Our model is similar in this respect to Rothkopf’s.⁹ The variance of our distribution may be the same or different from our opponents’ — depending on the relative quality of our information.

Model Input Data

Some believe that the input requirements for a competitive bidding model are quite severe — that reliable input is impossible to obtain. We do not think so. Unless one successfully engages in espionage, he is not going to know his opponent’s bid. But he does not need to. We have found that by studying the behavior of companies in past sales, we can get a fair clue as to what they will do in the future — close enough to make the model results meaningful.

Here is the information we think is necessary to make an intelligent bid. Keep in mind that each bit of input is an uncertain quantity. We treat it as uncertain by using probabilities and probability distributions. That, after all, is the way the world is.

We believe that the input data are best determined by a combination of historical data and the judgment of explorationists. To illustrate the use of our model, we will develop a set of input data for a purely hypothetical example.

What sort of data do we need? Primarily, we need information about the competition we are likely to face. We try to identify companies that are likely to bid on the parcel. This allows us to use any specific

knowledge we have about a competitor or his exploration activities. For each of the potential competitors, we then try to estimate the probability that he will bid. To the competitors specifically named, we can add some “other bidders” in order to make the expected number of bidders consistent with our beliefs:

Company	Probability of Bidding
A	0.8
B	0.7
C	0.5
Other bidder	0.5
Other bidder	0.5
Expected number of bidders	3.0

In this example, we expect three competitors, but we acknowledge that there could be as few as none or as many as five. In the simulation performed by our model, the number of competitors will vary, from trial to trial, from a low of zero to a high of five. The proportion of trials on which a given bidder appears will be approximately equal to the probability we have assigned above.

The next item we require is usually the most difficult to estimate: the bid level of each potential competitor. If he calculates a value of \$*X* for the property, what fraction of that value is he likely to bid? To further complicate the matter, we need to estimate this fraction as if the \$*X* value were based on our own rate of return criterion. In other words, the bid level is used to adjust for differences in evaluation criteria and for the fraction of value that a given competitor will bid.

We believe that historical data can be of help in estimating bid levels. We can go back to a previous sale or sales and compare a given competitor’s bids with the value estimates we made on the same tracts. At first we were tempted to compute the ratio of a competitor’s bid to our value on each tract and then average these ratios over all tracts. We discovered that under the assumptions of our model of the bidding process this gives a biased estimate of the competitor’s bid level. We can show that to get an unbiased estimate of his bid level on a tract we need to divide the ratio of his bid to our value by the quantity e^{σ^2} . Here σ^2 is the variance of the natural logarithm of our value estimate on the tract. (Our value estimate, remember, is considered a random variable. Estimates of σ^2 are not easy to come by, but again historical data can be of help.) We can then calculate an average bid level for the competitor from these unbiased estimates on all the tracts. This bid level estimate incorporates differences in evaluation criteria, as well as the fraction of value that the com-

TABLE 7—INPUT DATA FOR COMPETITION

Company	Probability Of Bidding	Bid Level	Variance
A	0.8	0.6	0.6
B	0.7	0.6	0.6
C	0.5	0.4	0.6
Other bidder	0.5	0.3	0.8
Other bidder	0.5	0.3	0.8

petitor bids, on average. We then modify this according to our explorationists' judgment about the current sale and the particular tract in question to add another column to our hypothetical input data:

Company	Bid Level
A	0.6
B	0.6
C	0.4
Other bidder	0.3
Other bidder	0.3

We also need to specify how much variation we think is possible in each competitor's bid. Even if we knew that the bid levels specified above were precisely correct, we still would be uncertain as to the actual bids because we do not know the value that each competitor places on the tract. We must try to estimate the variability in each competitor's value estimate. We do this by specifying the variance of the estimate. (Actually, we specify the variance of the natural logarithm of the estimate. Hereafter, when we mention variance, we will be referring to the variance of the logarithm of a quantity, because this is a useful parameter in the log-normal distribution.)

We can again get some help from data on past sales. On individual tracts about 1.2 has been the average variance of the bids.¹⁰ This includes more than just the variation in value estimates, though. It also includes differences in bid levels and evaluation criteria among competitors. The variance in value estimates for a single company would average something less — we have guessed about 0.6.

Another way to estimate this variance, if we assume it is constant over all tracts, is to compare an individual competitor's bids with our values on the tracts in a given sale. This should eliminate variation due to differences in evaluation criteria, assuming a company uses the same criterion in all of its evaluations. If we measure the variance of the ratio of a competitor's bid to our value, there are three components to this variance:

1. Variance of our value estimate (Y)
2. Variance of the competitor's value estimate (X)
3. Variance of the competitor's bid level (K) from tract to tract.

We can show that these components are additive. The variable whose variance we are measuring is $\log_e(KX/Y)$. We can write

$$\log_e(KX/Y) = \log_e(K) + \log_e(X) - \log_e(Y).$$

If K , X , and Y are independent,

$$\begin{aligned} \text{Var} [\log_e(KX/Y)] &= \text{Var} [\log_e(K)] \\ &+ \text{Var} [\log_e(X)] + \text{Var} [\log_e(Y)]. \end{aligned}$$

By assuming that the last two components are equal and the first is about 0.15, we calculated an average variance for our opponents' value estimates in several sales. The values were not far from the 0.6 estimated above.

We feel free to modify this estimate in accordance with the nature of the sale and the tract in question.

For example, we felt that the 1969 North Slope Sale was characterized by more uncertainty than the typical offshore Louisiana sale. Thus, we generally assigned higher variances to value estimates. In drainage situations, we use lower variances to reflect the fact that the value estimates should be closer to the true values. We also try to differentiate among competitors. Those we feel have better information about a tract are given lower variances and those with poorer information, higher variances. So we shall add another column to our input data:

Company	Variance
A	0.6
B	0.6
C	0.6
Other bidder	0.8
Other bidder	0.8

Table 7 shows a complete set of the input data on competition.

We add another component, $\text{Var} \log_e(K)$ mentioned above (usually about 0.15), to these variances to reflect our uncertainty about our competitors' bid levels. Finally, we estimate the variance in our value estimate by assessing the quality of our information relative to that of our opponents'.

Mechanics of the Model

The parameters for the log-normal distributions assigned to the value estimates of the various bidders (including us) come directly from the data given above. We usually run the model thousands of times to simulate the competitive and evaluation possibilities on a single tract. (See flow chart, Fig. 3.) On each trial, a value is drawn for each random variable, which results in a set of bids by the participating companies. The results of the "sale" are then recorded and the whole process is repeated. After enough trials have been run, the expected results are calculated and printed.

Model Output

The output of the model includes expected results for 15 different bid levels, from 0.1 to 1.5 times our value estimate. Results from our hypothetical example are shown in Table 8. The values in the first column indicate possible bidding levels as fractions of our value estimate. The second column gives the amount of our bid at each level. We have assumed that our estimate of the value of this tract is \$10 million. The next column shows the probability of winning, as calculated by the model, for each bidding level. This is useful in estimating the amount of acreage, reserves, etc., we expect to win. The expected amount of our expenditure is shown in the fourth column. In the next column we have the expected present worth for each bidding level. The last column indicates how high we can expect our value estimate to be if we win. If we bid full value (bid level of 1.0) and win on tracts such as this, our value estimate will, on the average, be 1.35 times the true value. It is again obvious that we have to bid less than full value just to break even.

Optimization of Bids

The expected present worth of the submitted bid we will designate as EPW_{Bid} . Given all our usual information about the tract and other bidders, what bid should we submit? What is our optimum bid for the example above?

We can use a graph of EPW_{Bid} vs bid level to consider this problem (Fig. 4). First, what happens if we do not bid? The bid level is zero. No expenditures will be made, and the EPW_{Bid} is zero. Second, what happens if we bid our estimate of the tract value? For the tracts we win, we tend to overestimate value. Hence, the average value of the tracts we win is less than our original estimates. Thus in the example we have a negative EPW_{Bid} of \$1.9 million. Third, what happens if we bid less than our estimate? This strategy really provides the only chance we have to get a positive EPW_{Bid} . We must bid somewhere between the one extreme of a very low bid (which means very low chance of winning a big positive value) and the other extreme of a very high bid (which means a high chance of winning a big negative value).

What then is the optimum bid? For the single tract illustration above, and for our investment criterion of maximizing the EPW_{Bid} rather than maximizing it: serves or some other goal, we would choose a bid level of 0.35. There may not always be a positive value of EPW_{Bid} , in which case we would not bid.

Usually, however, there is a positive maximum value. It is not always at the same bid level. The maximum shifts along the bid level axis with changes in the number of bidders, their bid levels, and the variances of their estimates.

Deviation from the optimum bid level in either direction will decrease the EPW_{Bid} . If someone "feels" we should bid higher or lower, we can show what this feel costs in terms of EPW. Any bid giving a positive EPW_{Bid} will, of course, give an expected IRR greater than the discount rate. Suppose the discount rate used is the marginal acceptable IRR. Going to a larger bid level than that giving maximum EPW gives a lower EPW. Therefore, that marginal increase in bid has a negative EPW associated with it. Look at Table 8. Going from a bid of 0.5 to 0.6 costs \$283 thousand in EPW. Taking an action that decreases the EPW is the same as taking an action that invests money at less than the acceptable IRR. According to the model, then, he who would go above his optimum bid level to gain probability of win advantage can expect to invest part of his money at a return lower than the minimum he said he would accept.

Before leaving the subject of bid optimization, we will comment on another frequently mentioned criterion. Under the existing conditions of uncertainty, there will be "money left on the table" (difference between the winning bid and second-high bid) and rightly so. We can minimize the money left on the

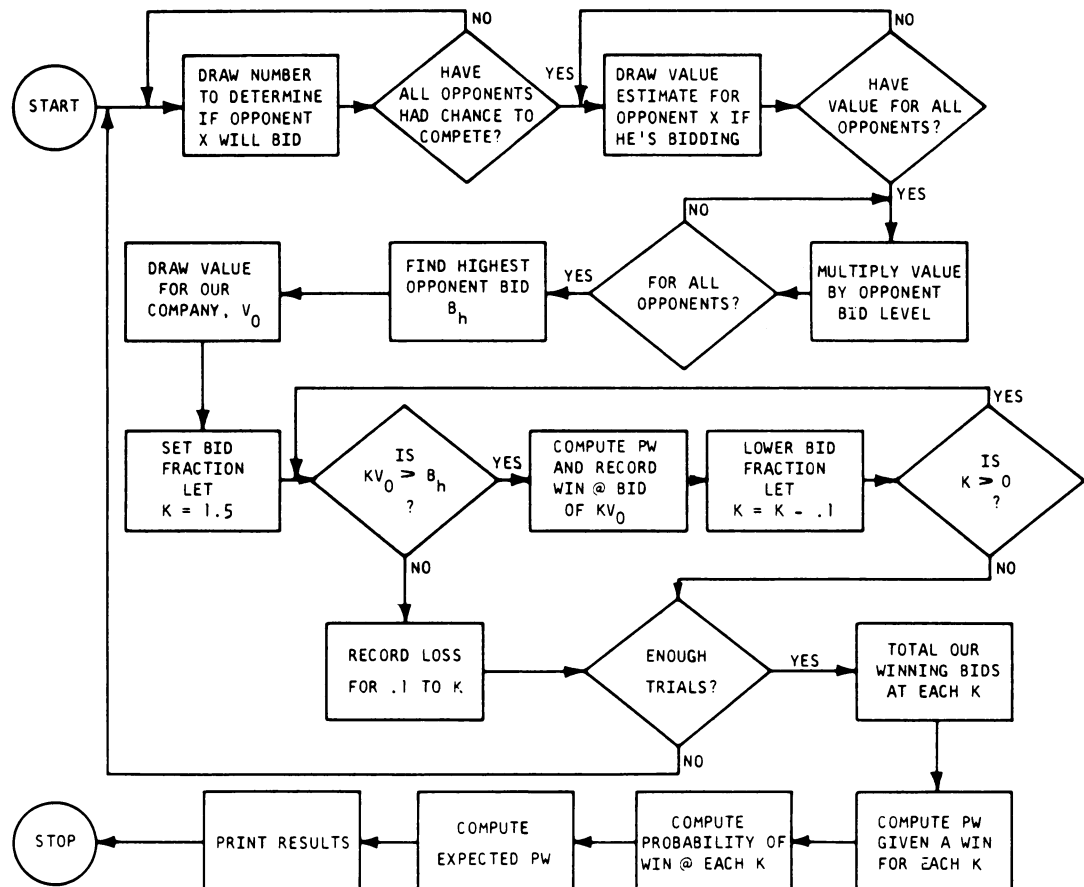


Fig. 3—A bidding model.

TABLE 8—MODEL OUTPUT

Bidding Level	Bid*	Probability of Winning	Expected Bonus Spent*	Expected PW of Bid*	Expected Ratio of Our Estimate To True Value, Given We Win
0.0	0	0	0	0	—
0.1	1,000	0.03	30	232	2.32
0.2	2,000	0.09	180	512	2.16
0.3	3,000	0.16	490	670	1.97
0.4	4,000	0.23	933	664	1.79
0.5	5,000	0.29	1,472	477	1.68
0.6	6,000	0.36	2,136	194	1.58
0.7	7,000	0.41	2,878	-212	1.50
0.8	8,000	0.46	3,675	-707	1.44
0.9	9,000	0.50	4,523	-1,288	1.39
1.0	10,000	0.54	5,407	-1,918	1.35
1.1	11,000	0.57	6,324	-2,607	1.32
1.2	12,000	0.61	7,321	-3,313	1.29
1.3	13,000	0.64	8,342	-4,076	1.26
1.4	14,000	0.67	9,330	-4,868	1.24
1.5	15,000	0.69	10,348	-5,682	1.22

*Thousands of dollars.

table by not bidding at all; or, if we have positive knowledge of others' bids, we can bid one cent more. Obviously both approaches are impractical. The fact is that there *will* be money left on the table, so we have chosen the objective of maximizing EPW — not minimizing money left on the table.

Possible Weaknesses in the Approach

If we thought there were any serious flaws in our ideas on bidding, we would not want to waste your time. On the other hand, while we have gathered considerable evidence to support our theory, some chance always remains that we have simply been fooled by the data.

The *log-normal distribution* does not look as if it can work very well as a model for describing uncertainty on tracts of small value. People who use discounted cash flows to arrive at present worths can get negative values even though they expect the tracts to make some, though small, positive rate of return. The log-normal probability distribution allows for no negative values. The effects of differences among company discount rates become highly magnified on the less valuable acreage.

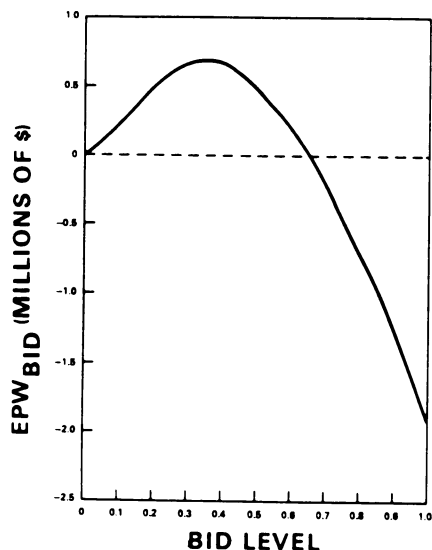


Fig. 4—EPW_{b,d} vs bid level.

The reason our model can suggest such a low bid level as a reasonable strategy is the *magnitude of the uncertainty* that we believe is associated with the reserves-value estimating process. We had occasion to compare our independent reserves estimates with those of a partner and found the disagreement to be quite large, though there was no bias by either party. We were as likely to be high as he. If you look at published bids, you can, indirectly, get the same results.

In Fig. 2 we showed that the highest estimator would be off, on the average, by a factor of 2.5 in his expected reserves estimates if he were competing against 11 other independent estimators. Anyone who feels his own reserves estimates are never off by more than 50 percent will feel severe pains swallowing our factor of 2.5.

Of course the amount of uncertainty is just an input parameter for the model. One can put in whatever he likes.

Another problem is our assumption that reserves and value as reflected in final bid estimates *tend to be unbiased*. If we did not make this assumption we would change our ways. No manager is going to submit a bid based on value estimates that he knows are too high or too low. He will enter a multiplier with the intention of being correct on the average. But that tactic does not necessarily guarantee he will be.

We have recognized another weakness without finding much of a solution. How do we account for the competitor who does not bid at all on a particular lease? Does he think it worthless? Has he no interest? Or has he run out of funds? One might argue forcefully that in a major sale he always faces 15 to 20 competitors, whether all of them bid or not.

Conclusions

It is still said that, after many years of exploration, many barrels of oil found, many cubic feet of gas found, and after much red ink, the outlook for future offshore potential is bright. Maybe it is.

Unexpectedly low rates of return, however, follow the industry into competitive lease sale environments year after year. This must mean that by and large industry is paying more for the property than it ultimately is worth. But each competitor thinks he is play-

ing a reasonable strategy. How can the industry go astray? Our sojourn into competitive bidding theory tells us to expect exactly what has happened. It is, then, a theory not only that is mathematically sound, but also that fits reality. Even though each bidder estimates his values properly on average, he tends to win at the worst times — namely when he most overestimates value. The error is not the fault of the explorationists. They are doing creditable work on a tough job. The problem is simply a quirk of the competitive bidding environment.

Acknowledgment

L. P. (Barney) Whorton, Manager of Atlantic Richfield's Production Research and Development Laboratories and former SPE president, deserves much of the credit for the work that led to this paper. We want to thank him for his unending support, encouragement, and open-minded criticisms of the research effort.

References

1. Friedman, Lawrence: "A Competitive Bidding Strategy", *Operations Research* (Feb., 1956) 104-112.
2. Barrow, T. S.: "Economics of Offshore Development", *Exploration and Economics of the Petroleum Industry*, Matthew Bender & Co., New York (1967) 5.
3. "Offshore Know-how Racing but Costs Look Formidable", *Oil and Gas J.* (May 26, 1969) 35-38.
4. "U. S. Offshore Long Way From Paying Off", *Oil and Gas J.* (July 6, 1970) 55-56.
5. Winzeler, A. R.: "Economics of Offshore Exploration — A Post-Appraisal of Recent Federal Sales", *Exploration and Economics of the Petroleum Industry*, Matthew Bender & Co., New York (1969) 7.
6. Friedman, Lawrence: oral communication, John D. Kettelle Corp. Seminar on Competitive Bidding, Philadelphia, Pa. (Dec., 1969).
7. Crawford, Paul B.: "Texas Offshore Bidding Patterns", *J. Pet. Tech.* (March, 1970) 283-289.
8. Davis, Louis F.: "Economic Judgment and Planning in North American Petroleum Exploration", *J. Pet. Tech.* (May, 1968) 467-474.
9. Rothkopf, Michael H.: "A Model of Rational Competitive Bidding", *Management Science* (March, 1969) 362-373.
10. Brown, Keith C.: "A Theoretical and Statistical Study of Decision-Making Under Uncertainty: Competitive Bidding for Leases on Offshore Petroleum Lands", unpublished PhD dissertation, Southern Methodist U., Dallas (1966).
11. Aitchison, J. and Brown, J. A. C.: *The Lognormal Distribution*, Cambridge U. Press, New York (1957).

APPENDIX

Mathematical arguments leave most people cold. On the other hand, it is nice to know that the logic of English has the solid support of mathematics — especially when we try to explain why a bid level that maximizes present worth should often go down as the number of competitors increases. Of course some of you may have learned long ago to beware of the English language and to trust naught but mathematical rigor. For you, we offer this Appendix.

In the main text we said that we could not carry out the necessary integrations if we used the log-normal distribution. We can, however, analyze a probability distribution that has properties similar to the log-normal. The exponential distribution is our candidate. It is properly skewed. It is defined on the interval 0 to ∞. And if we choose an exponential distribution whose mean is 1.0 (corresponding to true

value normalized to 1.0), we have a spread not too unlike that log-normal whose variance describes the kind of uncertainties faced in the Gulf of Mexico.

The level of mathematics we use is not difficult — a little calculus and a little probability theory. We want to derive an equation that will tell us EPW (Expected Present Worth) as a function of our bid level, the opponent's bid levels, and the number of opponents. By solving that equation, we will show that the bid level for which we get the largest EPW peaks out at two or three opponents and then falls.

$h(x) = \lambda e^{-\lambda x}$, probability density function for value estimate for each bidder

$\frac{1}{\lambda}$ = mean of value distribution for ourselves and our opponents, assumed equal to true value

c_0 = fraction of our value estimate we choose to bid

$g(x) = \frac{\lambda}{c_0} e^{-\lambda x/c_0}$, probability density function for our bid

c_i = fraction of his value estimate that Company i chooses to bid

$f_i(x) = \frac{\lambda}{c_i} e^{-\lambda x/c_i}$, probability density function for bid of Company i

$F_i(x) = 1 - e^{-\lambda x/c_i}$, cumulative bid distribution for Company i

n = number of bidding opponents

$\binom{n}{k} = \frac{n!}{k!(n-k)!}$, notation for combinations

$\prod_{i=1}^n F_i(x)$ = probability that all opponents will bid less than x , or the probability that we win if we bid x

To get our EPW we multiply 3 terms:

- PW if we bid x and win
- Probability of winning if we bid x
- Probability of our bidding x .

Then we integrate or sum up over all possible values of x .

$$EPW = \int_0^{\infty} \left(\frac{1}{\lambda} - x \right) \left[\prod_{i=1}^n F_i(x) \right] \frac{\lambda}{c_0} e^{-\lambda x/c_0} dx. \quad \dots \dots \dots (A-1)$$

Let us simplify by assuming that all opponents will use the same bid levels. Then

$$\prod_{i=1}^n F_i(x) = [F_i(x)]^n = (1 - e^{-\lambda x/c_i})^n,$$

which we expand binomially

$$= \sum_{k=0}^n (-1)^k \binom{n}{k} e^{-\lambda k x/c_i} \dots \dots (A-2)$$

Then

$$\begin{aligned}
 EPW &= \int_0^{\infty} \left(\frac{1}{\lambda} - x \right) \left[\sum_{k=0}^n (-1)^k \binom{n}{k} e^{-\lambda k x / c_i} \right] \\
 &\quad \cdot \frac{\lambda}{c_o} e^{-\lambda x / c_o} dx \\
 &= \int_0^{\infty} \left(\frac{1}{\lambda} - x \right) \frac{\lambda}{c_o} \left[\sum_{k=0}^n (-1)^k \binom{n}{k} e^{-\lambda x \left(\frac{k}{c_i} + \frac{1}{c_o} \right)} \right] dx \\
 &= \int_0^{\infty} \frac{1}{c_o} \left[\sum_{k=0}^n (-1)^k \binom{n}{k} e^{-\lambda x \left(\frac{k}{c_i} + \frac{1}{c_o} \right)} \right] dx \\
 &\quad - \int_0^{\infty} \frac{x \lambda}{c_o} \left[\sum_{k=0}^n (-1)^k \binom{n}{k} e^{-\lambda x \left(\frac{k}{c_i} + \frac{1}{c_o} \right)} \right] dx \\
 &= \frac{1}{\lambda} \sum_{k=0}^n \frac{(-1)^k \binom{n}{k}}{\left(\frac{k c_o}{c_i} + 1 \right)} - \frac{c_o}{\lambda} \sum_{k=0}^n \frac{(-1)^k \binom{n}{k}}{\left(\frac{k c_o}{c_i} + 1 \right)^2} \\
 &= \frac{1}{\lambda} \sum_{k=0}^n \left\{ \frac{(-1)^k \binom{n}{k}}{\left(\frac{k c_o}{c_i} + 1 \right)} \left[1 - \frac{c_o}{\left(\frac{k c_o}{c_i} + 1 \right)} \right] \right\} \quad \dots \dots \dots (A-3)
 \end{aligned}$$

In the computations we will normalize by setting

$$\frac{1}{\lambda} = 1.0.$$

Mathematics does not interpret anything. People have to do that. Look at Fig. 5, which shows the results of the computations. For purposes of this example, we have chosen to let all opponents use exactly the same strategy: each bids one-half of his particular value estimate. We consider that all opponents have information of equal quality and that the mean of the distribution from which their value estimate comes is the true tract value. We plot our optimum bid level (bid level that maximizes our EPW) vs the number of opponents we face.

At the left of the graph you see that for no opposition the mathematics says to bid a penny. That will be the highest bid and will win. In reality that will not work. The selling authority may set a minimum bid. It may also choose, for one reason or another, not to honor the highest bid. But then no one seriously proposes the use of a competitive bidding model when there is no competition.

The optimum bid level goes up (maximum of 0.28) until the number of opponents reaches two, whereupon it begins its descent. We interpret the curve to be saying that we should bid fairly low if the number of opponents is very small (like one) because there is a good chance that we will be able to pick up a bargain or two. The mathematics appears to be telling us that if we bid any higher, we will just be leaving money on the table. The more competitors we have,




Fig. 5—Optimum bid level vs number of opponents.

the less chance there is for bargains and the higher we must bid to get the property (make our investment). This is the kind of influence of increasing competition that most people see immediately. We call it competitive influence of the first kind.

But we see that after the second opponent the optimum bid level begins to fall. For 12 opponents it has dropped to only 0.15 — about half the maximum it achieved for two opponents. A counter-influence has begun to dominate. The tracts we win tend to be those on which we have overestimated value. The more opponents, the worse our error on the average when we win. We call this competitive influence of the second kind.

Both competitive influences are always present. They do not, however, always “weigh” the same. For most competitive situations, we think competitive influence of the second kind is more important.

The purist may be unhappy that we have drawn a curve through our computed points, giving values for such impossibilities as 3.33 opponents. In setting up a strategy, however, we are never certain of how many competitors we will face on a given parcel. If we thought there was a one-third chance of facing four opponents and a two-thirds chance of facing three opponents, then we would be justified in “expecting” 3.33 opponents. For actual computing with the formula just derived, we should be able to switch from factorials to gamma functions if we expect fractional opponents.

We would get somewhat different pictures if we altered the strategies of our opponents, but the principal characteristics that we used to illustrate the two kinds of competitive influences would remain.

Our simulations using the log-normal distribution show results similar to the ones in this analysis. That is not too surprising. As we pointed out earlier, the log-normal and the exponential have some important similarities. Furthermore, the simulation we carry out is really a numerical integration of the kinds of factors we have examined analytically in this Appendix.

JPT

Original manuscript received in Society of Petroleum Engineers office Aug. 7, 1970. Revised manuscript received March 31, 1971. Paper (SPE 2993) was presented at SPE 45th Annual Fall Meeting, held in Houston, Oct. 4-7, 1970. © Copyright 1971 American Institute of Mining, Metallurgical, and Petroleum Engineers, Inc.

DECISION ANALYSIS: PERSPECTIVES ON INFERENCE, DECISION, AND EXPERIMENTATION

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from Proceedings
of the IEEE, Vol. 58, No. 5, May 1970.
pp. 632-643.

Copyright © 1970 The Institute of Electrical
and Electronics Engineers, Inc.

Decision Analysis: Perspectives on Inference, Decision, and Experimentation

RONALD A. HOWARD, SENIOR MEMBER, IEEE

Invited Paper

Abstract—This paper illustrates by using a simple coin-tossing example how the new discipline of decision analysis sheds light on the perennial problems of inference, decision, and experimentation. The inference problem is first discussed from the classical viewpoints of maximum likelihood estimation and hypothesis testing, and then from the viewpoint of subjective probability and Bayesian updating. The problem is next placed in a decision setting to demonstrate how an estimate is related to the nature of the loss structure. Experimental possibilities are evaluated for the case where the size of the experiment must be determined a priori and for the case where experimentation can cease at any point. The decision-analysis philosophy allows consideration of all these problems within one philosophical and methodological framework.

I. INTRODUCTION

DECISION analysis is a term used to describe a new professional field that combines the viewpoints of decision theory and systems analysis. From decision theory, decision analysts have learned how to be rational in simple, but uncertain situations, while from systems analysis they have learned how to extend these insights to complex and dynamic problems. Decision analysis is usually performed on large, one-of-a-kind decisions, but it is equally applicable where decisions are repetitive. The last decade has seen decision analysis progress from a curiosity to a working tool of the largest private and public enterprises.

The communication engineer will find that the decision-analysis viewpoint shows him how to treat inference and decision problems when there is little or no experimental information available. Indeed, the cornerstones of decision analysis are that any information received can be incorporated in the process of inference and that a decision can be made given any state of information.

Many aspects of decision analysis have been discussed in a recent publication [1]. In particular, this publication described the professional aspects of the field, the question of how to treat large problems, and the nature of successful applications. Since the present readership is likely to be composed of individuals who are well aware of systems analysis, I would like to take this opportunity to illustrate how the decision theoretic side of decision analysis provides a broad conceptual framework for viewing the problems of inference, decision, and experimentation. The vehicle for this illustration will be an example based on the simplest of probabilistic phenomena, the Bernoulli trial.

Manuscript received March 3, 1970. This research was partially supported by the National Science Foundation under Grant NSF-GK-16125.

The author is with the Department of Engineering-Economic Systems, Stanford University, Stanford, Calif.

II. A COIN-TOSSING PROBLEM

Suppose that we are sitting with a group of our associates and that someone raises the question of the fairness of coins. In the course of this discussion a 1964 U. S. \$0.25 piece is offered for experimental purposes. It is examined and found to be free from visible defects and in possession of both a head and a tail. The coin is then tossed spinning into the air and allowed to fall on the floor. In a sequence of 100 tosses, 54 heads are observed. The question arises, how should we estimate the probability of heads for this coin.

When I have asked this question of a group, a majority say "0.50," a minority say "0.54," and a few give an answer like "0.51" or "0.52." The reason for their diverse answers seems to lie in the inferential systems they are using. We shall discuss a few of these systems and compare their advantages.

A. Hypothesis Testing

One viewpoint on inference is that of the hypothesis tester. In this case he might say, "All coins are fair unless proved otherwise. Unless observing 54 heads in 100 tosses is a *rare* result for a fair coin, I am going to say that this coin is fair." Taking his remarks at face value we might compute the probability that 100 tosses of a fair coin will produce exactly 54 heads. From the binomial distribution we find

$$\begin{aligned} \mathcal{P} \{r \text{ successes in } n \text{ Bernoulli trials with} \\ \text{probability of success } p \text{ on each trial}\} \\ = p_h(r|n, p) = \binom{n}{r} p^r (1-p)^{n-r}. \end{aligned} \quad (1)$$

Identifying a head with a success we have

$$p_h(54|100, 1/2) = 0.058. \quad (2)$$

When we point out to the hypothesis tester that this event is not very likely, he says that we misunderstood him; that if the actual number of heads is close enough to the expected number for a fair coin, he will accept the hypothesis that the coin is fair. He further explains that the observed number of heads will be considered close enough to the expected number if the probability that a fair coin would produce a number of heads deviating from the expected number by at least as much as the observed number is at least 0.05. Since the expected number and standard deviation of the number of heads in 100 tosses of a fair coin are 50 and 5, it turns out that accepting the hypothesis will require a number of heads between 40 and 60. That is, there is only a probability 0.05 that a fair coin would produce fewer than

40 or more than 60 heads. Since the coin in question produced a number of heads within the 40 to 60 range, the hypothesis tester would not reject the hypothesis that the coin is fair.

The choice of the probability that will distinguish the rare event is, of course, open to question. It is usually taken as 0.05 or 0.01, the latter choice broadening the range of outcomes that will not cause rejection of the hypothesis.

The group members who estimated 0.50 say that their estimate was based on the general idea of hypothesis testing.

B. Likelihood Maximization

Another viewpoint on inference is that of the likelihood maximizer. He might say, "The results of tossing the coin 100 times are all the information I have about it. I should estimate the probability of heads to be the value of that quantity that would make the experimental result most probable." Since the experimental result is specified by the number of tosses n and the number of successes r , he fixes these quantities in (1) and finds the value of p that makes the observed experimental result most probable. By differentiating and setting the derivative to zero, he finds that the maximizing value of p is r/n , the empirical fraction of successes. This value of p is called the maximum likelihood estimate and is 0.54 in the case of the coin. Note that it depends only on the observed data.

Group members who estimated 0.54 say that their reasoning was essentially that of the likelihood maximizer.

C. Characteristics of Classical Approaches

The classical approaches of hypothesis testing and likelihood maximization have difficulties as logical solutions to the inference problem. Sometimes these difficulties are revealed by supposing that we were tossing a thumbtack rather than a coin, with point-down corresponding to heads and point-up corresponding to tails. For example, in the case of hypothesis testing, we must ask, "How is the hypothesis to be constructed?" While this might be an obvious matter in the case of the coin, it is far from obvious for a thumbtack. Hypothesis testing also suffers from the need to specify the probability of a rare outcome. What is acceptable as a rare outcome in betting in the coffee pool may not be acceptable in matters of safety of life and limb. Finally, hypothesis testing becomes nonoperational if the hypothesis is rejected. Suppose, for example, that 61 heads were observed in the 100 tosses. The hypothesis of a fair coin is rejected—and now what? One can imagine a procedure for constructing new hypotheses *post hoc* until one is acceptable, but this procedure would have few defenders.

Maximum likelihood has difficulties when the experiment is small. For example, would seeing one head on one toss of a coin or thumbtack justify estimating the probability to be $r/n = 1$? It is easy to think of cases where information available before the experiment casts serious doubt on the maximum likelihood estimate.

Although hypothesis testing may be an acceptable estimation procedure when the sample size is small (because the hypothesis will almost surely not be rejected) and maximum likelihood may be acceptable when the sample size is large

(because its validity increases with sample size), we would prefer to have an inferential methodology that made logical sense for any size experiment, from no tosses to an endless stream. Fortunately, just such an approach is available, based on the concepts of probability developed by Bayes and Laplace.

III. INFERENCE

To study inference, and indeed, the allied problems of decision-making and experimental design, we shall need a convenient notation. We call this notation inferential notation and define it as follows.

A. Inferential Notation

The basic concept of inferential notation is that every probability assignment is conditional on some state of information, which we may describe generically by \mathcal{S} . If A is some event, we define $\{A|\mathcal{S}\}$ to be the probability of A given the state of information \mathcal{S} . If x is a random variable, then we define $\{x|\mathcal{S}\}$ to be the density function of x given \mathcal{S} . In most cases we can ignore the distinction between events and random variables up to the point of computation. By extension, if y is another random variable then $\{x, y|\mathcal{S}\}$ is the joint density function of x and y , $\{x|y, \mathcal{S}\}$ is the conditional density function of x given y , etc.

A particularly important state of information is the prior experience brought to the problem, which we define as \mathcal{E} . Any probability assignment conditional only on \mathcal{E} is called a prior assignment. Thus, $\{A|\mathcal{E}\}$ is the prior probability of the event A ; $\{x|\mathcal{E}\}$ is the prior probability density of the random variable x .

The notation extends to moments. We define $\langle x|\mathcal{S}\rangle$ to be the expectation of the random variable x given the state of information \mathcal{S} , computed from

$$\langle x|\mathcal{S}\rangle = \int_x x\{x|\mathcal{S}\} \tag{3}$$

where \int_x is a general summation operator. Then $\langle x^n|\mathcal{S}\rangle$ is the n th moment of x . We denoted the variance of x by

$$\mathcal{V}\langle x|\mathcal{S}\rangle = \langle x^2|\mathcal{S}\rangle - \langle x|\mathcal{S}\rangle^2 = \langle (x - \langle x|\mathcal{S}\rangle)^2|\mathcal{S}\rangle. \tag{4}$$

Two operations are especially important. The first is Bayes' theorem for determining the effect on knowledge of a variable x of observing some other variable y .

$$\{x|y, \mathcal{S}\} = \frac{\{y|x, \mathcal{S}\}\{x|\mathcal{S}\}}{\{y|\mathcal{S}\}}. \tag{5}$$

The other is expansion, which allows us to describe knowledge of one variable x in terms of knowledge of another variable y ,

$$\{x|\mathcal{S}\} = \int_y \{x|y, \mathcal{S}\}\{y|\mathcal{S}\}. \tag{6}$$

Notice that the expansion relation can be extended to expectations by multiplying both sides by x and summing over x ,

$$\int_x x \{x|\mathcal{S}\} = \int_y \int_x x \{x|y, \mathcal{S}\} \{y|\mathcal{S}\} \quad (7)$$

$$\langle x|\mathcal{S}\rangle = \int_y \langle x|y, \mathcal{S}\rangle \{y|\mathcal{S}\}.$$

All these relations can be considered simply as consequences of elementary probability theory. However, their application often provides immediate insight into inference problems.

B. The Prior Distribution

We now come to the question of how to balance experience and evidence in the case of the coin. We could directly assign a probability of heads on the first toss, but this mechanism would not let us update our state of knowledge as the sequence of tosses developed. Consequently, let ϕ be the fraction of heads that would be observed in a very large number of tosses of the coin. No matter how knowledgeable we are about coins in general or about this coin in particular we shall have some uncertainty about ϕ . We shall represent this uncertainty by assigning a prior density function on ϕ , $\{\phi|\mathcal{E}\}$, where we restrict ϕ to satisfy $0 \leq \phi \leq 1$. We can alternately represent this information in the form of a cumulative distribution $\{\phi \leq f|\mathcal{E}\}$ plotted as a function of f . Several techniques are available for the encoding of such distributions [2] that range from interview techniques to questionnaires to interactive computer programs; we shall not comment on them further here.

The end result, however, is a probability density for ϕ based solely on what the individual knows about the device before it was tossed. It is clear that the form of the prior information will have a major effect on the shape of the distribution. For example, in the case of a coin the belief that heads were as likely as tails in the long run would produce a density function symmetrical about 0.5. The density function would not have zero width, however, unless the individual believed that it was impossible for the coin to be biased in either direction. In the case of the thumbtack we would not be surprised to find the density function asymmetrical and quite broad.

To be specific we shall choose a particular prior density function for our example, as shown in Fig. 1. It is symmetric about 0.5 and has a standard deviation of about 0.05. Such a density function would be consistent with the belief that there was a 2/3 probability that the long-run fractional number of heads would fall between 0.45 and 0.55. It represents neither the degree of knowledge we would expect of a man who had spent his life observing coin tosses at the mint nor the degree possessed by someone culturally unfamiliar with the coin-tossing process. Yet it requires that the individual assign 0.5 as the probability that the next toss will fall heads.

To see this, let H represent the event of heads on the next toss and use the expansion result to expand H in terms of ϕ ,

$$\{H|\mathcal{S}\} = \int_{\phi} \{H|\phi, \mathcal{S}\} \{\phi|\mathcal{S}\}. \quad (8)$$




Fig. 1. A prior density function.

However, if ϕ were known to the individual, he would assign ϕ as the probability of H ,

$$\{H|\phi, \mathcal{S}\} = \phi. \quad (9)$$

Therefore,

$$\{H|\mathcal{S}\} = \int_{\phi} \phi \{\phi|\mathcal{S}\} = \langle \phi|\mathcal{S}\rangle, \quad (10)$$

and we have established that the probability of a head on the next toss is just the mean of the ϕ density function, or 0.5 in the present case.

C. Learning from Observation

Now the question arises as to how knowledge of ϕ is changed by the observation of tosses. Suppose that the individual has observed one head in addition to his information \mathcal{S} . What has he learned? From Bayes' theorem

$$\{\phi|H\} = \frac{\{H|\phi, \mathcal{S}\} \{\phi|\mathcal{S}\}}{\{H|\mathcal{S}\}}. \quad (11)$$

However, as we have just found this expression can be written as

$$\{\phi|H, \mathcal{S}\} = \frac{\phi \{\phi|\mathcal{S}\}}{\langle \phi|\mathcal{S}\rangle}. \quad (12)$$

The effect of seeing a head is to multiply the ϕ density function by ϕ and to divide it by a normalization factor that is equal to its mean. Similarly, if a tail T is observed, the effect on the ϕ distribution is given by

$$\{\phi|T, \mathcal{S}\} = \frac{(1 - \phi) \{\phi|\mathcal{S}\}}{1 - \langle \phi|\mathcal{S}\rangle}. \quad (13)$$

It is multiplied by $1 - \phi$ and divided by a normalization factor equal to one minus its mean. The normalization factor is always the integral of the numerator over all values of ϕ .

We could now design a computer to perform the updating job. It would start with the prior density function $\{\phi|\mathcal{E}\}$. If a head were observed, it would multiply this density function by ϕ , and renormalize it to have total area one. Observing a tail would cause multiplication by $1 - \phi$ and the same renormalization. Fig. 2 shows that the effect of observing a mixed sequence of heads and tails is to make the ϕ density function narrower and narrower until it becomes

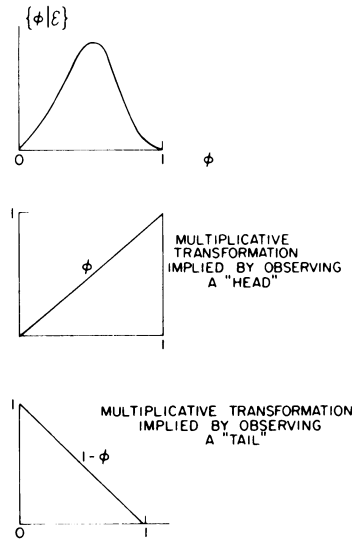


Fig. 2. The effect of observations.

in the limit a very narrow spike located at the empirical ratio of heads to total tosses.

To make this behavior clear, suppose that r heads are observed in n tosses. Then in view of the effect of observing heads and tails that we have discussed, the posterior distribution on ϕ , $\{\phi|r, n, \mathcal{E}\}$, will be given by

$$\{\phi|r, n, \mathcal{E}\} = k\phi^r(1 - \phi)^{n-r}\{\phi|\mathcal{E}\} \quad (14)$$

where k is whatever constant is required to make the posterior a density function.

D. Conjugate Distributions

While these results constitute a complete procedure for updating information, they do not have the analytical convenience necessary if we are to demonstrate decision-making and experimentation. Think how much easier our work would be if the prior and posterior density functions in the above equations could be members of the same family. This would be so if their dependence on ϕ were of the form $\phi^a(1 - \phi)^b$ where a and b are constants, for then both sides of the equation would have that form. In particular, if we choose a prior density function of the form

$$\{\phi|\mathcal{E}\} \sim \phi^{r'-1}(1 - \phi)^{n'-r'-1} \quad (15)$$

where r' and n' are constants, then the posterior density function will depend on ϕ as

$$\begin{aligned} \{\phi|r, n, \mathcal{E}\} &\sim \phi^r(1 - \phi)^{n-r}\phi^{r'-1}(1 - \phi)^{n'-r'-1} \\ &\sim \phi^{r+r'-1}(1 - \phi)^{n+n'-(r+r')-1} \\ &\sim \phi^{r''-1}(1 - \phi)^{n''-r''-1} \end{aligned} \quad (16)$$

where we have defined posterior parameters $r'' = r + r'$ and $n'' = n + n'$ to show that the posterior is, in fact, in the same family as the prior.

When we can find a family of distributions such that the prior and posterior belong to the family for some sampling process, we say that the family is a conjugate family with respect to the sampling process. We have found such a family for the Bernoulli sampling process and have shown

furthermore that the posterior parameters are obtained simply by adding the numbers that describe the experimental outcome to the prior parameters, a very simple procedure.

The conjugate family we have found is the family of beta distributions. The beta density function is defined by

$$f_\beta(\phi|r', n') = \frac{\Gamma(n')}{\Gamma(r')\Gamma(n' - r')} \phi^{r'-1}(1 - \phi)^{n'-r'-1} \quad (17)$$

$$0 \leq \phi \leq 1, 0 < r' < n'$$

with mean

$$m_\beta(r', n') = \frac{r'}{n'} \quad (18)$$

the ratio of the parameters, and variance

$$v_\beta(r', n') = \frac{r'}{n'} \left(1 - \frac{r'}{n'}\right) \frac{1}{n' + 1} \quad (19)$$

the mean times one minus the mean divided by the second parameter plus one.

Of course, it is entirely possible that the individual's state of knowledge will not be adequately described by a member of the beta family. In this case, all we can do is refer him to our general results. However, if the prior density function is well described by a beta distribution, then both the inferential and decision-making problems become much easier.

To show just how restrictive the beta family is, a representative set of beta density functions has been plotted in Fig. 3. The upper portion shows several betas with a parameter ratio and hence mean of 0.5. They range from the U-shaped (1/4, 1/2) to the quite peaked (60, 120) as the second parameter increases, and hence the variance decreases. The uniform density function is the beta (1, 2) while the two triangular density functions are the betas (1, 3) and (2, 3).

The lower portion of the figure shows a group of betas with a parameter ratio, and hence mean, of 0.05. Their mass is very much concentrated at the left end of the unit interval. As the second parameter increases and the variance decreases, the form of the beta progresses through several quite distinct forms.

Although one can imagine many prior density functions that do not fit within the beta family [for example, density functions with two peaks within the (0, 1) interval], the family does describe with ease a wide variety of possible states of information.

E. Solution of the Inference Problem

We are ready to solve the inference problem of the individual with the prior density of Fig. 1, who has now seen 54 heads in 100 tosses. We begin by determining whether his prior density function is adequately described by a beta distribution. Since there is nothing about its appearance to suggest the contrary, we might find the equivalent beta parameters by equating the mean and variance of the prior to their corresponding expressions in terms of beta parameters. The mean, 0.5, should equal r'/n' ; the variance 0.0025 should equal $(r'/n')(1 - (r'/n'))/(n' + 1)$. Solv-





Fig. 3. Beta distributions.

ing for the nearest convenient values, we find $r' = 50$, $n' = 100$,

$$\{\phi|\mathcal{E}\} = f_{\beta}(\phi|r' = 50, n' = 100), \quad (20)$$

with mean

$$\langle \phi|\mathcal{E} \rangle = \frac{r'}{n'} = 0.5, \quad (21)$$

and variance

$$\langle \phi|\mathcal{E} \rangle = \frac{r'}{n'} \left(1 - \frac{r'}{n'}\right) \frac{1}{n'+1} = 0.002475. \quad (22)$$

The ultimate test of whether this beta distribution adequately describes the original prior would be answered by plotting both it and the prior as a density or cumulative distribution and seeing if the difference is acceptable. In this case there is no problem.

Having encoded the original information as a beta distribution it is a simple matter to see how the state of information is affected by observing 54 heads in 100 tosses. We know that the posterior will also be a beta distribution and that its parameters will be numbers r'' and n'' obtained by adding the number of heads observed, r , to the prior parameter r' and the number of tosses, n , to the prior parameter n' .




Fig. 4. Prior and posterior density functions.

$$\begin{aligned} r'' &= r' + r = 50 + 54 = 104 \\ n'' &= n' + n = 100 + 100 = 200. \end{aligned} \quad (23)$$

The posterior density function is therefore a beta density function with parameters $r'' = 104$, $n'' = 200$,

$$\{\phi|r, n, \mathcal{E}\} = f_{\beta}(\phi|r'' = 104, n'' = 200). \quad (24)$$

The mean is

$$\langle \phi|r, n, \mathcal{E} \rangle = \frac{r''}{n''} = \frac{104}{200} = 0.52, \quad (25)$$

which means, of course, that if the individual were required to assign the probability of obtaining a head on the next toss, he would assign 0.52. The variance is

$$\begin{aligned} \langle \phi|r, n, \mathcal{E} \rangle &= \frac{r''}{n''} \left(1 - \frac{r''}{n''}\right) \frac{1}{n''+1} = 0.52(0.48) \frac{1}{201} \\ &= 0.001242, \end{aligned} \quad (26)$$

which is about 1/2 the prior variance. This shows that the posterior standard deviation and hence density function width is about 70 percent that of the prior. The prior and posterior density functions appear in Fig. 4.

Thus, the net effect of the observation has been to shift the prior slightly toward $\phi = 1$ and to narrow it. We see that there is a basic symmetry or duality between the roles of the prior and the experimental observations. Exactly the same posterior would have resulted if the prior had had parameters 54 and 100, and 50 heads had been observed in 100 tosses.

If an individual had a rather broad beta prior with parameters 5 and 10, the 54 head out of 100 toss result would have produced a beta posterior with parameters 59 and 110 having a mean of 0.536. He would have been more influenced by the same experiment. Conversely, an individual with a narrow beta prior with parameters (500, 1000) would have a beta posterior with parameters (554, 1100) as the result of the same experiment. Since the posterior mean would now be 0.504, he would have been far less influenced by the experimental result.

Table I shows that increasing the size of the experiment will increase its effect on the posterior. Very large experiments will result in the posterior being virtually a spike at the observed fractional number of heads.

TABLE I

POSTERIOR MEAN (PROBABILITY OF HEAD ON NEXT TOSS) AS A FUNCTION OF PRIOR PARAMETERS AND EXPERIMENTAL RESULTS

Prior Beta Parameters r' n'	Heads r Tosses n	54 100	540 1000
	5 10		0.5364
50 100		0.5200	0.5364
500 1000		0.5036	0.5200

Now that we have managed to describe an individual's state of information by the pair of numbers giving the current beta parameters, it is a simple matter to incorporate any further experimental evidence that he might see. For example, suppose that the individual whose posterior was the (104, 200) beta in Fig. 4 saw 100 additional tosses that produced 48 heads. His old posterior would become his new prior and his new posterior would be a beta distribution with parameters $104 + 48 = 152$ and $200 + 100 = 300$. Since the mean would then be $152/300 = 0.507$, the net effect of the latest data would be to shift the ϕ distribution slightly to the left, and to assign probability 0.507 to the event that the next toss will produce a head. This latter posterior would have a variance about 1/3 as great as the original beta (50, 100) prior. Any future experimental evidence would be treated in the same way.

F. Some Observations

We have now accomplished what we set out to do in the inference problem; namely, to find an inferential methodology that made as much sense for no data as it did for an overwhelming amount. No matter what size the experiment or how extensive the initial state of information, the inferential problem can be successfully treated.

While successful and complete treatment of inferential problems is a hallmark of decision analysis, the major advantages of the approach appear when there is a decision to be made. We shall show how the inferential concepts can be augmented to clarify the decision-making process.

IV. DECISION

Suppose that the coin under discussion, the 1964 \$0.25 piece, has never been tossed. A philanthropist enters the room and says that he would like to finance a game based on tossing the coin. He proposes that the coin be sent to a reputable research laboratory where it will be tossed, say, ten million times. (We might have to wait a year or two for the result.) The fractional number of heads observed in this experiment will be called ϕ . The task of the players is to guess what ϕ will be. One might think that we had already solved this problem in the last section in estimating the probability of heads. However, in that case there was no profit or loss associated with the estimate and consequently no incentive to examine how the guess should relate to the state of knowledge. Our present task is to examine this question very carefully.

A. Profit Relations

To provide us with incentive, the philanthropist announces that he will pay \$2500 to each player less an amount equal to \$1 000 000 times the square of the difference between ϕ and the player's guess g , the payment to be made when ϕ is reported. Thus, the player's profit π will be

$$\pi(\phi, g) = a - c(\phi - g)^2 \tag{27}$$

where $a = \$2500$ and $c = \$1\,000\,000$. If the player happens to guess ϕ exactly, he will earn \$2500. If he misses by 0.01 he will still earn \$2400. Missing by 0.05 causes him to break even; a miss greater than 0.05 causes him to lose money (pay the philanthropist). If he misses by 0.10, he will lose \$7500; by 0.2, \$32 500; and by 0.5, \$247 500. Finally, if the player has the bad luck to guess zero when ϕ is one, or vice versa, he will lose almost \$1 000 000. You can see that this can be an exciting game, particularly, if played with a thumbtack rather than a coin.

The first question that arises is whether one should play. While this question involves personal attitude toward risk, a subject we shall mention later, we shall simply observe that by increasing α the game can be made attractive to anyone. And so everyone must face the question, what should the guess g be.

B. Expected Profit Maximization with Quadratic Costs

Let us assume that the individual is an expected value decision maker—he desires to submit a guess g that will make his expected profit as high as possible. If we write the above equation in the form

$$\pi(\phi, g) = a - l(\phi, g), \tag{28}$$

we see that his goal is equivalent to minimizing the expected value of the loss function

$$l(\phi, g) = c(\phi - g)^2 = \langle l | \phi, g, \mathcal{E} \rangle. \tag{29}$$

The expected value of the loss function can be written using the expansion result for expectation as

$$\langle l | g, \mathcal{E} \rangle = \int_{\phi} \langle l | \phi, g, \mathcal{E} \rangle \{ \phi | g, \mathcal{E} \}. \tag{30}$$

Since we have no reason to assume that our guess will influence the final value of ϕ , $\{ \phi | g, \mathcal{E} \} = \{ \phi | \mathcal{E} \}$, we can write

$$\begin{aligned} \langle l | g, \mathcal{E} \rangle &= \int_{\phi} \langle l | \phi, g, \mathcal{E} \rangle \{ \phi | \mathcal{E} \} \\ &= \int_{\phi} c(\phi - g)^2 \{ \phi | \mathcal{E} \}. \end{aligned} \tag{31}$$

Now we can seek the value g^* for the guess that will minimize this expectation,

$$g^* = \min_g \langle l | g, \mathcal{E} \rangle. \tag{32}$$

To find it, we set the derivative of $\langle l | g, \mathcal{E} \rangle$ with respect to g to zero,

$$\begin{aligned} \frac{d}{dg} \langle l|g, \mathcal{E} \rangle &= 2c \int_{\phi} (\phi - g) \{\phi|\mathcal{E}\} = 0 \\ \int_{\phi} \phi \{\phi|\mathcal{E}\} - g \int_{\phi} \{\phi|\mathcal{E}\} &= 0 \end{aligned} \quad (33)$$

$$\langle \phi|\mathcal{E} \rangle - g = 0$$

or

$$g^* = \langle \phi|\mathcal{E} \rangle. \quad (34)$$

The best strategy for the expected value decision maker is to guess the mean of his prior. His expected loss in this case is

$$\begin{aligned} \langle l|g = g^* = \langle \phi|\mathcal{E} \rangle, \mathcal{E} \rangle &= c \int_{\phi} (\phi - \langle \phi|\mathcal{E} \rangle)^2 \{\phi|\mathcal{E}\} \\ &= c \langle (\phi - \langle \phi|\mathcal{E} \rangle)^2 | \mathcal{E} \rangle \end{aligned} \quad (35)$$

or

$$\langle l|g^*, \mathcal{E} \rangle = c \langle \phi|\mathcal{E} \rangle. \quad (36)$$

The expected loss is just c or \$1 000 000 times the prior variance. Of course, if the individual is certain of the value of ϕ , then this variance will be zero and he will expect to lose nothing. In general $c \langle \phi|\mathcal{E} \rangle$ will be the amount that the individual would pay for perfect information, or clairvoyance, on the value of ϕ .

The Coin-Tossing Decision: Suppose now that the expected value decision maker had the beta prior with parameters $r' = 50$, $n' = 100$ for the coin, and no other information. Since the mean and variance of this beta were earlier found to be 0.5 and 0.002475, this individual would guess $g = 0.5$ and would expect a loss of \$2475 to offset his fee for playing of \$2500. Consequently, the game would be a break-even proposition for him.

However, after he had seen 54 heads in 100 tosses, his beta distribution has parameters 104 and 200 with mean 0.52 and variance 0.001242. Now he would guess $g = 0.52$ and expect a loss of \$1242 to diminish his playing fee of \$2500. At this point he would have a \$1258 expected profit. If he observed succeeding tosses, he would make adjustments in his guess and anticipate that his expected profit would continually increase. Note that his guess is always equal to the probability that the next toss will fall heads.

C. Expected Profit Maximization with Bilinear Costs

To examine the generality of this result, suppose that the philanthropist changed the problem by making $l(\phi, g)$ bilinear rather than quadratic in the guessing error; that is,

$$l(\phi, g) = \langle l|\phi, g, \mathcal{E} \rangle = \begin{cases} \beta(g - \phi) & \phi < g \\ \alpha(\phi - g) & g < \phi \end{cases} \quad (37)$$

If α were twice β , for example, then it would cost twice as much to guess low by a given amount as it would to guess high by that amount. To determine the expected effect of a guess in this case we again use the expansion in terms of expectations and the assumption that the ultimate value of ϕ is independent of our guess,

$$\begin{aligned} \langle l|g, \mathcal{E} \rangle &= \int_{\phi} \langle l|\phi, g, \mathcal{E} \rangle \{\phi|\mathcal{E}\} \\ &= \int_{\phi=-\infty}^g \beta(g - \phi) \{\phi|\mathcal{E}\} + \int_{\phi=g}^{\infty} \alpha(\phi - g) \{\phi|\mathcal{E}\}. \end{aligned} \quad (38)$$

To find the guess that will maximize this quantity we differentiate and set the result to zero,

$$\begin{aligned} \beta \int_{\phi=-\infty}^g \{\phi|\mathcal{E}\} - \alpha \int_{\phi=g}^{\infty} \{\phi|\mathcal{E}\} &= 0 \\ \beta \{\phi < g|\mathcal{E}\} - \alpha \{\phi > g|\mathcal{E}\} &= 0 \end{aligned} \quad (39)$$

or

$$\frac{\{\phi < g|\mathcal{E}\}}{\{\phi > g|\mathcal{E}\}} = \frac{\alpha}{\beta}. \quad (40)$$

Since $\{\phi > g|\mathcal{E}\} = 1 - \{\phi < g|\mathcal{E}\}$, we readily convert this into

$$\{\phi < g|\mathcal{E}\} = \frac{\alpha}{\alpha + \beta}. \quad (41)$$

In other words, the best guess is the $\alpha/\alpha + \beta$ fractile of the prior distribution. If this result is substituted into the equation for $\langle l|g, \mathcal{E} \rangle$, we obtain a sum of two integrals that will specify the expected loss.

The Coin-Tossing Decision: If $\alpha = \beta$ the losses are symmetric in the error and the best guess is the median of the prior distribution rather than the mean. However, for the beta distribution (50, 100) there is no difference between the two numbers and the guess will be 0.5. Suppose, however, that α is three times β , that it costs three times as much to be in error by a certain amount if you guess low as it does if you guess high. Then the best guess is the 0.75 fractional of the prior distribution. For the beta (50, 100) this turns out to be about 0.536. Therefore, even in the absence of experimental information the guess will be above the prior mean, and, consequently, above the probability of a head on the next toss.

D. Other Cost Functions

The number to be guessed depends very strongly on both the profit structure and the state of information. We have seen how quadratic costs lead to guessing the mean and how bilinear costs lead to guessing a fractional. It is easy to see that if the cost structure placed no penalty on a guess within 0.1 percent of the actual ϕ and a very high penalty otherwise, then the best guess would be the mode (most likely value) of the prior distribution. (This is sometimes called the "William Tell" loss function.) Indeed any profit structure will imply some guesses or set of guesses that will maximize the expected profit. In this simple problem selecting the best guess specifies the decision.

E. Loss Lotteries

While we have maximized the expected profit in selecting a guess, it is not at all obvious what probability density function on profit is implied by this guess. Returning to the case of quadratic costs, we know that guessing the mean will cause a loss upon revelation of ϕ that is given by

$$\langle l|\phi, g^*, \mathcal{E} \rangle = \langle l|\phi, g^* = \langle \phi|\mathcal{E} \rangle, \mathcal{E} \rangle = c(\phi - \langle \phi|\mathcal{E} \rangle)^2. \quad (42)$$




Fig. 5. Density function of loss for the case of quadratic costs.

The loss is proportional to the square of the difference between ϕ and the prior mean. Since we have assigned a prior density function to ϕ , we have implied some density function on l given by

$$\{l|g^*, \mathcal{E}\} = \int_{\phi} \{l|\phi, g^*, \mathcal{E}\} \{\phi|\mathcal{E}\} \quad (43)$$

where $\{l|\phi, g^*, \mathcal{E}\}$ assigns probability one to the value of l specified by the deterministic relationship $\langle l|\phi, g^*, \mathcal{E}\rangle$. Consequently, finding $\{l|g^*, \mathcal{E}\}$ requires solving a change of variable problem.

A useful property of beta distributions is that they can be well approximated by a normal distribution when the second parameter is relatively large and the mean is not too close to zero or one. These conditions are met for the beta distributions in our examples. If we therefore assume that $\{\phi|\mathcal{E}\}$ is given by a normal distribution with the same mean and variance as the beta, our change of variable problem becomes one of finding the density function for the square of the difference between a normal variable and its mean. It is easy to show that this density function is the gamma density function,

$$f_l(l) = \frac{1}{\sigma\sqrt{2\pi c}} l^{-1/2} \exp\left[-\frac{l}{2c\sigma^2}\right], \quad (44)$$

where c is \$1 000 000 and σ^2 is the variance of the beta. The expected loss as derived from this density function is just $c\sigma^2$ in accordance with our earlier results. The variance of the loss is $2c^2\sigma^4$.

Fig. 5 shows the density function of loss for the case where the individual has seen the 54 heads in 100 tosses and has guessed $g=0.52$. His beta variance at this point is 0.001242 and so he expects a loss of \$1242. Of course, he will lose money out of pocket only if the loss exceeds his playing fee of \$2500. The area under the density function shows that the probability of this event is 0.155. The overall profit lottery would be obtained by creating a new ordinate at \$2500 and then reversing the abscissa.

It is interesting to note how far from a normal density function the actual loss density function is. Anyone who was accustomed to "bell-shaped" thinking and who was content to deal only with the mean (\$1242) and standard deviation (\$1756) of this lottery would have a very distorted picture of the situation he faced.

To summarize what we have learned about profit (or loss) lotteries, every decision, including the optimum one, im-

plies a deterministic relationship between the final profit and the uncertain variable. Hence, our prior on that variable implies a density function on profit that is derived as a change of variable problem. Displaying this profit lottery to the decision maker may require a combination of analytical and computational techniques.

V. EXPERIMENTATION

Perhaps the greatest advantage of the decision analysis approach appears in the analysis of experimental design and evaluation problems. The reason is that an experiment is worthwhile only if it increases the profitability of the decision more than it costs. Since the profitability of the decision depends on the state of information, the value of experimentation cannot be determined by any method that does not take explicit account of the information state. The only inferential procedure consistent with this need is the one we have described.

A. Fixed Size Experimentation

To fix ideas, suppose that the decision maker in the coin-tossing problem who has the beta (50,100) prior is given the opportunity to purchase the right to observe tosses of the coin before submitting his guess. If he chooses to purchase any samples, they will cost him \$500 plus \$4 per toss. The question is how many, if any, should he buy.

We recall that if this decision maker were required to submit a guess in the absence of sampling, he would guess 0.5 and have an expected loss of one million times his prior variance, or \$2475. Since the most he can reduce his loss by even a perfect experiment is \$2475, that number is an upper bound on the value of any experimental program. Because the opportunity he faces could cost less than \$2475 and still be informative, it is worth investigating.

If the experiment is performed the net profit to the decision maker will be

$$\pi = a - l - e \quad (45)$$

where a is the fee for playing, l is the loss from the decision, and e is the cost of the experiment. If the experiment is to take n samples, the expected profit is

$$\langle \pi|n, \mathcal{E}\rangle = a - \langle l|n, \mathcal{E}\rangle - \langle e|n, \mathcal{E}\rangle \quad (46)$$

where it is implicitly assumed that the guess will always be the mean of the beta prior on ϕ at the conclusion of the experiment. The problem is now to find the experiment size $n = n^*$ that will maximize $\langle \pi|n, \mathcal{E}\rangle$, or, equivalently, minimize the sum of the expected loss from the decision and the cost of the experiment,

$$\begin{aligned} n^* &= \max_n^{-1} \langle \pi|n, \mathcal{E}\rangle \\ &= \min_n^{-1} \langle l|n, \mathcal{E}\rangle + \langle e|n, \mathcal{E}\rangle. \end{aligned} \quad (47)$$

To evaluate $\langle l|n, \mathcal{E}\rangle$ we expand it in terms of r , the number of heads that the experiment might produce,

$$\langle l|n, \mathcal{E}\rangle = \int_{r=0}^n \langle l|r, n, \mathcal{E}\rangle \{r|n, \mathcal{E}\}. \quad (48)$$

The expected loss given that n samples have produced r heads, $\langle l|r, n, \mathcal{E} \rangle$, is equal to c times the variance of the beta posterior that will result in that case,

$$\begin{aligned} \langle l|r, n, \mathcal{E} \rangle &= c \langle \phi |r, n, \mathcal{E} \rangle \\ &= c \frac{r + r'}{n + n'} \left(1 - \frac{r + r'}{n + n'} \right) \frac{1}{n + n' + 1} \end{aligned} \quad (49)$$

To find the probability that n samples will produce r heads, $\{r|n, \mathcal{E}\}$, we expand in terms of ϕ ,

$$\{r|n, \mathcal{E}\} = \int_{\phi=0}^1 \{r|n, \phi, \mathcal{E}\} \{\phi|n, \mathcal{E}\} \quad (50)$$

The probability $\{r|n, \phi, \mathcal{E}\}$ that n Bernoulli trials will produce r heads when we know that the long-run fractional number of heads is ϕ is given by the binomial distribution,

$$\{r|n, \phi, \mathcal{E}\} = \frac{n!}{r!(n-r)!} \phi^r (1-\phi)^{n-r} \quad (51)$$

Since the mere fact that n samples are taken should not influence ϕ , $\{\phi|n, \mathcal{E}\} = \{\phi|\mathcal{E}\} = f_{\beta}(\phi|r', n')$. Consequently, we have

$$\begin{aligned} \{r|n, \mathcal{E}\} &= \int_{\phi=0}^1 \frac{n!}{r!(n-r)!} \phi^r (1-\phi)^{n-r} f_{\beta}(\phi|r', n') \\ &= \frac{n!}{r!(n-r)!} \frac{\Gamma(n')}{\Gamma(r')\Gamma(n'-r')} \int_0^1 d\phi \phi^r (1-\phi)^{n-r} \\ &\quad \cdot \phi^{r'-1} (1-\phi)^{n'-r'-1} \\ &= \frac{n!}{r!(n-r)!} \frac{\Gamma(n')}{\Gamma(r')\Gamma(n'-r')} \frac{\Gamma(r+r')\Gamma(n+n'-r-r')}{\Gamma(n+n')} \\ &= \frac{n!(n'-1)!(r+r'-1)!(n+n'-r-r'-1)!}{r!(n-r)!(r'-1)!(n'-r'-1)!(n+n'-1)!} \end{aligned} \quad (52)$$

$r=0, 1, 2, \dots, n$

where in the last step we have assumed that r' and n' are integers. The expression for $\{r|n, \mathcal{E}\}$ is called the beta-binomial distribution. It answers the question, if the probability of success for a sequence of Bernoulli trials is selected from a beta distribution with parameters r' and n' , what is the probability that r successes will be observed in n trials.

Two special cases may make this complex result more understandable. First, if $r' = 1, n' = 2$ the beta distribution is the uniform distribution on the $(0, 1)$ interval. In this case we find

$$\{r|n, \mathcal{E}\} = \frac{1}{n+1} \quad r = 0, 1, 2, \dots, n; \quad (53)$$

all numbers of heads between 0 and n are equally likely. Second, if r' and n' are general, but $n = 1$ so that we are interested in the probability of a head on one trial, we have

$$\{r|n = 1, \mathcal{E}\} = \begin{cases} \frac{r'}{n'} & r = 1 \\ 1 - \frac{r'}{n'} & r = 0 \end{cases}, \quad (54)$$




Fig. 6. Experimental design.

in accordance with our earlier result that the probability of a head should be the mean of the ϕ distribution.

We are now ready to compute $\langle l|n, \mathcal{E} \rangle$ by substituting the results of (52) and (49) into (48). The algebra involved is considerable, but the result is amazingly simple,

$$\begin{aligned} \langle l|n, \mathcal{E} \rangle &= c \frac{r'}{n'} \left(1 - \frac{r'}{n'} \right) \frac{1}{n' + 1} \cdot \frac{n'}{n + n'} \\ &= c \langle \phi | \mathcal{E} \rangle \cdot \frac{n'}{n + n'} \\ &= \langle l | \mathcal{E} \rangle \frac{n'}{n + n'} \end{aligned} \quad (55)$$

The expected loss from the decision is just the expected loss with no sampling times a factor $n'/(n+n')$. This result checks when $n=0$ and when n is very large because in that case the expected loss should be zero.

Experimental Design: Fig. 6 illustrates the process. Here we have plotted the expected loss from the decision, which, for the prior in question, is given by

$$\langle l|n, \mathcal{E} \rangle = 2475 \cdot \frac{100}{n+100} \quad (56)$$

Next, we plot the cost of purchasing n samples,

$$\langle e|n, \mathcal{E} \rangle = \begin{cases} K + kn & n = 1, 2, 3, \dots \\ 0 & n = 0 \end{cases} \quad (57)$$

where K is the fixed cost of sampling, \$500, and k is the cost per sample, \$4. The expected total cost from the decision is then the sum of the expected loss from the decision and the cost of sampling. Note that this expected total cost is \$2475 when $n=0$ and jumps to $2475 \cdot (100/101) + 500 + 4 = \2954 when $n=1$. The minimum expected total cost is incurred when $n=149$ samples. At this point the expected loss from the decision is $2475 \times (100/249) = \994 , the sampling cost is $500 + 149 \cdot 4 = \$1096$, for a total of \$2090. Thus, we have found that if an experiment is to be performed, 149 samples should be bought. But should the experiment be performed? Yes, because the expected loss from the process, \$2090, is less than the expected loss of \$2475 without experimentation; that is, the minimum of the expected total cost curve

TABLE II
EFFECT OF SAMPLE COST

Col. 1 Cost per Sample	Col. 2 Optimum Sample Size	Col. 3 Expected Loss from Decision $\langle l n, \mathcal{E} \rangle$	Col. 4 Variable Cost of Sampling	Col. 5 Total Expected Loss	Col. 6 Maximum Tolerable Fixed Sampling Cost	Col. 7 Total Cost of Experiment with Maximum Tolerable Fixed Sampling Cost
k	$n^* = \sqrt{\frac{247500}{k}} - 100$	$\frac{247500}{n^* + 100}$	kn^*	Col. 3 + Col. 4	$2475 - \text{Col. 5}$	Col. 4 + Col. 6
1	398	496	398	894	1581	1979
2	252	702	504	1206	1269	1773
4	149	994	596	1590	885	1481
9	66	1455	594	2049	426	1020
16	25	1980	400	2380	95	495
25	0	2475	0	2475	0	0

lies below \$2475. The fixed cost of sampling K would have to be increased by $\$2475 - \$2090 = \$385$ to $\$885$ before experimentation would become unprofitable. We see that if samples cost \$4 either 149 should be bought or none—the location of the minimum is not affected by changes in the fixed cost of experimentation.

It is important to observe that even if the cost of buying n tosses were not linear in n , the logical analysis we have described would be modified only by using a different curve for $\langle e|n, \mathcal{E} \rangle$.

Marginal Analysis: Important insights into the effect of experimentation can be obtained by examining the expected reduction in loss from the decision due to purchasing one additional sample, $\langle \Delta l|n, \mathcal{E} \rangle = \langle l|n, \mathcal{E} \rangle - \langle l|n+1, \mathcal{E} \rangle$. By using (55) and writing \bar{l} for $\langle l|\mathcal{E} \rangle$ we have

$$\begin{aligned} \langle \Delta l|n, \mathcal{E} \rangle &= \bar{l} \left[\frac{n'}{n+n'} - \frac{n'}{n+n'+1} \right] \\ &= \bar{l} \frac{n'}{(n+n')(n+n'+1)}. \end{aligned} \tag{58}$$

With $\bar{l} = 2475$, $n' = 100$,

$$\langle \Delta l|n, \mathcal{E} \rangle = 2475 \frac{100}{(n+100)(n+101)}. \tag{59}$$

This quantity is just the magnitude of the derivative of $\langle l|n, \mathcal{E} \rangle$ in Fig. 6. At $n=0$ it achieves its maximum value $\langle \Delta l|n=0, \mathcal{E} \rangle = \24.50 . Since it would be unprofitable to purchase a first sample unless it cost less than this amount, \$24.50 is the maximum cost of each sample before sampling becomes unprofitable.

As n increases, $\langle \Delta l|n, \mathcal{E} \rangle$ will decrease. The number of samples purchased should be increased until the cost of buying the next one is just equal to the expected savings in loss from the decision due to having it. That is, the optimum sample size n^* (assuming sampling is profitable) is the value of n that satisfies

$$\begin{aligned} \langle \Delta l|n, \mathcal{E} \rangle &= k \\ \bar{l} \frac{n'}{(n+n')(n+n'+1)} &= k. \end{aligned} \tag{60}$$

If we assume that $n+n' \gg 1$, we have

$$\begin{aligned} \bar{l} \frac{n'}{(n+n')^2} &= k \\ (n+n')^2 &= \frac{\bar{l}n'}{k} \\ n^* &= \sqrt{\frac{\bar{l}n'}{k}} - n'. \end{aligned} \tag{61}$$

This expression shows how to compute the optimum sample size from the variable cost per sample, the expected loss without sampling, and the prior n' parameter. With $k=4$, $\bar{l}=2475$, $n'=100$, we find $n^*=149$ in agreement with the graphical result. Of course, when $n=149$, $\langle \Delta l|n, \mathcal{E} \rangle$ is equal to the cost per sample, \$4.

Effect of Sample Cost: Table II shows how changing the cost per sample k changes the number of samples purchased, the total expected loss, and hence, the maximum tolerable fixed cost of sampling. The first column shows the cost per sample; the second, the optimum sample size computed from (61) with $\bar{l}=2475$, $n'=100$. The third column shows the expected loss from the decision from (56); the fourth, the variable cost of purchasing the samples, obtained by multiplying the numbers in columns 1 and 2. The fifth column shows the sum of the expected loss from the decision and the variable cost of sampling computed by adding columns 3 and 4. By subtracting the number in column 5 from $\bar{l}=2475$ we obtain in column 6 the maximum fixed cost of sampling above which no experiment will be performed. For example, when $k=4$ we confirm our earlier results that the optimum sample size is 149 and that the maximum tolerable fixed cost of sampling is 885. The seventh column shows the total cost of the experiment with the maximum tolerable fixed sampling cost as computed by adding the numbers in columns 4 and 6.

We observe that as the cost per sample increases fewer are bought. When the cost per sample reaches \$25, none are bought in correspondence with our earlier conclusion that \$24.50 is the maximum variable cost of samples. Increasing

the cost per sample increases the expected loss from the decision.

The maximum cost of sampling occurs between $k = 4$ and $k = 9$. To find the k that will achieve the maximum we write $kn^* = \sqrt{247500k - 100k}$, differentiate with respect to k , and set the result to zero. We learn that a cost per sample of \$6.19 will maximize variable sampling cost, and that the maximum will be \$619. This would be of interest to someone who wanted to sell this experiment to our decision maker, but who was constrained to charge no fixed cost of sampling.

$$v(r', n') = \min_{c,s} \begin{cases} S: \text{stop} & c \frac{r'}{n'} \left(1 - \frac{r'}{n'}\right) \frac{1}{n' + 1} \\ C: \text{continue} & k + \frac{r'}{n'} v(r' + 1, n' + 1) + \left(1 - \frac{r'}{n'}\right) v(r', n' + 1) \quad 0 < r' < n'. \end{cases} \quad (62)$$

We see from the sixth column that the maximum tolerable fixed cost of sampling decreases as the cost per sample increases. Column 7 shows, however, that if the seller of the experiment charges just slightly less than the maximum tolerable fixed sampling cost and thus ensures that the experiment is performed, his total revenue will decrease as the cost per sample increases. This makes sense, for if he gave the samples away free except for fixed cost, the decision maker would take a very large number, receive virtually perfect information, reduce his expected loss from the decision to zero, and consequently be willing to pay a fixed cost of almost \$2475 for the privilege of having free samples.

B. Sequential Sampling

Suppose that the decision maker does not have to commit himself to a batch of tosses, but rather can stop purchasing them whenever he likes; this procedure is called sequential sampling. The question we must answer is, when should the decision maker buy another sample and when should he stop and submit his guess.

We began by assuming that there is no fixed cost of sampling and that the cost per toss is \$ k as before. The analysis of sequential experimental plans is significantly simplified when we can describe the state of knowledge of the decision maker at any time by a few numbers. This is the case in the coin-tossing problem because the individual's state of knowledge is completely described by his beta parameters, r' and n' . When the reward structure of the problem has been specified, the profitability of his actions can depend only on these parameters. If at this point the individual is offered the chance to observe one toss at a cost k , he has only two choices: 1) he can refuse, stop the experiment, and make his decision, or 2) he can buy the toss, observe it, and then decide what to do.

If he chooses the first alternative, he will guess the mean, $g^* = r'/n'$, and by so doing will have an expected loss equal to c times the variance of his prior at that point, the quantity $I = c(r'/n')(1 - (r'/n'))/(n' + 1)$. However, if he decides to continue experimentation, he must first pay k for the sample and then will observe either a head or a tail. The probability that he will observe a head is r'/n' , the mean of his current

beta. If he does observe a head, his new beta parameters will be $(r' + 1, n' + 1)$. The probability that he will observe a tail is $1 - (r'/n')$; in this case his new beta parameters are $(r', n' + 1)$.

If we define $v(r', n')$ to be the minimum expected cost of a sequential sampling plan when the current beta parameters are r' and n' , then the expected cost from continuing sampling will be $k + (r'/n')v(r' + 1, n' + 1) + (1 - (r'/n'))v(r', n' + 1)$. Of course, to satisfy its definition $v(r', n')$ must be the minimum of the expected cost incurred by sampling or stopping. Thus, $v(r', n')$ must satisfy the recursive equation,

If we could solve this equation, we would know for any state of information described by (r', n') not only the expected cost from the best sequential experimental plan, but also whether or not a sample should be bought in that position. This knowledge would constitute a solution of the problem originally posed.

To develop the solution, we are going to evaluate $v(r', n')$ over an integral grid. We realize that if we could find the function $v(r', n' + 1)$ for all $0 < r' < n' + 1$ for some value n' , then we could use the recursive equation to find the function $v(r', n')$. The process then could be repeated for smaller and smaller values of the second parameter until the entire grid was evaluated.

The problem is how to evaluate $v(r', n')$ for some value of n' . The reasoning we use is that for some sufficiently large value of n' , say, n_m , the quantity $c(r'/n')(1 - (r'/n'))/(n' + 1)$ can be made less than k , the cost of taking another sample. At this point we are sure that the choice will be made to stop and that

$$v(r', n_m) = c \frac{r'}{n_m} \left(1 - \frac{r'}{n_m}\right) \frac{1}{n_m + 1}. \quad (63)$$

It is then a simple matter to solve for $v(\cdot, n_m - 1), v(\cdot, n_m - 2), \dots$ and we have generated the complete solution.

However, there is one further improvement we can make in this procedure. From (58), we know that the expected reduction in cost from taking one sample when the current beta parameters are (r', n') is

$$\langle \Delta I | n = 0, \mathcal{E} \rangle = I \frac{1}{n' + 1} = c \frac{r'}{n'} \left(1 - \frac{r'}{n'}\right) \frac{1}{(n' + 1)^2}. \quad (64)$$

We would expect to purchase such a sample as long as this expected reduction exceeds the cost of the sample k . Thus, for any value of n' , the values of r' that will separate the regions of continuing and stopping on the (r', n') grid are those that satisfy

$$c \frac{r'}{n'} \left(1 - \frac{r'}{n'}\right) \frac{1}{(n' + 1)^2} = k. \quad (65)$$




Fig. 7. Sequential sampling regions for coin-tossing problem.

Fig. 7 illustrates the continue and stop regions for the coin-tossing problem as calculated using these methods.¹ The three curves correspond to sampling costs k of \$3, \$4, and \$5. Each is the locus of (r', n') values that satisfy (65). We shall focus on $k = \$4$, the cost per sample analyzed at length in the study of batch experiments. To use the diagram, we first locate the point corresponding to the prior, $r' = 50$, $n' = 100$, in our case. Since this point is in the continue region, a sample would be purchased. If the sample were a tail, the state of information would move to the right to the point $(r' = 50, n' = 101)$. If the sample were a head, it would move to the right and up to the point $(r' = 51, n' = 101)$. Either point would be in the continue region, so another sample would be purchased. Every sample purchased moves the state of information point one unit to the right, and one unit up only if the sample is a head; consequently, the trajectories leading from any point must lie within a fan bounded by a line with slope zero and a line with slope one. It is clear that ultimately the state of information will be in the stop region and no further tosses will be purchased. We see that the longest experiments are those for which the state of information moves along the line $r' = n' / 2$. The asymmetric nature of this diagram would disappear if we plotted the regions in the coordinates $(r' - (n'/2), n')$ rather than in the coordinates (r', n') .

If we evaluated the expected cost $v(r', n')$ associated with entering upon a sequential sampling program with a beta prior with parameters (r', n') and if we subtracted it from the expected loss without sampling, $c(r'/n')(1 - (r'/n'))/(n' + 1)$, we would have obtained the maximum fixed cost of the sequential sampling program. However, everything must end, and so we shall conclude our discussion of experimentation with the thought that the average engineer today has at his fingertips more computational power than did a research statistician twenty years ago. Sequential sampling

¹ The computation of this diagram was performed by Dr. D. W. North.

programs can now be designed "to order" as the occasion requires by anyone familiar with the type of reasoning we have used in this example.

VI. CONCLUSION

The coin-tossing problem has given us the opportunity to explore in detail the philosophy and methodology of decision analysis in questions of inference, decision, and experimentation.

Perhaps the most serious omission in the presentation is lack of a discussion of the use of utility theory to describe and implement attitudes toward risk. The reason for the omission is that the subject is too broad to be covered in this paper and yet not so generally known that familiarity with it could be assumed. In essence, utility theory shows that if a decision maker subscribes to certain simple axioms, then his preference among uncertain propositions can be specified by a function called the utility function. A typical axiom is transitivity: If he likes prospect A better than prospect B and B better than C, then he must like A better than C. The utility function implied by these axioms has the property 1) that the decision maker will prefer one lottery to another if and only if its utility is higher and 2) that the utility of any lottery is the expected utility of its prices. The use of utility theory retains the advantages of expected value computation while still allowing the choice between lotteries to be made on a basis other than expected monetary value. Since we usually find experimentally that individuals confronting a lottery will forego expected profit in exchange for decreasing the probability of large deviations from the expected profit, this result is of great practical importance. The interested reader will find several discussions of the treatment of risk attitudes in [1]. We shall comment here only that adding the possibility of risk aversion does not change in any way the spirit of what we have done—it only makes the computational requirements greater and, in some cases, eliminates the possibility of compact representations.

Decision analysis provides the most comprehensive philosophy of the decision phenomenon that has been developed up to the present time. It will be a significant achievement when major resource allocations are routinely subjected to the same degree of logical analysis we have used in examining how to bet on coin tosses.

REFERENCES

- [1] Special Issue on Decision Analysis (R. A. Howard, Guest Ed.), *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-4, September 1968.
- [2] R. A. Howard, "Decision analysis: applied decision theory," *Proc. 4th Internat. Conf. on Operations Research*, Boston, Mass.: Wiley-Interscience, 1966, pp. 55-71.
- [3] —, "Bayesian decision models for system engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-1, pp. 36-40, November 1965.
- [4] —, "Information value theory," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 22-26, August 1966.
- [5] —, "Value of information lotteries," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-3, pp. 54-60, June 1967.
- [6] H. Raiffa and R. Schlaiffer, *Applied Statistical Decision Theory*, Cambridge, Mass.: Harvard University Press, 1961.
- [7] H. Raiffa, *Decision Analysis. Introductory Lectures on Choices under Uncertainty*, Reading Mass.: Addison-Wesley, 1968.
- [8] M. Tribus, *Rational Descriptions Decisions and Designs*, New York: Pergamon, 1969.

BAYESIAN DECISION MODELS FOR SYSTEMS ENGINEERING

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from IEEE Transactions on Systems Science and Cybernetics,
Vol. SSC-1, No. 1, November 1965. pp. 36-40

Copyright © 1965 by Institute of Electrical and Electronics Engineers, Inc.

Bayesian Decision Models for Systems Engineering

RONALD A. HOWARD, MEMBER, IEEE

Abstract—This paper shows how modern developments in statistical decision theory can be applied to a typical systems engineering problem. The problem is how to design an experiment to evaluate a reliability parameter for a device and then make a decision about whether to accept a contract for the development and maintenance of a system of these devices. We introduce the concept of subjective probability distribution to permit encoding prior knowledge about the uncertainty in the process. The expected value of clairvoyance is computed as an upper bound to the value of any experimental program. The structure of decision trees serves as a means for establishing the optimum size and type of experimentation and for acting on the basis of experimental results. The subjective probability approach to decision processes allows us to consider and solve problems that previously we could not even formulate.

INTRODUCTION

IN THE PAST few years we have experienced a renaissance of comprehensive statistical models for decision problems. In brief, these are the methods pioneered by

Manuscript received October 15, 1964.
The author is with Stanford University, Palo Alto, Calif.

Bayes and Laplace two centuries ago. However, the intervening period brought development of other statistical techniques based primarily, not on probability theory, but on certain other solution principles that had an ad hoc basis. The new developments, which go under the names of Bayesian statistics, decision theory, and subjective probability, allow us to consider questions in the analysis of probabilistic decision problems which we previously could not even discuss. Although adherents of these new, or rediscovered, techniques are still in a minority in the field of statistics, there is every reason to believe that they will ultimately cause a revolution in the way statistics is viewed and used by nonstatisticians. The purpose of this paper is to indicate the character of the new approach by applying it to a problem in reliability.

THE PROBLEM

A manufacturer is offered a fixed price contract to build and maintain a system of N devices for a period of T years. Every failure in the system during the T years must

be replaced by the manufacturer at a cost of C dollars. The system will cost k_0 dollars to establish, and the price of the contract is α .

The desirability of accepting the contract depends, of course, on the number of failures to be expected during the operating period. Therein lies the rub, the manufacturer does not know for sure what the failure characteristics of the device are. He believes that the failures will be Poisson with some rate λ , but he is very unsure about the value of λ . However, he can compute the expected cost of the contract when λ is known. The expected number of failures during the contract is $N\lambda T$, each of which must be replaced at a cost C . Therefore, the total expected repair cost is $CN\lambda T$, or $k_1\lambda$, where $k_1 = CNT$. The total expected cost of fulfilling the contract is then $k_0 + k_1\lambda$.

We show in Fig. 1 how the cost and price of the contract depends on the actual value of λ and also how the profit v from the contract depends on λ . We see that, if the contract is accepted, the expected profit is $\alpha - k_0 - k_1\lambda = k_2 - k_1\lambda$, where k_2 is the difference between the price of the contract α and the cost of establishing the system k_0 . Of course, if the contract is rejected, then the profit is zero regardless of λ . The figure shows that if the company knew λ it would choose to accept the contract if λ were less than $\lambda_0 = k_2/k_1$ and to reject it if λ were greater than this value. This decision assumes that the company wants to operate so as to maximize its expected profit.




Fig. 1. Economic implications as a function of failure rate λ .

We find it reassuring to know what to do if λ were known, but in fact λ is not known. Therefore, we must find some way of summarizing what the company knows about λ so that it can make the best decision in the face of its uncertainty.

INFERENCE NOTATION

To approach this goal we develop some notation. First we let $\{x|s\}$ be the probability density of a random variable x based on the state of information s . We use $\langle x|s \rangle$ to

denote the expectation of x with the same state of information s . If A is an event, we shall use $\{A|s\}$ to represent the probability of A assigned on the basis of the information s . A particularly important state of information is \mathcal{E} , the total prior experience we bring to the problem. It contains, among other things, the problem description and any other information we have that is relevant to assigning probabilities to the random variables in the problem. We call $\{x|\mathcal{E}\}$ the prior distribution of x because its probability assignment is based only on prior experience.

The crux of our reliability problem is to assign a prior $\{x|\mathcal{E}\}$ on λ that reflects everything that the engineering staff of the company can say about the failure rate. Suppose that after much soul searching the staff produces the prior density function $\{\lambda|\mathcal{E}\}$ that appears in Fig. 2. The area under this curve between any two points $\lambda = a$ and $\lambda = b$, $b \geq a$ represents the subjective or personal probability assigned by the staff to the event that if the device were produced its failure rate would lie between a and b . Drawing this density function will not be easy, but one can make a strong case that it is the only meaningful procedure for encoding the information possessed by the engineering department.

If the density function of Fig. 2 summarizes the company's state of information about λ , what is its best policy? We shall let D_A be the event that the company decides to accept the contract and D_R be the event that it decides to reject it. Then from Fig. 1 we have,

$$\begin{aligned} \langle v|\lambda D_A \mathcal{E} \rangle &= k_2 - k_1 \lambda \\ \langle v|\lambda D_R \mathcal{E} \rangle &= 0. \end{aligned} \tag{1}$$

We can now write

$$\begin{aligned} \langle v|D_A \mathcal{E} \rangle &= \int_{\lambda=0}^{\infty} d\lambda \langle v|\lambda D_A \mathcal{E} \rangle \{\lambda|\mathcal{E}\} \\ &= \int_{\lambda=0}^{\infty} d\lambda [k_2 - k_1 \lambda] \{\lambda|\mathcal{E}\} \\ &= k_2 - k_1 \int_{\lambda=0}^{\infty} d\lambda \cdot \lambda \cdot \{\lambda|\mathcal{E}\} \\ &= k_2 - k_1 \langle \lambda|\mathcal{E} \rangle \end{aligned} \tag{2}$$

and

$$\langle v|D_R \mathcal{E} \rangle = \int_{\lambda=0}^{\infty} d\lambda \langle v|\lambda D_R \mathcal{E} \rangle \{\lambda|\mathcal{E}\} = 0. \tag{3}$$

Equation (2) shows that when the contract is accepted the expected profit depends only on the mean of the prior distribution on λ and not on any other feature of that prior. The company prefers to accept the contract when the expected profit from this decision $k_2 - k_1 \langle \lambda|\mathcal{E} \rangle$ is greater than the zero profit to be expected when it is rejected. That is, when

$$k_2 - k_1 \langle \lambda|\mathcal{E} \rangle > 0$$

or

$$\langle \lambda|\mathcal{E} \rangle < k_2/k_1 = \lambda_0$$

then the company would accept the contract. The decision hinges on whether the mean of the prior distribution on λ is greater or less than the critical value $\lambda_0 = k_2/k_1$.

The prior on λ shown in Fig. 2 has a mean that is less than λ_0 . Therefore, on the basis of this λ distribution, the company should decide to accept the contract.

The Value of Clairvoyance

Although we have provided a basis for decision, the managers of the company would be most unusual if they did not feel somewhat unhappy because they do not know what λ will, in fact, result. However, we can measure their unhappiness by imagining a friendly mystic who truly possesses clairvoyance: he can foretell the value of λ that will govern the device failures. The question is, what is the most the company should be willing to pay the mystic for his service? We shall call this sum the value of clairvoyance. We are not surprised that the company is willing to pay for this information because of the chance that the mystic will predict a value of λ which is greater than λ_0 . In this case the company would benefit because it could reject the contract and avoid the loss it would otherwise incur by accepting the contract under these conditions.

We compute the value of clairvoyance by drawing a curve that shows the amount the company would save by using clairvoyance when λ turned out to be each possible value. If we let v_c be the value of clairvoyance, then $\langle v_c | \lambda \mathcal{E} \rangle$ would be the designation for this curve; it appears in Fig. 3. We observe that, since the company has decided to accept the contract, it will save nothing if the mystic predicts $\lambda < \lambda_0$ and hence that acceptance is a good idea. However, if a mystic predicts a λ that is larger than λ_0 the company will change its decision from acceptance to rejection and will save the loss of $k_1\lambda - k_2$ which it would have experienced under acceptance. Yet the mystic's pronouncements are not known in advance: the best the company can do is assign the probability distribution $\{\lambda | \mathcal{E}\}$ over what he will say. Therefore, the expected value of clairvoyance $\langle v_c | \mathcal{E} \rangle$ is just $\langle v_c | \lambda \mathcal{E} \rangle$ weighted with respect to $\{\lambda | \mathcal{E}\}$,

$$\begin{aligned} \langle v_c | \mathcal{E} \rangle &= \int_{\lambda=0}^{\infty} d\lambda \langle v_c | \lambda \mathcal{E} \rangle \{\lambda | \mathcal{E}\} \\ &= \int_{\lambda=\lambda_0}^{\infty} d\lambda \cdot k_1(\lambda - \lambda_0) \{\lambda | \mathcal{E}\} \\ &= k_1 \int_{\lambda=\lambda_0}^{\infty} (\lambda - \lambda_0) \{\lambda | \mathcal{E}\}. \end{aligned} \tag{4}$$

Equation (4) states that the expected value of clairvoyance is the area under the product of the two curves in Figs. 2 and 3. It is clearly positive and possibly quite large if the prior has much area beyond λ_0 and the slope k_1 is large.

The expected value of clairvoyance has a very important interpretation in any decision problem. It is the maximum amount that should be spent for any experimental program that attempts to provide information about the uncertainty in the problem. Since most experiments provide much less information than clairvoyance, their experi-




Fig. 2. The prior on λ .




Fig. 3. The value of perfect information.

mental cost should be considerably below $\langle v_c | \mathcal{E} \rangle$ if they are to be valuable.

EXPERIMENTAL DESIGN

The management of the company has now computed $\langle v_c | \mathcal{E} \rangle$ and learned the dollar value to be placed on its uncertainty about λ . Suppose now that it decides to make and test some of the devices to obtain a better estimate of λ . A typical test might be to place n devices in operation and note their times of failure. Since the failure process is Poisson, the number of devices placed on test n and the sum of their times to failure τ are the only statistics required to describe the outcome of the experiment. The problem we pose is this: If $c(n)$ is the cost of conducting such an experiment of size n , what value of n describes the most profitable size of experiment for the company?

The Decision Tree

To answer this question let us consider the general structure of experimental decisions. First we select some test T from a set of tests. Then we observe a result R out of some set of results for the test we selected. Next, based on the result of the test, we take some action A from a set of alternatives open to us in the decision problem. Finally, the problem has some outcome O from a set of outcomes and we are now more or less happy with our selection of test and action. This chronological sequence of happenings is best illustrated by a decision tree like that shown in Fig. 4. The crosses indicate actual decision points where a choice can be made by the decision maker. The dots represent points where the path is selected by chance, or nature. At the tips of the tree we write $\langle v | TRA O \mathcal{E} \rangle$ the expected profit for a particular test, result, action, and outcome based on prior experience \mathcal{E} . The problem for the decision maker is to select the test T and action A that will maximize his expected profit in the face of the chance moves R and O made by nature.

The expected profit the decision maker will achieve is

$$\langle v | \mathcal{E} \rangle = \sum_{T R A O} \langle v | TRA O \mathcal{E} \rangle \{ TRA O | \mathcal{E} \}, \tag{5}$$

where S is a general summation operator over the set or variable on which it operates. We can write (5) in the form

$$\langle v|\epsilon \rangle = \sum_T \{T|\epsilon\} \sum_R \{R|T\epsilon\} \sum_A \{A|TR\epsilon\} \sum_O \{O|TRA\epsilon\} \langle v|TRAO\epsilon \rangle \quad (6)$$

The numbers $\{T|\epsilon\}$, $\{R|T\epsilon\}$, $\{A|TR\epsilon\}$, and $\{O|TRA\epsilon\}$ are the probabilities to be assigned to the four branches in the decision tree. Of course, $\{T|\epsilon\}$ and $\{A|TR\epsilon\}$ are at the command of the decision maker. He will find it to his advantage to choose with certainty the test T and action A , given R and T that have the highest expected profit a priori.

Equation (6) indicates the order of the procedure. First, we calculate $\{O|TRA\epsilon\}$, the probability that should be assigned to each outcome when the test, its result, and the decision maker's action are known. In all noncompetitive decision problems $\{O|TRA\epsilon\} = \{O|TR\epsilon\}$, since the outcome is unaffected by the action the decision maker takes. We shall leave aside for the moment the question of how to compute $\{O|TR\epsilon\}$.

If we knew $\{O|TR\epsilon\}$ we could calculate

$$\begin{aligned} \langle v|TRA\epsilon \rangle &= \sum_O \{O|TRA\epsilon\} \langle v|TRAO\epsilon \rangle \\ &= \sum_O \{O|TR\epsilon\} \langle v|TRAO\epsilon \rangle. \end{aligned} \quad (7)$$

Equation (6) then becomes

$$\langle v|\epsilon \rangle = \sum_T \{T|\epsilon\} \sum_R \{R|T\epsilon\} \sum_A \{A|TR\epsilon\} \langle v|TRA\epsilon \rangle. \quad (8)$$

Now the decision maker would choose $\{A|TR\epsilon\}$, that is, choose the action A for each test T and result R that would maximize his expected profit. We shall let $\langle v|TRA^*\epsilon \rangle$ be the expected profit that results from a test T and result R when the best action A is selected. Then

$$\langle v|\epsilon \rangle = \sum_T \{T|\epsilon\} \sum_R \{R|T\epsilon\} \langle v|TRA^*\epsilon \rangle. \quad (9)$$

If we knew the probability to assign to each result of the test T , $\{R|T\epsilon\}$, we could find the expected profit from using test T , $\langle v|TA^*\epsilon \rangle$, from

$$\langle v|TA^*\epsilon \rangle = \sum_R \{R|T\epsilon\} \langle v|TRA^*\epsilon \rangle. \quad (10)$$

Note that this expected profit assumes that the actions have been assigned optimally to tests and results. Now the decision maker chooses $\{T|\epsilon\}$, namely, picks the test T that maximizes $\langle v|TA^*\epsilon \rangle$. He has then finally obtained

$$\langle v|\epsilon \rangle = \langle v|T^*A^*\epsilon \rangle \quad (11)$$

the expected profit for the best choice of test and the best choice of action given the test and its result.

In describing this procedure, we have left unspecified the method for calculating the probabilities $\{O|TR\epsilon\}$ and $\{R|T\epsilon\}$. We find these probabilities by constructing a tree we call nature's tree for the decision problem. This tree represents the steps that nature must take chronologically




Fig. 4. The decision tree.




Fig. 5. Nature's tree.

as shown in Fig. 5. First nature must select an outcome O , then the decision maker chooses a test T , next nature specifies the result R of that test, and finally the decision maker chooses his action A . The probabilities $\{O|\epsilon\}$ are assigned directly by the decision maker in accordance with his uncertainty about the outcome that nature will produce. The probabilities $\{R|OT\epsilon\}$ specify the characteristics of the test; they state the probability that should be assigned to each possible test result when the outcome is known. The decision maker again must specify these probabilities directly.

However, now we are in a position to evaluate the probabilities $\{O|TR\epsilon\}$ are $\{R|T\epsilon\}$. First we write

$$\{R|T\epsilon\} = \sum_O \{R|OT\epsilon\} \{O|T\epsilon\}. \quad (12)$$

We use the fact that in noncompetitive decision problems the outcome does not depend on the test selected,

$$\{O|T\epsilon\} = \{O|\epsilon\}, \quad (13)$$

to produce

$$\{R|T\epsilon\} = \sum_O \{R|OT\epsilon\} \{O|\epsilon\}, \quad (14)$$

and we have found the probabilities $\{R|T\epsilon\}$ in terms of the probability structure provided by the decision maker.

Next we compute $\{O|TR\epsilon\}$ from Bayes' theorem,

$$\{O|TR\epsilon\} = \frac{\{R|OT\epsilon\} \{O|T\epsilon\}}{\{R|T\epsilon\}} \quad (15)$$

where the denominator is found from (14) and the numerator from the decision maker's probability assignments.

We have thus established a formalism that allows the best test to be selected. To summarize, the decision maker first assigns probabilities $\{O|\epsilon\}$ and $\{R|OT\epsilon\}$ in nature's tree. Then he uses (14) and (15) to express the probabilities for the chance nodes in the decision tree in terms of these specified probabilities. He places values $\langle v|TRAO\epsilon \rangle$ at the tips of the tree in accordance with the economic structure of the problem. Finally he proceeds backwards

through the decision tree, taking expectations at chance nodes and maximizing at decision nodes. The result is to produce not only the best test plan and the expected profit associated with using it, but also the rule to be used for deciding what action to take when the result of the test becomes known. Sequential testing plans and other more sophisticated tests require only slight modifications of the procedure.

Experimental Design in the Reliability Problem

We can apply the decision tree principles directly to the reliability problem. The outcome distribution $\{O|\mathcal{E}\}$ in this case is the continuous distribution $\{\lambda|\mathcal{E}\}$. The test T is specified by the number of devices n placed on test. The result R is the sum τ of the times for all n devices to fail. Thus $\{R|OT\mathcal{E}\} = \{\tau|\lambda n\mathcal{E}\}$ is the probability distribution of summed failure time τ if n units are placed on test and the failure rate is λ . By convolution of the independent times to failure we have

$$\{R|OT\mathcal{E}\} = \{\tau|\lambda n\mathcal{E}\} = \frac{\lambda(\lambda\tau)^{n-1}e^{-\lambda\tau}}{(n-1)!}. \quad (16)$$

The expected profits $\langle v|TRAO\mathcal{E}\rangle$ to write at the tip of the tree in this case depend only on the outcome and the test, not on the test result or alternative,

$$\langle v|TRAO\mathcal{E}\rangle = \langle v|TO\mathcal{E}\rangle = \langle v|n\lambda\mathcal{E}\rangle. \quad (17)$$

Since $c(n)$ is the cost of performing a test of size n we can write

$$\begin{aligned} \langle v|TO\mathcal{E}\rangle &= \langle v|O\mathcal{E}\rangle - c(n) \\ \langle v|n\lambda\mathcal{E}\rangle &= \langle v|\lambda\mathcal{E}\rangle - c(n) \end{aligned} \quad (18)$$

where $\langle v|\lambda\mathcal{E}\rangle$ is the profit the company would expect if it knew the outcome λ . The quantity $\langle v|\lambda\mathcal{E}\rangle$ is just the maximum of the accept and reject profit curves of Fig. 1.

We can now write all the probabilities for the decision tree from (14) and (15) and carry out the expectation-maximization procedure to see which test T or sample size n is best. It is quite conceivable that the best experiment is of size zero: no experiment should be performed.

Analysis of Experimental Results

Suppose now that the experimental design has been completed and that an experiment of size n_0 has been chosen as the optimum experiment. The experiment is carried out and a summed failure time τ_0 is observed. What action should the company take? Although the answer to this question is contained in our discussion of the decision tree, the actual form of the answer is instructive. Figure 6 shows how to compute the company's new feelings about λ as the result of the experiment. We first calculate $\{R|OT\mathcal{E}\} = \{\tau_0|\lambda n_0\mathcal{E}\}$ from (16). This is the probability of the experimental result given the value of λ ; we call it the likelihood function. Next, according to (15), we multiply the prior $\{\lambda|\mathcal{E}\}$ by the likelihood function $\{\tau_0|\lambda n_0\mathcal{E}\}$ and divide




Fig. 6. Analysis of experimental data.

the resulting curve by its area to normalize it and produce $\{O|TR\mathcal{E}\} = \{\lambda|n_0\tau_0\mathcal{E}\}$, the posterior on λ .

The experimental result (n_0, τ_0) indicated by the likelihood function in Fig. 6 corresponds to a rather high failure rate. The effect of combining this likelihood function with the prior distribution is to produce a posterior with most of its area lying above $\lambda_0 = k_2/k_1$. In particular, the mean of the posterior $\langle \lambda|n_0\tau_0\mathcal{E}\rangle$ is greater than λ_0 . Therefore the result of this particular experiment is to cause the company to change its mind about accepting the contract and instead to reject it.

In general, the larger the experiment, the narrower will be the likelihood function and consequently the posterior. In the limit of an infinitely large experiment, the posterior distribution becomes an impulse at the value of λ indicated by the experiment. Note, however, that even in the other extreme of no experiment, the decision structure provides a rational basis for decision in this problem.

CONCLUSION

We have only been able to sketch the generality that decision theory brings to the analysis of decision problems. There is virtually no limit to the potential application of this theory at every level from power system operation to voltage divider design. Wherever the system engineer encounters decisions in the face of uncertainty, he can now enjoy the aid of a conceptually satisfying and practically powerful inference theory.

REFERENCES

[1] H. Raiffa and R. Schlaiffer, *Applied Statistical Decision Theory*. Cambridge, Mass.: Harvard University Press, 1961.

PROXIMAL DECISION ANALYSIS

Ronald A. Howard

Department of Engineering-Economic Systems

Stanford University

Reprinted with permission from Management Science,
Vol. 17, No. 9, May 1971. pp. 507-541

Copyright © 1971 by the Institute of Management Science.
All rights reserved.

PROXIMAL DECISION ANALYSIS*†

RONALD A. HOWARD

Stanford University

This paper presents simplified techniques for analyzing the effect of uncertainty in large decision problems.

Starting with the development of approximate expressions for the moments of a value lottery, we show that the probabilistic assessments of jointly related random variables necessary for these approximations are quite reasonable in number. The concepts of risk aversion, certain equivalent, and exponential utility function then permit writing useful approximations for the certain equivalent of the value lottery.

Deterministic sensitivity analyses are described first for the case when the decision variables are fixed and then for the case when they can be changed to compensate for variations in state variables. The approximate effect and value of clairvoyance (revelation of ultimate values of uncertain variables) is derived from the original probabilistic assessment and the results of the deterministic sensitivity analysis. We next determine the approximate value of wizardry (changing uncertain variables into decision variables). The amount by which decision variables must be adjusted to account for risk aversion is established from earlier results.

The final portion of the paper discusses a simple economic example that illustrates the application of the development.

1. Introduction

This paper presents simplified techniques for analyzing the effect of uncertainty in large decision problems.

The difficulties of treating problems of this type are many. First, since the number of uncertain variables is large, a logically complete solution would require that the decision-maker assign a joint probability distribution over many variables—a task he is hard-pressed to perform in all but the simplest cases. Second, even if we had the necessary probability assignments, we would have to perform many expensive simulation runs of the decision model to achieve stable results. As in any simulation, the insight generated by the runs may not be commensurate with their expense. Third, although the effect of risk aversion can be determined from these simulations, the computational procedure is unlikely to make the effect clear. Fourth, the uncertainty in overall outcome may be unbelievably large. An executive viewing the results may think: "It couldn't come out that bad because I would have done something about it." The analysis does not generally take into account the ability to compensate for ultimate state variable changes through adjustments in decision variables.

This paper discusses an approximate method of analyzing a large decision problem that:

- (1) Requires only a reasonable number of simple subjective assignments from the decision-maker;
- (2) Allows using deterministic sensitivity analysis to estimate the effect of uncertainty and of risk aversion, and
- (3) Permits estimating the effect and hence the value of making compensating decision variable changes in the light of state variable changes.

* Received April 1970; revised December 1970.

† This research was partially supported by the National Science Foundation under Grant NSF-GK-16125 and by the Office of Naval Research under Contract ONR N00014-67-A0112-0008.

A simple example will illustrate application of the methodology.

2. Notation

Throughout the paper we shall use inferential notation [2] defined as follows. [Readers familiar with inferential notation can proceed to §3.]

One Variable

If x is a random variable and \mathcal{S} our state of knowledge, then we use $\{x | \mathcal{S}\}$ to represent the probability density function of x given the state of knowledge \mathcal{S} . The mean of x is denoted by $\langle x | \mathcal{S} \rangle$ and defined by

$$(2.1) \quad \langle x | \mathcal{S} \rangle = \int_x x \{x | \mathcal{S}\},$$

where we consider \int to be a general summation operator. When no ambiguity about state of information can arise we may use \bar{x} as an abbreviation for $\langle x | \mathcal{S} \rangle$. The second moment of x is then $\langle x^2 | \mathcal{S} \rangle$ and defined by

$$(2.2) \quad \langle x^2 | \mathcal{S} \rangle = \int_x x^2 \{x | \mathcal{S}\} = \bar{x}^2.$$

The variance of x is denoted by $\checkmark \langle x | \mathcal{S} \rangle$ and defined as usual by

$$(2.3) \quad \checkmark \langle x | \mathcal{S} \rangle = \langle (x - \langle x | \mathcal{S} \rangle)^2 | \mathcal{S} \rangle = \langle x^2 | \mathcal{S} \rangle - \langle x | \mathcal{S} \rangle^2.$$

The variance $\checkmark \langle x | \mathcal{S} \rangle$ may be indicated by \check{x} when the state of information is obvious,

$$(2.4) \quad \check{x} = \bar{x}^2 - \bar{x}^2.$$

The standard deviation of x , ${}^\sigma \langle x | \mathcal{S} \rangle$, is, of course, given by

$$(2.5) \quad {}^\sigma \langle x | \mathcal{S} \rangle = \checkmark \langle x | \mathcal{S} \rangle^{1/2}.$$

Two Variables

If we are dealing with two random variables x and y , then their joint probability density function will be $\{x, y | \mathcal{S}\}$. The conditional distribution of x given y would be $\{x | y\mathcal{S}\}$. The mean of this distribution, the conditional mean of x given y , would be

$$(2.6) \quad \langle x | y\mathcal{S} \rangle = \int_x x \{x | y\mathcal{S}\}.$$

The covariance of x and y is given the symbol ${}^{\text{cov}} \langle x, y | \mathcal{S} \rangle$; it is defined by

$$(2.7) \quad {}^{\text{cov}} \langle x, y | \mathcal{S} \rangle = \langle (x - \langle x | \mathcal{S} \rangle)(y - \langle y | \mathcal{S} \rangle) | \mathcal{S} \rangle \\ = \langle xy | \mathcal{S} \rangle - \langle x | \mathcal{S} \rangle \langle y | \mathcal{S} \rangle = {}^{\text{cov}} \langle y, x | \mathcal{S} \rangle.$$

Where the state of information is clear we can write the covariance of x and y as $\text{cov}(x, y)$,

$$(2.8) \quad \text{cov}(x, y) = \overline{xy} - \bar{x}\bar{y}.$$

Note that

$$(2.9) \quad \text{cov}(x, x) = \check{x}$$

and that when a is a constant

$$(2.10) \quad \text{cov}(a, x) = 0.$$

Several Variables

If we are dealing with several random variables x_1, x_2, \dots, x_n , we can think of them as a vector \mathbf{x} and write their joint probability density function as $\{\mathbf{x} | \mathcal{S}\}$. The vector of marginal means of the variables would then be $\langle \mathbf{x} | \mathcal{S} \rangle$,

$$(2.11) \quad \langle \mathbf{x} | \mathcal{S} \rangle = \int_{\mathbf{x}} \mathbf{x} \{\mathbf{x} | \mathcal{S}\} = \bar{\mathbf{x}}.$$

We shall call $\bar{\mathbf{x}}$ the centroid of the joint density function $\{\mathbf{x} | \mathcal{S}\}$ because it establishes its center of mass in n -dimensional space. Any other notation can be developed by analogy.

Experience

A very special state of information is the experience available at the beginning of the problem, a state we symbolize by \mathcal{E} . Any probability density function or expectation conditional only on \mathcal{E} is allowed the adjective "prior." Thus $\{x | \mathcal{E}\}$ is the prior density function of x , $\langle x | \mathcal{E} \rangle$ is the prior variance of x , and $\langle \mathbf{x} | \mathcal{E} \rangle$ is the prior vector of marginal means of \mathbf{x} .

The inferential notation will be very helpful in keeping track of what we are talking about and how much we know about it.

3. The Decision Model

We envision a decision model like that shown in Figure 3.1 and discussed in reference [2]. The environment is described by a set of N state variables s_1, s_2, \dots, s_N denoted by the state vector \mathbf{s} . The decision-maker has control of M decision variables d_1, d_2, \dots, d_M , the decision vector \mathbf{d} . When the values of \mathbf{s} and \mathbf{d} are specified a unique value v is produced,

$$(3.1) \quad v = \langle v | \mathbf{s} \mathbf{d} \mathcal{E} \rangle = v(\mathbf{s}, \mathbf{d}).$$

Unfortunately, however, the state variables are uncertain. The prior knowledge about them is summarized as $\{\mathbf{s} | \mathcal{E}\}$. Therefore, for any decision vector \mathbf{d} the value v will be uncertain and described by $\{v | \mathbf{d} \mathcal{E}\}$, the value lottery. This quantity is computed from

$$(3.2) \quad \{v | \mathbf{d} \mathcal{E}\} = \int_{\mathbf{s}} \{v | \mathbf{s} \mathbf{d} \mathcal{E}\} \{\mathbf{s} | \mathcal{E}\},$$

where $\{v | \mathbf{s} \mathbf{d} \mathcal{E}\}$ is a unit impulse at $v = v(\mathbf{s}, \mathbf{d})$.

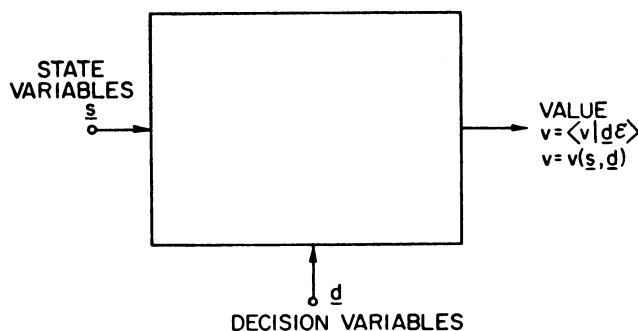


FIGURE 3.1. The decision model

The general problem we face is how to adjust the decision variables so that the distribution of v is "best" in some sense. We shall soon introduce the concept of risk aversion to make the concept of "best" precise. If the number of variables were small or the problem one of special structure we could readily derive the value lottery analytically. If $v(\mathbf{s}, \mathbf{d})$ could be quickly evaluated by a computer, then simulation would produce the value lottery very cheaply. However, it is far more common in large decision analysis problems that a single evaluation of $v(\mathbf{s}, \mathbf{d})$ will cost a significant amount of money. For example, evaluating $v(\mathbf{s}, \mathbf{d})$ may mean the deterministic simulation of a power system over many years, or perhaps running some detailed marketing model. In these cases of expensive $v(\mathbf{s}, \mathbf{d})$ evaluation we cannot afford the luxury of simulation, and must develop new methods.

The approach we shall present should be viewed in the context of expensive $v(\mathbf{s}, \mathbf{d})$ evaluation. However, we shall find that it has side benefits in terms of probability assignment and of insight.

4. An Expansion

Let us perform a Taylor series expansion of $v(\mathbf{s}, \mathbf{d})$ about the centroid $\bar{\mathbf{s}}$. Rejecting terms of higher than second degree, we have

$$(4.1) \quad v(\mathbf{s}, \mathbf{d}) \approx v(\bar{\mathbf{s}}, \mathbf{d}) + \sum_i \left. \frac{\partial v}{\partial s_i} \right|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i) + \frac{1}{2} \sum_{i,j} \left. \frac{\partial^2 v}{\partial s_i \partial s_j} \right|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i)(s_j - \bar{s}_j).$$

If we take the expectation of this equation with respect to $\{\mathbf{s} \mid \mathcal{E}\}$ and recall that the expectation of the difference between a variable and its expectation is zero, we have

$$(4.2) \quad \begin{aligned} \langle v(\mathbf{s}, \mathbf{d}) \mid \mathcal{E} \rangle &= \langle v \mid \mathcal{E} \rangle \approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_{i,j} \left. \frac{\partial^2 v}{\partial s_i \partial s_j} \right|_{\bar{\mathbf{s}}} \overline{(s_i - \bar{s}_i)(s_j - \bar{s}_j)} \\ &\approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_{i,j} \left. \frac{\partial^2 v}{\partial s_i \partial s_j} \right|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j). \end{aligned}$$

This equation shows that the mean of the value lottery is approximated by the value at the centroid plus a correction factor involving the covariance. Since the expansion in Equation (4.1) would be exact for a quadratic surface, the result of Equation (4.2) is also exact for a quadratic surface and a good approximation if the surface is almost quadratic.

If we square Equation (4.1) and retain only terms through second degree, we have

$$(4.3) \quad \begin{aligned} v^2(\mathbf{s}, \mathbf{d}) &\approx v^2(\bar{\mathbf{s}}, \mathbf{d}) + \sum_{i,j} \left. \frac{\partial v}{\partial s_i} \right|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i) \left. \frac{\partial v}{\partial s_j} \right|_{\bar{\mathbf{s}}} (s_j - \bar{s}_j) \\ &+ 2v(\bar{\mathbf{s}}, \mathbf{d}) \sum_i \left. \frac{\partial v}{\partial s_i} \right|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i) \\ &+ v(\bar{\mathbf{s}}, \mathbf{d}) \sum_{i,j} \left. \frac{\partial^2 v}{\partial s_i \partial s_j} \right|_{\bar{\mathbf{s}}} (s_i - \bar{s}_i)(s_j - \bar{s}_j). \end{aligned}$$

By taking the expectation of this equation with respect to $\{\mathbf{s} \mid \mathcal{E}\}$ we obtain an approximation for the second moment of the value lottery,

$$(4.4) \quad \begin{aligned} \langle v^2 \mid \mathcal{E} \rangle &\approx v^2(\bar{\mathbf{s}}, \mathbf{d}) + \sum_{i,j} \left. \frac{\partial v}{\partial s_i} \right|_{\bar{\mathbf{s}}} \left. \frac{\partial v}{\partial s_j} \right|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j) \\ &+ v(\bar{\mathbf{s}}, \mathbf{d}) \sum_{i,j} \left. \frac{\partial^2 v}{\partial s_i \partial s_j} \right|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j). \end{aligned}$$

Since Equation (4.3) is exact only for a planar surface $v(\mathbf{s}, d)$, this expansion is also exact only in that circumstance.

The variance of the value lottery $\langle v | \mathbf{d}\mathcal{E} \rangle$ may now be approximated by combining the results of Equations (4.2) and (4.4),

$$(4.5) \quad \begin{aligned} \langle v | \mathbf{d}\mathcal{E} \rangle &= \langle v^2 | \mathbf{d}\mathcal{E} \rangle - \langle v | \mathbf{d}\mathcal{E} \rangle^2 \\ &\approx \sum_{i, j} \frac{\partial v}{\partial s_i} \bigg|_{\bar{\mathbf{s}}} \frac{\partial v}{\partial s_j} \bigg|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j), \end{aligned}$$

where we have dropped the term in the square of the covariance. Equation (4.5) will be exact only if the surface is planar. While this might seem to be an important limitation, the alternative of extending the variance approximation to be exact for the quadratic case is not attractive because it requires assignment of third and fourth order covariances like

$$\overline{(s_i - \bar{s}_i)(s_j - \bar{s}_j)(s_k - \bar{s}_k)}$$

and

$$\overline{(s_i - \bar{s}_i)(s_j - \bar{s}_j)(s_k - \bar{s}_k)(s_l - \bar{s}_l)},$$

a difficult task for anyone no matter how great his experience in the environment of the problem.

It might seem that persuading decision-makers to assign the covariances needed in Equations (4.2) and (4.5) for the mean and variance would be a formidable task in itself. However, it is not too difficult. We recall

$$(4.6) \quad \text{cov}(s_i, s_j) = \overline{s_i s_j} - \bar{s}_i \bar{s}_j.$$

Once we have the decision-maker assign all the marginal distributions $\{s_i | \mathcal{E}\}$, $i = 1, 2, \dots, N$, we can compute the marginal means \bar{s}_i from

$$(4.7) \quad \bar{s}_i = \langle s_i | \mathcal{E} \rangle = \int_{s_i} s_i \{s_i | \mathcal{E}\}.$$

To find the covariance, we still need the expectation of the product, $\overline{s_i s_j}$. To obtain it we write

$$(4.8) \quad \overline{s_i s_j} = \langle s_i s_j | \mathcal{E} \rangle = \int_{s_i} \langle s_i s_j | s_i \mathcal{E} \rangle \{s_i | \mathcal{E}\} = \int_{s_i} s_i \langle s_j | s_i \mathcal{E} \rangle \{s_i | \mathcal{E}\},$$

which shows us how to compute the expectation of the product $s_i s_j$ from the prior $\{s_i | \mathcal{E}\}$ and the conditional expectation $\langle s_j | s_i \mathcal{E} \rangle$. Thus all we have to do is ask the decision-maker to draw a picture of how the mean of state variable s_j depends upon the value of state variable s_i . Furthermore, since the covariance is symmetric in s_i and s_j and since $\text{cov}(s_i, s_i) = \bar{s}_i$, a quantity computable from $\{s_i | \mathcal{E}\}$, we only need the conditional means $\langle s_j | s_i \mathcal{E} \rangle$ for $j < i$. When there are N state variables, this will require that the decision-maker draw $\frac{1}{2}N(N - 1)$ conditional mean pictures. Since he has already drawn N priors on the state variables, we are requiring a total of $\frac{1}{2}N(N + 1)$ graphical contributions from him. In a ten-state variable problem, this would mean 55 inputs. However, in any well-formulated problem, most of the state variables will be independent, and so the amount of effort required of the decision-maker would be considerably reduced. The important thing is to see just how much return he achieves on his effort.

5. Risk Aversion

As it is well known that most decision-makers, individual or organizational, are not willing to make decisions on the basis of expected values, it is important to have a place in our methodology for risk aversion. In this section we shall develop the notation and terminology necessary for the treatment of risk preference.

If an individual subscribes to certain preference axioms, then his preferences can be encoded in terms of a utility function on value that we shall denote by $\langle u | v\mathcal{E} \rangle$. The incorporation of \mathcal{E} in the notation makes clear the subjective nature of the evaluation; when this is not necessary, the simpler representation $u(v)$ can be used. If the individual is faced with a lottery L that will produce a value v then the expected utility for the lottery is computed from

$$(5.1) \quad \langle u | L\mathcal{E} \rangle = \int_v \langle u | v\mathcal{E} \rangle \{v | L\mathcal{E}\}.$$

If faced with a choice among lotteries, he will select the one with highest utility. In our model the decision vector \mathbf{d} determines the value lottery; therefore the utility of a particular vector can be derived from

$$(5.2) \quad \langle u | \mathbf{d}\mathcal{E} \rangle = \int_v \langle u | v\mathcal{E} \rangle \{v | \mathbf{d}\mathcal{E}\}.$$

The best decision vector is then the one that generates the highest utility.

Certain Equivalent

An important concept in our work is that of the "certain equivalent" of a lottery. The certain equivalent of a lottery L on v described by $\{v | L\mathcal{E}\}$ is denoted by $\sim\langle v | L\mathcal{E} \rangle$ and is the value of v that has the same utility as the lottery:

$$(5.3) \quad \langle u | v = \sim\langle v | L\mathcal{E} \rangle, \mathcal{E} \rangle = \langle u | L\mathcal{E} \rangle.$$

In the decision model the choice of a vector \mathbf{d} would produce some certain equivalent $\sim\langle v | \mathbf{d}\mathcal{E} \rangle$ defined by substituting \mathbf{d} for L in this equation. Since the utility curve is monotonically increasing in v , the decision vector that maximizes utility will maximize certain equivalent, and vice versa. Therefore we can formulate our problem just as well by saying that we seek the setting of the decision vector that will produce the value lottery having highest certain equivalent.

Approximation

A good approximation to the certain equivalent is obtained from the mean and variance of a lottery using the following equation [3].

$$(5.4) \quad \sim\langle v | L\mathcal{E} \rangle \approx \langle v | L\mathcal{E} \rangle - \frac{1}{2}r(\langle v | L\mathcal{E} \rangle) \langle v | L\mathcal{E} \rangle,$$

where $r(v)$ is called the risk aversion coefficient and is defined by

$$(5.5) \quad r(v) = -u''(v)/u'(v),$$

the negative ratio of second to first derivatives of the utility curve. Note that the approximation requires evaluating the risk aversion coefficient at the mean of the lottery.

The difference,

$$(5.6) \quad \langle v | L\mathcal{E} \rangle - \sim\langle v | L\mathcal{E} \rangle \approx \frac{1}{2}r(\langle v | L\mathcal{E} \rangle) \langle v | L\mathcal{E} \rangle,$$

is called the risk premium. Another quantity of note is the reciprocal of the risk aversion coefficient; it is called the risk tolerance and given the symbol $\rho(v)$,

$$(5.7) \quad \rho(v) = 1/r(v).$$

Exponential Utility

Suppose that the decision-maker will receive an amount Δ to augment the outcome of the lottery $\{v | L\mathcal{E}\}$ regardless of the value of the outcome. This creates a new lottery $\{v + \Delta | L\mathcal{E}\}$. Some decision-makers feel that their certain equivalent for the new lottery should be just Δ greater than their certain equivalent for the original lottery,

$$(5.8) \quad \tilde{\langle}v + \Delta | L\mathcal{E}\rangle = \tilde{\langle}v | L\mathcal{E}\rangle + \Delta.$$

We say that such decision-makers satisfy the “delta property.” It can be shown that adoption of the delta property severely limits the choice of utility function to either a straight line (and thus implies expected value decision-making) or to an exponential [4],

$$(5.9) \quad \langle u | v\mathcal{E}\rangle \sim e^{-\gamma v}.$$

If we normalize to the form

$$(5.10) \quad \langle u | v\mathcal{E}\rangle = u(v) = (1 - e^{-\gamma v}) / (1 - e^{-\gamma}),$$

then the straight line is included as the limiting case where $\gamma = 0$. Furthermore, this expression has the useful properties that $u(0) = 0$, $u(1) = 1$.

By applying Equation (5.5) we see that the risk aversion coefficient for such a utility curve is γ , independent of v . The approximation of Equation (5.4) then becomes

$$(5.11) \quad \tilde{\langle}v | L\mathcal{E}\rangle \approx \langle v | L\mathcal{E}\rangle - \frac{1}{2}\gamma \check{\langle}v | L\mathcal{E}\rangle$$

for the exponential case. The approximation is good if

$$(5.12) \quad \check{\langle}v | L\mathcal{E}\rangle \ll 1/\gamma^2,$$

$$(5.13) \quad \text{\textcircled{v}}\langle v | L\mathcal{E}\rangle \ll 1/\gamma;$$

that is, when the standard deviation of the lottery is small compared to the risk tolerance. Furthermore, the approximation is exact when $\{v | L\mathcal{E}\}$ is a normal distribution regardless of the variance (Appendix A).

We shall use the exponential utility curve in our developments not because the delta property is always persuasive, but because the exponential will fit adequately almost any reasonable utility curve in the region of interest. We must add the qualifying word “reasonable” because it is easy to construct a hypothetical utility function that fits the axioms and yet has a most unusual shape, for example, one that looks like the cross-section of a staircase.

However, the utility functions assessed by actual decision-makers seldom have this property [5]. They are usually smooth functions that are concave downward and representable by an exponential at least over a limited range of monetary outcomes.

Where this is not possible, the utility function can still be bounded by exponential utility functions having risk aversion coefficients that are the maximum and minimum values of risk aversion coefficient assumed by the actual utility function over the same range. The certain equivalents developed for these exponential utility functions will bound the certain equivalent for the actual utility functions over this range [3].

Risk Aversion in the Decision Model

We are now at a point where we can introduce risk aversion into the decision model. If risk aversion is described by an exponential utility curve with risk aversion coefficient γ , then Equation (5.11) shows that the certain equivalent value produced by

the decision model will be given by the approximate equation

$$(5.14) \quad \tilde{\langle v | \mathbf{d}\mathcal{E} \rangle} \approx \langle v | \mathbf{d}\mathcal{E} \rangle - \frac{1}{2} \gamma \langle v | \mathbf{d}\mathcal{E} \rangle.$$

In view of Equations (4.2) and (4.5) for the approximate mean and variance of value in the decision model, we can write the further approximate equation

$$(5.15) \quad \begin{aligned} \tilde{\langle v | \mathbf{d}\mathcal{E} \rangle} &\approx \langle v | \mathbf{d}\mathcal{E} \rangle - \frac{1}{2} \gamma \langle v | \mathbf{d}\mathcal{E} \rangle \\ &\approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \Big|_{\bar{\mathbf{s}}} \text{cov}(s_i, s_j) - \frac{1}{2} \gamma \sum_{i,j} \frac{\partial v}{\partial s_i} \Big|_{\bar{\mathbf{s}}} \frac{\partial v}{\partial s_j} \Big|_{\bar{\mathbf{s}}} \\ &\quad \cdot \text{cov}(s_i, s_j) \\ &\approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_{i,j} \left[\frac{\partial^2 v}{\partial s_i \partial s_j} \Big|_{\bar{\mathbf{s}}} - \gamma \frac{\partial v}{\partial s_i} \Big|_{\bar{\mathbf{s}}} \frac{\partial v}{\partial s_j} \Big|_{\bar{\mathbf{s}}} \right] \text{cov}(s_i, s_j). \end{aligned}$$

Note that the certain equivalent is approximated by the value at the centroid plus $\frac{1}{2}$ the sum of the covariances weighted by a factor representing the difference between the effect of nonlinearity on the mean and the effect of risk aversion. If all state variables were independent, this expression would reduce to

$$(5.16) \quad \tilde{\langle v | \mathbf{d}\mathcal{E} \rangle} \approx v(\bar{\mathbf{s}}, \mathbf{d}) + \frac{1}{2} \sum_i \left[\frac{\partial^2 v}{\partial s_i^2} \Big|_{\bar{\mathbf{s}}} - \gamma \left(\frac{\partial v}{\partial s_i} \Big|_{\bar{\mathbf{s}}} \right)^2 \right] \xi_i.$$

In most practical problems maximization of the approximate certain equivalent of Equation (5.15) will serve as well as maximizing the exact certain equivalent, and will be considerably easier.

6. Deterministic Sensitivity—Open Loop

An important question is how are the derivatives necessary in these expressions to be computed. Of course, we shall have to assume that the surface is reasonably behaved if all the derivatives are to exist. Suppose that the decision and state variables have been set to some values, ${}^0\mathbf{s}$ and ${}^0\mathbf{d}$ and that we increment them by amounts $\Delta\mathbf{s}$ and $\Delta\mathbf{d}$; that is,

$$(6.1) \quad \begin{aligned} s_i &= {}^0s_i + \Delta s_i, & i &= 1, 2, \dots, N, \\ d_k &= {}^0d_k + \Delta d_k, & k &= 1, 2, \dots, M. \end{aligned}$$

The value v will then be the value at the original setting, 0v , plus some increment Δv . By Taylor series expansion through squared terms we can relate the increase in v , Δv , to the elements of $\Delta\mathbf{s}$ and $\Delta\mathbf{d}$ by the equation

$$(6.2) \quad \begin{aligned} \Delta v &= \sum_i \frac{\partial v}{\partial s_i} \Delta s_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \Delta s_i \Delta s_j + \\ &\quad \sum_k \frac{\partial v}{\partial d_k} \Delta d_k + \frac{1}{2} \sum_{k,m} \frac{\partial^2 v}{\partial d_k \partial d_m} \Delta d_k \Delta d_m + \\ &\quad \sum_{i,k} \frac{\partial^2 v}{\partial s_i \partial d_k} \Delta s_i \Delta d_k, \quad [({}^0\mathbf{s}, {}^0\mathbf{d})]. \end{aligned}$$

The notation $[({}^0\mathbf{s}, {}^0\mathbf{d})]$ at the right means that all derivatives in this equation are evaluated at $({}^0\mathbf{s}, {}^0\mathbf{d})$, the original operating point. Therefore to evaluate derivatives at that point, all we have to do is observe in the model the Δv that results from certain combinations of small increments $\Delta\mathbf{s}$ and $\Delta\mathbf{d}$, substitute the results into Equation

(6.2) and solve for the derivatives. For example, if all but Δs_i are set equal to zero, we have

$$(6.3) \quad \Delta v = \frac{\partial v}{\partial s_i} \Delta s_i + \frac{1}{2} \frac{\partial^2 v}{\partial s_i^2} (\Delta s_i)^2, \quad [({}^0 \mathbf{s}, {}^0 \mathbf{d})].$$

By giving Δs_i a positive and a negative increment, the equation can be solved for the first and second derivatives of v with respect to s_i at ${}^0 s_i$.

Similarly, if all but one decision variable increment Δd_k is set to zero then Equation (6.2) becomes

$$(6.4) \quad \Delta v = \frac{\partial v}{\partial d_k} \Delta d_k + \frac{1}{2} \frac{\partial^2 v}{\partial d_k^2} (\Delta d_k)^2, \quad [({}^0 \mathbf{s}, {}^0 \mathbf{d})],$$

and the same procedure can be used. Now if two increments Δs_i and Δs_j are the only nonzero increments, Equation (6.2) shows that in addition to terms in derivatives already evaluated there will appear a new term in $\partial^2 v / \partial s_i \partial s_j$. The measured value of Δv at this point will then provide just the information required to evaluate $\partial^2 v / \partial s_i \partial s_j$. Similar procedures apply to estimating $\partial^2 v / \partial d_k \partial d_m$ and $\partial^2 v / \partial s_i \partial d_k$.

Computing the derivatives necessary for the approximations is a simple matter if we have a convenient means of evaluating the decision model. Appendix B describes the procedure in more detail.

Now suppose we use the centroid values $\bar{\mathbf{s}}$ for ${}^0 \mathbf{s}$ and that we adjust the decision vector \mathbf{d} until $v(\bar{\mathbf{s}}, \mathbf{d})$ is a maximum. This would be accomplished by increasing d_k whenever $\partial v / \partial d_k$ was positive, and vice versa. The resulting optimizing decision vector \mathbf{d} we designate as \mathbf{d}^* ,

$$(6.5) \quad \mathbf{d}^* = \text{Max}_{\mathbf{d}}^{-1} v(\bar{\mathbf{s}}, \mathbf{d}).$$

At the maximum,

$$(6.6) \quad \partial v / \partial d_k = 0, \quad k = 1, 2, \dots, M, \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)].$$

If we set all state variables to their mean values and determined the effect on Δv of changing one decision variable d_k , we would obtain a curve like that shown in Figure




FIGURE 6.1. Deterministic sensitivity to decision variable d_k

6.1; it would represent a standard deterministic sensitivity analysis on a decision variable. This curve is the one approximated by Equation (6.4).

If we set all incremental state variables but Δs_i to zero, maintained $\mathbf{d} = \mathbf{d}^*$, and then explored the Δv that would result from various magnitudes of Δs_i , we would generate a curve like that shown in Figure 6.2. This curve is the one approximated in Equation (6.3), and is the typical deterministic sensitivity analysis to one state variable performed in a decision problem.

Note that its curvature shows the effect of uncertainty in Δs_i upon the expected value of v . Taking the expectation of Equation (6.3) yields

$$(6.7) \quad \overline{\Delta v} = \frac{1}{2} \frac{\partial^2 v}{\partial s_i^2} \overline{(\Delta s_i)^2} = \frac{1}{2} \frac{\partial^2 v}{\partial s_i^2} \bar{s}_i, \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)].$$

If the curve is concave upward the expected value will be increased by the uncertainty; if concave downward, decreased. Of course, taking the expectation of Equation (6.2) with $\Delta \mathbf{d} = 0$ shows the general effect of uncertainty in the state variables upon the mean

$$(6.8) \quad \overline{\Delta v} = \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \overline{\Delta s_i \Delta s_j} = \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \text{cov}(s_i, s_j), \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)],$$

in accordance with Equation (4.2).

We can think of these deterministic sensitivity analyses as representing what the decision-maker would pay (or ask to be paid) to have the state and decision variables depart from their norms $\bar{\mathbf{s}}$ and \mathbf{d}^* if there were no uncertainty in the state variables. However, we realize that if the state variables should be known to be different from their means, then the setting \mathbf{d}^* may no longer be optimum. Since we do not allow such compensating changes in \mathbf{d} we call this sensitivity analysis "open loop" in control theory terminology.




FIGURE 6.2. Deterministic sensitivity to state variable s_i .

7. Deterministic Sensitivity—Closed Loop

Suppose that when state variables are found to depart from their expectations, the decision variables can be adjusted in light of the change; we call this “closed loop” sensitivity analysis. The first question is: How should the decision variables be changed in response to a change $\Delta \mathbf{s} = \mathbf{s} - \bar{\mathbf{s}}$ in the state variables?

If we consider the surface approximated by Equation (6.2), we would like to maximize Δv with respect to $\Delta \mathbf{d}$ when $\Delta \mathbf{s}$ is fixed. We do this by differentiating Equation (6.2) with respect to d_k and setting the result to zero. We discard all terms of greater than second degree and obtain,

$$(7.1) \quad \frac{\partial \Delta v}{\partial d_k} = 0 = \sum_i \frac{\partial^2 v}{\partial s_i \partial d_k} \Delta s_i + \frac{\partial^2 v}{\partial d_k^2} \Delta d_k, \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)].$$

Thus,

$$(7.2) \quad \Delta d_k = - \left(\frac{1}{\partial^2 v / \partial d_k^2} \right) \sum_i \frac{\partial^2 v}{\partial s_i \partial d_k} \Delta s_i, \quad k = 1, 2, \dots, N, \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)].$$

The adjustment in the decision variable d_k is the negative sum of the increments Δs_i weighted by the cross derivative between s_i and d_k and divided by the second derivative of v with respect to d_k . Note that we have already evaluated all these derivatives in the open loop sensitivity analysis.

We can now substitute the changes in decision variables shown in Equation (7.2) and Equation (6.2) to determine the net effect of the state variable changes. We recall that $\partial v / \partial d_k |_{\bar{\mathbf{s}}, \mathbf{d}^*} = 0$ and find

$$(7.3) \quad \begin{aligned} \Delta v = & \sum_i \frac{\partial v}{\partial s_i} \Delta s_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \Delta s_i \Delta s_j + \\ & \frac{1}{2} \sum_{k,m} \frac{\partial^2 v}{\partial d_k \partial d_m} \left(- \frac{1}{\partial^2 v / \partial d_k^2} \sum_i \frac{\partial^2 v}{\partial s_i \partial d_k} \Delta s_i \right) \left(- \frac{1}{\partial^2 v / \partial d_m^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_m} \Delta s_j \right) + \\ & \sum_{i,k} \frac{\partial^2 v}{\partial s_i \partial d_k} \Delta s_i \left(- \frac{1}{\partial^2 v / \partial d_k^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_k} \Delta s_j \right) \quad [(\bar{\mathbf{s}}, \mathbf{d}^*)] \end{aligned}$$

or

$$(7.4) \quad \begin{aligned} \Delta v = & \sum_i \frac{\partial v}{\partial s_i} \Delta s_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \Delta s_i \Delta s_j & \left. \begin{array}{l} \text{Open} \\ \text{loop} \\ \text{Sensi-} \\ \text{tivity} \end{array} \right\} & \text{CLOSED} \\ & + \frac{1}{2} \sum_{k,m} \frac{\partial^2 v / \partial d_k \partial d_m}{(\partial^2 v / \partial d_k^2)(\partial^2 v / \partial d_m^2)} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \cdot \frac{\partial^2 v}{\partial s_j \partial d_m} \Delta s_i \Delta s_j & \left. \begin{array}{l} \text{Effect} \\ \text{of} \\ \text{Com-} \\ \text{pensa-} \\ \text{tion} \end{array} \right\} & \text{LOOP} \\ & - \sum_{i,k} \frac{\partial^2 v / \partial s_i \partial d_k}{\partial^2 v / \partial d_k^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_k} \Delta s_i \Delta s_j & & \text{SENSI-} \\ & & & \text{TIVITY} \\ & & & [(\bar{\mathbf{s}}, \mathbf{d}^*)]. \end{aligned}$$

This equation shows that the closed loop sensitivity is composed of terms representing the open loop sensitivity to state variables plus terms that show the effect of compensation. Since we have already computed all the derivatives necessary in this equation, it is a simple matter to provide insight on the effect of compensation.

Because Equation (7.4) is complex, it may be helpful to write it for the case where

there is only one state variable s and one decision variable d ,

$$(7.5) \quad \underbrace{\Delta v = \frac{\partial v}{\partial s} \Delta s + \frac{1}{2} \frac{\partial^2 v}{\partial s^2} (\Delta s)^2}_{\text{Open Loop Sensitivity}} - \underbrace{\frac{1}{2} \frac{(\partial^2 v / \partial s \partial d)^2}{\partial^2 v / \partial d^2} (\Delta s)^2}_{\text{Effect of Compensation}} \quad [(\bar{s}, \mathbf{d})].$$

CLOSED LOOP SENSITIVITY

Note that there will be an effect of compensation only when $\partial^2 v / \partial s \partial d \neq 0$, that is, when there is a dependence of v on joint values of s and d .

When decision variables can be compensated for changes in state variables, the expected effect on the mean value due to uncertainty in the state variables is computed by taking the expectation of Equation (7.4),

$$(7.6) \quad \begin{aligned} \bar{\Delta v} = & \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \text{cov}(s_i, s_j) \\ & + \frac{1}{2} \sum_{k,m} \frac{\partial^2 v / \partial d_k \partial d_m}{(\partial^2 v / \partial d_k^2) (\partial^2 v / \partial d_m^2)} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial^2 v}{\partial s_j \partial d_k} \text{cov}(s_i, s_j) \\ & - \sum_{i,k} \frac{\partial^2 v / \partial s_i \partial d_k}{\partial^2 v / \partial d_k^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_k} \text{cov}(s_i, s_j), \quad [(\bar{s}, \mathbf{d}^*)]. \end{aligned}$$

8. Clairvoyance

We are now ready to discuss the value of clairvoyance. Clairvoyance is the opportunity to have revealed the actual values of the state variables that will appear in the problem. The practical impossibility of clairvoyance in most problems does not preclude its usefulness as a concept in evaluating information-gathering alternatives.

Clairvoyance has economic value only when it might change the decisions to be made; otherwise it merely satisfies curiosity. Its value lies in the difference between the value that can be attained with it and without it. For example, a decision-maker who makes decisions on the basis of expected value ($\gamma = 0$) would measure the value of clairvoyance by the increase in expected value that it brings. Equation (7.6) shows the expected increment in value that will result if the decision variables can be optimized after the state variables are observed. Equation (6.8) shows the increment in expected value due to uncertainty when clairvoyance is not available. The difference is the expected value of clairvoyance on the state vector \mathbf{s} , written $\langle v_{c_s} | \mathcal{E} \rangle$,

$$(8.1) \quad \begin{aligned} \langle v_{c_s} | \mathcal{E} \rangle = & \frac{1}{2} \sum_{k,m} \frac{\partial^2 v / \partial d_k \partial d_m}{(\partial^2 v / \partial d_k^2) (\partial^2 v / \partial d_m^2)} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial^2 v}{\partial s_j \partial d_m} \text{cov}(s_i, s_j) \\ & - \sum_{i,k} \frac{\partial^2 v / \partial s_i \partial d_k}{\partial^2 v / \partial d_k^2} \sum_j \frac{\partial^2 v}{\partial s_j \partial d_k} \text{cov}(s_i, s_j), \quad [(\bar{s}, \mathbf{d}^*)]. \end{aligned}$$

Recall, once again, that we have measured all the derivatives necessary to evaluate this expression. A succinct description of the expected value of clairvoyance is that it is the expected effect of compensation.

An important specialization occurs if all decision variables are independent, $\partial^2 v / \partial d_k \partial d_m = 0$, $k \neq m$,

$$(8.2) \quad \langle v_{c_s} | \mathcal{E} \rangle = - \frac{1}{2} \sum_k \frac{1}{\partial^2 v / \partial d_k^2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial^2 v}{\partial s_j \partial d_k} \text{cov}(s_i, s_j), \quad [(\bar{s}, \mathbf{d}^*)].$$

If, in addition, all state variables are uncorrelated, $\text{cov}(s_i, s_j) = 0$, $i \neq j$, then

$$(8.3) \quad \langle v_{c_s} | \mathcal{E} \rangle = - \frac{1}{2} \sum_k \frac{1}{\partial^2 v / \partial d_k^2} \sum_i \left(\frac{\partial^2 v}{\partial s_i \partial d_k} \right)^2 \bar{s}_i, \quad [(\bar{s}, \mathbf{d}^*)].$$

Finally, if there is only one state variable s and one decision variable d , we have

$$(8.4) \quad \langle v_{c_s} | \varepsilon \rangle = -\frac{1}{2} \frac{(\partial^2 v / \partial s \partial d)^2}{\partial^2 v / \partial d^2} \check{s}, \quad [(\bar{s}, d^*)].$$

This last form is particularly easy to interpret. The quantity $(\partial^2 v / \partial s \partial d)^2$ is, of course, nonnegative as is the variance \check{s} . Since $\partial^2 v / \partial d^2$ will be negative as we can see from Figure 6.1, the expected value of clairvoyance cannot be negative. It will be 0 if \check{s} is 0 and hence there is no uncertainty in the state variable s , or if $\partial^2 v / \partial s \partial d$ is 0 because the state and decision variables are independent.

If the decision-maker has a risk sensitive exponential utility curve ($\gamma \neq 0$), then the value of clairvoyance to him will be the increase in certain equivalent that he can achieve as a result of the clairvoyance. Since the clairvoyance does not, to a first approximation, change the variance of the value lottery, we see from Equation (5.4) that the value of clairvoyance to the risk sensitive decision-maker is substantially the same as that given by Equation (8.1) for the risk indifferent decision-maker.

9. Wizardry

A wizard is another construction: an individual who will set any or all state variables to any values we like. In essence, he makes decision variables out of state variables. A wizard is different from a clairvoyant because he eliminates uncertainty, rather than simply resolving it in advance.

The most direct way to evaluate wizardry is just to make the state variables on which the wizard will operate into decision variables and then see what increase in value would result if the settings are changed. The results of the deterministic sensitivity analysis are usually quite helpful here.

Suppose that we have the opportunity to have one state variable, say s_r , set to any value we like by the wizard. On an open loop basis where we cannot change the decision, the effect on the value will be that illustrated in Figure 6.2.

However, knowing s_i in advance will also affect the certain equivalent through the corrections of nonlinearity and risk aversion. The reason is that the uncertainty in variable s_i has been removed, and so all covariances with this variable must be equal to zero. Suppose we write the combined term for nonlinearity and risk aversion in the form

$$(9.1) \quad \sum_{i,j} f_{ij} \text{cov}(s_i, s_j),$$

where as we see from Equation (5.15)

$$(9.2) \quad f_{ij} = \frac{1}{2} \left[\frac{\partial^2 v}{\partial s_i \partial s_j} - \gamma \frac{\partial v}{\partial s_i} \frac{\partial v}{\partial s_j} \right], \quad [(\bar{s}, d^*)].$$

Since $\text{cov}(s_i, s_r) = 0$ for $i \neq r$, we can write the new correction term as

$$(9.3) \quad \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) (1 - \delta_{ir})(1 - \delta_{jr})$$

where $\delta_{ij} = 1$ if $i = j$ and zero otherwise. Then the new correction term will be

$$(9.4) \quad \begin{aligned} & \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) (1 - \delta_{ir})(1 - \delta_{jr}) \\ &= \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) - \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) \delta_{ir} \\ & \quad - \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) \delta_{jr} + \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) \delta_{ir} \delta_{jr} \\ &= \sum_{i,j} f_{ij} \text{cov}(s_i, s_j) - 2 \sum_j f_{rj} \text{cov}(s_r, s_j) + f_{rr} \check{s}_r. \end{aligned}$$

Therefore the reduction in the effects of nonlinearity and risk aversion due to wizardry

about s_r will be

$$(9.5) \quad 2 \sum_j f_{rj} \text{cov}(s_r, s_j) - f_{rr} \bar{s}_r,$$

and if s_r is uncorrelated with all other state variables, just $f_{rr} \bar{s}_r$. Thus from this calculation and the deterministic open loop sensitivity we can obtain an excellent idea of the effect of wizardry on any state variable.

Similar considerations apply if we are allowed to change the decision variables to take advantage of wizardry. The direct effect on value will be given by the closed loop sensitivity of Equation (7.4). However, the effect on the contributions of nonlinearity and risk aversion must also be considered. This is done by including in f_{ij} not only the terms in Equation (9.2), but also those new terms appearing in Equation (7.6). Thus

$$(9.6) \quad f_{ij} = \frac{1}{2} \left[\frac{\partial^2 v}{\partial s_i \partial s_j} - \gamma \frac{\partial v}{\partial s_i} \frac{\partial v}{\partial s_j} \right] + \frac{1}{2} \sum_{k,m} \frac{\partial^2 v / \partial d_k \partial d_m}{(\partial^2 v / \partial d_k^2) (\partial^2 v / \partial d_m^2)} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial^2 v}{\partial s_j \partial d_m} - \sum_k \frac{\partial^2 v / \partial s_i \partial d_k}{\partial^2 v / \partial d_k^2} \frac{\partial^2 v}{\partial s_j \partial d_k}.$$

Equations (9.4) and (9.5) then apply directly.

10. Correction of Decision Variables for Risk Aversion and Nonlinearity

Up to this time our setting of the decision vector has been the vector $\mathbf{d} = \mathbf{d}^*$ that maximizes $v(\mathbf{s}, \mathbf{d})$, as indicated in Equation (6.5). However, it is very possible that the setting to maximize certain equivalent will not be the same as \mathbf{d}^* . Let \mathbf{d}^{**} be the setting of \mathbf{d} that maximizes the certain equivalent $\bar{v}(\mathbf{d} | \mathbf{d}\mathcal{E})$,

$$(10.1) \quad \mathbf{d}^{**} = \text{Max}_{\mathbf{d}}^{-1} \bar{v}(\mathbf{d} | \mathbf{d}\mathcal{E}).$$

The question we now ask is how to find a vector $\Delta \mathbf{d}^*$ such that

$$(10.2) \quad \mathbf{d}^{**} = \mathbf{d}^* + \Delta \mathbf{d}^*.$$

In other words, how can we correct the decision vector that maximizes $v(\bar{\mathbf{s}}, \mathbf{d})$ to produce a vector that will maximize $\bar{v}(\mathbf{d} | \mathbf{d}\mathcal{E})$?

Equation (5.15) shows the relationship between $\bar{v}(\mathbf{d} | \mathbf{d}\mathcal{E})$ and $v(\bar{\mathbf{s}}, \mathbf{d})$. Writing \bar{v} for $\bar{v}(\mathbf{d} | \mathbf{d}\mathcal{E})$ we have

$$(10.3) \quad \bar{v} = v(\bar{\mathbf{s}}, \mathbf{d}) + \underbrace{\frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j} \text{cov}(s_i, s_j)}_{\text{Nonlinearity Effect}} - \underbrace{\frac{1}{2} \gamma \sum_{i,j} \frac{\partial v}{\partial s_i} \frac{\partial v}{\partial s_j} \text{cov}(s_i, s_j)}_{\text{Risk Premium}}.$$

When $\mathbf{d} = \mathbf{d}^{**}$ we shall have

$$(10.4) \quad \partial \bar{v} / \partial d_k |_{\mathbf{d}^{**}} = 0, \quad k = 1, 2, \dots, M,$$

or from Equation (10.3),

$$(10.5) \quad 0 = \frac{\partial \bar{v}}{\partial d_k} \Big|_{\mathbf{d}^{**}} = \frac{\partial v}{\partial d_k} \Big|_{\mathbf{d}^{**}} + \frac{1}{2} \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial s_j \partial d_k} \text{cov}(s_i, s_j) - \frac{1}{2} \gamma \sum_{i,j} 2 \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial v}{\partial s_j} \text{cov}(s_i, s_j).$$

However,

$$(10.6) \quad \frac{\partial v}{\partial d_k} \Big|_{\mathbf{d}^{**}} = \frac{\partial v}{\partial d_k} \Big|_{\mathbf{d}^*} + \Delta d_k^* \frac{\partial^2 v}{\partial d_k^2} \Big|_{\mathbf{d}^*}$$

where we know that $\partial v / \partial d_k |_{\mathbf{d}^*} = 0$ because \mathbf{d}^* is the decision vector that maximizes v . When we substitute this result into Equation (10.5) we find

$$(10.7) \quad 0 = \Delta d_k^* \frac{\partial^2 v}{\partial d_k^2} \Big|_{\mathbf{d}^*} + \frac{1}{2} \sum_{i,j} \frac{\partial^3 v}{\partial s_i \partial s_j \partial d_k} \text{cov}(s_i, s_j) - \gamma \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial v}{\partial s_j} \text{cov}(s_i, s_j).$$

If we assume that all higher order derivatives are equal at \mathbf{d}^* and \mathbf{d}^{**} we have

$$(10.8) \quad \Delta d_k^* = \frac{-1}{\partial^2 v / \partial d_k^2} \left[\frac{1}{2} \sum_{i,j} \frac{\partial^3 v}{\partial s_i \partial s_j \partial d_k} \text{cov}(s_i, s_j) - \gamma \sum_{i,j} \frac{\partial^2 v}{\partial s_i \partial d_k} \frac{\partial v}{\partial s_j} \text{cov}(s_i, s_j) \right], \quad [(\bar{s}, \mathbf{d}^*)].$$

This equation shows how the decision vector must be adjusted to take into account nonlinearity and risk aversion. The correction for nonlinearity involves terms of the form $\partial^3 v / \partial s_i \partial s_j \partial d_k$, derivatives that we have not previously measured. However, in practice this correction is usually so small that it may be neglected. We include it only for logical completeness.

The adjustment in decision vector is thus mainly due to risk aversion and is given by

$$(10.9) \quad \Delta d_k^* = \gamma \sum_{i,j} \frac{(\partial^2 v / \partial s_i \partial d_k) \cdot (\partial v / \partial s_j)}{\partial^2 v / \partial d_k^2} \text{cov}(s_i, s_j), \quad [(\bar{s}, \mathbf{d}^*)].$$

Since all derivatives in this expression are evaluated at the point (\bar{s}, \mathbf{d}^*) , they have previously been computed. Note that the magnitude of the correction is proportional to the risk aversion coefficient.

11. An Example—The Entrepreneur's Problem

To illustrate the concepts we have developed, let us apply them to the problem of an entrepreneur trying to decide on a price for his product. The problem is described pictorially in Figure 11.1. When the entrepreneur selects a price he determines a quantity q that he will sell from the demand curve $q(p)$. This quantity q will have a cost of




FIGURE 11.1. A pictorial representation of the entrepreneur's problem

manufacture c given by the total cost curve $c(q)$. The entrepreneur's profit π will then be the difference between his revenue pq and his cost c or

$$(11.1) \quad \pi(p) = pq(p) - c(q(p)).$$

The entrepreneur desires to find the price p that will maximize this profit.

This problem would be very simple if the demand curve and total cost curve were known with certainty, but that is seldom the case. We shall assume that the quantity $q(p)$ determined from the demand curve is only a nominal value and that the actual quantity sold will be $q(p) + \Delta q$, where Δq is a random variable. Furthermore, producing this quantity $q(p) + \Delta q$ will not cost $c(q(p) + \Delta q)$, but rather

$$c(q(p) + \Delta q) + \Delta c,$$

where Δc is another random variable. This modification of the problem appears diagrammatically as Figure 11.2. Note that the profit as a function of Δq , Δc , and p is now

$$(11.2) \quad \pi(\Delta q, \Delta c, p) = p[q(p) + \Delta q] - c(q(p) + \Delta q) - \Delta c.$$

We assume further that the random variables Δq and Δc are independent,

$$(11.3) \quad \{\Delta q, \Delta c \mid \varepsilon\} = \{\Delta q \mid \varepsilon\}\{\Delta c \mid \varepsilon\}$$

and that their prior means are both equal to zero,

$$(11.4) \quad \langle \Delta q \mid \varepsilon \rangle = \langle \Delta c \mid \varepsilon \rangle = 0.$$

Functional Forms

Our primary goal is to determine the effect of these uncertainties upon the solution of the problem. However, to pursue this goal we shall first have to solve the nominal problem originally posed. We must therefore become specific both with respect to functional forms and numerical values. We shall imagine that the product is an expensive one like a commercial air transport. We shall think of the quantity in units and the monetary amounts in millions of dollars.

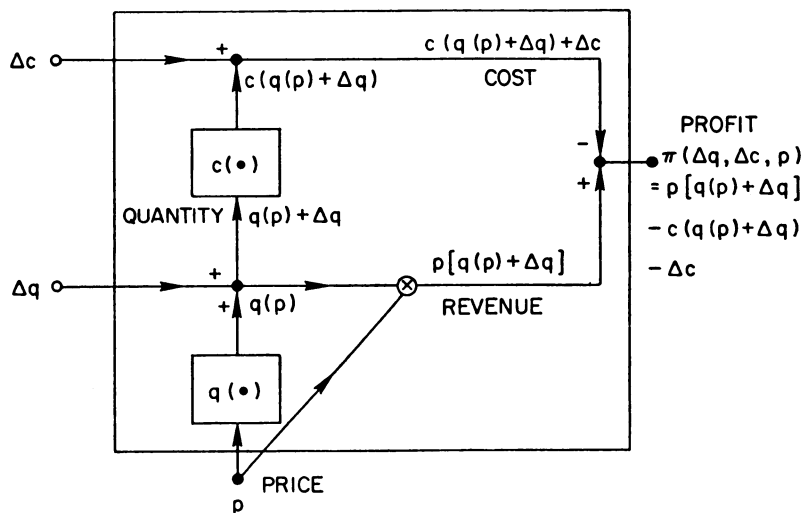


FIGURE 11.2. Uncertain perturbations in the entrepreneur's problem

PROXIMAL DECISION ANALYSIS




FIGURE 11.3. The demand curve

Demand. For the demand curve we shall choose the functional form

$$(11.5) \quad q(p) = \frac{1}{\beta} [\ln \alpha - \ln p], \quad 0 < p \leq \alpha,$$

with the constants given by $\alpha = 50$, $\beta = 1/80$. This function is plotted in Figure 11.3. We observe that the quantity falls monotonically with price and that none will be sold when the price reaches 50. For future reference, the derivatives of the demand curve are

$$(11.6) \quad q'(p) = -1/\beta p, \quad q''(p) = 1/\beta p^2.$$

Of course, we can also solve for the price as a function of the quantity

$$(11.7) \quad p(q) = \alpha e^{-\beta q}, \quad 0 \leq q,$$

to show that the selling price is an exponentially decreasing function of the quantity sold. The derivative in this form is

$$(11.8) \quad p'(q) = -\alpha \beta e^{-\beta q} = -\beta p(q).$$

The function $p(q)$ is tabulated in Table 11.1 along with several other functions that we shall define.

TABLE 11.1
Evaluation of Certain Functions in the Entrepreneur's Problem

Quantity q	Price $p(q)$	Revenue $r(q)$	Marginal Revenue $r'(q)$	Total Cost $c(q)$	Marginal Cost $c'(q)$	Profit $\pi(q)$	Marginal Profit $\pi'(q)$
0	50.00	0	50.00	700.0	12.00	-700.0	38.00
10	44.12	441.2	38.60	812.5	10.55	-371.3	28.05
20	38.94	778.8	29.20	911.9	9.36	-133.1	19.84
30	34.37	1031.0	21.47	1000.5	8.39	29.5	13.08
40	30.33	1213.0	15.17	1080.3	7.59	132.7	7.58
50	26.76	1338.0	10.01	1152.8	6.94	185.2	3.07
60	23.62	1417.2	5.92	1219.5	6.41	197.7	-0.49
70	20.84	1458.8	2.57	1281.4	5.97	177.4	-3.40
80	18.40	1471.6	0	1339.2	5.62	132.4	-5.62
90	16.23	1460.7	-2.04	1393.9	5.32	66.8	-7.36
100	14.33	1432.5	-3.58	1445.9	5.08	-13.4	-8.66
110	12.64	1390.4	-4.74	1495.7	4.89	-105.3	-9.63
120	11.16	1338.6	-5.53	1543.7	4.73	-205.1	-10.26
130	9.85	1279.9	-6.15	1590.3	4.59	-310.4	-10.74
140	8.69	1216.6	-6.57	1635.7	4.49	-419.1	-11.06
150	7.67	1149.8	-6.74	1680.0	4.40	-530.2	-11.14
160	6.77	1082.4	-6.84	1723.7	4.33	-641.3	-11.17

Revenue. The revenue $r(q)$ associated with selling q is just

$$(11.9) \quad r(q) = qp(q) = \alpha qe^{-\beta q},$$

with derivative

$$(11.10) \quad \begin{aligned} r'(q) &= p(q) + qp'(q) \\ &= p(q)[1 - \beta q]. \end{aligned}$$

The revenue function is plotted in Figure 11.4.

Cost. For the total cost function we assume the form

$$(11.11) \quad c(q) = k_0 + k_1q + k_2(1 - \exp(-k_3q))$$

with constants

$$(11.12) \quad k_0 = 700, \quad k_1 = 4, \quad k_2 = 400, \quad k_3 = \frac{1}{50}.$$

The plot of this function in Figure 11.4 shows that total costs are continually increasing, as we would expect. The marginal cost

$$(11.13) \quad c'(q) = k_1 + k_2k_3 \exp(-k_3q)$$

approaches $k_1 = 4$ as the quantity becomes large. The second derivative

$$(11.14) \quad c''(q) = -k_2k_3^2 \exp(-k_3q)$$

shows that the marginal cost is always decreasing.

Profit. Once the revenue and cost at a given quantity are known, the profit at that quantity is just the difference

$$(11.15) \quad \pi(q) = r(q) - c(q),$$

a difference plotted in Figure 11.5. We see that a maximum profit of 198 will occur when the quantity sold is 58.5. The venture appears to be profitable for sales between about 28 and 97.




FIGURE 11.4. Revenue and total cost curves

We gain more insight into the solution by considering marginal quantities. By differentiating Equation (11.15) we have

$$(11.16) \quad \pi'(q) = r'(q) - c'(q),$$

where all three terms appear as curves in Figure 11.6. The maximum profit occurs when the marginal profit is zero; namely, at $q = 58.5$. This point is also the one where marginal revenue and marginal cost curves intersect, because we wish to increase sales until the increase in revenue from selling the unit is just equal to the increase in the cost of making it. The curve $p(q)$, the demand curve in alternate form, shows that the price at which the optimum quantity will be sold is 24.1.

An important quantity to a businessman is his margin m , the difference between the selling price and the marginal cost,

$$(11.17) \quad m(q) = p(q) - c'(q).$$

We see that at the optimum quantity, the margin will be 17.6. Our later analysis will show why margin is important to businessmen although it has not traditionally been important to economists.

Since price p is the entrepreneur's decision variable, it is appropriate to show how profit depends directly on price as indicated in Equation (11.1). This function $\pi(p)$




FIGURE 11.5. Profit

is plotted in Figure 11.7. We verify that the maximum profit of 198 is achieved at a price of 24.1 and furthermore find that prices between about 15 and 35 will lead to positive profits. Note, however, that profit is very sensitive to price when the price differs from 24.1 by even a few percent.

Deterministic Sensitivity—Open Loop

We have now established that in the absence of uncertainty, the price should be set at $p = p^* = 24.1$, the quantity will be given by $q = q^* = 58.5$, and the profit will be $\pi = \pi^* = 198$. The question now arises as to how uncertainty in the state variables and changes in the decision variable p could affect the result. We define $\Delta p = p - p^*$ and recall that Δq and Δc are in fact the state variables in the problem. Then in view of Equation (11.2) we can write

$$(11.18) \quad \begin{aligned} &\pi(\Delta q, \Delta c, \Delta p) \\ &= [p^* + \Delta p][q(p^* + \Delta p) + \Delta q] - c(q(p^* + \Delta p) + \Delta q) - \Delta c. \end{aligned}$$




FIGURE 11.6. Marginal quantities

This equation shows how the profit depends on changes in all variables with respect to the centroid; it therefore forms the basis for deterministic sensitivity analysis. If we define $\Delta\pi = \pi - \pi^*$,

$$\begin{aligned}
 \Delta\pi(\Delta q, \Delta c, \Delta p) &= \pi(\Delta q, \Delta c, \Delta p) - \pi(\Delta q = 0, \Delta c = 0, \Delta p = 0) \\
 (11.19) \quad &= [p^* + \Delta p][q(p^* + \Delta p) + \Delta q] - c(q(p^* + \Delta p) + \Delta q) \\
 &\quad - \Delta c - p^*q(p^*) + c(q(p^*)),
 \end{aligned}$$

we can examine changes in profit due to changes in the variables. We could, of course, evaluate the sensitivity by computation, but direct analytical evaluation will be easier here because of the simple functional forms we have chosen.

Sensitivity to Δq . To find sensitivity to Δq given that Δc and Δp are zero, we write

$$\begin{aligned}
 \Delta\pi(\Delta q, 0, 0) &= \Delta\pi(\Delta q) \\
 (11.20) \quad &= p^*[q(p^*) + \Delta q] - c(q(p^*) + \Delta q) - p^*q(p^*) - c(q(p^*)) \\
 &= p^*\Delta q - c(q(p^*) + \Delta q) + c(q(p^*)).
 \end{aligned}$$




FIGURE 11.7. Profit as a function of price

Expansion in linear and squared terms in Δq then shows

$$\begin{aligned}
 \Delta\pi(\Delta q) &\approx \left. \frac{\partial\Delta\pi}{\partial\Delta q} \right|_{\Delta q=0} \Delta q + \frac{1}{2} \left. \frac{\partial^2\Delta\pi}{\partial(\Delta q)^2} \right|_{\Delta q=0} (\Delta q)^2 \\
 &= [p^* - c'(q(p^*))] \Delta q + \frac{1}{2} [-c''(q(p^*))] (\Delta q)^2,
 \end{aligned}
 \tag{11.21}$$

and insertion of numerical values produces

$$\Delta\pi(\Delta q) = 17.58 \Delta q + 0.0249 (\Delta q)^2.
 \tag{11.22}$$

Notice that the coefficient of the linear term is just the margin at the optimum point. Figure 11.8 shows graphically the effect of unanticipated changes in quantity sold. The curve is slightly convex upward although the curvature is not too apparent. We note, for example, that an increase of 10 units in the quantity sold will increase profits by 178.3, almost doubling them, while a decrease by 10 units will reduce profits by 173.3.

Sensitivity to Δc . Sensitivity to Δc is readily apparent when we use Equation (11.19)

$$\Delta\pi(0, \Delta c, 0) = \Delta\pi(\Delta c) = -\Delta c.
 \tag{11.23}$$

Any change in total cost has exactly the opposite effect on profit: if total cost increases by 100, total profit decreases by 100. Figure 11.9 makes this very clear.




FIGURE 11.8. Deterministic sensitivity to Δq




FIGURE 11.9. Deterministic sensitivity to Δc

Sensitivity to Δp . Sensitivity to the decision variable change Δp also follows from Equation (11.19)

$$(11.24) \quad \begin{aligned} \Delta\pi(0, 0, \Delta p) &= \Delta\pi(\Delta p) \\ &= [p^* + \Delta p]q(p^* + \Delta p) - c(q(p^* + \Delta p)) - p^*q(p^*) + c(q(p^*)). \end{aligned}$$

If we desire to expand $\Delta\pi(\Delta p)$ in the quadratic form,

$$(11.25) \quad \Delta\pi(\Delta p) = \left. \frac{\partial \Delta\pi}{\partial \Delta p} \right|_{\Delta p=0} \Delta p + \frac{1}{2} \left. \frac{\partial^2 \Delta\pi}{\partial (\Delta p)^2} \right|_{\Delta p=0} (\Delta p)^2,$$

we must evaluate the derivatives $\partial \Delta\pi / \partial \Delta p |_{\Delta p=0}$ and $\partial^2 \Delta\pi / \partial (\Delta p)^2 |_{\Delta p=0}$. We find

$$(11.26) \quad \begin{aligned} \left. \frac{\partial \Delta\pi}{\partial \Delta p} \right|_{\Delta p=0} &= [q(p^* + \Delta p) + (p^* + \Delta p)q'(p^* + \Delta p) \\ &\quad - c'(q(p^* + \Delta p))q'(p^* + \Delta p)]_{\Delta p=0} \\ &= q(p^*) + p^*q'(p^*) - c'(q(p^*))q'(p^*) \\ &= \frac{d}{dp} [r(q(p)) - c(q(p))]_{p=p^*} \\ &= [r'(q(p^*)) - c'(q(p^*))]q'(p^*) \\ &= 0, \end{aligned}$$

since marginal revenue and marginal cost are equal at the optimum. Of course, we found $p = p^*$ by requiring that $\partial \Delta\pi / \partial \Delta p |_{\Delta p=0} = 0$, so this result is not surprising.

To evaluate the second derivative we write

$$\begin{aligned}
 \frac{\partial^2 \Delta \pi}{\partial (\Delta p)^2} \Big|_{\Delta p=0} &= \{ [2q'(p^* + \Delta p) + (p^* + \Delta p)q''(p^* + \Delta p) \\
 &\quad - c''(q(p^* + \Delta p)) [q'(p^* + \Delta p)]^2 \\
 &\quad - c'(q(p^* + \Delta p)) q''(p^* + \Delta p)] \} \Big|_{\Delta p=0} \\
 (11.27) \qquad \qquad \qquad &= 2q'(p^*) + p^* q''(p^*) - c''(q(p^*)) [q'(p^*)]^2 \\
 &\qquad \qquad \qquad \qquad \qquad \qquad - c'(q(p^*)) q''(p^*).
 \end{aligned}$$

This expression shows why we evaluated certain derivatives of our demand and cost functions in Equations (11.6), (11.13), and (11.14). When

$$p = p^*, \quad \partial^2 \Delta \pi / \partial (\Delta p)^2 \Big|_{\Delta p=0} = -3.670$$

and thus Equation (11.25) becomes

$$(11.28) \quad \Delta \pi(\Delta p) = \frac{1}{2} \frac{\partial^2 \Delta \pi}{\partial (\Delta p)^2} \Big|_{\Delta p=0} (\Delta p)^2 = -1.835 (\Delta p)^2.$$

Figure 11.10 illustrates the sensitivity of profit to changes in price from the optimum value. Obviously, any change in price will decrease profit. If the price is increased or decreased by 2, for example, the profit will decrease by about 7.3.

Open Loop Effect of Uncertainty

We can now evaluate the effect on expected profit of uncertainty in the state variables using Equation (6.8). We identify Δs_1 with Δq , Δs_2 with Δc . Then Equation (6.8) becomes

$$(11.29) \quad \overline{\Delta \pi} = \frac{1}{2} \frac{\partial^2 \pi}{\partial q^2} \Delta \vee q + \frac{1}{2} \frac{\partial^2 \pi}{\partial c^2} \Delta \vee c + \frac{\partial^2 \pi}{\partial q \partial c} \overline{\Delta q \Delta c}, \quad [(\Delta q = \Delta c = \Delta p = 0)].$$

Because Δq and Δc are independent and have zero means, the last term is zero. We shall assume that the variances of Δq and Δc are 100 and 10000,

$$(11.30) \quad \Delta \vee q = 100, \quad \Delta \vee c = 10000.$$

The corresponding standard deviations 10 and 100 represent significant uncertainty in quantity and cost in the range of interest.

The derivative $\partial^2 \pi / \partial q^2$ has already been evaluated in Equations (11.21) and (11.22) as

$$(11.31) \quad \partial^2 \pi / \partial q^2 = -c''(q(p^*)) = 0.0497.$$




FIGURE 11.10. Deterministic sensitivity to Δp

The derivative $\partial^2\pi/\partial c^2 = 0$, as is clear from the linear sensitivity to Δc ; to a first approximation, the uncertainty in Δc does not affect the expected profit. Thus,

$$(11.32) \quad \begin{aligned} \overline{\Delta\pi} &= \frac{1}{2} \frac{\partial^2\pi}{\partial q^2} \Delta^2 q = \frac{1}{2} (0.0497) (10000) \\ &= 2.49. \end{aligned}$$

The effect of uncertainty in q is to increase the expected profit by 2.5 from 198 to 200.5. The reason for the increase is the upward curvature of Figure 11.8. This means that a decision-maker who is indifferent to risk would be willing to pay 2.5 to preserve the uncertainty in q rather than have a wizard assure that $\Delta q = 0$. Uncertainty can increase expected profit.

Deterministic Sensitivity—Closed Loop

Suppose that a clairvoyant should reveal Δq and Δc in advance. How would the price p be adjusted to compensate for this information, and how much would the information be worth? Equation (7.2) shows how much the price should be changed in response to Δq and Δc ,

$$(11.33) \quad \Delta p = - \left(\frac{\partial^2\pi}{\partial(\Delta p)^2} \right)^{-1} \left[\frac{\partial^2\pi}{\partial\Delta q\partial\Delta p} \Delta q + \frac{\partial^2\pi}{\partial\Delta c\partial\Delta q} \Delta c \right].$$

It is clear from Equation (11.23) that $\partial^2\pi/\partial\Delta c\partial\Delta p = 0$ and therefore that knowing Δc will not affect Δp ; consequently, clairvoyance on Δc has no value.

We evaluate $\partial^2\pi/\partial\Delta q\partial\Delta p$ from Equation (11.18) as

$$(11.34) \quad \begin{aligned} \partial^2\pi/\partial\Delta q\partial\Delta p_{\Delta q=\Delta c=\Delta p=0} &= 1 - c''(q(p^*))q'(p^*) \\ &= 0.835. \end{aligned}$$

Since we have already found that $\partial^2\pi/\partial(\Delta p)^2 = -3.670$, we write

$$(11.35) \quad \Delta p = (0.835/3.670)\Delta q = 0.228 \Delta q.$$

Thus if Δq should be 10, Δp would be increased by 2.28.

Altering Δp in response to Δq will change the sensitivity of profit to Δq from open-loop to closed loop sensitivity. From Equation (7.5) we have

$$(11.36) \quad \begin{aligned} \Delta\pi &= \frac{\partial\pi}{\partial q} \Delta q + \frac{1}{2} \frac{\partial^2\pi}{\partial q^2} (\Delta q)^2 - \frac{1}{2} \frac{(\partial^2\pi/\partial q\partial p)^2}{(\partial^2\pi/\partial p^2)} (\Delta q)^2 \\ &= 17.58 \Delta q + 0.0249 (\Delta q)^2 + \frac{1}{2} \frac{(0.835)^2}{-3.670} (\Delta q)^2 \\ &= \underbrace{17.58 \Delta q + 0.0249 (\Delta q)^2}_{\text{Open loop Sensitivity}} + \underbrace{0.095 (\Delta q)^2}_{\text{Effect of Compensation}} \\ &\quad \underbrace{\hspace{10em}}_{\text{Closed loop Sensitivity}} \end{aligned}$$

The open and closed loop sensitivity to Δq is shown in exaggerated form in Figure 11.11. The effect of the compensation is to make the sensitivity function even more curved than it was before. We found earlier that if Δq were +10, profit would increase by 178.3 without changing p . If Δp is made 2.28 as Equation (11.35) would indicate,




FIGURE 11.11. Open and closed loop sensitivity to Δq

then Equation (11.36) shows that $\Delta\pi$ would increase by another 9.5 units as a result of the compensation.

Clairvoyance

The expected value of clairvoyance on Δq is given by Equation (8.4) as

$$(11.37) \quad \langle v_{e_{\Delta q}} | \varepsilon \rangle = -\frac{1}{2} \frac{(\partial^2 \pi / \partial q \partial p)^2}{(\partial^2 \pi / \partial p^2)} \Delta^* q$$

which is just the expected value of the effect of the compensation term in Equation (11.36). Since $\Delta^* q = 100$,

$$(11.38) \quad \langle v_{e_{\Delta q}} | \varepsilon \rangle = 9.5.$$

The risk indifferent decision-maker would pay 9.5 units to know Δq in advance, about five per cent of his expected profit.

We can now summarize the effect of uncertainty and clairvoyance with respect to Δq . We have established that the uncertainty in q increases expected profit by 2.5 on an open-loop basis and that the expected effect of compensation after clairvoyance is 9.5 units. Therefore the total expected increase in profit due to uncertainty on a closed loop basis is 12.0.

This means that a risk indifferent decision-maker would pay 9.5 units to a clairvoyant for revealing Δq . Furthermore, if the clairvoyance were free, he would pay up to 12 units to avoid having a wizard set $\Delta q = 0$. Free clairvoyants can never be a bad deal, but free wizards can.

Risk Aversion and Wizardry. Up to this point we have assumed that the entrepreneur is interested only in expected profit. Now we shall introduce risk aversion. Suppose that the decision-maker is indifferent between participating in a lottery with equal probability of winning 0 and 100 on the one hand and receiving 45 for certain on the other. His certain equivalent for the lottery is therefore 45. If we assume that the decision-maker has a constant risk aversion coefficient γ , then we can find γ using Equation (5.14) and the observation that the lottery has a mean of 50 and a variance of 2500,

$$(11.39) \quad 45 = 50 - \frac{1}{2} \gamma (2500), \quad \gamma = 0.004.$$

This risk aversion coefficient implies that the entrepreneur has a risk tolerance of $1/\gamma = 250$.

Suppose that the entrepreneur leaves the price fixed at $p = p^* = 24.1$ and is interested in determining the certain equivalent $\bar{\pi}(p)$ of his uncertain profit. Equa-

tion (5.15) shows how to compute it, but we have already performed part of the calculation. The certain equivalent $\bar{\pi}(p)$ is approximately given by

$$(11.40) \quad \bar{\pi}(p) = \bar{\pi}(p) - \frac{1}{2}\gamma\check{\pi}(p),$$

and we have found that $\bar{\pi}(p)$ at $p = p^*$ is 200.5, the sum of $\pi(p^*) = 198$ and the nonlinearity correction of 2.5.

The variance $\check{\pi}(p)$ is obtained from Equation (4.5) as

$$(11.41) \quad \check{\pi}(p) = \left(\frac{\partial\pi}{\partial q}\right)^2 \Delta \check{q} + \left(\frac{\partial\pi}{\partial c}\right)^2 \Delta \check{c}$$

since the covariance of Δq and Δc is zero. From the deterministic sensitivity studies we have already found that the derivative $\partial\pi/\partial q$ is the margin, 17.58 at the optimum point, while $\partial\pi/\partial c = -1$. Therefore,

$$(11.42) \quad \check{\pi}(p) = (17.58)^2 \Delta \check{q} + (-1)^2 \Delta \check{c}.$$

It is interesting that the margin computed by the businessman is just the quantity needed to determine the variance in profit produced by a variance in quantity. With the numerical values of the variances we have

$$(11.43) \quad \begin{aligned} \check{\pi}(p) &= (17.58)^2(100) + (10000) \\ &= 30913 + 10000 \\ &= 40913. \end{aligned}$$

Returning to Equation (11.40) we find

$$(11.44) \quad \begin{aligned} \bar{\pi}(p) &= 200.5 - \frac{1}{2}(0.004)(40913) \\ &= 200.5 - 81.8 \\ &= 118.7 \end{aligned}$$

The certain equivalent is 118.7, an amount less than the expected value of 200.5 by the risk premium of 81.8.

The composition of the risk premium is itself interesting. As we see from Equations (11.40) and (11.41) the contribution to the risk premium of uncertainty in q is

$$-\frac{1}{2}\gamma(\partial\pi/\partial q)^2 \Delta \check{q} = 61.8,$$

while the contribution of uncertainty in c is $-\frac{1}{2}\gamma(\partial\pi/\partial c)^2 \Delta \check{c} = 20$. We can therefore determine the value of wizardry concerning each variable.

Suppose that a wizard offers to set $\Delta c = 0$. Since Δc does not contribute to the nonlinearity term, the value of this offer is just the saving in risk premium, 20. If the wizard offers to set $\Delta q = 0$, the saving in risk premium is 61.8. However, setting Δq to zero will cost 2.5 in expected value because of the nonlinear effect of Δq . Therefore the net value of wizardry that sets $\Delta q = 0$ is $61.8 - 2.5 = 59.3$.

All these computations rest on the approximation of Equation (11.40). As stated in Equation (5.13), the approximation is good if the standard deviation of the lottery is small compared to the risk tolerance, or, in the present case, if

$$(11.45) \quad (\check{\pi}(p))^{1/2} \ll \frac{1}{\gamma} = 250.$$

Figure 11.12 shows the standard deviation of profit computed according to Equation (11.41) as a function of price. We see first that the standard deviation of profit in-




FIGURE 11.12. Comparison of standard deviation of profit with risk tolerance

creases with price and second that it is considerably less than the risk tolerance at the price $p^* = 24.1$. Consequently, we can feel confident that our approximations are useful at the present operating point.

Now that we know the variance of profit as a function of p it is a simple matter to use Equation (11.40) to determine how certain equivalent depends on p . The result is plotted in Figure 11.13. We see that while nominal profit is maximized at the price $p = p^* = 24.1$, the certain equivalent is in fact maximized at $p = p^{**} = 22.5$. This maximum certain equivalent is 122.4, 3.7 units greater than the certain equivalent of 118.7 achieved at p^* .

Adjustment for Risk Aversion

We would like to know how to change p when risk aversion is introduced without performing the entire evaluation of Figure 11.13. We can use Equation (10.8) for this purpose. However, computation of the third derivative terms arising from non-linearity shows that they can be neglected. Thus the simpler form of Equation (10.9) is adequate,

$$(11.46) \quad \Delta p^* = p^{**} - p^* = \frac{\gamma}{\partial^2 \pi / \partial (\Delta p)^2} \left[\frac{\partial^2 \pi}{\partial \Delta q \partial \Delta p} \frac{\partial \pi}{\partial \Delta q} \Delta^v q + \frac{\partial^2 \pi}{\partial \Delta c \partial \Delta p} \frac{\partial \pi}{\partial \Delta p} \Delta^v c \right]$$

However, as we have already seen, $\partial^2 \pi / \partial \Delta c \partial \Delta p = 0$. Therefore we can use previously

PROXIMAL DECISION ANALYSIS




FIGURE 11.13. Comparison of nominal and certain equivalent profits as a function of price

evaluated derivatives to write

$$\begin{aligned}
 (11.47) \quad \Delta p^* &= \frac{0.004}{-3.670} [(0.835)(17.58)100] \\
 &= -1.6.
 \end{aligned}$$

This result shows that the optimum price when risk aversion is introduced will be $p^{**} = p^* + \Delta p^* = 24.1 + (-1.6) = 22.5$, in accordance with the plot of Figure 11.13.

Once the value of p^{**} is known, changes in other problem variables are easy to ascertain. For example, the change in q due to this change in p is given by

$$\begin{aligned}
 (11.48) \quad \Delta q &= \Delta p q'(p) = (-1.6)(-3.3) \\
 &= 5.
 \end{aligned}$$

Decreasing the price by 1.6 will increase the quantity sold by 5.

To find the certain equivalent of the profit at $p = p^{**}$, we write

$$\begin{aligned}
 \hat{\pi}(p^{**}) &= \pi(p^{**}) + \frac{1}{2} \frac{\partial^2 \pi}{\partial(\Delta q)^2} \Big|_{p^{**}} \Delta \nu q - \frac{1}{2} \gamma \hat{\pi}(p^{**}) \\
 (11.49) \qquad &= 193.0 + 2.2 - \frac{1}{2} \gamma (36406) \\
 &= 195.2 - 18203\gamma.
 \end{aligned}$$

Since $\gamma = 0.004$ the risk premium is $18203(0.004) = 72.8$, which is 9 less than it was at $p^* = 24.1$. The certain equivalent at p^{**} is then

$$\begin{aligned}
 (11.50) \qquad \hat{\pi}(p^{**}) &= 195.2 - 72.8 \\
 &= 122.4
 \end{aligned}$$

which again agrees with Figure 11.13. It would be a simple matter to recompute the values of clairvoyance and wizardry when $p = p^{**}$; the earlier results are not substantially changed.

Risk Sensitivity Profile

If people with several different attitudes toward risk must evaluate the decision it becomes worthwhile to show how the certain equivalent of the profit will depend on the risk aversion coefficient. We call a plot of certain equivalent versus risk aversion coefficient the risk sensitivity profile. The profile for the entrepreneur's problem is shown in Figure 11.14. It is drawn with the understanding that $p = p^{**} = 22.5$ using Equation (11.49). The certain equivalent of the 0-100 lottery is also shown for comparative purposes. We verify, for example, that when $\gamma = 0.004$, the profit has a certain equivalent of 122.4, while the 0-100 lottery has a certain equivalent of 45.

When $\gamma = 0$ the certain equivalent is the expected value; consequently the $\gamma = 0$ axis indicates the expected values of the profit and the 0-100 lottery. It is clear that a risk indifferent decision-maker would prefer the entrepreneur's problem to the 0-100




FIGURE 11.14. Risk sensitivity profile

lottery. This preference would exist as the risk aversion coefficient increases until it reaches a value of 0.00856, at which point the two propositions have the same certain equivalent. When γ exceeds 0.00856, the 0-100 lottery is preferred to the entrepreneur's problem. The risk sensitivity profile thus provides a basis for comparing the situation at hand with other standard options.

12. Conclusion

We have now shown how a feasible number of probabilistic assessments and a limited amount of computation can lead to extensive insights into a decision problem. Even where the approximate formulation is not appropriate for a final analysis, it can still provide a convenient and rapid procedure for guiding the evolution and evaluation of a more exact model. In short, it is a first step to overcoming the computational impasse that is so often encountered in complex decision models.

Appendix A

Certain Equivalents of Arbitrary Lotteries Under Exponential Risk Sensitivity

Consider a lottery on a random variable v described by the density function $f_v(\cdot)$. The exponential transform of this variable written in the transform variable s would then be

$$(A1) \quad f_v^e(s) = \int_{-\infty}^{\infty} dv_0 \exp(-sv_0) f_v(v_0) = \overline{\exp(-sv)}.$$

Thus we can interpret the exponential transform as the expectation of e raised to the negative product of transform variable and random variable.

Assume that the utility function of the decision-maker is the exponential function $u(\cdot)$ of Equation (5.10). By definition, the utility of the certain equivalent of the lottery must be the expected utility of the lottery,

$$u(\bar{v}) = u(v),$$

or,

$$\begin{aligned} \frac{1 - e^{-\gamma\bar{v}}}{1 - e^{-\gamma}} &= \int_{-\infty}^{\infty} dv_0 \left(\frac{1 - e^{-\gamma v_0}}{1 - e^{-\gamma}} \right) f_v(v_0) \\ &= \left(1 - \int_{-\infty}^{\infty} dv_0 e^{-\gamma v_0} f_v(v_0) \right) / (1 - e^{-\gamma}). \end{aligned}$$

Therefore,

$$\begin{aligned} e^{-\gamma\bar{v}} &= \int_{-\infty}^{\infty} dv_0 e^{-\gamma v_0} f_v(v_0) \\ &= f_v^e(\gamma), \end{aligned}$$

and

$$(A2) \quad \bar{v} = -\frac{1}{\gamma} \ln f_v^e(\gamma).$$

When the decision-maker has an exponential utility function with risk aversion coefficient γ , his certain equivalent for any lottery is one over γ times the natural logarithm of the exponential transform of the lottery evaluated at γ .

For example, the normal lottery with mean μ and standard deviation σ as defined by the density function

$$f_v(v_0) = \exp(-(v_0 - \mu)^2/2\sigma^2)/\sigma(2\pi)^{1/2}, \quad -\infty < v_0 < \infty,$$

has the exponential transform

$$f_v^e(s) = \exp(-\mu s + \sigma^2 s^2/2).$$

In view of Equation (A2), the certain equivalent \bar{v} of this lottery is then

$$\begin{aligned} \bar{v} &= -\frac{1}{\gamma} \ln \exp(-\mu\gamma + \sigma^2\gamma^2/2) \\ (A3) \quad &= \mu - \frac{1}{2}\gamma\sigma^2. \end{aligned}$$

The certain equivalent of a normal lottery is given exactly by the mean-variance approximation of Equation (5.11).

It is interesting to note that Equation (A2) can serve as a basis for approximating the certain equivalent to any desired degree of accuracy. To demonstrate this, let $g(s)$ be the natural logarithm of the exponential transform of the random variable v ,

$$(A4) \quad g(s) = \ln f_v^e(s).$$

We can perform a power series expansion of $g(s)$ about $s = 0$ according to

$$g(s) = \sum_{k=0}^{\infty} \frac{g^{(k)}(0)}{k!} s^k,$$

where $g^{(k)}(s)$ is the k th derivative of $g(s)$. Since $g(0) = 0$, Equation (A4) becomes

$$(A5) \quad g(s) = \sum_{k=1}^{\infty} \frac{g^{(k)}(0)}{k!} s^k = \sum_{k=1}^{\infty} \frac{(-1)^k g^{(k)}(0)}{k!} (-s)^k.$$

Let ${}^k v = (-1)^k g^{(k)}(0)$; then,

$$(A6) \quad g(s) = \sum_{k=1}^{\infty} \frac{{}^k v}{k!} (-s)^k.$$

The quantity ${}^k v$ is called the k th cumulant of the random variable v . By direct evaluation of the derivatives or from reference [1], we have

$$\begin{aligned} (A7) \quad {}^1 v &= \bar{v} && \text{(the mean of } v) \\ {}^2 v &= \overline{(v - \bar{v})^2} = \check{v} && \text{(the variance of } v) \\ {}^3 v &= \overline{(v - \bar{v})^3} && \text{(the third central moment of } v) \\ {}^4 v &= \overline{(v - \bar{v})^4} - 3\overline{(v - \bar{v})^2}^2. \end{aligned}$$

Since the exponential transform of the sum of independent random variables is the product of their exponential transforms, the logarithm of the exponential transform of the sum must be the sum of the logarithms of the individual exponential transforms. In view of Equations (A4) and (A6), it follows that the cumulants of the sum of independent random variables are the sum of corresponding cumulants of those variables.

We use cumulants to approximate the certain equivalent by combining the results

of Equations (A2), (A4), and (A5),

$$\begin{aligned}
 \bar{v} &= -\frac{1}{\gamma} \ln f_v^e(\gamma) = -\frac{1}{\gamma} g(\gamma) \\
 (A8) \quad &= -\frac{1}{\gamma} \sum_{k=1}^{\infty} \frac{k v}{k!} (-\gamma)^k \\
 &= \sum_{k=1}^{\infty} \frac{k v}{k!} (-\gamma)^{k-1}.
 \end{aligned}$$

The certain equivalent of any lottery under exponential risk sensitivity can be represented by a power series in γ whose coefficients are just the cumulants of the lottery divided by corresponding factorials.

By truncating the series, we can obtain any degree of approximation we wish. For example, if only the first two terms in the summation are considered, we obtain $\bar{v} \approx {}^1v - \frac{1}{2}\gamma^2 v = \bar{v} - \frac{1}{2}\gamma \check{v}$, the approximation of Equation (5.11). Considering the first three terms produces

$$\begin{aligned}
 \bar{v} &\approx {}^1v - \frac{1}{2}\gamma^2 v + \frac{1}{6}\gamma^3 v \\
 &\approx \bar{v} - \frac{1}{2}\gamma \check{v} + \frac{1}{6}\gamma^2 (\bar{v} - \bar{v})^3,
 \end{aligned}$$

showing that lotteries with positive third central moments (skewed to the right) tend to have certain equivalents higher than symmetric lotteries with the same mean and variance.

Appendix B

Numerical Evaluation of Derivatives

The numerical calculation of derivatives can be readily performed using deterministic sensitivity results. Suppose that the value v is a function of two variables x and y , $v = v(x, y)$. We know that if we have measured $v_0 = v(x_0, y_0)$ at some nominal point (x_0, y_0) , then the series expansion of v about that point can be written in terms of the distributions of v at the point and the change in the values of the variables. We shall use v_x to represent $\partial v / \partial x |_{(x_0, y_0)}$, v_{xy} to represent $\partial^2 v / \partial x \partial y |_{x_0, y_0}$, etc., and define $\Delta x = x - x_0$, $\Delta y = y - y_0$, and $\Delta v(\Delta x, \Delta y) = v(x, y) - v_0$. Then we write,

$$(B1) \quad \Delta v(\Delta x, \Delta y) \approx v_x \Delta x + v_y \Delta y + v_{xx} (\Delta x)^2 + v_{yy} (\Delta y)^2 + v_{xy} \Delta x \Delta y.$$

To evaluate v_x and v_{xx} , we set $\Delta y = 0$ and measure $\Delta v(\Delta x, 0)$ and $\Delta v(-\Delta x, 0)$; that is, leave y at its nominal value and determine the change in v that results from increasing x by the amount Δx and decreasing x by the amount Δx . Then from Equation (B1) we have

$$\begin{aligned}
 (B2) \quad \Delta v(\Delta x, 0) &\approx v_x \Delta x + v_{xx} (\Delta x)^2 \\
 \Delta v(-\Delta x, 0) &\approx -v_x \Delta x + v_{xx} (\Delta x)^2.
 \end{aligned}$$

If we subtract these equations we produce

$$\Delta v(\Delta x, 0) - \Delta v(-\Delta x, 0) \approx 2v_x \Delta x$$

or

$$(B3) \quad v_x \approx (\Delta v(\Delta x, 0) - \Delta v(-\Delta x, 0)) / 2\Delta x.$$

If we add the equations we find

$$\Delta v(\Delta x, 0) + \Delta v(-\Delta x, 0) \approx 2v_{xx}(\Delta x)^2$$

or

$$(B4) \quad v_{xx} \approx (\Delta v(\Delta x, 0) + \Delta v(-\Delta x, 0))/2(\Delta x)^2.$$

Thus Equation (B3) shows that the first partial derivative of v with respect to x at the nominal point is approximately equal to the difference in the changes in v divided by the total excursion of x , $2\Delta x$. Equation (B4) shows that the second partial derivative of v with respect to x at the nominal point is approximately equal to the sum of the changes in v divided by twice the square of the x increment. Of course, similar equations serve to determine v_y and v_{yy} from the measurements $v(0, \Delta y)$ and $v(0, -\Delta y)$. Since such one-variable-at-a-time changes are usually performed in the course of sensitivity analysis, measurement of single variable first and second partial derivatives is quite straightforward.

When we come to the cross derivative v_{xy} , we have a choice. One approach is to combine our knowledge of v_x , v_{xx} , v_y , and v_{yy} with the measured $\Delta v(\Delta x, \Delta y)$ to solve Equation (B1) for v_{xy} . While such a procedure may be economic where evaluations are very expensive, it can usually be improved upon in cases where all four of the measurements $\Delta v(\Delta x, \Delta y)$, $\Delta v(\Delta x, -\Delta y)$, $\Delta v(-\Delta x, \Delta y)$, and $\Delta v(-\Delta x, -\Delta y)$ have already been evaluated in a deterministic joint sensitivity analysis. The improvement is likely because these quantities allow us to approximate v_{xy} in a manner independent of our calculations of the one-variable derivatives. We proceed by writing Equation (B1) at the four points of measurement:

$$(B5) \quad \begin{aligned} \Delta v(\Delta x, \Delta y) &= v_x \Delta x + v_y \Delta y + v_{xx}(\Delta x)^2 + v_{yy}(\Delta y)^2 + v_{xy} \Delta x \Delta y, \\ \Delta v(\Delta x, -\Delta y) &= v_x \Delta x - v_y \Delta y + v_{xx}(\Delta x)^2 + v_{yy}(\Delta y)^2 - v_{xy} \Delta x \Delta y, \\ \Delta v(-\Delta x, \Delta y) &= -v_x \Delta x + v_y \Delta y + v_{xx}(\Delta x)^2 + v_{yy}(\Delta y)^2 - v_{xy} \Delta x \Delta y, \\ \Delta v(-\Delta x, -\Delta y) &= -v_x \Delta x - v_y \Delta y + v_{xx}(\Delta x)^2 + v_{yy}(\Delta y)^2 + v_{xy} \Delta x \Delta y. \end{aligned}$$

If we subtract the sum of the two middle equations from the sum of the first and last equations, the only term remaining on the right side is $4v_{xy}\Delta x\Delta y$. Therefore, we obtain

$$(B6) \quad v_{xy} = \frac{1}{4\Delta x\Delta y} (\Delta v(\Delta x, \Delta y) - \Delta v(\Delta x, -\Delta y) - \Delta v(-\Delta x, \Delta y) + \Delta v(-\Delta x, -\Delta y)),$$

and we have evaluated the cross partial derivative we need directly from the four measurements.

While we have considered only symmetric changes in each variable, the results can be readily extended to the asymmetric case for situations where the convenience of symmetry is not appropriate.

Evaluations of the derivatives for the example of §11 were performed using this method on a time-shared computer. The increments for Δp , Δq , and Δc were 2, 10, and 100, corresponding to the standard deviations of the variables in the latter two cases. The results were numerically indistinguishable from those in §11.

References

1. CRAMER, HAROLD, *Mathematical Method of Statistics*, Princeton University Press, 1946, Section 15.10.

PROXIMAL DECISION ANALYSIS

2. HOWARD, RONALD A., "The Foundations of Decision Analysis," *IEEE Transactions on Systems Science and Cybernetics*, Vol. SSC-4, No. 3 (September 1968).
3. PRATT, JOHN W., "Risk Aversion in the Small and in the Large," *Econometrica*, Vol. 32, No. 1-2 (January-April, 1964).
4. RAIFFA, HOWARD, *Decision Analysis, Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley, 1968, p. 201.
5. SPETZLER, CARL S., "The Development of a Corporate Risk Policy for Capital Investment Decisions," *IEEE Transactions on Systems Science and Cybernetics*, Vol. SSC-4, No. 3 (September 1968).

RISK-SENSITIVE MARKOV DECISION PROCESSES

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted from Management Science, Vol. 18, No. 7,
March 1972, pp. 356-369.

Copyright © 1972 The Institute of Management Sciences.
All rights reserved.

RISK-SENSITIVE MARKOV DECISION PROCESSES*

RONALD A. HOWARD† AND JAMES E. MATHESON‡§

This paper considers the maximization of certain equivalent reward generated by a Markov decision process with constant risk sensitivity. First, value iteration is used to optimize possibly time-varying processes of finite duration. Then a policy iteration procedure is developed to find the stationary policy with highest certain equivalent gain for the infinite duration case. A simple example demonstrates both procedures.

1. Introduction

An important limitation of previous analyses of Markov reward and decision processes [1] is that there has been no provision for incorporating risk sensitivity. The present paper shows how risk sensitivity may be treated in such processes if the utility function is exponential in form (constant risk aversion).

2. Risk Sensitivity

If a decision maker subscribes to certain arguments regarding risky propositions [3], then his risk preference may be represented by a utility function that assigns a real number to each possible outcome. Furthermore, his preference ranking of these uncertain propositions, called "lotteries," will be in accordance with the expectation of these numbers. We shall call this expectation "the utility of the lottery." Thus if v is the real-valued outcome of a lottery, $u(v)$ is the utility to be assigned to the outcome v . We assume that larger values of v are preferred, and therefore that $u(\cdot)$ is monotonically increasing.

Certain Equivalent

An important concept of our discussion will be that of certain equivalent. The certain equivalent of a lottery is the outcome whose utility is the same as the utility of the lottery. We use the symbol \bar{v} for the certain equivalent of a lottery on an outcome v with utility $u(v)$,

$$\overline{u(v)} = u(\bar{v}).$$

Exponential Utility

In many situations a decision maker is willing to accept what we call the "delta property": if all prizes in a lottery are increased by the same amount Δ , then he wants his certain equivalent for the lottery to increase by Δ ,

$$\overline{(\bar{v} + \Delta)} = \Delta + \bar{v}.$$

A decision maker who accepts the delta property is saying that his certain equivalent for any proposed new lottery is independent of his current wealth. While few decision makers would accept the delta property in all circumstances, it can be a very useful approximation in practical problems.

* Received February 1971; revised April 1971.

† Stanford University.

‡ Stanford Research Institute.

§ The authors express their appreciation to Arthur F. Veinott, Jr. for the many helpful suggestions he has made regarding the presentation of this paper.

It is easy to show [4] that the utility function of someone accepting the delta property must be either linear or exponential. The linear case implies risk indifference; it was treated in [1]. The exponential case may be described by writing the utility function in the form

$$(2.1) \quad u(v) = -(\operatorname{sgn} \gamma)e^{-\gamma v}$$

with inverse

$$(2.2) \quad u^{-1}(x) = -\frac{1}{\gamma} \ln(-(\operatorname{sgn} \gamma)x),$$

where γ is the risk aversion coefficient, and $\operatorname{sgn} \gamma$ denotes the sign of γ . The exponential utility function implies

$$(2.3) \quad u(v + \Delta) = -(\operatorname{sgn} \gamma)e^{-\gamma(v+\Delta)} = e^{-\gamma\Delta}u(v);$$

adding a constant Δ to all prizes in a lottery causes their utilities to be multiplied by $e^{-\gamma\Delta}$.

A positive risk aversion coefficient implies risk aversion: establishing a certain equivalent for a lottery that is less than its expected value. A negative risk aversion coefficient implies risk preference, the contrary behavior. We shall characterize any risk attitude that is not risk indifferent as "risk sensitive."

3. A Time-Varying Markov Reward Process

Consider an N -state time-varying Markov process that has transition probability matrix $P(n)$ at a time when n transitions (stages) remain. The ij th element of this matrix $p_{ij}(n)$ is the probability that the process will make its next transition to state j if it currently occupies state i and has n transitions remaining. A transition from state i to state j on the n th transition pays a reward $r_{ij}(n)$, positive or negative. The reward structure for any transition n is therefore summarized by a reward matrix $R(n)$ with elements $r_{ij}(n)$. We are interested in analyzing this reward process in the risk sensitive case.

Suppose that the reward process is to be allowed to continue for $n + 1$ transitions and that the process is currently in state i . The total reward the process will generate before termination we define to be $v_i(n + 1)$. If the decision maker satisfies the delta property, this uncertain reward will have a certain equivalent $\bar{v}_i(n + 1)$ that is independent of his wealth, and thus independent of rewards he has received previously. This quantity represents the amount that he would be willing to take for certain instead of receiving the reward generated by the Markov process.

To compute this quantity, consider what would happen on the next transition. If the process makes its next transition from state i to state j , it will earn a reward $r_{ij}(n + 1)$ and place the decision maker in a position where he has n transitions remaining. This position will have a certain equivalent $\bar{v}_j(n)$ that is independent of $r_{ij}(n + 1)$ as well as of previous rewards because the decision maker satisfies the delta property. With probability $p_{ij}(n + 1)$ the next transition will be to state j , whereupon the decision maker will use the delta property to assign the certain equivalent $r_{ij}(n + 1) + \bar{v}_j(n)$. Consistency requires that

$$(3.1) \quad u(\bar{v}_i(n + 1)) = \sum_{j=1}^N p_{ij}(n + 1)u(r_{ij}(n + 1) + \bar{v}_j(n)), \quad n = 0, 1, 2, \dots$$

The utility of accepting the certain equivalent must be the same as the utility of continuing. The quantities $\bar{v}_j(0)$ can be assigned directly by the decision maker.

Using the property of Equation (2.5), we can write Equation (3.1) as

$$(3.2) \quad u(\bar{v}_i(n+1)) = \sum_{j=1}^N p_{ij}(n+1)e^{-\gamma r_{ij}(n+1)}u(\bar{v}_j(n)), \quad n = 0, 1, 2, \dots$$

Note that this substitution has allowed us to write an equation relating the utility of the Markov reward process lottery at two successive stages. This happy development is directly traceable to the fact that the risk attitude of the decision maker is independent of his wealth.

We define the utility of the reward process when it occupies state j with n transitions remaining as $u_j(n)$,

$$(3.3) \quad u_j(n) = u(\bar{v}_j(n)) = -(\text{sgn } \gamma)e^{-\gamma \bar{r}_j(n)}, \quad n = 0, 1, 2, \dots$$

Equation (3.2) becomes

$$(3.4) \quad u_i(n+1) = \sum_{j=1}^N p_{ij}(n+1)e^{-\gamma r_{ij}(n+1)}u_j(n), \quad n = 0, 1, 2, \dots$$

We can directly interpret the terms in Equation (3.4). Thus in the risk-averse case of positive γ , $e^{-\gamma r_{ij}(n)}$, which we shall call $e_{ij}(n)$, is the negative utility or "disutility" of the reward r_{ij} associated with the transition from state i to state j at transition time n . The term "disutility" will hereafter be applied to $e_{ij}(n)$ regardless of the sign of γ . We shall let $q_{ij}(n)$ be the symbol for the product of transition probability and disutility,

$$q_{ij}(n) = p_{ij}(n)e^{-\gamma r_{ij}(n)} = p_{ij}(n)e_{ij}(n).$$

We shall call it the "disutility contribution" of the transition from state i to state j at time n , and define the disutility contribution matrix $Q(n)$ with elements $q_{ij}(n)$. It is clear that all elements of $Q(n)$ are nonnegative.

Now we can write Equation (3.4) in the simple form,

$$(3.5) \quad u_i(n+1) = \sum_{j=1}^N q_{ij}(n+1)u_j(n), \quad n = 0, 1, 2, \dots$$

This equation or Equation (3.4) provides a recursive relation for computing the successive utilities of the process. To find the certain equivalents of the process we then use the result implied by Equation (3.3),

$$(3.6) \quad \bar{v}_i(n) = -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_i(n)].$$

4. The Stationary Markov Reward Process

Our relationships assume important special form when the transition probabilities and rewards are the same for all transitions:

$$p_{ij}(n) = p_{ij}, \quad r_{ij}(n) = r_{ij}, \quad n = 0, 1, 2, \dots$$

In this case the process is completely specified by the transition probability matrix P and the reward matrix R or, equivalently, by the disutility contribution matrix Q with elements defined by

$$(4.1) \quad q_{ij} = p_{ij}e^{-\gamma r_{ij}}.$$

Example

To illustrate the recursive evaluation of certain equivalents and all other computations in this paper, we shall use the taxicab example of [1]. This example describes the behavior of a taxicab driver who carries out his business in three towns that we identify

with states 1, 2, and 3. His trips between towns are governed by transition probabilities; each trip entails a corresponding reward. The data for the example appear in Table 4.1. In towns 1 and 3, the driver has three alternatives: to cruise, to go to a taxi stand, or to wait for a radio call. In town 2, only the first two alternatives are available. The variable k is used to index the alternatives in each state.

We shall now investigate the policy of going to a stand, the policy composed of alternative 2 in each state. It is shown in the reference that this policy produces the highest possible average reward per transition for the process. The transition probabilities and rewards for this policy are given by:

$$(4.2) \quad P = \begin{bmatrix} 1/16 & 3/4 & 3/16 \\ 1/16 & 7/8 & 1/16 \\ 1/8 & 3/4 & 1/8 \end{bmatrix}, \quad R = \begin{bmatrix} 8 & 2 & 4 \\ 8 & 16 & 8 \\ 6 & 4 & 2 \end{bmatrix}.$$

We shall use the risk aversion coefficient $\gamma = 1.0$. For this policy and γ , the matrix Q defined by Equation (4.1) becomes:

$$(4.3) \quad Q = \begin{bmatrix} 2.09664 \times 10^{-5} & 1.01501 \times 10^{-1} & 3.43418 \times 10^{-3} \\ 2.09664 \times 10^{-5} & 9.84683 \times 10^{-8} & 2.09664 \times 10^{-5} \\ 3.09844 \times 10^{-4} & 1.37367 \times 10^{-2} & 1.69169 \times 10^{-2} \end{bmatrix}.$$

Table 4.2 shows the results of the computation using Equations (3.5) and (3.6). It provides the certain equivalent $\bar{v}_i(n)$ for each state i and for a range of number of stages remaining n . It also shows the differences in the certain equivalent of each state on successive stages under the convenient assumption that $\bar{v}_1(0) = \bar{v}_2(0) = \bar{v}_3(0) = 0$.

Notice that this difference approaches the constant 4.07438 for all states as the number of transitions remaining grows. We can interpret this constant difference as the amount that a person with a risk aversion coefficient of 1 should be just willing to pay to increase the number of stages available to him by one when he already has several available. We shall call it the "certain equivalent gain" of the process.

Notice also that the differences between the certain equivalents of states at the same stage seem to approach a constant as n increases. For example, $\bar{v}_1(n) - \bar{v}_3(n)$ approaches 1.55518. We interpret this number as the amount that a person with a risk evaluation coefficient of 1 should be just willing to pay to be in state 1 rather than in

TABLE 4.1
Taxicab Example

State i	Alternative k	Transition Probabilities p_{ij}^k			Transition Rewards r_{ij}^k		
		$j = 1$	$j = 2$	$j = 3$	$j = 1$	$j = 2$	$j = 3$
1	1 (Cruise)	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	10	4	8
	2 (Stand)	$\frac{1}{16}$	$\frac{3}{4}$	$\frac{3}{16}$	8	2	4
	3 (Call)	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{5}{8}$	4	6	4
2	1 (Cruise)	$\frac{1}{2}$	0	$\frac{1}{2}$	14	0	18
	2 (Stand)	$\frac{1}{16}$	$\frac{7}{8}$	$\frac{1}{16}$	8	16	8
3	1 (Cruise)	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{2}$	10	2	8
	2 (Stand)	$\frac{1}{8}$	$\frac{3}{4}$	$\frac{1}{8}$	6	4	2
	3 (Call)	$\frac{3}{4}$	$\frac{1}{16}$	$\frac{3}{16}$	4	0	8

TABLE 4.2
Recursive Evaluation of Taxicab Problem with Risk Aversion under Policy of Going to a Stand in Every Town

Stage	Certain Equivalent			Change in Certain Equivalent with Stage		
	State			State		
	1	2	3	1	2	3
1	2.25421	10.07710	3.47495			
2	9.08988	12.76828	7.49312	6.83567	2.69118	4.01817
3	13.02518	18.08124	11.56472	3.93530	5.31296	4.07160
4	17.19456	22.12856	15.63873	4.16938	4.04732	4.07401
5	21.26756	26.21985	19.71309	4.07300	4.09129	4.07436
6	25.34264	30.29397	23.78746	4.07508	4.07412	4.07437
7	29.41700	34.36847	27.86183	4.07436	4.07450	4.07437
8	33.49138	38.44284	31.93621	4.07438	4.07437	4.07438
9	37.56575	42.51722	36.01058	4.07437	4.07438	4.07437
10	41.64013	46.59159	40.08495	4.07438	4.07437	4.07437

state 3 with a large number of transitions remaining. For further reference we also note that $\bar{v}_2(n) - \bar{v}_3(n)$ approaches 6.50664.

We shall now investigate the nature of the certain equivalent gain and the convergence of certain equivalent differences in more detail.

Certain Equivalent Gain

The nature of the utility transformation implied by Equation (3.5) for this case becomes evident if we define the column vector $\mathbf{u}(n)$ with components $u_j(n)$, recall that $Q(n) = Q$ for the stationary case, and write the equation in the form

$$\mathbf{u}(n + 1) = Q\mathbf{u}(n), \quad n = 0, 1, 2, \dots,$$

with solution

$$\mathbf{u}(n) = Q^n \mathbf{u}(0), \quad n = 0, 1, 2, \dots$$

In the terminology of Appendix A, if the Markov transition probability is irreducible and acyclic, then the disutility contribution matrix Q is irreducible and primitive. Appendix A shows that for this case

$$(4.4) \quad \lim_{n \rightarrow \infty} \left(\frac{1}{\lambda^n} \right) Q^n \mathbf{u}(0) = \lim_{n \rightarrow \infty} \left(\frac{1}{\lambda^n} \right) \mathbf{u}(n) = k\mathbf{u},$$

where λ is the largest eigenvalue of Q and \mathbf{u} is the corresponding eigenvector with k chosen so that $u_N = -\text{sgn } \gamma$. In other words, for large n , the utility of any state will be multiplied by λ at each successive stage.

To find the implications for certain equivalents, we apply to the component form of this equation the transformation indicated in Equation (3.6),

$$-\frac{1}{\gamma} \ln \left[-(\text{sgn } \gamma) \lim_{n \rightarrow \infty} \left(\frac{1}{\lambda^n} \right) u_i(n) \right] = -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)ku_i],$$

$$\lim_{n \rightarrow \infty} \left\{ -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_i(n)] - n \left(-\frac{1}{\gamma} \ln \lambda \right) \right\} = -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_i] - \frac{1}{\gamma} \ln k,$$

or

$$(4.5) \quad \lim_{n \rightarrow \infty} \left[\bar{v}_i(n) - n \left(-\frac{1}{\gamma} \ln \lambda \right) \right] = \bar{v}_i + c,$$

where $c = -\ln k/\gamma$, and \bar{v}_i is defined to be the relative certain equivalent of state i ,

$$(4.6) \quad \bar{v}_i = -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_i].$$

Note that the normalization of u_N causes \bar{v}_N to be zero.

We find that the certain equivalent of the process will grow linearly with stage at a rate $-\ln \lambda/\gamma$ that we have called the certain equivalent gain. We designate it by the symbol \bar{g} ,

$$(4.7) \quad \bar{g} = -\frac{1}{\gamma} \ln \lambda.$$

Then the asymptotic form of $\bar{v}_i(n)$ can be written as $n\bar{g} + \bar{v}_i + c$.

It is easy to show that the certain equivalent gain must be bounded by the smallest and largest transition rewards. From Equation (A2) of Appendix A with the matrix A replaced by the matrix Q , we have

$$\min_i \sum_j q_{ij} \leq \lambda \leq \max_i \sum_j q_{ij}$$

or

$$\min_i \sum_j p_{ij} e^{-\gamma r_{ij}} \leq \lambda \leq \max_i \sum_j p_{ij} e^{-\gamma r_{ij}}.$$

Since each row sum above is a weighted average of the disutilities, we can bound each row sum by the largest and smallest disutilities for that row, i.e.,

$$\min_i \min_j e^{-\gamma r_{ij}} \leq \min_i \sum_j p_{ij} e^{-\gamma r_{ij}} \leq \lambda \leq \max_i \sum_j p_{ij} e^{-\gamma r_{ij}} \leq \max_i \max_j e^{-\gamma r_{ij}}.$$

This inequality implies

$$\exp(-\max_{i,j} \gamma r_{ij}) \leq \lambda \leq \exp(-\min_{i,j} \gamma r_{ij}).$$

By taking the natural logarithm and dividing by minus γ , and using Equation (4.7), we find

$$\min_{i,j} r_{ij} \leq \bar{g} \leq \max_{i,j} r_{ij}$$

regardless of the sign of γ .

We can obtain more insight into the nature of the certain equivalent gain by writing Equation (3.5) for the stationary case,

$$u_i(n+1) = \sum_{j=1}^N q_{ij} u_j(n),$$

dividing by λ^n , letting $n \rightarrow \infty$, and using Equation (4.4),

$$(4.8) \quad \lambda u_i = \sum_{j=1}^N q_{ij} u_j.$$

We recall from Equation (4.6) that we can write u_i as

$$u_i = -(\text{sgn } \gamma) e^{-\gamma \bar{v}_i}$$

and from Equation (4.7) that λ is related to the certain equivalent gain by $\lambda = e^{-\gamma \bar{g}}$. Then we can write Equation (4.8) as

$$e^{-\gamma(\bar{g} + \bar{v}_i)} = \sum_{j=1}^N q_{ij} e^{-\gamma \bar{v}_j}$$

or, using the explicit form of the disutility contribution,

$$(4.9) \quad e^{-\gamma(\bar{g} + \bar{v}_i)} = \sum_{j=1}^N p_{ij} e^{-\gamma(r_{ij} + \bar{v}_j)}.$$

As γ approaches zero and therefore the decision maker approaches risk indifference, these equations imply

$$\tilde{g} + \bar{v}_i = \sum_{j=1}^N p_{ij}(r_{ij} + \bar{v}_j),$$

which are the usual relative value equations for a Markov reward process. In this case the certain equivalent gain becomes the ordinary gain or expected reward per transition of the process, and the \bar{v}_i 's become the relative values of the states.

Thus we can think of Equation (4.9) as the analog of the usual relative value equations for the case of exponential risk sensitivity. Note that they share with the usual equations the property that adding a constant to all \bar{v}_i 's leaves the equations unchanged; consequently, we may solve for the \bar{v}_i 's only to within a constant. We shall therefore call \bar{v}_i the relative certain equivalent of state i . When the relative certain equivalent of one state, say state N , is set equal to zero, then we can solve Equation (4.9) for the certain equivalent gain and the relative certain equivalents of the other states for any irreducible acyclic Markov process.

Inspection of Equation (4.9) shows that if we add the same constant Δ to all rewards r_{ij} , the new solution of the equations will have the same relative certain equivalents, but a certain equivalent gain increased by Δ : Increasing the rewards by a constant increases the certain equivalent gain by the same constant.

Let us apply the results of this section to the policy formed by the second alternative in each state in the taxicab example. The transition probabilities and rewards for this policy appear in Equation (4.2); the Q matrix, in Equation (4.3). We find that the largest eigenvalue of Q is $\lambda = 0.0170027$. The certain equivalent gain is then computed from Equation (4.7) to be 4.07438. Note that this is the same value for this quantity indicated by the iterative solution.

We find that the eigenvalue associated with this eigenvector is proportional to (0.17413, 0.00123, 0.82464). If we wish to set the relative certain equivalent value of a state to be zero, we see from Equation (4.6) that the corresponding component of the eigenvector must be set to $-(\text{sgn } \gamma)$. Thus if we wish to set $\bar{v}_3 = 0$ in this example, we must normalize \mathbf{u} so that $u_3 = -1$. With this normalization, the vector \mathbf{u} becomes $(-0.21115, -0.0014935, -1)$. We then use Equation (4.6) to find:

$$\begin{aligned} \bar{v}_1 &= -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_1] & \bar{v}_2 &= -\frac{1}{\gamma} \ln [-(\text{sgn } \gamma)u_2] \\ &= -\ln (0.21115) & &= -\ln (1.49349) \\ &= 1.55517, & &= 6.50664. \end{aligned}$$

We note that these relative certain equivalents are the ones observed earlier in the recursive calculation of certain equivalent values for this policy.

5. A Time-Varying Markov Decision Process

Suppose that at any transition n (measured from the end of the process) different transition probability and reward matrices can be used to govern the process. In any state i at transition n there is a choice among various alternatives k that specify the transition probability $p_{ij}^k(n)$ and rewards $r_{ij}^k(n)$ that will characterize the next transition of the process. The number of alternatives available may be different from transition to transition and from state to state. Our problem is to find which alternative should be used at each stage and state in order to maximize the utility or, equivalently, the certain equivalent of the reward subsequently generated by the process.

Let $d_i(n)$ be the number of the best alternative to use in state i when n transitions remain. When we have specified the column vector $\mathbf{d}(n)$ for all n , we have solved the problem. We call $\mathbf{d}(n)$ the optimum policy at time n . We can find the optimum policy by applying the principle of optimality to Equation (3.4). We define $u_i(n)$ to be the highest utility achievable from the process when it occupies state i , and n transitions remain. Then

$$(5.1) \quad u_i(n + 1) = \max_k \sum_{j=1}^N p_{ij}^k(n + 1) e^{-\gamma r_{ij}^k(n)} u_j(n), \quad n = 0, 1, 2, \dots,$$

where $d_i(n + 1)$ is the maximizing value of k . This equation allows us to compute the optimum policy for all stages as well as the utility of the process under this policy.

If we define $\bar{v}_i(n)$ as the certain equivalent of the lottery implied by being in state i with n stages remaining under the optimum policy, then we can find $\bar{v}_i(n)$ from $u_i(n)$ by using Equation (3.6).

Example

Table 5.1 shows the results of applying this procedure to the taxicab decision problem whose data appear in Table 4.1. The risk aversion coefficient used is $\gamma = 1.0$; rewards are assumed to be zero after all available transitions have been made. We see that when only one transition remains, the best policy is to use the first alternative in both states 1 and 2, and the second alternative in state 3. For all greater numbers of transitions remaining, the best policy is to use the first alternative in each state.

The table also shows the change in the certain equivalent of each state as the number of transitions increases. Note that it approaches approximately 8.48 for all states. As we shall see, this is the certain equivalent gain of the policy composed of the first alternative in each state.

Table 5.2 shows the same procedure with the sole change that the risk aversion coefficient has been changed to $\gamma = -1.0$ to illustrate a case of risk preference. When one transition remains, the best policy is to use the first alternative in each state. However, when more than one transition remains, the best policy is to use the third alternative in state 1 and the second alternative in states 2 and 3. Note that the differences in certain equivalent for any state appear to approach approximately 15.87 when more than a few transitions remain. We shall see that this is the certain equivalent gain of

TABLE 5.1
Value Iteration in Taxicab Problem with Risk Aversion Coefficient $\gamma = 1.0$

Stage	Policy			Certain Equivalent			Change in Certain Equivalent with Stage		
	State			State			State		
	1	2	3	1	2	3	1	2	3
1	1	1	2	5.36329	14.67500	3.47495			
2	1	1	1	12.82038	19.94218	12.15518	7.45709	5.26718	8.68023
3	1	1	1	21.39147	27.47853	20.73632	8.57109	7.53635	8.58114
4	1	1	1	29.93613	36.04996	29.18795	8.53466	8.57143	8.45163
5	1	1	1	38.40430	44.59130	37.66343	8.46817	8.54134	8.47548
6	1	1	1	46.88206	53.05974	46.14960	8.47776	8.46844	8.48617
7	1	1	1	55.36650	61.53781	54.63268	8.48444	8.47807	8.48308
8	1	1	1	63.84950	70.02220	63.11497	8.48300	8.48439	8.48229
9	1	1	1	72.33197	78.50518	71.59763	8.48247	8.48298	8.48266
10	1	1	1	80.81462	86.98765	80.08033	8.48265	8.48247	8.48270

TABLE 5.2
Value Iteration in Taxicab Problem with Risk Aversion Coefficient $\gamma = -1.0$

Stage	Policy			Certain Equivalent			Change in Certain Equivalent with Stage		
	State			State			State		
	1	2	3	1	2	3	1	2	3
1	1	1	1	9.37349	17.32500	8.85351			
2	3	2	2	21.24580	33.19147	21.03776	11.87231	15.86647	12.18425
3	3	2	2	37.11204	49.05794	36.90380	15.86624	15.86647	15.86604
4	3	2	2	52.97850	64.92441	52.77027	15.86646	15.86647	15.86647
5	3	2	2	68.84497	80.79088	68.63673	15.86647	15.86647	15.86646
6	3	2	2	84.71144	96.65735	84.50320	15.86647	15.86647	15.86647
7	3	2	2	100.57791	112.52381	100.36967	15.86647	15.86646	15.86647
8	3	2	2	116.44438	128.39028	116.23614	15.86647	15.86647	15.86647
9	3	2	2	132.31085	144.25675	132.10261	15.86647	15.86647	15.86647
10	3	2	2	148.17732	160.12322	147.96908	15.86647	15.86647	15.86647

the policy of using the third alternative in state 1, the second alternative in states 2 and 3 for every transition.

6. The Stationary Markov Decision Process

Suppose again the individual has various alternatives available for operating the system. However, unlike the earlier case, whatever alternative is selected in a state must be used for all transitions. The alternative k in state i therefore specifies transition probabilities p_{ij}^k and transition rewards r_{ij}^k that will govern the system whenever state i is entered. We describe the policy for the system by a vector \mathbf{d} whose i th element is the decision in state i , the number of the alternative to be used in state i . We seek the policy that will maximize the certain equivalent gain of the system.

We can find the optimum policy by a procedure analogous to policy iteration; it appears in Figure 6.1.¹ First we select an arbitrary policy and solve the relative certain equivalent Equations (4.9) to find the certain equivalent gain and relative certain equivalents corresponding to it. Then we perform a policy improvement by selecting as the new decision in each state the alternative k that maximizes

$$\tilde{V}_i^k = -\frac{1}{\gamma} \ln \left[\sum_{j=1}^N p_{ij}^k e^{-\gamma(r_{ij}^k + \tilde{v}_j)} \right]$$

using the relative certain equivalents of the previous policy. When this has been done for all states, we have a new policy which we evaluate and attempt to improve in the same manner. When no change is possible in the alternative selected in any state, we have found the optimum policy. The proof of the optimality of this procedure appears in Appendix B.

The method can also be used when strictly periodic and hence deterministic processes are possible. In this case $p_{ij} = 1$ for only one $j, j \neq i$, and otherwise $p_{ij} = 0$. The procedure of Figure 6.1 then reduces to the usual policy iteration of [1] and [2] involving the solution of linear simultaneous equations.

The procedure of Figure 6.1 may also be viewed in the utility formulation as shown

¹ We thank Arthur F. Veinott, Jr. for pointing out that the existence of a solution to an equation with the form of Equation (5.1) for the stationary infinite-horizon case is proved by Richard Bellman in the book *Dynamic Programming*, Princeton University Press, Princeton, N.J., 1957, p. 329.

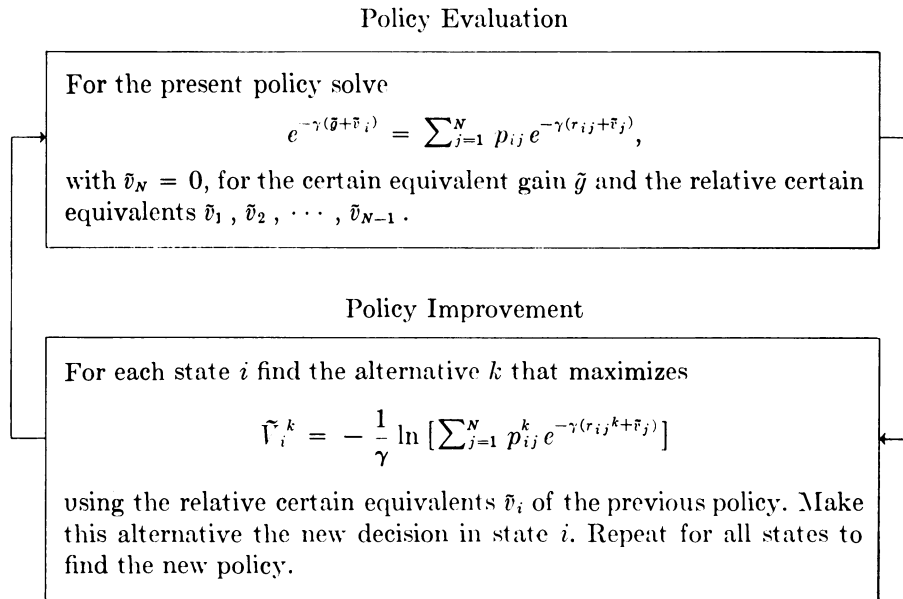


FIGURE 6.1. The Policy Iteration Cycle with Risk Sensitivity—Certain Equivalent Form

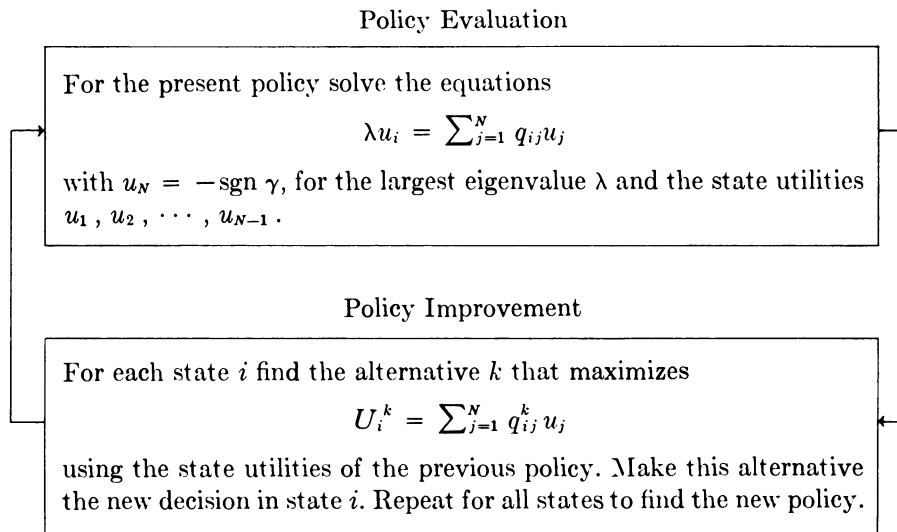


FIGURE 6.2. The Policy Iteration Cycle with Risk Sensitivity—Utility Form.

in Figure 6.2. The quantity q_{ij}^k is defined to be $p_{ij}^k e^{-\gamma r_{ij}^k}$ in agreement with the definition of q_{ij} . The nature of the iteration is clear from the figure.

Example

Table 6.1 brings the results of applying this policy iteration procedure to the taxicab example of Table 4.1. For a range of risk aversion coefficients γ from -1.0 to 1.0 , the table shows for each γ the optimum stationary policy, the certain equivalent gain, and the relative certain equivalents of states 1 and 2 relative to state 3. In each case the

TABLE 6.1
Results of Policy Iteration in Taxicab Problem for a Range of Risk Aversion Coefficients

Risk Aversion Coefficient γ	Policy			Certain Equivalent Gain \bar{x}	Relative Certain Equivalents		No. of Iterations
	d				\bar{v}_1	\bar{v}_2	
-1.0	3	2	2	15.86647	0.20824	12.15414	2
-0.7	3	2	2	15.80924	-0.55936	12.22008	3
-0.50	3	2	2	15.73295	-1.57543	12.30717	3
-0.45	3	2	2	15.70329	-1.96342	12.34054	3
-0.44	2	2	2	15.69655	-1.99156	12.34803	3
-0.3	2	2	2	15.55569	-1.96311	12.49427	3
-0.2	2	2	2	15.34197	-1.89234	12.66973	3
-0.1	2	2	2	14.82655	-1.68692	12.88752	3
-0.01	2	2	2	13.56641	-1.24330	12.94136	3
-0.001	2	2	2	13.36751	-1.18326	12.66499	3
-0.0001	2	2	2	13.34678	-1.17715	12.65638	3
0	2	2	2	13.34454	-1.17647	12.65546	3*
0.0001	2	2	2	13.34216	-1.17579	12.65445	3
0.001	2	2	2	13.32137	-1.16966	12.64571	3
0.01	2	2	2	13.10536	-1.10755	12.54814	3
0.09	2	2	2	10.88344	-0.57796	11.00438	3
0.1	1	2	2	10.62679	-0.47318	10.76876	2
0.16	1	2	2	9.56203	1.03740	9.59114	3
0.17	1	2	1	9.42216	1.20493	9.39993	2
0.2	1	2	1	9.21697	1.21201	8.47294	2
0.24	1	2	1	8.98541	1.22502	7.50767	2
0.25	1	1	1	8.95762	1.22286	7.41749	1
0.3	1	1	1	8.91227	1.18909	7.39437	1
0.5	1	1	1	8.75025	1.04280	7.26815	1
0.7	1	1	1	8.62182	0.90394	7.11748	1
1.0	1	1	1	8.48267	0.73431	6.90733	3

* Run by usual procedure presented in [1].

procedure was started using the policy derived by setting $\bar{v}_i = 0$ for all i and then entering the policy improvement box. The number of iterations required to attain the optimum policy appear in the last column of the table.

The table includes all values of γ at which policy changes occur. Note that only 5 of the 18 possible policies are optimum for any value of γ . As γ increases from -1 , the policy $\mathbf{d} = [3\ 2\ 2]$ is optimum for γ through -0.45 . Then $\mathbf{d} = [2\ 2\ 2]$ becomes optimum until γ reaches 0.1 . At this point $\mathbf{d} = [1\ 2\ 2]$ is best and remains so until $\mathbf{d} = [1\ 2\ 1]$ becomes optimum at $\gamma = 0.17$. When γ reaches 0.25 , the optimum policy changes permanently to $\mathbf{d} = [1\ 1\ 1]$.

As γ ranges from -1 to 1 , the certain equivalent gain decreases from 15.86647 to 8.48267 . The number of iterations required for convergence is never more than 3 .

Figure 6.3 is a plot of certain equivalent gain of the optimum policy versus risk aversion coefficient for the range covered by the table. The optimum policy regions are indicated at the top of the figure. We see that most of the dependence of gain on risk aversion coefficient occurs for γ in the range $-0.4 \leq \gamma \leq 0.4$. The figure shows quite clearly just how sensitive the optimum policy and certain equivalent gain are to changes in risk attitude:

RISK-SENSITIVE MARKOV DECISION PROCESSES




FIGURE 6.3. Certain Equivalent Gain of the Optimum Policy as a Function of Risk Aversion Coefficient.

7. Conclusion

The ability to extend Markov decision process analysis to the case of risk sensitivity should have important application in many areas. We have found the model exceptionally useful in considering optimum buying and selling strategies for a commodity market. Employing this approach in such traditional applications as the optimization of replacement and investment systems will provide interesting new insight into the robustness of maximum-expected-value-per-transition policies.

Appendix A—Matrix Results Utilized in This Paper

The results in this paper are derived from the theory of matrices with nonnegative elements. An excellent discussion of this subject can be found in F. R. Gantmacher, *The Theory of Matrices*, Vol. 2, Chelsea, 1960, Chapter XIII. We shall now summarize the important properties of irreducible nonnegative matrices used in the paper.

A matrix A is called reducible if there is a permutation that can place it in the form

$$\begin{pmatrix} B & 0 \\ C & D \end{pmatrix},$$

where B and D are square matrices; otherwise, A is called irreducible. (An irreducible Markov transition probability matrix represents a process in which all states communicate.)

An irreducible nonnegative matrix A always has a positive eigenvalue λ that is a simple root of the characteristic equation. The moduli of all the other eigenvalues do not exceed λ . To this "maximal" eigenvalue there corresponds an eigenvector with positive components having a unique direction. The matrix A is called "primitive" if some power of A has all elements positive. (A primitive Markov transition probability matrix is called "acyclic.") If A is primitive, then the moduli of all other eigenvalues are strictly less than λ . This means that asymptotically

$$(A1) \quad \lim_{n \rightarrow \infty} \left(\frac{1}{\lambda^n} \right) A^n \mathbf{x} = \mathbf{v},$$

where \mathbf{x} is any vector with nonnegative elements, and \mathbf{v} is an appropriately normalized eigenvector (of unique direction) corresponding to λ .

If A is an irreducible matrix with nonnegative elements a_{ij} and maximal eigenvalue λ , and \mathbf{x} is a vector with positive components x_i , then the following inequalities hold:

$$(A2) \quad \min_i \sum_j a_{ij} \leq \lambda \leq \max_i \sum_j a_{ij},$$

$$(A3) \quad \min_i \frac{\sum_j a_{ij} x_j}{x_i} \leq \lambda \leq \max_i \frac{\sum_j a_{ij} x_j}{x_i}.$$

Equation (A3) holds with the equality signs if and only if \mathbf{x} is an eigenvector corresponding to λ .

Appendix B—Proof of Convergence of the Policy Iteration Procedure to an Optimal Policy

The proof will be carried out only for the case of risk aversion, i.e., positive risk aversion coefficient and negative utilities. The same conclusions can be reached for risk preference with appropriate inequality changes.

Assume that the policy iteration has converged on policy B and that policy A is any other policy. Then from the policy improvement step we know that

$$(B1) \quad \sum_j q_{ij}^A u_j^B \leq \sum_j q_{ij}^B u_j^B,$$

where \mathbf{u}^B is an eigenvector of Q^B corresponding to the maximal eigenvalue λ^B ,

$$(B2) \quad \sum_j q_{ij}^B u_j^B = \lambda^B u_i^B.$$

Thus, we have

$$(B3) \quad \sum_j q_{ij}^A u_j^B \leq \lambda^B u_i^B.$$

Since all components u_j^B are negative, we can write

$$(B4) \quad \frac{\sum_j q_{ij}^A |u_j^B|}{|u_i^B|} \geq \lambda^B, \quad i = 1, 2, \dots, N.$$

The left inequality of Equation (A3) provides the condition

$$\min_i \frac{\sum_j q_{ij}^A |u_j^B|}{|u_i^B|} \leq \lambda^A,$$

where λ^A is the maximal eigenvalue of Q^A . Therefore we obtain

$$(B5) \quad \lambda^A \geq \min_i \frac{\sum_j q_{ij}^A |u_j^B|}{|u_i^B|} \geq \lambda^B.$$

To find the implication for certain equivalent gains of both policies, we note that (B5) implies

$$-\frac{1}{\gamma} \ln \lambda^A \leq -\frac{1}{\gamma} \ln \lambda^B$$

and $\tilde{g}^A \leq \tilde{g}^B$. Thus policy iteration can only converge on an optimum policy.

It now remains to be shown that the policy iteration procedure will converge on an optimum policy. Assume that the procedure has improved policy A to arrive at policy B . Then from the policy improvement step we know that

$$(B6) \quad \sum_j q_{ij}^B u_j^A \geq \sum_j q_{ij}^A u_j^A$$

with inequality for some i .

Recognizing \mathbf{u}^A as an eigenvector of Q^A corresponding to λ^A , we can write

$$(B7) \quad \sum_j q_{ij}^B u_j^A \geq \lambda^A u_i^A.$$

If \mathbf{u}^A happens also to be an eigenvector of Q^B then the above equation implies directly that $\lambda^B u_i^A > \lambda^A u_i^A$ for some i , which because the u_i^A are negative leads to $\lambda^B < \lambda^A$.

If \mathbf{u}^A is not an eigenvector of Q^B then we rewrite (B7) as

$$(B8) \quad \frac{\sum_j q_{ij}^B |u_j^A|}{|u_i^A|} \leq \lambda^A.$$

In view of Equation (A3), we have

$$(B9) \quad \lambda^B < \max_i \frac{\sum_j q_{ij}^B |u_j^A|}{|u_i^A|} \leq \lambda^A.$$

Note the strict inequality on the left occurs because we have assumed that \mathbf{u}^A is not an eigenvector of Q^B .

In terms of certain equivalent gain these results imply

$$(B10) \quad \tilde{g}^A < \tilde{g}^B.$$

Thus at each iteration the certain equivalent gain must increase.

Since there is only a finite number of possible policies, the procedure will converge in a finite number of iterations.

References

1. HOWARD, RONALD A., *Dynamic Programming and Markov Processes*, The M.I.T. Press, Cambridge, 1960.
2. —, *Dynamic Probabilistic Systems* (two volumes), John F. Wiley & Sons, Inc., New York, 1971.
3. LUCE, ROBERT D. AND RAIFFA, HOWARD, *Games and Decisions*, John F. Wiley & Sons, Inc., New York, 1957.
4. RAIFFA, HOWARD, *Decision Analysis: Introductory Lectures on Choices under Uncertainty*, Addison Wesley, Menlo Park, 1968.

CONTRIBUTIONS FROM PSYCHOLOGY

Preface

The following three papers are especially relevant to decision analysis. They either help us refine our assessment procedures or demonstrate the need for decision analysis in important decisions

"Judgment Under Uncertainty: Heuristics and Biases" clearly demonstrates several biases that occur when people express their judgments numerically. This paper catalogs these biases by defining the underlying empirical judgments that seem to be producing them. The probability assessment procedure presented earlier in this collection was motivated by this work.

"Prospect Theory: An Analysis of Decision Under Risk" presents utility theory as a descriptive model of decision-making under risk and develops an alternate model -- prospect theory. The paper identifies consistent biases people exhibit in decision-making situations and shows when and how they operate.

"The Framing of Decisions and the Psychology of Choice" shows that the "decision frame" adopted by the decision-maker explains many decision biases, that the frame is not often consciously controlled, and that alternate frames account for seeming reversals of preference. This paper questions what we really mean by rational choice.

JUDGMENT UNDER UNCERTAINTY:
HEURISTICS AND BIASES

Amos Tversky
Stanford University

Daniel Kahneman
University of British Columbia

Reprinted from Science, September 27, 1974, Volume 185, pp. 1124-1131.

Copyright © 1974 by the American Association
for the Advancement of Science.

Judgment under Uncertainty: Heuristics and Biases

Biases in judgments reveal some heuristics of
thinking under uncertainty.

Amos Tversky and Daniel Kahneman

Many decisions are based on beliefs concerning the likelihood of uncertain events such as the outcome of an election, the guilt of a defendant, or the future value of the dollar. These beliefs are usually expressed in statements such as "I think that . . .," "chances are . . .," "it is unlikely that . . .," and so forth. Occasionally, beliefs concerning uncertain events are expressed in numerical form as odds or subjective probabilities. What determines such beliefs? How do people assess the probability of an uncertain event or the value of an uncertain quantity? This article shows that people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.

The subjective assessment of probability resembles the subjective assessment of physical quantities such as distance or size. These judgments are all based on data of limited validity, which are processed according to heuristic rules. For example, the apparent distance of an object is determined in part by its clarity. The more sharply the object is seen, the closer it appears to be. This rule has some validity, because in any given scene the more distant objects are seen less sharply than nearer objects. However, the reliance on this rule leads to systematic errors in the estimation of distance. Specifically, distances are often overestimated when visibility is poor because the contours of objects are blurred. On the other hand, distances are often underesti-

mated when visibility is good because the objects are seen sharply. Thus, the reliance on clarity as an indication of distance leads to common biases. Such biases are also found in the intuitive judgment of probability. This article describes three heuristics that are employed to assess probabilities and to predict values. Biases to which these heuristics lead are enumerated, and the applied and theoretical implications of these observations are discussed.

Representativeness

Many of the probabilistic questions with which people are concerned belong to one of the following types: What is the probability that object A belongs to class B? What is the probability that event A originates from process B? What is the probability that process B will generate event A? In answering such questions, people typically rely on the representativeness heuristic, in which probabilities are evaluated by the degree to which A is representative of B, that is, by the degree to which A resembles B. For example, when A is highly representative of B, the probability that A originates from B is judged to be high. On the other hand, if A is not similar to B, the probability that A originates from B is judged to be low.

For an illustration of judgment by representativeness, consider an individual who has been described by a former neighbor as follows: "Steve is very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail." How do people assess the probability that Steve is engaged in a particular

occupation from a list of possibilities (for example, farmer, salesman, airline pilot, librarian, or physician)? How do people order these occupations from most to least likely? In the representativeness heuristic, the probability that Steve is a librarian, for example, is assessed by the degree to which he is representative of, or similar to, the stereotype of a librarian. Indeed, research with problems of this type has shown that people order the occupations by probability and by similarity in exactly the same way (1). This approach to the judgment of probability leads to serious errors, because similarity, or representativeness, is not influenced by several factors that should affect judgments of probability.

Insensitivity to prior probability of outcomes. One of the factors that have no effect on representativeness but should have a major effect on probability is the prior probability, or base-rate frequency, of the outcomes. In the case of Steve, for example, the fact that there are many more farmers than librarians in the population should enter into any reasonable estimate of the probability that Steve is a librarian rather than a farmer. Considerations of base-rate frequency, however, do not affect the similarity of Steve to the stereotypes of librarians and farmers. If people evaluate probability by representativeness, therefore, prior probabilities will be neglected. This hypothesis was tested in an experiment where prior probabilities were manipulated (1). Subjects were shown brief personality descriptions of several individuals, allegedly sampled at random from a group of 100 professionals—engineers and lawyers. The subjects were asked to assess, for each description, the probability that it belonged to an engineer rather than to a lawyer. In one experimental condition, subjects were told that the group from which the descriptions had been drawn consisted of 70 engineers and 30 lawyers. In another condition, subjects were told that the group consisted of 30 engineers and 70 lawyers. The odds that any particular description belongs to an engineer rather than to a lawyer should be higher in the first condition, where there is a majority of engineers, than in the second condition, where there is a majority of lawyers. Specifically, it can be shown by applying Bayes' rule that the ratio of these odds should be $(.7/.3)^2$, or 5.44, for each description. In a sharp violation of Bayes' rule, the subjects in the two conditions produced essen-

The authors are members of the department of psychology at the Hebrew University, Jerusalem, Israel.

tially the same probability judgments. Apparently, subjects evaluated the likelihood that a particular description belonged to an engineer rather than to a lawyer by the degree to which this description was representative of the two stereotypes, with little or no regard for the prior probabilities of the categories.

The subjects used prior probabilities correctly when they had no other information. In the absence of a personality sketch, they judged the probability that an unknown individual is an engineer to be .7 and .3, respectively, in the two base-rate conditions. However, prior probabilities were effectively ignored when a description was introduced, even when this description was totally uninformative. The responses to the following description illustrate this phenomenon:

Dick is a 30 year old man. He is married with no children. A man of high ability and high motivation, he promises to be quite successful in his field. He is well liked by his colleagues.

This description was intended to convey no information relevant to the question of whether Dick is an engineer or a lawyer. Consequently, the probability that Dick is an engineer should equal the proportion of engineers in the group, as if no description had been given. The subjects, however, judged the probability of Dick being an engineer to be .5 regardless of whether the stated proportion of engineers in the group was .7 or .3. Evidently, people respond differently when given no evidence and when given worthless evidence. When no specific evidence is given, prior probabilities are properly utilized; when worthless evidence is given, prior probabilities are ignored (1).

Insensitivity to sample size. To evaluate the probability of obtaining a particular result in a sample drawn from a specified population, people typically apply the representativeness heuristic. That is, they assess the likelihood of a sample result, for example, that the average height in a random sample of ten men will be 6 feet (180 centimeters), by the similarity of this result to the corresponding parameter (that is, to the average height in the population of men). The similarity of a sample statistic to a population parameter does not depend on the size of the sample. Consequently, if probabilities are assessed by representativeness, then the judged probability of a sample statistic will be essentially independent of

sample size. Indeed, when subjects assessed the distributions of average height for samples of various sizes, they produced identical distributions. For example, the probability of obtaining an average height greater than 6 feet was assigned the same value for samples of 1000, 100, and 10 men (2). Moreover, subjects failed to appreciate the role of sample size even when it was emphasized in the formulation of the problem. Consider the following question:

A certain town is served by two hospitals. In the larger hospital about 45 babies are born each day, and in the smaller hospital about 15 babies are born each day. As you know, about 50 percent of all babies are boys. However, the exact percentage varies from day to day. Sometimes it may be higher than 50 percent, sometimes lower.

For a period of 1 year, each hospital recorded the days on which more than 60 percent of the babies born were boys. Which hospital do you think recorded more such days?

- ▶ The larger hospital (21)
- ▶ The smaller hospital (21)
- ▶ About the same (that is, within 5 percent of each other) (53)

The values in parentheses are the number of undergraduate students who chose each answer.

Most subjects judged the probability of obtaining more than 60 percent boys to be the same in the small and in the large hospital, presumably because these events are described by the same statistic and are therefore equally representative of the general population. In contrast, sampling theory entails that the expected number of days on which more than 60 percent of the babies are boys is much greater in the small hospital than in the large one, because a large sample is less likely to stray from 50 percent. This fundamental notion of statistics is evidently not part of people's repertoire of intuitions.

A similar insensitivity to sample size has been reported in judgments of posterior probability, that is, of the probability that a sample has been drawn from one population rather than from another. Consider the following example:

Imagine an urn filled with balls, of which $\frac{2}{3}$ are of one color and $\frac{1}{3}$ of another. One individual has drawn 5 balls from the urn, and found that 4 were red and 1 was white. Another individual has drawn 20 balls and found that 12 were red and 8 were white. Which of the two individuals should feel more confident that the urn contains $\frac{2}{3}$ red balls and $\frac{1}{3}$ white balls, rather than the opposite? What odds should each individual give?

In this problem, the correct posterior odds are 8 to 1 for the 4 : 1 sample and 16 to 1 for the 12 : 8 sample, assuming equal prior probabilities. However, most people feel that the first sample provides much stronger evidence for the hypothesis that the urn is predominantly red, because the proportion of red balls is larger in the first than in the second sample. Here again, intuitive judgments are dominated by the sample proportion and are essentially unaffected by the size of the sample, which plays a crucial role in the determination of the actual posterior odds (2). In addition, intuitive estimates of posterior odds are far less extreme than the correct values. The underestimation of the impact of evidence has been observed repeatedly in problems of this type (3, 4). It has been labeled "conservatism."

Misconceptions of chance. People expect that a sequence of events generated by a random process will represent the essential characteristics of that process even when the sequence is short. In considering tosses of a coin for heads or tails, for example, people regard the sequence H-T-H-T-T-H to be more likely than the sequence H-H-H-T-T-T, which does not appear random, and also more likely than the sequence H-H-H-H-T-H, which does not represent the fairness of the coin (2). Thus, people expect that the essential characteristics of the process will be represented, not only globally in the entire sequence, but also locally in each of its parts. A locally representative sequence, however, deviates systematically from chance expectation: it contains too many alternations, and too few runs. Another consequence of the belief in local representativeness is the well-known gambler's fallacy. After observing a long run of red on the roulette wheel, for example, most people erroneously believe that black is now due, presumably because the occurrence of black will result in a more representative sequence than the occurrence of an additional red. Chance is commonly viewed as a self-correcting process in which a deviation in one direction induces a deviation in the opposite direction to restore the equilibrium. In fact, deviations are not "corrected" as a chance process unfolds, they are merely diluted.

Misconceptions of chance are not limited to naive subjects. A study of the statistical intuitions of experienced research psychologists (5) revealed a lingering belief in what may be called the "law of small numbers," according to which even small samples are highly

representative of the populations from which they are drawn. The responses of these investigators reflected the expectation that a valid hypothesis about a population will be represented by a statistically significant result in a sample—with little regard for its size. As a consequence, the researchers put too much faith in the results of small samples and grossly overestimated the replicability of such results. In the actual conduct of research, this bias leads to the selection of samples of inadequate size and to overinterpretation of findings.

Insensitivity to predictability. People are sometimes called upon to make such numerical predictions as the future value of a stock, the demand for a commodity, or the outcome of a football game. Such predictions are often made by representativeness. For example, suppose one is given a description of a company and is asked to predict its future profit. If the description of the company is very favorable, a very high profit will appear most representative of that description; if the description is mediocre, a mediocre performance will appear most representative. The degree to which the description is favorable is unaffected by the reliability of that description or by the degree to which it permits accurate prediction. Hence, if people predict solely in terms of the favorableness of the description, their predictions will be insensitive to the reliability of the evidence and to the expected accuracy of the prediction.

This mode of judgment violates the normative statistical theory in which the extremeness and the range of predictions are controlled by considerations of predictability. When predictability is nil, the same prediction should be made in all cases. For example, if the descriptions of companies provide no information relevant to profit, then the same value (such as average profit) should be predicted for all companies. If predictability is perfect, of course, the values predicted will match the actual values and the range of predictions will equal the range of outcomes. In general, the higher the predictability, the wider the range of predicted values.

Several studies of numerical prediction have demonstrated that intuitive predictions violate this rule, and that subjects show little or no regard for considerations of predictability (*I*). In one of these studies, subjects were presented with several paragraphs, each describing the performance of a stu-

dent teacher during a particular practice lesson. Some subjects were asked to *evaluate* the quality of the lesson described in the paragraph in percentile scores, relative to a specified population. Other subjects were asked to *predict*, also in percentile scores, the standing of each student teacher 5 years after the practice lesson. The judgments made under the two conditions were identical. That is, the prediction of a remote criterion (success of a teacher after 5 years) was identical to the evaluation of the information on which the prediction was based (the quality of the practice lesson). The students who made these predictions were undoubtedly aware of the limited predictability of teaching competence on the basis of a single trial lesson 5 years earlier; nevertheless, their predictions were as extreme as their evaluations.

The illusion of validity. As we have seen, people often predict by selecting the outcome (for example, an occupation) that is most representative of the input (for example, the description of a person). The confidence they have in their prediction depends primarily on the degree of representativeness (that is, on the quality of the match between the selected outcome and the input) with little or no regard for the factors that limit predictive accuracy. Thus, people express great confidence in the prediction that a person is a librarian when given a description of his personality which matches the stereotype of librarians, even if the description is scanty, unreliable, or outdated. The unwarranted confidence which is produced by a good fit between the predicted outcome and the input information may be called the illusion of validity. This illusion persists even when the judge is aware of the factors that limit the accuracy of his predictions. It is a common observation that psychologists who conduct selection interviews often experience considerable confidence in their predictions, even when they know of the vast literature that shows selection interviews to be highly fallible. The continued reliance on the clinical interview for selection, despite repeated demonstrations of its inadequacy, amply attests to the strength of this effect.

The internal consistency of a pattern of inputs is a major determinant of one's confidence in predictions based on these inputs. For example, people express more confidence in predicting the final grade-point average of a student

whose first-year record consists entirely of B's than in predicting the grade-point average of a student whose first-year record includes many A's and C's. Highly consistent patterns are most often observed when the input variables are highly redundant or correlated. Hence, people tend to have great confidence in predictions based on redundant input variables. However, an elementary result in the statistics of correlation asserts that, given input variables of stated validity, a prediction based on several such inputs can achieve higher accuracy when they are independent of each other than when they are redundant or correlated. Thus, redundancy among inputs decreases accuracy even as it increases confidence, and people are often confident in predictions that are quite likely to be off the mark (*J*).

Misconceptions of regression. Suppose a large group of children has been examined on two equivalent versions of an aptitude test. If one selects ten children from among those who did best on one of the two versions, he will usually find their performance on the second version to be somewhat disappointing. Conversely, if one selects ten children from among those who did worst on one version, they will be found, on the average, to do somewhat better on the other version. More generally, consider two variables *X* and *Y* which have the same distribution. If one selects individuals whose average *X* score deviates from the mean of *X* by *k* units, then the average of their *Y* scores will usually deviate from the mean of *Y* by less than *k* units. These observations illustrate a general phenomenon known as regression toward the mean, which was first documented by Galton more than 100 years ago.

In the normal course of life, one encounters many instances of regression toward the mean, in the comparison of the height of fathers and sons, of the intelligence of husbands and wives, or of the performance of individuals on consecutive examinations. Nevertheless, people do not develop correct intuitions about this phenomenon. First, they do not expect regression in many contexts where it is bound to occur. Second, when they recognize the occurrence of regression, they often invent spurious causal explanations for it (*I*). We suggest that the phenomenon of regression remains elusive because it is incompatible with the belief that the predicted outcome should be maximally

representative of the input, and, hence, that the value of the outcome variable should be as extreme as the value of the input variable.

The failure to recognize the import of regression can have pernicious consequences, as illustrated by the following observation (1). In a discussion of flight training, experienced instructors noted that praise for an exceptionally smooth landing is typically followed by a poorer landing on the next try, while harsh criticism after a rough landing is usually followed by an improvement on the next try. The instructors concluded that verbal rewards are detrimental to learning, while verbal punishments are beneficial, contrary to accepted psychological doctrine. This conclusion is unwarranted because of the presence of regression toward the mean. As in other cases of repeated examination, an improvement will usually follow a poor performance and a deterioration will usually follow an outstanding performance, even if the instructor does not respond to the trainee's achievement on the first attempt. Because the instructors had praised their trainees after good landings and admonished them after poor ones, they reached the erroneous and potentially harmful conclusion that punishment is more effective than reward.

Thus, the failure to understand the effect of regression leads one to overestimate the effectiveness of punishment and to underestimate the effectiveness of reward. In social interaction, as well as in training, rewards are typically administered when performance is good, and punishments are typically administered when performance is poor. By regression alone, therefore, behavior is most likely to improve after punishment and most likely to deteriorate after reward. Consequently, the human condition is such that, by chance alone, one is most often rewarded for punishing others and most often punished for rewarding them. People are generally not aware of this contingency. In fact, the elusive role of regression in determining the apparent consequences of reward and punishment seems to have escaped the notice of students of this area.

Availability

There are situations in which people assess the frequency of a class or the probability of an event by the ease with

which instances or occurrences can be brought to mind. For example, one may assess the risk of heart attack among middle-aged people by recalling such occurrences among one's acquaintances. Similarly, one may evaluate the probability that a given business venture will fail by imagining various difficulties it could encounter. This judgmental heuristic is called availability. Availability is a useful clue for assessing frequency or probability, because instances of large classes are usually recalled better and faster than instances of less frequent classes. However, availability is affected by factors other than frequency and probability. Consequently, the reliance on availability leads to predictable biases, some of which are illustrated below.

Biases due to the retrievability of instances. When the size of a class is judged by the availability of its instances, a class whose instances are easily retrieved will appear more numerous than a class of equal frequency whose instances are less retrievable. In an elementary demonstration of this effect, subjects heard a list of well-known personalities of both sexes and were subsequently asked to judge whether the list contained more names of men than of women. Different lists were presented to different groups of subjects. In some of the lists the men were relatively more famous than the women, and in others the women were relatively more famous than the men. In each of the lists, the subjects erroneously judged that the class (sex) that had the more famous personalities was the more numerous (6).

In addition to familiarity, there are other factors, such as salience, which affect the retrievability of instances. For example, the impact of seeing a house burning on the subjective probability of such accidents is probably greater than the impact of reading about a fire in the local paper. Furthermore, recent occurrences are likely to be relatively more available than earlier occurrences. It is a common experience that the subjective probability of traffic accidents rises temporarily when one sees a car overturned by the side of the road.

Biases due to the effectiveness of a search set. Suppose one samples a word (of three letters or more) at random from an English text. Is it more likely that the word starts with r or that r is the third letter? People approach this problem by recalling words that

begin with r (road) and words that have r in the third position (car) and assess the relative frequency by the ease with which words of the two types come to mind. Because it is much easier to search for words by their first letter than by their third letter, most people judge words that begin with a given consonant to be more numerous than words in which the same consonant appears in the third position. They do so even for consonants, such as r or k, that are more frequent in the third position than in the first (6).

Different tasks elicit different search sets. For example, suppose you are asked to rate the frequency with which abstract words (thought, love) and concrete words (door, water) appear in written English. A natural way to answer this question is to search for contexts in which the word could appear. It seems easier to think of contexts in which an abstract concept is mentioned (love in love stories) than to think of contexts in which a concrete word (such as door) is mentioned. If the frequency of words is judged by the availability of the contexts in which they appear, abstract words will be judged as relatively more numerous than concrete words. This bias has been observed in a recent study (7) which showed that the judged frequency of occurrence of abstract words was much higher than that of concrete words, equated in objective frequency. Abstract words were also judged to appear in a much greater variety of contexts than concrete words.

Biases of imaginability. Sometimes one has to assess the frequency of a class whose instances are not stored in memory but can be generated according to a given rule. In such situations, one typically generates several instances and evaluates frequency or probability by the ease with which the relevant instances can be constructed. However, the ease of constructing instances does not always reflect their actual frequency, and this mode of evaluation is prone to biases. To illustrate, consider a group of 10 people who form committees of k members, $2 \leq k \leq 8$. How many different committees of k members can be formed? The correct answer to this problem is given by the binomial coefficient $\binom{10}{k}$ which reaches a maximum of 252 for $k = 5$. Clearly, the number of committees of k members equals the number of committees of $(10 - k)$ members, because any committee of k

members defines a unique group of $(10 - k)$ nonmembers.

One way to answer this question without computation is to mentally construct committees of k members and to evaluate their number by the ease with which they come to mind. Committees of few members, say 2, are more available than committees of many members, say 8. The simplest scheme for the construction of committees is a partition of the group into disjoint sets. One readily sees that it is easy to construct five disjoint committees of 2 members, while it is impossible to generate even two disjoint committees of 8 members. Consequently, if frequency is assessed by imaginability, or by availability for construction, the small committees will appear more numerous than larger committees, in contrast to the correct bell-shaped function. Indeed, when naive subjects were asked to estimate the number of distinct committees of various sizes, their estimates were a decreasing monotonic function of committee size (6). For example, the median estimate of the number of committees of 2 members was 70, while the estimate for committees of 8 members was 20 (the correct answer is 45 in both cases).

Imaginability plays an important role in the evaluation of probabilities in real-life situations. The risk involved in an adventurous expedition, for example, is evaluated by imagining contingencies with which the expedition is not equipped to cope. If many such difficulties are vividly portrayed, the expedition can be made to appear exceedingly dangerous, although the ease with which disasters are imagined need not reflect their actual likelihood. Conversely, the risk involved in an undertaking may be grossly underestimated if some possible dangers are either difficult to conceive of, or simply do not come to mind.

Illusory correlation. Chapman and Chapman (8) have described an interesting bias in the judgment of the frequency with which two events co-occur. They presented naive judges with information concerning several hypothetical mental patients. The data for each patient consisted of a clinical diagnosis and a drawing of a person made by the patient. Later the judges estimated the frequency with which each diagnosis (such as paranoia or suspiciousness) had been accompanied by various features of the drawing (such as peculiar eyes). The subjects markedly overestimated the frequency of co-occurrence of

natural associates, such as suspiciousness and peculiar eyes. This effect was labeled illusory correlation. In their erroneous judgments of the data to which they had been exposed, naive subjects "rediscovered" much of the common, but unfounded, clinical lore concerning the interpretation of the draw-a-person test. The illusory correlation effect was extremely resistant to contradictory data. It persisted even when the correlation between symptom and diagnosis was actually negative, and it prevented the judges from detecting relationships that were in fact present.

Availability provides a natural account for the illusory-correlation effect. The judgment of how frequently two events co-occur could be based on the strength of the associative bond between them. When the association is strong, one is likely to conclude that the events have been frequently paired. Consequently, strong associates will be judged to have occurred together frequently. According to this view, the illusory correlation between suspiciousness and peculiar drawing of the eyes, for example, is due to the fact that suspiciousness is more readily associated with the eyes than with any other part of the body.

Lifelong experience has taught us that, in general, instances of large classes are recalled better and faster than instances of less frequent classes; that likely occurrences are easier to imagine than unlikely ones; and that the associative connections between events are strengthened when the events frequently co-occur. As a result, man has at his disposal a procedure (the availability heuristic) for estimating the numerosity of a class, the likelihood of an event, or the frequency of co-occurrences, by the ease with which the relevant mental operations of retrieval, construction, or association can be performed. However, as the preceding examples have demonstrated, this valuable estimation procedure results in systematic errors.

Adjustment and Anchoring

In many situations, people make estimates by starting from an initial value that is adjusted to yield the final answer. The initial value, or starting point, may be suggested by the formulation of the problem, or it may be the result of a partial computation. In either case, adjustments are typically insufficient (4).

That is, different starting points yield different estimates, which are biased toward the initial values. We call this phenomenon anchoring.

Insufficient adjustment. In a demonstration of the anchoring effect, subjects were asked to estimate various quantities, stated in percentages (for example, the percentage of African countries in the United Nations). For each quantity, a number between 0 and 100 was determined by spinning a wheel of fortune in the subjects' presence. The subjects were instructed to indicate first whether that number was higher or lower than the value of the quantity, and then to estimate the value of the quantity by moving upward or downward from the given number. Different groups were given different numbers for each quantity, and these arbitrary numbers had a marked effect on estimates. For example, the median estimates of the percentage of African countries in the United Nations were 25 and 45 for groups that received 10 and 65, respectively, as starting points. Payoffs for accuracy did not reduce the anchoring effect.

Anchoring occurs not only when the starting point is given to the subject, but also when the subject bases his estimate on the result of some incomplete computation. A study of intuitive numerical estimation illustrates this effect. Two groups of high school students estimated, within 5 seconds, a numerical expression that was written on the blackboard. One group estimated the product

$$8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$$

while another group estimated the product

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$$

To rapidly answer such questions, people may perform a few steps of computation and estimate the product by extrapolation or adjustment. Because adjustments are typically insufficient, this procedure should lead to underestimation. Furthermore, because the result of the first few steps of multiplication (performed from left to right) is higher in the descending sequence than in the ascending sequence, the former expression should be judged larger than the latter. Both predictions were confirmed. The median estimate for the ascending sequence was 512, while the median estimate for the descending sequence was 2,250. The correct answer is 40,320.

Biases in the evaluation of conjunctive and disjunctive events. In a recent

study by Bar-Hillel (9) subjects were given the opportunity to bet on one of two events. Three types of events were used: (i) simple events, such as drawing a red marble from a bag containing 50 percent red marbles and 50 percent white marbles; (ii) conjunctive events, such as drawing a red marble seven times in succession, with replacement, from a bag containing 90 percent red marbles and 10 percent white marbles; and (iii) disjunctive events, such as drawing a red marble at least once in seven successive tries, with replacement, from a bag containing 10 percent red marbles and 90 percent white marbles. In this problem, a significant majority of subjects preferred to bet on the conjunctive event (the probability of which is .48) rather than on the simple event (the probability of which is .50). Subjects also preferred to bet on the simple event rather than on the disjunctive event, which has a probability of .52. Thus, most subjects bet on the less likely event in both comparisons. This pattern of choices illustrates a general finding. Studies of choice among gambles and of judgments of probability indicate that people tend to overestimate the probability of conjunctive events (10) and to underestimate the probability of disjunctive events. These biases are readily explained as effects of anchoring. The stated probability of the elementary event (success at any one stage) provides a natural starting point for the estimation of the probabilities of both conjunctive and disjunctive events. Since adjustment from the starting point is typically insufficient, the final estimates remain too close to the probabilities of the elementary events in both cases. Note that the overall probability of a conjunctive event is lower than the probability of each elementary event, whereas the overall probability of a disjunctive event is higher than the probability of each elementary event. As a consequence of anchoring, the overall probability will be overestimated in conjunctive problems and underestimated in disjunctive problems.

Biases in the evaluation of compound events are particularly significant in the context of planning. The successful completion of an undertaking, such as the development of a new product, typically has a conjunctive character: for the undertaking to succeed, each of a series of events must occur. Even when each of these events is very likely, the overall probability of success can be quite low if the number of events is

large. The general tendency to overestimate the probability of conjunctive events leads to unwarranted optimism in the evaluation of the likelihood that a plan will succeed or that a project will be completed on time. Conversely, disjunctive structures are typically encountered in the evaluation of risks. A complex system, such as a nuclear reactor or a human body, will malfunction if any of its essential components fails. Even when the likelihood of failure in each component is slight, the probability of an overall failure can be high if many components are involved. Because of anchoring, people will tend to underestimate the probabilities of failure in complex systems. Thus, the direction of the anchoring bias can sometimes be inferred from the structure of the event. The chain-like structure of conjunctions leads to overestimation, the funnel-like structure of disjunctions leads to underestimation.

Anchoring in the assessment of subjective probability distributions. In decision analysis, experts are often required to express their beliefs about a quantity, such as the value of the Dow-Jones average on a particular day, in the form of a probability distribution. Such a distribution is usually constructed by asking the person to select values of the quantity that correspond to specified percentiles of his subjective probability distribution. For example, the judge may be asked to select a number, X_{90} , such that his subjective probability that this number will be higher than the value of the Dow-Jones average is .90. That is, he should select the value X_{90} so that he is just willing to accept 9 to 1 odds that the Dow-Jones average will not exceed it. A subjective probability distribution for the value of the Dow-Jones average can be constructed from several such judgments corresponding to different percentiles.

By collecting subjective probability distributions for many different quantities, it is possible to test the judge for proper calibration. A judge is properly (or externally) calibrated in a set of problems if exactly Π percent of the true values of the assessed quantities falls below his stated values of X_{Π} . For example, the true values should fall below X_{01} for 1 percent of the quantities and above X_{99} for 1 percent of the quantities. Thus, the true values should fall in the confidence interval between X_{01} and X_{99} on 98 percent of the problems.

Several investigators (11) have ob-

tained probability distributions for many quantities from a large number of judges. These distributions indicated large and systematic departures from proper calibration. In most studies, the actual values of the assessed quantities are either smaller than X_{01} or greater than X_{99} for about 30 percent of the problems. That is, the subjects state overly narrow confidence intervals which reflect more certainty than is justified by their knowledge about the assessed quantities. This bias is common to naive and to sophisticated subjects, and it is not eliminated by introducing proper scoring rules, which provide incentives for external calibration. This effect is attributable, in part at least, to anchoring.

To select X_{90} for the value of the Dow-Jones average, for example, it is natural to begin by thinking about one's best estimate of the Dow-Jones and to adjust this value upward. If this adjustment—like most others—is insufficient, then X_{90} will not be sufficiently extreme. A similar anchoring effect will occur in the selection of X_{10} , which is presumably obtained by adjusting one's best estimate downward. Consequently, the confidence interval between X_{10} and X_{90} will be too narrow, and the assessed probability distribution will be too tight. In support of this interpretation it can be shown that subjective probabilities are systematically altered by a procedure in which one's best estimate does not serve as an anchor.

Subjective probability distributions for a given quantity (the Dow-Jones average) can be obtained in two different ways: (i) by asking the subject to select values of the Dow-Jones that correspond to specified percentiles of his probability distribution and (ii) by asking the subject to assess the probabilities that the true value of the Dow-Jones will exceed some specified values. The two procedures are formally equivalent and should yield identical distributions. However, they suggest different modes of adjustment from different anchors. In procedure (i), the natural starting point is one's best estimate of the quantity. In procedure (ii), on the other hand, the subject may be anchored on the value stated in the question. Alternatively, he may be anchored on even odds, or 50-50 chances, which is a natural starting point in the estimation of likelihood. In either case, procedure (ii) should yield less extreme odds than procedure (i).

To contrast the two procedures, a set of 24 quantities (such as the air dis-

tance from New Delhi to Peking) was presented to a group of subjects who assessed either X_{10} or X_{90} for each problem. Another group of subjects received the median judgment of the first group for each of the 24 quantities. They were asked to assess the odds that each of the given values exceeded the true value of the relevant quantity. In the absence of any bias, the second group should retrieve the odds specified to the first group, that is, 9 : 1. However, if even odds or the stated value serve as anchors, the odds of the second group should be less extreme, that is, closer to 1 : 1. Indeed, the median odds stated by this group, across all problems, were 3 : 1. When the judgments of the two groups were tested for external calibration, it was found that subjects in the first group were too extreme, in accord with earlier studies. The events that they defined as having a probability of .10 actually obtained in 24 percent of the cases. In contrast, subjects in the second group were too conservative. Events to which they assigned an average probability of .34 actually obtained in 26 percent of the cases. These results illustrate the manner in which the degree of calibration depends on the procedure of elicitation.

Discussion

This article has been concerned with cognitive biases that stem from the reliance on judgmental heuristics. These biases are not attributable to motivational effects such as wishful thinking or the distortion of judgments by payoffs and penalties. Indeed, several of the severe errors of judgment reported earlier occurred despite the fact that subjects were encouraged to be accurate and were rewarded for the correct answers (2, 6).

The reliance on heuristics and the prevalence of biases are not restricted to laymen. Experienced researchers are also prone to the same biases—when they think intuitively. For example, the tendency to predict the outcome that best represents the data, with insufficient regard for prior probability, has been observed in the intuitive judgments of individuals who have had extensive training in statistics (1, 5). Although the statistically sophisticated avoid elementary errors, such as the gambler's fallacy, their intuitive judgments are liable to similar fallacies in more intricate and less transparent problems.

It is not surprising that useful heuristics such as representativeness and availability are retained, even though they occasionally lead to errors in prediction or estimation. What is perhaps surprising is the failure of people to infer from lifelong experience such fundamental statistical rules as regression toward the mean, or the effect of sample size on sampling variability. Although everyone is exposed, in the normal course of life, to numerous examples from which these rules could have been induced, very few people discover the principles of sampling and regression on their own. Statistical principles are not learned from everyday experience because the relevant instances are not coded appropriately. For example, people do not discover that successive lines in a text differ more in average word length than do successive pages, because they simply do not attend to the average word length of individual lines or pages. Thus, people do not learn the relation between sample size and sampling variability, although the data for such learning are abundant.

The lack of an appropriate code also explains why people usually do not detect the biases in their judgments of probability. A person could conceivably learn whether his judgments are externally calibrated by keeping a tally of the proportion of events that actually occur among those to which he assigns the same probability. However, it is not natural to group events by their judged probability. In the absence of such grouping it is impossible for an individual to discover, for example, that only 50 percent of the predictions to which he has assigned a probability of .9 or higher actually came true.

The empirical analysis of cognitive biases has implications for the theoretical and applied role of judged probabilities. Modern decision theory (12, 13) regards subjective probability as the quantified opinion of an idealized person. Specifically, the subjective probability of a given event is defined by the set of bets about this event that such a person is willing to accept. An internally consistent, or coherent, subjective probability measure can be derived for an individual if his choices among bets satisfy certain principles, that is, the axioms of the theory. The derived probability is subjective in the sense that different individuals are allowed to have different probabilities for the same event. The major contribution of this approach is that it provides a rigorous

subjective interpretation of probability that is applicable to unique events and is embedded in a general theory of rational decision.

It should perhaps be noted that, while subjective probabilities can sometimes be inferred from preferences among bets, they are normally not formed in this fashion. A person bets on team A rather than on team B because he believes that team A is more likely to win; he does not infer this belief from his betting preferences. Thus, in reality, subjective probabilities determine preferences among bets and are not derived from them, as in the axiomatic theory of rational decision (12).

The inherently subjective nature of probability has led many students to the belief that coherence, or internal consistency, is the only valid criterion by which judged probabilities should be evaluated. From the standpoint of the formal theory of subjective probability, any set of internally consistent probability judgments is as good as any other. This criterion is not entirely satisfactory, because an internally consistent set of subjective probabilities can be incompatible with other beliefs held by the individual. Consider a person whose subjective probabilities for all possible outcomes of a coin-tossing game reflect the gambler's fallacy. That is, his estimate of the probability of tails on a particular toss increases with the number of consecutive heads that preceded that toss. The judgments of such a person could be internally consistent and therefore acceptable as adequate subjective probabilities according to the criterion of the formal theory. These probabilities, however, are incompatible with the generally held belief that a coin has no memory and is therefore incapable of generating sequential dependencies. For judged probabilities to be considered adequate, or rational, internal consistency is not enough. The judgments must be compatible with the entire web of beliefs held by the individual. Unfortunately, there can be no simple formal procedure for assessing the compatibility of a set of probability judgments with the judge's total system of beliefs. The rational judge will nevertheless strive for compatibility, even though internal consistency is more easily achieved and assessed. In particular, he will attempt to make his probability judgments compatible with his knowledge about the subject matter, the laws of probability, and his own judgmental heuristics and biases.

Summary

This article described three heuristics that are employed in making judgments under uncertainty: (i) representativeness, which is usually employed when people are asked to judge the probability that an object or event A belongs to class or process B; (ii) availability of instances or scenarios, which is often employed when people are asked to assess the frequency of a class or the plausibility of a particular development; and (iii) adjustment from an anchor, which is usually employed in numerical prediction when a relevant value is available. These heuristics are highly economical

and usually effective, but they lead to systematic and predictable errors. A better understanding of these heuristics and of the biases to which they lead could improve judgments and decisions in situations of uncertainty.

References and Notes

1. D. Kahneman and A. Tversky, *Psychol. Rev.* **80**, 237 (1973).
2. ———, *Cognitive Psychol.* **3**, 430 (1972).
3. W. Edwards, in *Formal Representation of Human Judgment*, B. Kleinmuntz, Ed. (Wiley, New York, 1968), pp. 17–52.
4. P. Slovic and S. Lichtenstein, *Organ. Behav. Hum. Performance* **6**, 649 (1971).
5. A. Tversky and D. Kahneman, *Psychol. Bull.* **76**, 105 (1971).
6. ———, *Cognitive Psychol.* **5**, 207 (1973).
7. R. C. Galbraith and B. J. Underwood, *Mem. Cognition* **1**, 56 (1973).
8. L. J. Chapman and J. P. Chapman, *J. Abnorm. Psychol.* **73**, 193 (1967); *ibid.*, **74**, 271 (1969).
9. M. Bar-Hillel, *Organ. Behav. Hum. Performance* **9**, 396 (1973).
10. J. Cohen, E. I. Chesnick, D. Haran, *Br. J. Psychol.* **63**, 41 (1972).
11. M. Alpert and H. Raiffa, unpublished manuscript; C. A. S. von Holstein, *Acta Psychol.* **35**, 478 (1971); R. L. Winkler, *J. Am. Stat. Assoc.* **62**, 776 (1967).
12. L. J. Savage, *The Foundations of Statistics* (Wiley, New York, 1954).
13. B. De Finetti, in *International Encyclopedia of the Social Sciences*, D. E. Sills, Ed. (Macmillan, New York, 1968), vol. 12, pp. 496–504.
14. This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under contract N00014-73-C-0438 to the Oregon Research Institute, Eugene. Additional support for this research was provided by the Research and Development Authority of the Hebrew University, Jerusalem, Israel.

PROSPECT THEORY: AN ANALYSIS
OF DECISION UNDER RISK

Daniel Kahneman
University of British Columbia

Amos Tversky
Stanford University

PROSPECT THEORY: AN ANALYSIS OF DECISION UNDER RISK

BY DANIEL KAHNEMAN AND AMOS TVERSKY¹

This paper presents a critique of expected utility theory as a descriptive model of decision making under risk, and develops an alternative model, called prospect theory. Choices among risky prospects exhibit several pervasive effects that are inconsistent with the basic tenets of utility theory. In particular, people underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty. This tendency, called the certainty effect, contributes to risk aversion in choices involving sure gains and to risk seeking in choices involving sure losses. In addition, people generally discard components that are shared by all prospects under consideration. This tendency, called the isolation effect, leads to inconsistent preferences when the same choice is presented in different forms. An alternative theory of choice is developed, in which value is assigned to gains and losses rather than to final assets and in which probabilities are replaced by decision weights. The value function is normally concave for gains, commonly convex for losses, and is generally steeper for losses than for gains. Decision weights are generally lower than the corresponding probabilities, except in the range of low probabilities. Overweighting of low probabilities may contribute to the attractiveness of both insurance and gambling.

1. INTRODUCTION

EXPECTED UTILITY THEORY has dominated the analysis of decision making under risk. It has been generally accepted as a normative model of rational choice [24], and widely applied as a descriptive model of economic behavior, e.g. [15, 4]. Thus, it is assumed that all reasonable people would wish to obey the axioms of the theory [47, 36], and that most people actually do, most of the time.

The present paper describes several classes of choice problems in which preferences systematically violate the axioms of expected utility theory. In the light of these observations we argue that utility theory, as it is commonly interpreted and applied, is not an adequate descriptive model and we propose an alternative account of choice under risk.

2. CRITIQUE

Decision making under risk can be viewed as a choice between prospects or gambles. A prospect $(x_1, p_1; \dots; x_n, p_n)$ is a contract that yields outcome x_i with probability p_i , where $p_1 + p_2 + \dots + p_n = 1$. To simplify notation, we omit null outcomes and use (x, p) to denote the prospect $(x, p; 0, 1 - p)$ that yields x with probability p and 0 with probability $1 - p$. The (riskless) prospect that yields x with certainty is denoted by (x) . The present discussion is restricted to prospects with so-called objective or standard probabilities.

The application of expected utility theory to choices between prospects is based on the following three tenets.

- (i) Expectation: $U(x_1, p_1; \dots; x_n, p_n) = p_1 u(x_1) + \dots + p_n u(x_n)$.

¹ This work was supported in part by grants from the Harry F. Guggenheim Foundation and from the Advanced Research Projects Agency of the Department of Defense and was monitored by Office of Naval Research under Contract N00014-78-C-0100 (ARPA Order No. 3469) under Subcontract 78-072-0722 from Decisions and Designs, Inc. to Perceptronics, Inc. We also thank the Center for Advanced Study in the Behavioral Sciences at Stanford for its support.

That is, the overall utility of a prospect, denoted by U , is the expected utility of its outcomes.

(ii) Asset Integration: $(x_1, p_1; \dots; x_n, p_n)$ is acceptable at asset position w iff $U(w + x_1, p_1; \dots; w + x_n, p_n) > u(w)$.

That is, a prospect is acceptable if the utility resulting from integrating the prospect with one's assets exceeds the utility of those assets alone. Thus, the domain of the utility function is final states (which include one's asset position) rather than gains or losses.

Although the domain of the utility function is not limited to any particular class of consequences, most applications of the theory have been concerned with monetary outcomes. Furthermore, most economic applications introduce the following additional assumption.

(iii) Risk Aversion: u is concave ($u'' < 0$).

A person is risk averse if he prefers the certain prospect (x) to any risky prospect with expected value x . In expected utility theory, risk aversion is equivalent to the concavity of the utility function. The prevalence of risk aversion is perhaps the best known generalization regarding risky choices. It led the early decision theorists of the eighteenth century to propose that utility is a concave function of money, and this idea has been retained in modern treatments (Pratt [33], Arrow [4]).

In the following sections we demonstrate several phenomena which violate these tenets of expected utility theory. The demonstrations are based on the responses of students and university faculty to hypothetical choice problems. The respondents were presented with problems of the type illustrated below.

Which of the following would you prefer?

- A: 50% chance to win 1,000, B: 450 for sure.
50% chance to win nothing;

The outcomes refer to Israeli currency. To appreciate the significance of the amounts involved, note that the median net monthly income for a family is about 3,000 Israeli pounds. The respondents were asked to imagine that they were actually faced with the choice described in the problem, and to indicate the decision they would have made in such a case. The responses were anonymous, and the instructions specified that there was no 'correct' answer to such problems, and that the aim of the study was to find out how people choose among risky prospects. The problems were presented in questionnaire form, with at most a dozen problems per booklet. Several forms of each questionnaire were constructed so that subjects were exposed to the problems in different orders. In addition, two versions of each problem were used in which the left-right position of the prospects was reversed.

The problems described in this paper are selected illustrations of a series of effects. Every effect has been observed in several problems with different outcomes and probabilities. Some of the problems have also been presented to groups of students and faculty at the University of Stockholm and at the

University of Michigan. The pattern of results was essentially identical to the results obtained from Israeli subjects.

The reliance on hypothetical choices raises obvious questions regarding the validity of the method and the generalizability of the results. We are keenly aware of these problems. However, all other methods that have been used to test utility theory also suffer from severe drawbacks. Real choices can be investigated either in the field, by naturalistic or statistical observations of economic behavior, or in the laboratory. Field studies can only provide for rather crude tests of qualitative predictions, because probabilities and utilities cannot be adequately measured in such contexts. Laboratory experiments have been designed to obtain precise measures of utility and probability from actual choices, but these experimental studies typically involve contrived gambles for small stakes, and a large number of repetitions of very similar problems. These features of laboratory gambling complicate the interpretation of the results and restrict their generality.

By default, the method of hypothetical choices emerges as the simplest procedure by which a large number of theoretical questions can be investigated. The use of the method relies on the assumption that people often know how they would behave in actual situations of choice, and on the further assumption that the subjects have no special reason to disguise their true preferences. If people are reasonably accurate in predicting their choices, the presence of common and systematic violations of expected utility theory in hypothetical problems provides presumptive evidence against that theory.

Certainty, Probability, and Possibility

In expected utility theory, the utilities of outcomes are weighted by their probabilities. The present section describes a series of choice problems in which people's preferences systematically violate this principle. We first show that people overweight outcomes that are considered certain, relative to outcomes which are merely probable—a phenomenon which we label the *certainty effect*.

The best known counter-example to expected utility theory which exploits the certainty effect was introduced by the French economist Maurice Allais in 1953 [2]. Allais' example has been discussed from both normative and descriptive standpoints by many authors [28, 38]. The following pair of choice problems is a variation of Allais' example, which differs from the original in that it refers to moderate rather than to extremely large gains. The number of respondents who answered each problem is denoted by N , and the percentage who choose each option is given in brackets.

PROBLEM 1: Choose between

A: 2,500 with probability	.33,	B: 2,400 with certainty.
2,400 with probability	.66,	
0 with probability	.01;	
$N = 72$	[18]	[82]*

PROSPECT THEORY

.8 to .2. The following pair of choice problems illustrates the certainty effect with non-monetary outcomes.

PROBLEM 5:

A: 50% chance to win a three-week tour of England, France, and Italy;	B: A one-week tour of England, with certainty.
$N = 72$ [22]	[78]*

PROBLEM 6:

C: 5% chance to win a three-week tour of England, France, and Italy;	D: 10% chance to win a one-week tour of England.
$N = 72$ [67]*	[33]

The certainty effect is not the only type of violation of the substitution axiom. Another situation in which this axiom fails is illustrated by the following problems.

PROBLEM 7:

A: (6,000, .45),	B: (3,000, .90).
$N = 66$ [14]	[86]*

PROBLEM 8:

C: (6,000, .001),	D: (3,000, .002).
$N = 66$ [73]*	[27]

Note that in Problem 7 the probabilities of winning are substantial (.90 and .45), and most people choose the prospect where winning is more probable. In Problem 8, there is a *possibility* of winning, although the probabilities of winning are minuscule (.002 and .001) in both prospects. In this situation where winning is possible but not probable, most people choose the prospect that offers the larger gain. Similar results have been reported by MacCrimmon and Larsson [28].

The above problems illustrate common attitudes toward risk or chance that cannot be captured by the expected utility model. The results suggest the following empirical generalization concerning the manner in which the substitution axiom is violated. If (y, pq) is equivalent to (x, p) , then (y, pqr) is preferred to (x, pr) , $0 < p, q, r < 1$. This property is incorporated into an alternative theory, developed in the second part of the paper.

The Reflection Effect

The previous section discussed preferences between positive prospects, i.e., prospects that involve no losses. What happens when the signs of the outcomes are reversed so that gains are replaced by losses? The left-hand column of Table I displays four of the choice problems that were discussed in the previous section, and the right-hand column displays choice problems in which the signs of the outcomes are reversed. We use $-x$ to denote the loss of x , and $>$ to denote the prevalent preference, i.e., the choice made by the majority of subjects.

TABLE I
PREFERENCES BETWEEN POSITIVE AND NEGATIVE PROSPECTS

	Positive prospects		Negative prospects
Problem 3:	$(4,000, .80) < (3,000)$	Problem 3':	$(-4,000, .80) > (-3,000)$
$N = 95$	[20]	$N = 95$	[92]*
Problem 4:	$(4,000, .20) > (3,000, .25)$	Problem 4':	$(-4,000, .20) < (-3,000, .25)$
$N = 95$	[65]*	$N = 95$	[42]
Problem 7:	$(3,000, .90) > (6,000, .45)$	Problem 7':	$(-3,000, .90) < (-6,000, .45)$
$N = 66$	[86]*	$N = 66$	[8]
Problem 8:	$(3,000, .002) < (6,000, .001)$	Problem 8':	$(-3,000, .002) > (-6,000, .001)$
$N = 66$	[27]	$N = 66$	[70]*

In each of the four problems in Table I the preference between negative prospects is the mirror image of the preference between positive prospects. Thus, the reflection of prospects around 0 reverses the preference order. We label this pattern the *reflection effect*.

Let us turn now to the implications of these data. First, note that the reflection effect implies that risk aversion in the positive domain is accompanied by risk seeking in the negative domain. In Problem 3', for example, the majority of subjects were willing to accept a risk of .80 to lose 4,000, in preference to a sure loss of 3,000, although the gamble has a lower expected value. The occurrence of risk seeking in choices between negative prospects was noted early by Markowitz [29]. Williams [48] reported data where a translation of outcomes produces a dramatic shift from risk aversion to risk seeking. For example, his subjects were indifferent between $(100, .65; -100, .35)$ and (0) , indicating risk aversion. They were also indifferent between $(-200, .80)$ and (-100) , indicating risk seeking. A recent review by Fishburn and Kochenberger [14] documents the prevalence of risk seeking in choices between negative prospects.

Second, recall that the preferences between the positive prospects in Table I are inconsistent with expected utility theory. The preferences between the corresponding negative prospects also violate the expectation principle in the same manner. For example, Problems 3' and 4', like Problems 3 and 4, demonstrate that outcomes which are obtained with certainty are overweighted relative to uncertain outcomes. In the positive domain, the certainty effect contributes to a risk averse preference for a sure gain over a larger gain that is merely probable. In the negative domain, the same effect leads to a risk seeking preference for a loss

PROSPECT THEORY

that is merely probable over a smaller loss that is certain. The same psychological principle—the overweighting of certainty—favors risk aversion in the domain of gains and risk seeking in the domain of losses.

Third, the reflection effect eliminates aversion for uncertainty or variability as an explanation of the certainty effect. Consider, for example, the prevalent preferences for (3,000) over (4,000, .80) and for (4,000, .20) over (3,000, .25). To resolve this apparent inconsistency one could invoke the assumption that people prefer prospects that have high expected value and small variance (see, e.g., Allais [2]; Markowitz [30]; Tobin [41]). Since (3,000) has no variance while (4,000, .80) has large variance, the former prospect could be chosen despite its lower expected value. When the prospects are reduced, however, the difference in variance between (3,000, .25) and (4,000, .20) may be insufficient to overcome the difference in expected value. Because (−3,000) has both higher expected value and lower variance than (−4,000, .80), this account entails that the sure loss should be preferred, contrary to the data. Thus, our data are incompatible with the notion that certainty is generally desirable. Rather, it appears that certainty increases the aversiveness of losses as well as the desirability of gains.

Probabilistic Insurance

The prevalence of the purchase of insurance against both large and small losses has been regarded by many as strong evidence for the concavity of the utility function for money. Why otherwise would people spend so much money to purchase insurance policies at a price that exceeds the expected actuarial cost? However, an examination of the relative attractiveness of various forms of insurance does not support the notion that the utility function for money is concave everywhere. For example, people often prefer insurance programs that offer limited coverage with low or zero deductible over comparable policies that offer higher maximal coverage with higher deductibles—contrary to risk aversion (see, e.g., Fuchs [16]). Another type of insurance problem in which people's responses are inconsistent with the concavity hypothesis may be called probabilistic insurance. To illustrate this concept, consider the following problem, which was presented to 95 Stanford University students.

PROBLEM 9: Suppose you consider the possibility of insuring some property against damage, e.g., fire or theft. After examining the risks and the premium you find that you have no clear preference between the options of purchasing insurance or leaving the property uninsured.

It is then called to your attention that the insurance company offers a new program called *probabilistic insurance*. In this program you pay half of the regular premium. In case of damage, there is a 50 per cent chance that you pay the other half of the premium and the insurance company covers all the losses; and there is a 50 per cent chance that you get back your insurance payment and suffer all the losses. For example, if an accident occurs on an odd day of the month, you pay the other half of the regular premium and your losses are covered; but if the accident

occurs on an even day of the month, your insurance payment is refunded and your losses are not covered.

Recall that the premium for full coverage is such that you find this insurance barely worth its cost.

Under these circumstances, would you purchase probabilistic insurance:

	Yes,	No.
$N = 95$	[20]	[80]*

Although Problem 9 may appear contrived, it is worth noting that probabilistic insurance represents many forms of protective action where one pays a certain cost to reduce the probability of an undesirable event—without eliminating it altogether. The installation of a burglar alarm, the replacement of old tires, and the decision to stop smoking can all be viewed as probabilistic insurance.

The responses to Problem 9 and to several other variants of the same question indicate that probabilistic insurance is generally unattractive. Apparently, reducing the probability of a loss from p to $p/2$ is less valuable than reducing the probability of that loss from $p/2$ to 0.

In contrast to these data, expected utility theory (with a concave u) implies that probabilistic insurance is superior to regular insurance. That is, if at asset position w one is just willing to pay a premium y to insure against a probability p of losing x , then one should definitely be willing to pay a smaller premium ry to reduce the probability of losing x from p to $(1-r)p$, $0 < r < 1$. Formally, if one is indifferent between $(w-x, p; w, 1-p)$ and $(w-y)$, then one should prefer probabilistic insurance $(w-x, (1-r)p; w-y, rp; w-ry, 1-p)$ over regular insurance $(w-y)$.

To prove this proposition, we show that

$$pu(w-x) + (1-p)u(w) = u(w-y)$$

implies

$$(1-r)pu(w-x) + rpu(w-y) + (1-p)u(w-ry) > u(w-y).$$

Without loss of generality, we can set $u(w-x) = 0$ and $u(w) = 1$. Hence, $u(w-y) = 1-p$, and we wish to show that

$$rp(1-p) + (1-p)u(w-ry) > 1-p \quad \text{or} \quad u(w-ry) > 1-rp$$

which holds if and only if u is concave.

This is a rather puzzling consequence of the risk aversion hypothesis of utility theory, because probabilistic insurance appears intuitively riskier than regular insurance, which entirely eliminates the element of risk. Evidently, the intuitive notion of risk is not adequately captured by the assumed concavity of the utility function for wealth.

The aversion for probabilistic insurance is particularly intriguing because all insurance is, in a sense, probabilistic. The most avid buyer of insurance remains vulnerable to many financial and other risks which his policies do not cover. There appears to be a significant difference between probabilistic insurance and what may be called contingent insurance, which provides the certainty of coverage for a

PROSPECT THEORY

specified type of risk. Compare, for example, probabilistic insurance against all forms of loss or damage to the contents of your home and contingent insurance that eliminates all risk of loss from theft, say, but does not cover other risks, e.g., fire. We conjecture that contingent insurance will be generally more attractive than probabilistic insurance when the probabilities of unprotected loss are equated. Thus, two prospects that are equivalent in probabilities and outcomes could have different values depending on their formulation. Several demonstrations of this general phenomenon are described in the next section.

The Isolation Effect

In order to simplify the choice between alternatives, people often disregard components that the alternatives share, and focus on the components that distinguish them (Tversky [44]). This approach to choice problems may produce inconsistent preferences, because a pair of prospects can be decomposed into common and distinctive components in more than one way, and different decompositions sometimes lead to different preferences. We refer to this phenomenon as the *isolation effect*.

PROBLEM 10: Consider the following two-stage game. In the first stage, there is a probability of .75 to end the game without winning anything, and a probability of .25 to move into the second stage. If you reach the second stage you have a choice between

(4,000, .80) and (3,000).

Your choice must be made before the game starts, i.e., before the outcome of the first stage is known.

Note that in this game, one has a choice between $.25 \times .80 = .20$ chance to win 4,000, and a $.25 \times 1.0 = .25$ chance to win 3,000. Thus, in terms of final outcomes and probabilities one faces a choice between (4,000, .20) and (3,000, .25), as in Problem 4 above. However, the dominant preferences are different in the two problems. Of 141 subjects who answered Problem 10, 78 per cent chose the latter prospect, contrary to the modal preference in Problem 4. Evidently, people ignored the first stage of the game, whose outcomes are shared by both prospects, and considered Problem 10 as a choice between (3,000) and (4,000, .80), as in Problem 3 above.

The standard and the sequential formulations of Problem 4 are represented as decision trees in Figures 1 and 2, respectively. Following the usual convention, squares denote decision nodes and circles denote chance nodes. The essential difference between the two representations is in the location of the decision node. In the standard form (Figure 1), the decision maker faces a choice between two risky prospects, whereas in the sequential form (Figure 2) he faces a choice between a risky and a riskless prospect. This is accomplished by introducing a dependency between the prospects without changing either probabilities or




FIGURE 1.—The representation of Problem 4 as a decision tree (standard formulation).




FIGURE 2.—The representation of Problem 10 as a decision tree (sequential formulation).

outcomes. Specifically, the event 'not winning 3,000' is included in the event 'not winning 4,000' in the sequential formulation, while the two events are independent in the standard formulation. Thus, the outcome of winning 3,000 has a certainty advantage in the sequential formulation, which it does not have in the standard formulation.

The reversal of preferences due to the dependency among events is particularly significant because it violates the basic supposition of a decision-theoretical analysis, that choices between prospects are determined solely by the probabilities of final states.

It is easy to think of decision problems that are most naturally represented in one of the forms above rather than in the other. For example, the choice between two different risky ventures is likely to be viewed in the standard form. On the other hand, the following problem is most likely to be represented in the sequential form. One may invest money in a venture with some probability of losing one's capital if the venture fails, and with a choice between a fixed agreed return and a percentage of earnings if it succeeds. The isolation effect implies that the contingent certainty of the fixed return enhances the attractiveness of this option, relative to a risky venture with the same probabilities and outcomes.

3. THEORY

The preceding discussion reviewed several empirical effects which appear to invalidate expected utility theory as a descriptive model. The remainder of the paper presents an alternative account of individual decision making under risk, called prospect theory. The theory is developed for simple prospects with monetary outcomes and stated probabilities, but it can be extended to more involved choices. Prospect theory distinguishes two phases in the choice process: an early phase of editing and a subsequent phase of evaluation. The editing phase consists of a preliminary analysis of the offered prospects, which often yields a simpler representation of these prospects. In the second phase, the edited prospects are evaluated and the prospect of highest value is chosen. We next outline the editing phase, and develop a formal model of the evaluation phase.

The function of the editing phase is to organize and reformulate the options so as to simplify subsequent evaluation and choice. Editing consists of the application of several operations that transform the outcomes and probabilities associated with the offered prospects. The major operations of the editing phase are described below.

Coding. The evidence discussed in the previous section shows that people normally perceive outcomes as gains and losses, rather than as final states of wealth or welfare. Gains and losses, of course, are defined relative to some neutral reference point. The reference point usually corresponds to the current asset position, in which case gains and losses coincide with the actual amounts that are received or paid. However, the location of the reference point, and the consequent coding of outcomes as gains or losses, can be affected by the formulation of the offered prospects, and by the expectations of the decision maker.

Combination. Prospects can sometimes be simplified by combining the probabilities associated with identical outcomes. For example, the prospect $(200, .25; 200, .25)$ will be reduced to $(200, .50)$, and evaluated in this form.

Segregation. Some prospects contain a riskless component that is segregated from the risky component in the editing phase. For example, the prospect $(300, .80; 200, .20)$ is naturally decomposed into a sure gain of 200 and the risky prospect $(100, .80)$. Similarly, the prospect $(-400, .40; -100, .60)$ is readily seen to consist of a sure loss of 100 and of the prospect $(-300, .40)$.

The preceding operations are applied to each prospect separately. The following operation is applied to a set of two or more prospects.

Cancellation. The essence of the isolation effects described earlier is the discarding of components that are shared by the offered prospects. Thus, our respondents apparently ignored the first stage of the sequential game presented in Problem 10, because this stage was common to both options, and they evaluated the prospects with respect to the results of the second stage (see Figure 2). Similarly, they neglected the common bonus that was added to the prospects in Problems 11 and 12. Another type of cancellation involves the discarding of common constituents, i.e., outcome-probability pairs. For example, the choice

PROSPECT THEORY

between (200, .20; 100, .50; -50, .30) and (200, .20; 150, .50; -100, .30) can be reduced by cancellation to a choice between (100, .50; -50, .30) and (150, .50; -100, .30).

Two additional operations that should be mentioned are simplification and the detection of dominance. The first refers to the simplification of prospects by rounding probabilities or outcomes. For example, the prospect (101, .49) is likely to be recoded as an even chance to win 100. A particularly important form of simplification involves the discarding of extremely unlikely outcomes. The second operation involves the scanning of offered prospects to detect dominated alternatives, which are rejected without further evaluation.

Because the editing operations facilitate the task of decision, it is assumed that they are performed whenever possible. However, some editing operations either permit or prevent the application of others. For example, (500, .20; 101, .49) will appear to dominate (500, .15; 99, .51) if the second constituents of both prospects are simplified to (100, .50). The final edited prospects could, therefore, depend on the sequence of editing operations, which is likely to vary with the structure of the offered set and with the format of the display. A detailed study of this problem is beyond the scope of the present treatment. In this paper we discuss choice problems where it is reasonable to assume either that the original formulation of the prospects leaves no room for further editing, or that the edited prospects can be specified without ambiguity.

Many anomalies of preference result from the editing of prospects. For example, the inconsistencies associated with the isolation effect result from the cancellation of common components. Some intransitivities of choice are explained by a simplification that eliminates small differences between prospects (see Tversky [43]). More generally, the preference order between prospects need not be invariant across contexts, because the same offered prospect could be edited in different ways depending on the context in which it appears.

Following the editing phase, the decision maker is assumed to evaluate each of the edited prospects, and to choose the prospect of highest value. The overall value of an edited prospect, denoted V , is expressed in terms of two scales, π and v .

The first scale, π , associates with each probability p a decision weight $\pi(p)$, which reflects the impact of p on the over-all value of the prospect. However, π is not a probability measure, and it will be shown later that $\pi(p) + \pi(1-p)$ is typically less than unity. The second scale, v , assigns to each outcome x a number $v(x)$, which reflects the subjective value of that outcome. Recall that outcomes are defined relative to a reference point, which serves as the zero point of the value scale. Hence, v measures the value of deviations from that reference point, i.e., gains and losses.

The present formulation is concerned with simple prospects of the form $(x, p; y, q)$, which have at most two non-zero outcomes. In such a prospect, one receives x with probability p , y with probability q , and nothing with probability $1-p-q$, where $p+q \leq 1$. An offered prospect is strictly positive if its outcomes are all positive, i.e., if $x, y > 0$ and $p+q = 1$; it is strictly negative if its outcomes

are all negative. A prospect is regular if it is neither strictly positive nor strictly negative.

The basic equation of the theory describes the manner in which π and v are combined to determine the over-all value of regular prospects.

If $(x, p; y, q)$ is a regular prospect (i.e., either $p + q < 1$, or $x \geq 0 \geq y$, or $x \leq 0 \leq y$), then

$$(1) \quad V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y)$$

where $v(0) = 0$, $\pi(0) = 0$, and $\pi(1) = 1$. As in utility theory, V is defined on prospects, while v is defined on outcomes. The two scales coincide for sure prospects, where $V(x, 1.0) = V(x) = v(x)$.

Equation (1) generalizes expected utility theory by relaxing the expectation principle. An axiomatic analysis of this representation is sketched in the Appendix, which describes conditions that ensure the existence of a unique π and a ratio-scale v satisfying equation (1).

The evaluation of strictly positive and strictly negative prospects follows a different rule. In the editing phase such prospects are segregated into two components: (i) the riskless component, i.e., the minimum gain or loss which is certain to be obtained or paid; (ii) the risky component, i.e., the additional gain or loss which is actually at stake. The evaluation of such prospects is described in the next equation.

If $p + q = 1$ and either $x > y > 0$ or $x < y < 0$, then

$$(2) \quad V(x, p; y, q) = v(y) + \pi(p)[v(x) - v(y)].$$

That is, the value of a strictly positive or strictly negative prospect equals the value of the riskless component plus the value-difference between the outcomes, multiplied by the weight associated with the more extreme outcome. For example, $V(400, .25; 100, .75) = v(100) + \pi(.25)[v(400) - v(100)]$. The essential feature of equation (2) is that a decision weight is applied to the value-difference $v(x) - v(y)$, which represents the risky component of the prospect, but not to $v(y)$, which represents the riskless component. Note that the right-hand side of equation (2) equals $\pi(p)v(x) + [1 - \pi(p)]v(y)$. Hence, equation (2) reduces to equation (1) if $\pi(p) + \pi(1 - p) = 1$. As will be shown later, this condition is not generally satisfied.

Many elements of the evaluation model have appeared in previous attempts to modify expected utility theory. Markowitz [29] was the first to propose that utility be defined on gains and losses rather than on final asset positions, an assumption which has been implicitly accepted in most experimental measurements of utility (see, e.g., [7, 32]). Markowitz also noted the presence of risk seeking in preferences among positive as well as among negative prospects, and he proposed a utility function which has convex and concave regions in both the positive and the negative domains. His treatment, however, retains the expectation principle; hence it cannot account for the many violations of this principle; see, e.g., Table I.

The replacement of probabilities by more general weights was proposed by Edwards [9], and this model was investigated in several empirical studies (e.g.,

[3, 42]). Similar models were developed by Fellner [12], who introduced the concept of decision weight to explain aversion for ambiguity, and by van Dam [46] who attempted to scale decision weights. For other critical analyses of expected utility theory and alternative choice models, see Allais [2], Coombs [6], Fishburn [13], and Hansson [22].

The equations of prospect theory retain the general bilinear form that underlies expected utility theory. However, in order to accommodate the effects described in the first part of the paper, we are compelled to assume that values are attached to changes rather than to final states, and that decision weights do not coincide with stated probabilities. These departures from expected utility theory must lead to normatively unacceptable consequences, such as inconsistencies, intransitivities, and violations of dominance. Such anomalies of preference are normally corrected by the decision maker when he realizes that his preferences are inconsistent, intransitive, or inadmissible. In many situations, however, the decision maker does not have the opportunity to discover that his preferences could violate decision rules that he wishes to obey. In these circumstances the anomalies implied by prospect theory are expected to occur.

The Value Function

An essential feature of the present theory is that the carriers of value are changes in wealth or welfare, rather than final states. This assumption is compatible with basic principles of perception and judgment. Our perceptual apparatus is attuned to the evaluation of changes or differences rather than to the evaluation of absolute magnitudes. When we respond to attributes such as brightness, loudness, or temperature, the past and present context of experience defines an adaptation level, or reference point, and stimuli are perceived in relation to this reference point [23]. Thus, an object at a given temperature may be experienced as hot or cold to the touch depending on the temperature to which one has adapted. The same principle applies to non-sensory attributes such as health, prestige, and wealth. The same level of wealth, for example, may imply abject poverty for one person and great riches for another—depending on their current assets.

The emphasis on changes as the carriers of value should not be taken to imply that the value of a particular change is independent of initial position. Strictly speaking, value should be treated as a function in two arguments: the asset position that serves as reference point, and the magnitude of the change (positive or negative) from that reference point. An individual's attitude to money, say, could be described by a book, where each page presents the value function for changes at a particular asset position. Clearly, the value functions described on different pages are not identical: they are likely to become more linear with increases in assets. However, the preference order of prospects is not greatly altered by small or even moderate variations in asset position. The certainty equivalent of the prospect (1,000, .50), for example, lies between 300 and 400 for most people, in a wide range of asset positions. Consequently, the representation

of value as a function in one argument generally provides a satisfactory approximation.

Many sensory and perceptual dimensions share the property that the psychological response is a concave function of the magnitude of physical change. For example, it is easier to discriminate between a change of 3° and a change of 6° in room temperature, than it is to discriminate between a change of 13° and a change of 16°. We propose that this principle applies in particular to the evaluation of monetary changes. Thus, the difference in value between a gain of 100 and a gain of 200 appears to be greater than the difference between a gain of 1,100 and a gain of 1,200. Similarly, the difference between a loss of 100 and a loss of 200 appears greater than the difference between a loss of 1,100 and a loss of 1,200, unless the larger loss is intolerable. Thus, we hypothesize that the value function for changes of wealth is normally concave above the reference point ($v''(x) < 0$, for $x > 0$) and often convex below it ($v''(x) > 0$, for $x < 0$). That is, the marginal value of both gains and losses generally decreases with their magnitude. Some support for this hypothesis has been reported by Galanter and Pliner [17], who scaled the perceived magnitude of monetary and non-monetary gains and losses.

The above hypothesis regarding the shape of the value function was based on responses to gains and losses in a riskless context. We propose that the value function which is derived from risky choices shares the same characteristics, as illustrated in the following problems.

PROBLEM 13:

$$(6,000, .25), \quad \text{or} \quad (4,000, .25; 2,000, .25).$$

$$N = 68 \quad [18] \qquad \qquad \qquad [82]^*$$

PROBLEM 13':

$$(-6,000, .25), \quad \text{or} \quad (-4,000, .25; -2,000, .25).$$

$$N = 64 \quad [70]^* \qquad \qquad \qquad [30]$$

Applying equation 1 to the modal preference in these problems yields

$$\pi(.25)v(6,000) < \pi(.25)[v(4,000) + v(2,000)] \quad \text{and}$$

$$\pi(.25)v(-6,000) > \pi(.25)[v(-4,000) + v(-2,000)].$$

Hence, $v(6,000) < v(4,000) + v(2,000)$ and $v(-6,000) > v(-4,000) + v(-2,000)$. These preferences are in accord with the hypothesis that the value function is concave for gains and convex for losses.

Any discussion of the utility function for money must leave room for the effect of special circumstances on preferences. For example, the utility function of an individual who needs \$60,000 to purchase a house may reveal an exceptionally steep rise near the critical value. Similarly, an individual's aversion to losses may increase sharply near the loss that would compel him to sell his house and move to

PROSPECT THEORY

a less desirable neighborhood. Hence, the derived value (utility) function of an individual does not always reflect "pure" attitudes to money, since it could be affected by additional consequences associated with specific amounts. Such perturbations can readily produce convex regions in the value function for gains and concave regions in the value function for losses. The latter case may be more common since large losses often necessitate changes in life style.

A salient characteristic of attitudes to changes in welfare is that losses loom larger than gains. The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount [17]. Indeed, most people find symmetric bets of the form $(x, .50; -x, .50)$ distinctly unattractive. Moreover, the aversiveness of symmetric fair bets generally increases with the size of the stake. That is, if $x > y \geq 0$, then $(y, .50; -y, .50)$ is preferred to $(x, .50; -x, .50)$. According to equation (1), therefore,

$$v(y) + v(-y) > v(x) + v(-x) \quad \text{and} \quad v(-y) - v(-x) > v(x) - v(y).$$

Setting $y = 0$ yields $v(x) < -v(-x)$, and letting y approach x yields $v'(x) < v'(-x)$, provided v' , the derivative of v , exists. Thus, the value function for losses is steeper than the value function for gains.

In summary, we have proposed that the value function is (i) defined on deviations from the reference point; (ii) generally concave for gains and commonly convex for losses; (iii) steeper for losses than for gains. A value function which satisfies these properties is displayed in Figure 3. Note that the proposed S-shaped value function is steepest at the reference point, in marked contrast to the utility function postulated by Markowitz [29] which is relatively shallow in that region.




FIGURE 3.—A hypothetical value function.

Although the present theory can be applied to derive the value function from preferences between prospects, the actual scaling is considerably more complicated than in utility theory, because of the introduction of decision weights. For example, decision weights could produce risk aversion and risk seeking even with a linear value function. Nevertheless, it is of interest that the main properties ascribed to the value function have been observed in a detailed analysis of von Neumann–Morgenstern utility functions for changes of wealth (Fishburn and Kochenberger [14]). The functions had been obtained from thirty decision makers in various fields of business, in five independent studies [5, 18, 19, 21, 40]. Most utility functions for gains were concave, most functions for losses were convex, and only three individuals exhibited risk aversion for both gains and losses. With a single exception, utility functions were considerably steeper for losses than for gains.

The Weighting Function

In prospect theory, the value of each outcome is multiplied by a decision weight. Decision weights are inferred from choices between prospects much as subjective probabilities are inferred from preferences in the Ramsey-Savage approach. However, decision weights are not probabilities: they do not obey the probability axioms and they should not be interpreted as measures of degree or belief.

Consider a gamble in which one can win 1,000 or nothing, depending on the toss of a fair coin. For any reasonable person, the probability of winning is .50 in this situation. This can be verified in a variety of ways, e.g., by showing that the subject is indifferent between betting on heads or tails, or by his verbal report that he considers the two events equiprobable. As will be shown below, however, the decision weight $\pi(.50)$ which is derived from choices is likely to be smaller than .50. Decision weights measure the impact of events on the desirability of prospects, and not merely the perceived likelihood of these events. The two scales coincide (i.e., $\pi(p) = p$) if the expectation principle holds, but not otherwise.

The choice problems discussed in the present paper were formulated in terms of explicit numerical probabilities, and our analysis assumes that the respondents adopted the stated values of p . Furthermore, since the events were identified only by their stated probabilities, it is possible in this context to express decision weights as a function of stated probability. In general, however, the decision weight attached to an event could be influenced by other factors, e.g., ambiguity [10, 11].

We turn now to discuss the salient properties of the weighting function π , which relates decision weights to stated probabilities. Naturally, π is an increasing function of p , with $\pi(0) = 0$ and $\pi(1) = 1$. That is, outcomes contingent on an impossible event are ignored, and the scale is normalized so that $\pi(p)$ is the ratio of the weight associated with the probability p to the weight associated with the certain event.

We first discuss some properties of the weighting function for small probabilities. The preferences in Problems 8 and 8' suggest that for small values of p , π

PROSPECT THEORY

is a subadditive function of p , i.e., $\pi(rp) > r\pi(p)$ for $0 < r < 1$. Recall that in Problem 8, (6,000, .001) is preferred to (3,000, .002). Hence

$$\frac{\pi(.001)}{\pi(.002)} > \frac{v(3,000)}{v(6,000)} > \frac{1}{2} \quad \text{by the concavity of } v.$$

The reflected preferences in Problem 8' yield the same conclusion. The pattern of preferences in Problems 7 and 7', however, suggests that subadditivity need not hold for large values of p .

Furthermore, we propose that very low probabilities are generally overweighted, that is, $\pi(p) > p$ for small p . Consider the following choice problems.

PROBLEM 14:

$$\begin{array}{lll} (5,000, .001), & \text{or} & (5). \\ N = 72 & [72]^* & [28] \end{array}$$

PROBLEM 14':

$$\begin{array}{lll} (-5,000, .001), & \text{or} & (-5). \\ N = 72 & [17] & [83]^* \end{array}$$

Note that in Problem 14, people prefer what is in effect a lottery ticket over the expected value of that ticket. In Problem 14', on the other hand, they prefer a small loss, which can be viewed as the payment of an insurance premium, over a small probability of a large loss. Similar observations have been reported by Markowitz [29]. In the present theory, the preference for the lottery in Problem 14 implies $\pi(.001)v(5,000) > v(5)$, hence $\pi(.001) > v(5)/v(5,000) > .001$, assuming the value function for gains is concave. The readiness to pay for insurance in Problem 14' implies the same conclusion, assuming the value function for losses is convex.

It is important to distinguish overweighting, which refers to a property of decision weights, from the overestimation that is commonly found in the assessment of the probability of rare events. Note that the issue of overestimation does not arise in the present context, where the subject is assumed to adopt the stated value of p . In many real-life situations, overestimation and overweighting may both operate to increase the impact of rare events.

Although $\pi(p) > p$ for low probabilities, there is evidence to suggest that, for all $0 < p < 1$, $\pi(p) + \pi(1-p) < 1$. We label this property subcertainty. It is readily seen that the typical preferences in any version of Allias' example (see, e.g., Problems 1 and 2) imply subcertainty for the relevant value of p . Applying

equation (1) to the prevalent preferences in Problems 1 and 2 yields, respectively,

$$\begin{aligned} v(2,400) &> \pi(.66)v(2,400) + \pi(.33)v(2,500), \quad \text{i.e.,} \\ [1 - \pi(.66)]v(2,400) &> \pi(.33)v(2,500) \quad \text{and} \\ \pi(.33)v(2,500) &> \pi(.34)v(2,400); \quad \text{hence,} \\ 1 - \pi(.66) &> \pi(.34), \quad \text{or} \quad \pi(.66) + \pi(.34) < 1. \end{aligned}$$

Applying the same analysis to Allais' original example yields $\pi(.89) + \pi(.11) < 1$, and some data reported by MacCrimmon and Larsson [28] imply subcertainty for additional values of p .

The slope of π in the interval $(0, 1)$ can be viewed as a measure of the sensitivity of preferences to changes in probability. Subcertainty entails that π is regressive with respect to p , i.e., that preferences are generally less sensitive to variations of probability than the expectation principle would dictate. Thus, subcertainty captures an essential element of people's attitudes to uncertain events, namely that the sum of the weights associated with complementary events is typically less than the weight associated with the certain event.

Recall that the violations of the substitution axiom discussed earlier in this paper conform to the following rule: If (x, p) is equivalent to (y, pq) then (x, pr) is not preferred to (y, pqr) , $0 < p, q, r \leq 1$. By equation (1),

$$\begin{aligned} \pi(p)v(x) = \pi(pq)v(y) \quad \text{implies} \quad \pi(pr)v(x) \leq \pi(pqr)v(y); \quad \text{hence,} \\ \frac{\pi(pq)}{\pi(p)} \leq \frac{\pi(pqr)}{\pi(pr)}. \end{aligned}$$

Thus, for a fixed ratio of probabilities, the ratio of the corresponding decision weights is closer to unity when the probabilities are low than when they are high. This property of π , called subproportionality, imposes considerable constraints on the shape of π : it holds if and only if $\log \pi$ is a convex function of $\log p$.

It is of interest to note that subproportionality together with the overweighting of small probabilities imply that π is subadditive over that range. Formally, it can be shown that if $\pi(p) > p$ and subproportionality holds, then $\pi(rp) > r\pi(p)$, $0 < r < 1$, provided π is monotone and continuous over $(0, 1)$.

Figure 4 presents a hypothetical weighting function which satisfies overweighting and subadditivity for small values of p , as well as subcertainty and subproportionality. These properties entail that π is relatively shallow in the open interval and changes abruptly near the end-points where $\pi(0) = 0$ and $\pi(1) = 1$. The sharp drops or apparent discontinuities of π at the endpoints are consistent with the notion that there is a limit to how small a decision weight can be attached to an event, if it is given any weight at all. A similar quantum of doubt could impose an upper limit on any decision weight that is less than unity. This quantal effect may reflect the categorical distinction between certainty and uncertainty. On the other hand, the simplification of prospects in the editing phase can lead the individual to discard events of extremely low probability and to treat events of extremely high probability as if they were certain. Because people are limited in

PROSPECT THEORY

their ability to comprehend and evaluate extreme probabilities, highly unlikely events are either ignored or overweighted, and the difference between high probability and certainty is either neglected or exaggerated. Consequently, π is not well-behaved near the end-points.




FIGURE 4.—A hypothetical weighting function.

The following example, due to Zeckhauser, illustrates the hypothesized nonlinearity of π . Suppose you are compelled to play Russian roulette, but are given the opportunity to purchase the removal of one bullet from the loaded gun. Would you pay as much to reduce the number of bullets from four to three as you would to reduce the number of bullets from one to zero? Most people feel that they would be willing to pay much more for a reduction of the probability of death from $1/6$ to zero than for a reduction from $4/6$ to $3/6$. Economic considerations would lead one to pay more in the latter case, where the value of money is presumably reduced by the considerable probability that one will not live to enjoy it.

An obvious objection to the assumption that $\pi(p) \neq p$ involves comparisons between prospects of the form $(x, p; x, q)$ and $(x, p'; x, q')$, where $p + q = p' + q' < 1$. Since any individual will surely be indifferent between the two prospects, it could be argued that this observation entails $\pi(p) + \pi(q) = \pi(p') + \pi(q')$, which in turn implies that π is the identity function. This argument is invalid in the present theory, which assumes that the probabilities of identical outcomes are combined in the editing of prospects. A more serious objection to the nonlinearity of π involves potential violations of dominance. Suppose $x > y > 0$, $p > p'$, and $p + q = p' + q' < 1$; hence, $(x, p; y, q)$ dominates $(x, p'; y, q')$. If preference obeys

dominance, then

$$\pi(p)v(x) + \pi(q)v(y) > \pi(p')v(x) + \pi(q')v(y),$$

or

$$\frac{\pi(p) - \pi(p')}{\pi(q') - \pi(q)} > \frac{v(y)}{v(x)}.$$

Hence, as y approaches x , $\pi(p) - \pi(p')$ approaches $\pi(q') - \pi(q)$. Since $p - p' = q' - q$, π must be essentially linear, or else dominance must be violated.

Direct violations of dominance are prevented, in the present theory, by the assumption that dominated alternatives are detected and eliminated prior to the evaluation of prospects. However, the theory permits indirect violations of dominance, e.g., triples of prospects so that A is preferred to B , B is preferred to C , and C dominates A . For an example, see Raiffa [34, p. 75].

Finally, it should be noted that the present treatment concerns the simplest decision task in which a person chooses between two available prospects. We have not treated in detail the more complicated production task (e.g., bidding) where the decision maker generates an alternative that is equal in value to a given prospect. The asymmetry between the two options in this situation could introduce systematic biases. Indeed, Lichtenstein and Slovic [27] have constructed pairs of prospects A and B , such that people generally prefer A over B , but bid more for B than for A . This phenomenon has been confirmed in several studies, with both hypothetical and real gambles, e.g., Grether and Plott [20]. Thus, it cannot be generally assumed that the preference order of prospects can be recovered by a bidding procedure.

Because prospect theory has been proposed as a model of choice, the inconsistency of bids and choices implies that the measurement of values and decision weights should be based on choices between specified prospects rather than on bids or other production tasks. This restriction makes the assessment of v and π more difficult because production tasks are more convenient for scaling than pair comparisons.

4. DISCUSSION

In the final section we show how prospect theory accounts for observed attitudes toward risk, discuss alternative representations of choice problems induced by shifts of reference point, and sketch several extensions of the present treatment.

Risk Attitudes

The dominant pattern of preferences observed in Allais' example (Problems 1 and 2) follows from the present theory iff

$$\frac{\pi(.33)}{\pi(.34)} > \frac{v(2,400)}{v(2,500)} > \frac{\pi(.33)}{1 - \pi(.66)}.$$

PROSPECT THEORY

Hence, the violation of the independence axiom is attributed in this case to subcertainty, and more specifically to the inequality $\pi(.34) < 1 - \pi(.66)$. This analysis shows that an Allais-type violation will occur whenever the v -ratio of the two non-zero outcomes is bounded by the corresponding π -ratios.

Problems 3 through 8 share the same structure, hence it suffices to consider one pair, say Problems 7 and 8. The observed choices in these problems are implied by the theory iff

$$\frac{\pi(.001)}{\pi(.002)} > \frac{v(3,000)}{v(6,000)} > \frac{\pi(.45)}{\pi(.90)}.$$

The violation of the substitution axiom is attributed in this case to the subproportionality of π . Expected utility theory is violated in the above manner, therefore, whenever the v -ratio of the two outcomes is bounded by the respective π -ratios. The same analysis applies to other violations of the substitution axiom, both in the positive and in the negative domain.

We next prove that the preference for regular insurance over probabilistic insurance, observed in Problem 9, follows from prospect theory—provided the probability of loss is overweighted. That is, if $(-x, p)$ is indifferent to $(-y)$, then $(-y)$ is preferred to $(-x, p/2; -y, p/2; -y/2, 1-p)$. For simplicity, we define for $x \geq 0$, $f(x) = -v(-x)$. Since the value function for losses is convex, f is a concave function of x . Applying prospect theory, with the natural extension of equation 2, we wish to show that

$$\begin{aligned} \pi(p)f(x) = f(y) \quad &\text{implies} \\ f(y) &\leq f(y/2) + \pi(p/2)[f(y) - f(y/2)] + \pi(p/2)[f(x) - f(y/2)] \\ &= \pi(p/2)f(x) + \pi(p/2)f(y) + [1 - 2\pi(p/2)]f(y/2). \end{aligned}$$

Substituting for $f(x)$ and using the concavity of f , it suffices to show that

$$f(y) \leq \frac{\pi(p/2)}{\pi(p)} f(y) + \pi(p/2)f(y) + f(y)/2 - \pi(p/2)f(y)$$

or

$$\pi(p)/2 \leq \pi(p/2), \quad \text{which follows from the subadditivity of } \pi.$$

According to the present theory, attitudes toward risk are determined jointly by v and π , and not solely by the utility function. It is therefore instructive to examine the conditions under which risk aversion or risk seeking are expected to occur. Consider the choice between the gamble (x, p) and its expected value (px) . If $x > 0$, risk seeking is implied whenever $\pi(p) > v(px)/v(x)$, which is greater than p if the value function for gains is concave. Hence, overweighting ($\pi(p) > p$) is necessary but not sufficient for risk seeking in the domain of gains. Precisely the same condition is necessary but not sufficient for risk aversion when $x < 0$. This analysis restricts risk seeking in the domain of gains and risk aversion in the domain of losses to small probabilities, where overweighting is expected to hold.

Indeed these are the typical conditions under which lottery tickets and insurance policies are sold. In prospect theory, the overweighting of small probabilities favors both gambling and insurance, while the S-shaped value function tends to inhibit both behaviors.

Although prospect theory predicts both insurance and gambling for small probabilities, we feel that the present analysis falls far short of a fully adequate account of these complex phenomena. Indeed, there is evidence from both experimental studies [37], survey research [26], and observations of economic behavior, e.g., service and medical insurance, that the purchase of insurance often extends to the medium range of probabilities, and that small probabilities of disaster are sometimes entirely ignored. Furthermore, the evidence suggests that minor changes in the formulation of the decision problem can have marked effects on the attractiveness of insurance [37]. A comprehensive theory of insurance behavior should consider, in addition to pure attitudes toward uncertainty and money, such factors as the value of security, social norms of prudence, the aversiveness of a large number of small payments spread over time, information and misinformation regarding probabilities and outcomes, and many others. Some effects of these variables could be described within the present framework, e.g., as changes of reference point, transformations of the value function, or manipulations of probabilities or decision weights. Other effects may require the introduction of variables or concepts which have not been considered in this treatment.

Shifts of Reference

So far in this paper, gains and losses were defined by the amounts of money that are obtained or paid when a prospect is played, and the reference point was taken to be the status quo, or one's current assets. Although this is probably true for most choice problems, there are situations in which gains and losses are coded relative to an expectation or aspiration level that differs from the status quo. For example, an unexpected tax withdrawal from a monthly pay check is experienced as a loss, not as a reduced gain. Similarly, an entrepreneur who is weathering a slump with greater success than his competitors may interpret a small loss as a gain, relative to the larger loss he had reason to expect.

The reference point in the preceding examples corresponded to an asset position that one had expected to attain. A discrepancy between the reference point and the current asset position may also arise because of recent changes in wealth to which one has not yet adapted [29]. Imagine a person who is involved in a business venture, has already lost 2,000 and is now facing a choice between a sure gain of 1,000 and an even chance to win 2,000 or nothing. If he has not yet adapted to his losses, he is likely to code the problem as a choice between $(-2,000, .50)$ and $(-1,000)$ rather than as a choice between $(2,000, .50)$ and $(1,000)$. As we have seen, the former representation induces more adventurous choices than the latter.

A change of reference point alters the preference order for prospects. In particular, the present theory implies that a negative translation of a choice

PROSPECT THEORY

problem, such as arises from incomplete adaptation to recent losses, increases risk seeking in some situations. Specifically, if a risky prospect $(x, p; -y, 1-p)$ is just acceptable, then $(x-z, p; -y-z, 1-p)$ is preferred over $(-z)$ for $x, y, z > 0$, with $x > z$.

To prove this proposition, note that

$$V(x, p; y, 1-p) = 0 \quad \text{iff} \quad \pi(p)v(x) = -\pi(1-p)v(-y).$$

Furthermore,

$$\begin{aligned} & V(x-z, p; -y-z, 1-p) \\ &= \pi(p)v(x-z) + \pi(1-p)v(-y-z) \\ &> \pi(p)v(x) - \pi(p)v(z) + \pi(1-p)v(-y) \\ &\quad + \pi(1-p)v(-z) \quad \text{by the properties of } v, \\ &= -\pi(1-p)v(-y) - \pi(p)v(z) + \pi(1-p)v(-y) \\ &\quad + \pi(1-p)v(-z) \quad \text{by substitution,} \\ &= -\pi(p)v(z) + \pi(1-p)v(-z) \\ &> v(-z)[\pi(p) + \pi(1-p)] \quad \text{since } v(-z) < -v(z), \\ &> v(-z) \quad \text{by subcertainty.} \end{aligned}$$

This analysis suggests that a person who has not made peace with his losses is likely to accept gambles that would be unacceptable to him otherwise. The well known observation [31] that the tendency to bet on long shots increases in the course of the betting day provides some support for the hypothesis that a failure to adapt to losses or to attain an expected gain induces risk seeking. For another example, consider an individual who expects to purchase insurance, perhaps because he has owned it in the past or because his friends do. This individual may code the decision to pay a premium y to protect against a loss x as a choice between $(-x+y, p; y, 1-p)$ and (0) rather than as a choice between $(-x, p)$ and $(-y)$. The preceding argument entails that insurance is likely to be more attractive in the former representation than in the latter.

Another important case of a shift of reference point arises when a person formulates his decision problem in terms of final assets, as advocated in decision analysis, rather than in terms of gains and losses, as people usually do. In this case, the reference point is set to zero on the scale of wealth and the value function is likely to be concave everywhere [39]. According to the present analysis, this formulation essentially eliminates risk seeking, except for gambling with low probabilities. The explicit formulation of decision problems in terms of final assets is perhaps the most effective procedure for eliminating risk seeking in the domain of losses.

Many economic decisions involve transactions in which one pays money in exchange for a desirable prospect. Current decision theories analyze such problems as comparisons between the status quo and an alternative state which includes the acquired prospect minus its cost. For example, the decision whether to pay 10 for the gamble (1,000, .01) is treated as a choice between (990, .01; -10, .99) and (0). In this analysis, readiness to purchase the positive prospect is equated to willingness to accept the corresponding mixed prospect.

The prevalent failure to integrate riskless and risky prospects, dramatized in the isolation effect, suggests that people are unlikely to perform the operation of subtracting the cost from the outcomes in deciding whether to buy a gamble. Instead, we suggest that people usually evaluate the gamble and its cost separately, and decide to purchase the gamble if the combined value is positive. Thus, the gamble (1,000, .01) will be purchased for a price of 10 if $\pi(.01)v(1,000) + v(-10) > 0$.

If this hypothesis is correct, the decision to pay 10 for (1,000, .01), for example, is no longer equivalent to the decision to accept the gamble (990, .01; -10, .99). Furthermore, prospect theory implies that if one is indifferent between $(x(1-p), p; -px, 1-p)$ and (0) then one will not pay px to purchase the prospect (x, p) . Thus, people are expected to exhibit more risk seeking in deciding whether to accept a fair gamble than in deciding whether to purchase a gamble for a fair price. The location of the reference point, and the manner in which choice problems are coded and edited emerge as critical factors in the analysis of decisions.

Extensions

In order to encompass a wider range of decision problems, prospect theory should be extended in several directions. Some generalizations are immediate; others require further development. The extension of equations (1) and (2) to prospects with any number of outcomes is straightforward. When the number of outcomes is large, however, additional editing operations may be invoked to simplify evaluation. The manner in which complex options, e.g., compound prospects, are reduced to simpler ones is yet to be investigated.

Although the present paper has been concerned mainly with monetary outcomes, the theory is readily applicable to choices involving other attributes, e.g., quality of life or the number of lives that could be lost or saved as a consequence of a policy decision. The main properties of the proposed value function for money should apply to other attributes as well. In particular, we expect outcomes to be coded as gains or losses relative to a neutral reference point, and losses to loom larger than gains.

The theory can also be extended to the typical situation of choice, where the probabilities of outcomes are not explicitly given. In such situations, decision weights must be attached to particular events rather than to stated probabilities, but they are expected to exhibit the essential properties that were ascribed to the weighting function. For example, if A and B are complementary events and neither is certain, $\pi(A) + \pi(B)$ should be less than unity—a natural analogue to subcertainty.

PROSPECT THEORY

The decision weight associated with an event will depend primarily on the perceived likelihood of that event, which could be subject to major biases [45]. In addition, decision weights may be affected by other considerations, such as ambiguity or vagueness. Indeed, the work of Ellsberg [10] and Fellner [12] implies that vagueness reduces decision weights. Consequently, subcertainty should be more pronounced for vague than for clear probabilities.

The present analysis of preference between risky options has developed two themes. The first theme concerns editing operations that determine how prospects are perceived. The second theme involves the judgmental principles that govern the evaluation of gains and losses and the weighting of uncertain outcomes. Although both themes should be developed further, they appear to provide a useful framework for the descriptive analysis of choice under risk.

*The University of British Columbia
and
Stanford University*

Manuscript received November, 1977; final revision received March, 1978.

APPENDIX²

In this appendix we sketch an axiomatic analysis of prospect theory. Since a complete self-contained treatment is long and tedious, we merely outline the essential steps and exhibit the key ordinal properties needed to establish the bilinear representation of equation (1). Similar methods could be extended to axiomatize equation (2).

Consider the set of all regular prospects of the form $(x, p; y, q)$ with $p + q < 1$. The extension to regular prospects with $p + q = 1$ is straightforward. Let \geq denote the relation of preference between prospects that is assumed to be connected, symmetric and transitive, and let \approx denote the associated relation of indifference. Naturally, $(x, p; y, q) \approx (y, q; x, p)$. We also assume, as is implicit in our notation, that $(x, p; 0, q) \approx (x, p; 0, r)$, and $(x, p; y, 0) \approx (x, p; z, 0)$. That is, the null outcome and the impossible event have the property of a multiplicative zero.

Note that the desired representation (equation (1)) is additive in the probability-outcome pairs. Hence, the theory of additive conjoint measurement can be applied to obtain a scale V which preserves the preference order, and interval scales f and g in two arguments such that

$$V(x, p; y, q) = f(x, p) + g(y, q).$$

The key axioms used to derive this representation are:

Independence: $(x, p; y, q) \geq (x, p; y'q')$ iff $(x', p'; y, q) \geq (x', p'; y', q')$.

Cancellation: If $(x, p; y'q') \geq (x', p'; y, q)$ and $(x', p'; y'', q'') \geq (x'', p''; y', q')$, then $(x, p; y'', q'') \geq (x'', p''; y, q)$.

Solvability: If $(x, p; y, q) \geq (z, r) \geq (x, p; y'q')$ for some outcome z and probability r , then there exist y'', q'' such that

$$(x, p; y''q'') = (z, r).$$

It has been shown that these conditions are sufficient to construct the desired additive representation, provided the preference order is Archimedean [8, 25]. Furthermore, since $(x, p; y, q) \approx (y, q; x, p)$, $f(x, p) + g(y, q) = f(y, q) + g(x, p)$, and letting $q = 0$ yields $f = g$.

Next, consider the set of all prospects of the form (x, p) with a single non-zero outcome. In this case, the bilinear model reduces to $V(x, p) = \pi(p)v(x)$. This is the multiplicative model, investigated in [35] and [25]. To construct the multiplicative representation we assume that the ordering of the probability-outcome pairs satisfies independence, cancellation, solvability, and the Archimedean axiom. In addition, we assume sign dependence [25] to ensure the proper multiplication of signs. It should be noted that the solvability axiom used in [35] and [25] must be weakened because the probability factor permits only bounded solvability.

² We are indebted to David H. Krantz for his help in the formulation of this section.

Combining the additive and the multiplicative representations yields

$$V(x, p; y, q) = f[\pi(p)v(x)] + f[\pi(q)v(y)].$$

Finally, we impose a new distributivity axiom:

$$(x, p; y, p) \approx (z, p) \quad \text{iff} \quad (x, q; y, q) \approx (z, q).$$

Applying this axiom to the above representation, we obtain

$$f[\pi(p)v(x)] + f[\pi(p)v(y)] = f[\pi(p)v(z)]$$

implies

$$f[\pi(q)v(x)] + f[\pi(q)v(y)] = f[\pi(q)v(z)].$$

Assuming, with no loss of generality, that $\pi(q) < \pi(p)$, and letting $\alpha = \pi(p)v(x)$, $\beta = \pi(p)v(y)$, $\gamma = \pi(p)v(z)$, and $\theta = \pi(q)/\pi(p)$, yields $f(\alpha) + f(\beta) = f(\gamma)$ implies $f(\theta\alpha) + f(\theta\beta) = f(\theta\gamma)$ for all $0 < \theta < 1$.

Because f is strictly monotonic we can set $\gamma = f^{-1}[f(\alpha) + f(\beta)]$. Hence, $\theta\gamma = \theta f^{-1}[f(\alpha) + f(\beta)] = f^{-1}[f(\theta\alpha) + f(\theta\beta)]$.

The solution to this functional equation is $f(\alpha) = k\alpha^c$ [1]. Hence, $V(x, p; y, q) = k[\pi(p)v(x)]^c + k[\pi(q)v(y)]^c$, for some $k, c > 0$. The desired bilinear form is obtained by redefining the scales π, v , and V so as to absorb the constants k and c .

REFERENCES

- [1] ACZÉL, J.: *Lectures on Functional Equations and Their Applications*. New York: Academic Press, 1966.
- [2] ALLAIS, M.: "Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, 21 (1953), 503-546.
- [3] ANDERSON, N. H., AND J. C. SHANTEAU: "Information Integration in Risky Decision Making," *Journal of Experimental Psychology*, 84 (1970), 441-451.
- [4] ARROW, K. J.: *Essays in the Theory of Risk-Bearing*. Chicago: Markham, 1971.
- [5] BARNES, J. D., AND J. E. REINMUTH: "Comparing Imputed and Actual Utility Functions in a Competitive Bidding Setting," *Decision Sciences*, 7 (1976), 801-812.
- [6] COOMBS, C. H.: "Portfolio Theory and the Measurement of Risk," in *Human Judgment and Decision Processes*, ed. by M. F. Kaplan and S. Schwartz. New York: Academic Press, 1975, pp. 63-85.
- [7] DAVIDSON, D., P. SUPPES, AND S. SIEGEL: *Decision-making: An Experimental Approach*. Stanford: Stanford University Press, 1957.
- [8] DEBREU, G.: "Topological Methods in Cardinal Utility Theory," *Mathematical Methods in the Social Sciences*, ed. by K. J. Arrow, S. Karlin, and P. Suppes. Stanford: Stanford University Press, 1960, pp. 16-26.
- [9] EDWARDS, W.: "Subjective Probabilities Inferred from Decisions," *Psychological Review*, 69 (1962), 109-135.
- [10] ELLSBERG, D.: "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics*, 75 (1961), 643-669.
- [11] FELLNER, W.: "Distortion of Subjective Probabilities as a Reaction to Uncertainty," *Quarterly Journal of Economics*, 75 (1961), 670-690.
- [12] ———: *Probability and Profit—A Study of Economic Behavior Along Bayesian Lines*. Homewood, Illinois: Richard D. Irwin, 1965.
- [13] FISHBURN, P. C.: "Mean-Risk Analysis with Risk Associated with Below-Target Returns," *American Economic Review*, 67 (1977), 116-126.
- [14] FISHBURN, P. C., AND G. A. KOCHENBERGER: "Two-Piece von Neumann-Morgenstern Utility Functions," forthcoming.
- [15] FRIEDMAN, M., AND L. J. SAVAGE: "The Utility Analysis of Choices Involving Risks," *Journal of Political Economy*, 56 (1948), 279-304.
- [16] FUCHS, V. R.: "From Bismark to Woodcock: The 'Irrational' Pursuit of National Health Insurance," *Journal of Law and Economics*, 19 (1976), 347-359.
- [17] GALANTER, E., AND P. PLINER: "Cross-Modality Matching of Money Against Other Continua," in *Sensation and Measurement*, ed. by H. R. Moskowitz et al. Dordrecht, Holland: Reidel, 1974, pp. 65-76.

PROSPECT THEORY

- [18] GRAYSON, C. J.: *Decisions under Uncertainty: Drilling Decisions by Oil and Gas Operators*. Cambridge, Massachusetts: Graduate School of Business, Harvard University, 1960.
- [19] GREEN, P. E.: "Risk Attitudes and Chemical Investment Decisions," *Chemical Engineering Progress*, 59 (1963), 35-40.
- [20] GREETHER, D. M., AND C. R. PLOTT: "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review*, forthcoming.
- [21] HALTER, A. N., AND G. W. DEAN: *Decisions under Uncertainty*. Cincinnati: South Western Publishing Co., 1971.
- [22] HANSSON, B.: "The Appropriateness of the Expected Utility Model," *Erkenntnis*, 9 (1975), 175-194.
- [23] HELSON, H.: *Adaptation-Level Theory*. New York: Harper, 1964.
- [24] KEENEY, R. L., AND H. RAIFFA: *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley, 1976.
- [25] KRANTZ, D. H., D. R. LUCE, P. SUPPES, AND A. TVERSKY: *Foundations of Measurement*. New York: Academic Press, 1971.
- [26] KUNREUTHER, H., R. GINSBERG, L. MILLER, P. SAGI, P. SLOVIC, B. BORKAN, AND N. KATZ: *Disaster Insurance Protection: Public Policy Lessons*. New York: Wiley, 1978.
- [27] LICHTENSTEIN, S., AND P. SLOVIC: "Reversal of Preference Between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology*, 89 (1971), 46-55.
- [28] MACCRIMMON, K. R., AND S. LARSSON: "Utility Theory: Axioms versus Paradoxes," in *Expected Utility Hypothesis and the Allais Paradox*, ed. by M. Allais and O. Hagen, forthcoming in *Theory and Decision*.
- [29] MARKOWITZ, H.: "The Utility of Wealth," *Journal of Political Economy*, 60 (1952), 151-158.
- [30] ———: *Portfolio Selection*. New York: Wiley, 1959.
- [31] MCGLOTHLIN, W. H.: "Stability of Choices among Uncertain Alternatives," *American Journal of Psychology*, 69 (1956), 604-615.
- [32] MOSTELLER, F., AND P. NOGEE: "An Experimental Measurement of Utility," *Journal of Political Economy*, 59 (1951), 371-404.
- [33] PRATT, J. W.: "Risk Aversion in the Small and in the Large," *Econometrica*, 32 (1964), 122-136.
- [34] RAIFFA H.: *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Reading, Massachusetts: Addison-Wesley, 1968.
- [35] ROSKIES, R.: "A Measurement Axiomatization for an Essentially Multiplicative Representation of Two Factors," *Journal of Mathematical Psychology*, 2 (1965), 266-276.
- [36] SAVAGE, L. J.: *The Foundations of Statistics*. New York: Wiley, 1954.
- [37] SLOVIC, P., B. FISCHHOFF, S. LICHTENSTEIN, B. CORRIGAN, AND B. COOMBS: "Preference for Insuring Against Probable Small Losses: Insurance Implications," *Journal of Risk and Insurance*, 44 (1977), 237-258.
- [38] SLOVIC, P., AND A. TVERSKY: "Who Accepts Savage's Axiom?," *Behavioral Science*, 19 (1974), 368-373.
- [39] SPETZLER, C. S.: "The Development of Corporate Risk Policy for Capital Investment Decisions," *IEEE Transactions on Systems Science and Cybernetics*, SSC-4 (1968), 279-300.
- [40] SWALM, R. O.: "Utility Theory—Insights into Risk Taking," *Harvard Business Review*, 44 (1966), 123-136.
- [41] TOBIN, J.: "Liquidity Preferences as Behavior Towards Risk," *Review of Economic Studies*, 26 (1958), 65-86.
- [42] TVERSKY, A.: "Additivity, Utility, and Subjective Probability," *Journal of Mathematical Psychology*, 4 (1967), 175-201.
- [43] ———: "Intransitivity of Preferences," *Psychological Review*, 76 (1969), 31-48.
- [44] ———: "Elimination by Aspects: A Theory of Choice," *Psychological Review*, 79 (1972), 281-299.
- [45] TVERSKY, A., AND D. KAHNEMAN: "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185 (1974), 1124-1131.
- [46] VAN DAM, C.: "Another Look at Inconsistency in Financial Decision-Making," presented at the Seminar on Recent Research in Finance and Monetary Economics, Cergy-Pontoise, March, 1975.
- [47] VON NEUMANN, J., AND O. MORGENSTERN, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 1944.
- [48] WILLIAMS, A. C.: "Attitudes toward Speculative Risks as an Indicator of Attitudes toward Pure Risks," *Journal of Risk and Insurance*, 33 (1966), 577-586.

THE FRAMING OF DECISIONS
AND THE PSYCHOLOGY OF CHOICE

Amos Tversky
Stanford University

Daniel Kahneman
University of British Columbia

Reprinted from Science, Vol. 211, pp. 453-458, January 30, 1981.

Copyright © 1981 by the American Association
for the Advancement of Science.

The Framing of Decisions and the Psychology of Choice

Amos Tversky and Daniel Kahneman

Explanations and predictions of people's choices, in everyday life as well as in the social sciences, are often founded on the assumption of human rationality. The definition of rationality has been much debated, but there is general agreement that rational choices should satisfy some elementary requirements of consistency and coherence. In this article

tional choice requires that the preference between options should not reverse with changes of frame. Because of imperfections of human perception and decision, however, changes of perspective often reverse the relative apparent size of objects and the relative desirability of options.

We have obtained systematic rever-

Summary. The psychological principles that govern the perception of decision problems and the evaluation of probabilities and outcomes produce predictable shifts of preference when the same problem is framed in different ways. Reversals of preference are demonstrated in choices regarding monetary outcomes, both hypothetical and real, and in questions pertaining to the loss of human lives. The effects of frames on preferences are compared to the effects of perspectives on perceptual appearance. The dependence of preferences on the formulation of decision problems is a significant concern for the theory of rational choice.

we describe decision problems in which people systematically violate the requirements of consistency and coherence, and we trace these violations to the psychological principles that govern the perception of decision problems and the evaluation of options.

A decision problem is defined by the acts or options among which one must choose, the possible outcomes or consequences of these acts, and the contingencies or conditional probabilities that relate outcomes to acts. We use the term "decision frame" to refer to the decision-maker's conception of the acts, outcomes, and contingencies associated with a particular choice. The frame that a decision-maker adopts is controlled partly by the formulation of the problem and partly by the norms, habits, and personal characteristics of the decision-maker.

It is often possible to frame a given decision problem in more than one way. Alternative frames for a decision problem may be compared to alternative perspectives on a visual scene. Veridical perception requires that the perceived relative height of two neighboring mountains, say, should not reverse with changes of vantage point. Similarly, ra-

sals of preference by variations in the framing of acts, contingencies, or outcomes. These effects have been observed in a variety of problems and in the choices of different groups of respondents. Here we present selected illustrations of preference reversals, with data obtained from students at Stanford University and at the University of British Columbia who answered brief questionnaires in a classroom setting. The total number of respondents for each problem is denoted by N , and the percentage who chose each option is indicated in brackets.

The effect of variations in framing is illustrated in problems 1 and 2.

Problem 1 [$N = 152$]: Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved. [72 percent]

If Program B is adopted, there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved. [28 percent]

Which of the two programs would you favor?

The majority choice in this problem is risk averse: the prospect of certainly saving 200 lives is more attractive than a risky prospect of equal expected value, that is, a one-in-three chance of saving 600 lives.

A second group of respondents was given the cover story of problem 1 with a different formulation of the alternative programs, as follows:

Problem 2 [$N = 155$]:

If Program C is adopted 400 people will die. [22 percent]

If Program D is adopted there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die. [78 percent]

Which of the two programs would you favor?

The majority choice in problem 2 is risk taking: the certain death of 400 people is less acceptable than the two-in-three chance that 600 will die. The preferences in problems 1 and 2 illustrate a common pattern: choices involving gains are often risk averse and choices involving losses are often risk taking. However, it is easy to see that the two problems are effectively identical. The only difference between them is that the outcomes are described in problem 1 by the number of lives saved and in problem 2 by the number of lives lost. The change is accompanied by a pronounced shift from risk aversion to risk taking. We have observed this reversal in several groups of respondents, including university faculty and physicians. Inconsistent responses to problems 1 and 2 arise from the conjunction of a framing effect with contradictory attitudes toward risks involving gains and losses. We turn now to an analysis of these attitudes.

The Evaluation of Prospects

The major theory of decision-making under risk is the expected utility model. This model is based on a set of axioms, for example, transitivity of preferences, which provide criteria for the rationality of choices. The choices of an individual who conforms to the axioms can be described in terms of the utilities of various outcomes for that individual. The utility of a risky prospect is equal to the expected utility of its outcomes, obtained by weighting the utility of each possible outcome by its probability. When faced with a choice, a rational decision-maker will prefer the prospect that offers the highest expected utility ($I, 2$).

Dr. Tversky is a professor of psychology at Stanford University, Stanford, California 94305, and Dr. Kahneman is a professor of psychology at the University of British Columbia, Vancouver, Canada V6T 1W5.

As will be illustrated below, people exhibit patterns of preference which appear incompatible with expected utility theory. We have presented elsewhere (3) a descriptive model, called prospect theory, which modifies expected utility theory so as to accommodate these observations. We distinguish two phases in the choice process: an initial phase in which acts, outcomes, and contingencies are framed, and a subsequent phase of evaluation (4). For simplicity, we restrict the formal treatment of the theory to choices involving stated numerical probabilities and quantitative outcomes, such as money, time, or number of lives.

Consider a prospect that yields outcome x with probability p , outcome y with probability q , and the status quo with probability $1 - p - q$. According to prospect theory, there are values $v(\cdot)$ associated with outcomes, and decision weights $\pi(\cdot)$ associated with probabilities, such that the overall value of the prospect equals $\pi(p)v(x) + \pi(q)v(y)$. A slightly different equation should be applied if all outcomes of a prospect are on the same side of the zero point (5).

In prospect theory, outcomes are expressed as positive or negative deviations (gains or losses) from a neutral reference point and convex below it, as illustrated in Fig. 1. For example, the difference in subjective value between gains of \$10 and \$20 is greater than the subjective difference between gains of \$110 and \$120. The same relation between value differences holds for the corresponding losses. Another property of the value function is that the response to losses is more extreme than the response to gains. The displeasure associated with losing a sum of money is generally greater than the pleasure associated with winning the same amount, as is reflected in people's reluctance to accept fair bets on a toss of a coin. Several studies of decision (3, 6) and judgment (7) have confirmed these properties of the value function (8).

The second major departure of prospect theory from the expected utility model involves the treatment of probabilities. In expected utility theory the utility of an uncertain outcome is weighted by its probability; in prospect theory the value of an uncertain outcome is multiplied by a decision weight $\pi(p)$, which is a monotonic function of p but is not a probability. The weighting function π




Fig. 1. A hypothetical value function.

has the following properties. First, impossible events are discarded, that is, $\pi(0) = 0$, and the scale is normalized so that $\pi(1) = 1$, but the function is not well behaved near the endpoints. Second, for low probabilities $\pi(p) > p$, but $\pi(p) + \pi(1 - p) \leq 1$. Thus low probabilities are overweighted, moderate and high probabilities are underweighted, and the latter effect is more pronounced than the former. Third, $\pi(pq)/\pi(p) < \pi(pqr)/\pi(pr)$ for all $0 < p, q, r \leq 1$. That is, for any fixed probability ratio q , the ratio of decision weights is closer to unity when the probabilities are low than when they are high, for example, $\pi(.1)/\pi(.2) > \pi(.4)/\pi(.8)$. A hypothetical weighting function which satisfies these properties is shown in Fig. 2. The major qualitative properties of decision weights can be extended to cases in which the probabilities of outcomes are subjectively assessed rather than explicitly given. In these situations, however, decision weights may also be affected by other characteristics of an event, such as ambiguity or vagueness (9).

Prospect theory, and the scales illustrated in Figs. 1 and 2, should be viewed as an approximate, incomplete, and simplified description of the evaluation of risky prospects. Although the properties of v and π summarize a common pattern of choice, they are not universal: the preferences of some individuals are not well described by an S-shaped value function and a consistent set of decision weights. The simultaneous measurement of values and decision weights involves serious experimental and statistical difficulties (10).

If π and v were linear throughout, the preference order between options would be independent of the framing of acts, outcomes, or contingencies. Because of the characteristic nonlinearities of π and v , however, different frames can lead to different choices. The following three sections describe reversals of preference caused by variations in the framing of acts, contingencies, and outcomes.

The Framing of Acts

Problem 3 [$N = 150$]: Imagine that you face the following pair of concurrent decisions. First examine both decisions, then indicate the options you prefer.

Decision (i). Choose between:

- A. a sure gain of \$240 [84 percent]
- B. 25% chance to gain \$1000, and 75% chance to gain nothing [16 percent]

Decision (ii). Choose between:

- C. a sure loss of \$750 [13 percent]
- D. 75% chance to lose \$1000, and 25% chance to lose nothing [87 percent]

The majority choice in decision (i) is risk averse: a riskless prospect is preferred to a risky prospect of equal or greater expected value. In contrast, the majority choice in decision (ii) is risk taking: a risky prospect is preferred to a riskless prospect of equal expected value. This pattern of risk aversion in choices involving gains and risk seeking in choices involving losses is attributable to the properties of v and π . Because the value function is S-shaped, the value associated with a gain of \$240 is greater than 24 percent of the value associated with a gain of \$1000, and the (negative) value associated with a loss of \$750 is smaller than 75 percent of the value associated with a loss of \$1000. Thus the shape of the value function contributes to risk aversion in decision (i) and to risk seeking in decision (ii). Moreover, the underweighting of moderate and high probabilities contributes to the relative attractiveness of the sure gain in (i) and to the relative aversiveness of the sure loss in (ii). The same analysis applies to problems 1 and 2.

Because (i) and (ii) were presented together, the respondents had in effect to choose one prospect from the set: A and C, B and C, A and D, B and D. The most common pattern (A and D) was chosen by 73 percent of respondents, while the least popular pattern (B and C) was chosen by only 3 percent of respondents. However, the combination of B and C is definitely superior to the combination A and D, as is readily seen in problem 4.

Problem 4 [$N = 86$]. Choose between:

- A & D. 25% chance to win \$240, and 75% chance to lose \$760. [0 percent]
- B & C. 25% chance to win \$250, and 75% chance to lose \$750. [100 percent]

When the prospects were combined and the dominance of the second option became obvious, all respondents chose the superior option. The popularity of the inferior option in problem 3 implies that this problem was framed as a pair of

separate choices. The respondents apparently failed to entertain the possibility that the conjunction of two seemingly reasonable choices could lead to an untenable result.

The violations of dominance observed in problem 3 do not disappear in the presence of monetary incentives. A different group of respondents who answered a modified version of problem 3, with real payoffs, produced a similar pattern of choices (11). Other authors have also reported that violations of the rules of rational choice, originally observed in hypothetical questions, were not eliminated by payoffs (12).

We suspect that many concurrent decisions in the real world are framed independently, and that the preference order would often be reversed if the decisions were combined. The respondents in problem 3 failed to combine options, although the integration was relatively simple and was encouraged by instructions (13). The complexity of practical problems of concurrent decisions, such as portfolio selection, would prevent people from integrating options without computational aids, even if they were inclined to do so.

The Framing of Contingencies

The following triple of problems illustrates the framing of contingencies. Each problem was presented to a different group of respondents. Each group was told that one participant in ten, pre-selected at random, would actually be playing for money. Chance events were realized, in the respondents' presence, by drawing a single ball from a bag containing a known proportion of balls of the winning color, and the winners were paid immediately.

Problem 5 [$N = 77$]: Which of the following options do you prefer?

- A. a sure win of \$30 [78 percent]
- B. 80% chance to win \$45 [22 percent]

Problem 6 [$N = 85$]: Consider the following two-stage game. In the first stage, there is a 75% chance to end the game without winning anything, and a 25% chance to move into the second stage. If you reach the second stage you have a choice between:

- C. a sure win of \$30 [74 percent]
- D. 80% chance to win \$45 [26 percent]

Your choice must be made before the game starts, i.e., before the outcome of the first stage is known. Please indicate the option you prefer.

Problem 7 [$N = 81$]: Which of the following options do you prefer?

- E. 25% chance to win \$30 [42 percent]
- F. 20% chance to win \$45 [58 percent]




Fig. 2. A hypothetical weighting function.

Let us examine the structure of these problems. First, note that problems 6 and 7 are identical in terms of probabilities and outcomes, because prospect C offers a .25 chance to win \$30 and prospect D offers a probability of $.25 \times .80 = .20$ to win \$45. Consistency therefore requires that the same choice be made in problems 6 and 7. Second, note that problem 6 differs from problem 5 only by the introduction of a preliminary stage. If the second stage of the game is reached, then problem 6 reduces to problem 5; if the game ends at the first stage, the decision does not affect the outcome. Hence there seems to be no reason to make a different choice in problems 5 and 6. By this logical analysis, problem 6 is equivalent to problem 7 on the one hand and problem 5 on the other. The participants, however, responded similarly to problems 5 and 6 but differently to problem 7. This pattern of responses exhibits two phenomena of choice: the certainty effect and the pseudocertainty effect.

The contrast between problems 5 and 7 illustrates a phenomenon discovered by Allais (14), which we have labeled the certainty effect: a reduction of the probability of an outcome by a constant factor has more impact when the outcome was initially certain than when it was merely probable. Prospect theory attributes this effect to the properties of π . It is easy to verify, by applying the equation of prospect theory to problems 5 and 7, that people for whom the value ratio $v(30)/v(45)$ lies between the weight ratios $\pi(.20)/\pi(.25)$ and $\pi(.80)/\pi(1.0)$ will prefer A to B and F to E, contrary to expected utility theory. Prospect theory does not predict a reversal of preference for every individual in problems 5 and 7. It only requires that an individual who has no preference between A and B prefer F to E. For group data, the theory predicts the observed directional shift of preference between the two problems.

The first stage of problem 6 yields the same outcome (no gain) for both acts. Consequently, we propose, people evaluate the options conditionally, as if the second stage had been reached. In this framing, of course, problem 6 reduces to problem 5. More generally, we suggest that a decision problem is evaluated conditionally when (i) there is a state in which all acts yield the same outcome, such as failing to reach the second stage of the game in problem 6, and (ii) the stated probabilities of other outcomes are conditional on the nonoccurrence of this state.

The striking discrepancy between the responses to problems 6 and 7, which are identical in outcomes and probabilities, could be described as a pseudocertainty effect. The prospect yielding \$30 is relatively more attractive in problem 6 than in problem 7, as if it had the advantage of certainty. The sense of certainty associated with option C is illusory, however, since the gain is in fact contingent on reaching the second stage of the game (15).

We have observed the certainty effect in several sets of problems, with outcomes ranging from vacation trips to the loss of human lives. In the negative domain, certainty exaggerates the aversiveness of losses that are certain relative to losses that are merely probable. In a question dealing with the response to an epidemic, for example, most respondents found "a sure loss of 75 lives" more aversive than "80% chance to lose 100 lives" but preferred "10% chance to lose 75 lives" over "8% chance to lose 100 lives," contrary to expected utility theory.

We also obtained the pseudocertainty effect in several studies where the description of the decision problems favored conditional evaluation. Pseudocertainty can be induced either by a sequential formulation, as in problem 6, or by the introduction of causal contingencies. In another version of the epidemic problem, for instance, respondents were told that risk to life existed only in the event (probability .10) that the disease was carried by a particular virus. Two alternative programs were said to yield "a sure loss of 75 lives" or "80% chance to lose 100 lives" if the critical virus was involved, and no loss of life in the event (probability .90) that the disease was carried by another virus. In effect, the respondents were asked to choose between 10 percent chance of losing 75 lives and 8 percent chance of losing 100 lives, but their preferences were the same as when the choice was

between a sure loss of 75 lives and 80 percent chance of losing 100 lives. A conditional framing was evidently adopted in which the contingency of the noncritical virus was eliminated, giving rise to a pseudocertainty effect. The certainty effect reveals attitudes toward risk that are inconsistent with the axioms of rational choice, whereas the pseudocertainty effect violates the more fundamental requirement that preferences should be independent of problem description.

Many significant decisions concern actions that reduce or eliminate the probability of a hazard, at some cost. The shape of π in the range of low probabilities suggests that a protective action which reduces the probability of a harm from 1 percent to zero, say, will be valued more highly than an action that reduces the probability of the same harm from 2 percent to 1 percent. Indeed, probabilistic insurance, which reduces the probability of loss by half, is judged to be worth less than half the price of regular insurance that eliminates the risk altogether (3).

It is often possible to frame protective action in either conditional or unconditional form. For example, an insurance policy that covers fire but not flood could be evaluated either as full protection against the specific risk of fire or as a reduction in the overall probability of property loss. The preceding analysis suggests that insurance should appear more attractive when it is presented as the elimination of risk than when it is described as a reduction of risk. P. Slovic, B. Fischhoff, and S. Lichtenstein, in an unpublished study, found that a hypothetical vaccine which reduces the probability of contracting a disease from .20 to .10 is less attractive if it is described as effective in half the cases than if it is presented as fully effective against one of two (exclusive and equiprobable) virus strains that produce identical symptoms. In accord with the present analysis of pseudocertainty, the respondents valued full protection against an identified virus more than probabilistic protection against the disease.

The preceding discussion highlights the sharp contrast between lay responses to the reduction and the elimination of risk. Because no form of protective action can cover all risks to human welfare, all insurance is essentially probabilistic: it reduces but does not eliminate risk. The probabilistic nature of insurance is commonly masked by formulations that emphasize the completeness of protection against identified harms, but the sense of security that such formulations

provide is an illusion of conditional framing. It appears that insurance is bought as protection against worry, not only against risk, and that worry can be manipulated by the labeling of outcomes and by the framing of contingencies. It is not easy to determine whether people value the elimination of risk too much or the reduction of risk too little. The contrasting attitudes to the two forms of protective action, however, are difficult to justify on normative grounds (16).

The Framing of Outcomes

Outcomes are commonly perceived as positive or negative in relation to a reference outcome that is judged neutral. Variations of the reference point can therefore determine whether a given outcome is evaluated as a gain or as a loss. Because the value function is generally concave for gains, convex for losses, and steeper for losses than for gains, shifts of reference can change the value difference between outcomes and thereby reverse the preference order between options (6). Problems 1 and 2 illustrated a preference reversal induced by a shift of reference that transformed gains into losses.

For another example, consider a person who has spent an afternoon at the race track, has already lost \$140, and is considering a \$10 bet on a 15:1 long shot in the last race. This decision can be framed in two ways, which correspond to two natural reference points. If the status quo is the reference point, the outcomes of the bet are framed as a gain of \$140 and a loss of \$10. On the other hand, it may be more natural to view the present state as a loss of \$140, for the betting day, and accordingly frame the last bet as a chance to return to the reference point or to increase the loss to \$150. Prospect theory implies that the latter frame will produce more risk seeking than the former. Hence, people who do not adjust their reference point as they lose are expected to take bets that they would normally find unacceptable. This analysis is supported by the observation that bets on long shots are most popular on the last race of the day (17).

Because the value function is steeper for losses than for gains, a difference between options will loom larger when it is framed as a disadvantage of one option rather than as an advantage of the other option. An interesting example of such an effect in a riskless context has been noted by Thaler (18). In a debate on a proposal to pass to the consumer some of the costs associated with the process-

ing of credit-card purchases, representatives of the credit-card industry requested that the price difference be labeled a cash discount rather than a credit-card surcharge. The two labels induce different reference points by implicitly designating as normal reference the higher or the lower of the two prices. Because losses loom larger than gains, consumers are less willing to accept a surcharge than to forego a discount. A similar effect has been observed in experimental studies of insurance: the proportion of respondents who preferred a sure loss to a larger probable loss was significantly greater when the former was called an insurance premium (19, 20).

These observations highlight the lability of reference outcomes, as well as their role in decision-making. In the examples discussed so far, the neutral reference point was identified by the labeling of outcomes. A diversity of factors determine the reference outcome in everyday life. The reference outcome is usually a state to which one has adapted; it is sometimes set by social norms and expectations; it sometimes corresponds to a level of aspiration, which may or may not be realistic.

We have dealt so far with elementary outcomes, such as gains or losses in a single attribute. In many situations, however, an action gives rise to a compound outcome, which joins a series of changes in a single attribute, such as a sequence of monetary gains and losses, or a set of concurrent changes in several attributes. To describe the framing and evaluation of compound outcomes, we use the notion of a psychological account, defined as an outcome frame which specifies (i) the set of elementary outcomes that are evaluated jointly and the manner in which they are combined and (ii) a reference outcome that is considered neutral or normal. In the account that is set up for the purchase of a car, for example, the cost of the purchase is not treated as a loss nor is the car viewed as a gift. Rather, the transaction as a whole is evaluated as positive, negative, or neutral, depending on such factors as the performance of the car and the price of similar cars in the market. A closely related treatment has been offered by Thaler (18).

We propose that people generally evaluate acts in terms of a minimal account, which includes only the direct consequences of the act. The minimal account associated with the decision to accept a gamble, for example, includes the money won or lost in that gamble and excludes other assets or the outcome of

previous gambles. People commonly adopt minimal accounts because this mode of framing (i) simplifies evaluation and reduces cognitive strain, (ii) reflects the intuition that consequences should be causally linked to acts, and (iii) matches the properties of hedonic experience, which is more sensitive to desirable and undesirable changes than to steady states.

There are situations, however, in which the outcomes of an act affect the balance in an account that was previously set up by a related act. In these cases, the decision at hand may be evaluated in terms of a more inclusive account, as in the case of the bettor who views the last race in the context of earlier losses. More generally, a sunk-cost effect arises when a decision is referred to an existing account in which the current balance is negative. Because of the nonlinearities of the evaluation process, the minimal account and a more inclusive one often lead to different choices.

Problems 8 and 9 illustrate another class of situations in which an existing account affects a decision:

Problem 8 [$N = 183$]: Imagine that you have decided to see a play where admission is \$10 per ticket. As you enter the theater you discover that you have lost a \$10 bill.

Would you still pay \$10 for a ticket for the play?

Yes [88 percent] No [12 percent]

Problem 9 [$N = 200$]: Imagine that you have decided to see a play and paid the admission price of \$10 per ticket. As you enter the theater you discover that you have lost the ticket. The seat was not marked and the ticket cannot be recovered.

Would you pay \$10 for another ticket?

Yes [46 percent] No [54 percent]

The marked difference between the responses to problems 8 and 9 is an effect of psychological accounting. We propose that the purchase of a new ticket in problem 9 is entered in the account that was set up by the purchase of the original ticket. In terms of this account, the expense required to see the show is \$20, a cost which many of our respondents apparently found excessive. In problem 8, on the other hand, the loss of \$10 is not linked specifically to the ticket purchase and its effect on the decision is accordingly slight.

The following problem, based on examples by Savage (2, p. 103) and Thaler (18), further illustrates the effect of embedding an option in different accounts. Two versions of this problem were presented to different groups of subjects. One group ($N = 93$) was given the values that appear in parentheses, and the

other group ($N = 88$) the values shown in brackets.

Problem 10: Imagine that you are about to purchase a jacket for (\$125) [\$15], and a calculator for (\$15) [\$125]. The calculator salesman informs you that the calculator you wish to buy is on sale for (\$10) [\$120] at the other branch of the store, located 20 minutes drive away. Would you make the trip to the other store?

The response to the two versions of problem 10 were markedly different: 68 percent of the respondents were willing to make an extra trip to save \$5 on a \$15 calculator; only 29 percent were willing to exert the same effort when the price of the calculator was \$125. Evidently the respondents do not frame problem 10 in the minimal account, which involves only a benefit of \$5 and a cost of some inconvenience. Instead, they evaluate the potential saving in a more inclusive account, which includes the purchase of the calculator but not of the jacket. By the curvature of v , a discount of \$5 has a greater impact when the price of the calculator is low than when it is high.

A closely related observation has been reported by Pratt, Wise, and Zeckhauser (21), who found that the variability of the prices at which a given product is sold by different stores is roughly proportional to the mean price of that product. The same pattern was observed for both frequently and infrequently purchased items. Overall, a ratio of 2:1 in the mean price of two products is associated with a ratio of 1.86:1 in the standard deviation of the respective quoted prices. If the effort that consumers exert to save each dollar on a purchase, for instance by a phone call, were independent of price, the dispersion of quoted prices should be about the same for all products. In contrast, the data of Pratt *et al.* (21) are consistent with the hypothesis that consumers hardly exert more effort to save \$15 on a \$150 purchase than to save \$5 on a \$50 purchase (18). Many readers will recognize the temporary devaluation of money which facilitates extra spending and reduces the significance of small discounts in the context of a large expenditure, such as buying a house or a car. This paradoxical variation in the value of money is incompatible with the standard analysis of consumer behavior.

Discussion

In this article we have presented a series of demonstrations in which seemingly inconsequential changes in the formulation of choice problems caused significant shifts of preference. The in-

consistencies were traced to the interaction of two sets of factors: variations in the framing of acts, contingencies, and outcomes, and the characteristic nonlinearities of values and decision weights. The demonstrated effects are large and systematic, although by no means universal. They occur when the outcomes concern the loss of human lives as well as in choices about money; they are not restricted to hypothetical questions and are not eliminated by monetary incentives.

Earlier we compared the dependence of preferences on frames to the dependence of perceptual appearance on perspective. If while traveling in a mountain range you notice that the apparent relative height of mountain peaks varies with your vantage point, you will conclude that some impressions of relative height must be erroneous, even when you have no access to the correct answer. Similarly, one may discover that the relative attractiveness of options varies when the same decision problem is framed in different ways. Such a discovery will normally lead the decision-maker to reconsider the original preferences, even when there is no simple way to resolve the inconsistency. The susceptibility to perspective effects is of special concern in the domain of decision-making because of the absence of objective standards such as the true height of mountains.

The metaphor of changing perspective can be applied to other phenomena of choice, in addition to the framing effects with which we have been concerned here (19). The problem of self-control is naturally construed in these terms. The story of Ulysses' request to be bound to the mast of the ship in anticipation of the irresistible temptation of the Sirens' call is often used as a paradigm case (22). In this example of precommitment, an action taken in the present renders inoperative an anticipated future preference. An unusual feature of the problem of intertemporal conflict is that the agent who views a problem from a particular temporal perspective is also aware of the conflicting views that future perspectives will offer. In most other situations, decision-makers are not normally aware of the potential effects of different decision frames on their preferences.

The perspective metaphor highlights the following aspects of the psychology of choice. Individuals who face a decision problem and have a definite preference (i) might have a different preference in a different framing of the same problem, (ii) are normally unaware of alternative frames and of their potential effects on the relative attractiveness of options,

(iii) would wish their preferences to be independent of frame, but (iv) are often uncertain how to resolve detected inconsistencies (23). In some cases (such as problems 3 and 4 and perhaps problems 8 and 9) the advantage of one frame becomes evident once the competing frames are compared, but in other cases (problems 1 and 2 and problems 6 and 7) it is not obvious which preferences should be abandoned.

These observations do not imply that preference reversals, or other errors of choice or judgment (24), are necessarily irrational. Like other intellectual limitations, discussed by Simon (25) under the heading of "bounded rationality," the practice of acting on the most readily available frame can sometimes be justified by reference to the mental effort required to explore alternative frames and avoid potential inconsistencies. However, we propose that the details of the phenomena described in this article are better explained by prospect theory and by an analysis of framing than by ad hoc appeals to the notion of cost of thinking.

The present work has been concerned primarily with the descriptive question of how decisions are made, but the psychology of choice is also relevant to the normative question of how decisions ought to be made. In order to avoid the difficult problem of justifying values, the modern theory of rational choice has adopted the coherence of specific preferences as the sole criterion of rationality. This approach enjoins the decision-maker to resolve inconsistencies but offers no guidance on how to do so. It implicitly assumes that the decision-maker who carefully answers the question "What do I really want?" will eventually achieve coherent preferences. However, the susceptibility of preferences to variations of framing raises doubt about the feasibility and adequacy of the coherence criterion.

Consistency is only one aspect of the lay notion of rational behavior. As noted by March (26), the common conception of rationality also requires that preferences or utilities for particular outcomes should be predictive of the experiences of satisfaction or displeasure associated with their occurrence. Thus, a man could be judged irrational either because his preferences are contradictory or because his desires and aversions do not reflect his pleasures and pains. The predictive criterion of rationality can be applied to resolve inconsistent preferences and to improve the quality of decisions. A pre-

dictive orientation encourages the decision-maker to focus on future experience and to ask "What will I feel then?" rather than "What do I want now?" The former question, when answered with care, can be the more useful guide in difficult decisions. In particular, predictive considerations may be applied to select the decision frame that best represents the hedonic experience of outcomes.

Further complexities arise in the normative analysis because the framing of an action sometimes affects the actual experience of its outcomes. For example, framing outcomes in terms of overall wealth or welfare rather than in terms of specific gains and losses may attenuate one's emotional response to an occasional loss. Similarly, the experience of a change for the worse may vary if the change is framed as an uncompensated loss or as a cost incurred to achieve some benefit. The framing of acts and outcomes can also reflect the acceptance or rejection of responsibility for particular consequences, and the deliberate manipulation of framing is commonly used as an instrument of self-control (22). When framing influences the experience of consequences, the adoption of a decision frame is an ethically significant act.

References and Notes

1. J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton Univ. Press, Princeton, N.J., 1947); H. Raiffa, *Decision Analysis: Lectures on Choices Under Uncertainty* (Addison-Wesley, Reading, Mass., 1968); P. Fishburn, *Utility Theory for Decision Making* (Wiley, New York, 1970).
2. L. J. Savage, *The Foundations of Statistics* (Wiley, New York, 1954).
3. D. Kahneman and A. Tversky, *Econometrica* 47, 263 (1979).
4. The framing phase includes various editing operations that are applied to simplify prospects, for example by combining events or outcomes or by discarding negligible components (β).
5. If $p + q = 1$ and either $x > y > 0$ or $x < y < 0$, the equation in the text is replaced by $v(y) + \pi(p)[v(x) - v(y)]$, so that decision weights are not applied to sure outcomes.
6. P. Fishburn and G. Kochenberger, *Decision Sci.* 10, 503 (1979); D. J. Laughunn, J. W. Payne, R. Crum, *Manage. Sci.*, in press; J. W. Payne, D. J. Laughunn, R. Crum, *ibid.*, in press; S. A. Eraker and H. C. Sox, *Med. Decision Making*, in press. In the last study several hundred clinic patients made hypothetical choices between drug therapies for severe headaches, hypertension, and chest pain. Most patients were risk averse when the outcomes were described as positive (for example, reduced pain or increased life expectancy) and risk taking when the outcomes were described as negative (increased pain or reduced life expectancy). No significant differences were found between patients who actually suffered from the ailments described and patients who did not.
7. E. Galanter and P. Pliner, in *Sensation and Measurement*, H. R. Moskowitz et al., Eds. (Reidel, Dordrecht, 1974), pp. 65-76.
8. The extension of the proposed value function to multiattribute options, with or without risk, deserves careful analysis. In particular, indifference curves between dimensions of loss may be concave upward, even when the value functions for the separate losses are both convex, because of marked subadditivity between dimensions.
9. D. Ellsberg, *Q. J. Econ.* 75, 643 (1961); W. Fellner, *Probability and Profit—A Study of Economic Behavior Along Bayesian Lines* (Irwin, Homewood, Ill., 1965).
10. The scaling of v and π by pair comparisons requires a large number of observations. The procedure of pricing gambles is more convenient for scaling purposes, but it is subject to a severe anchoring bias: the ordering of gambles by their cash equivalents diverges systematically from the preference order observed in direct comparisons [S. Lichtenstein and P. Slovic, *J. Exp. Psychol.* 89, 46 (1971)].
11. A new group of respondents ($N = 126$) was presented with a modified version of problem 3, in which the outcomes were reduced by a factor of 50. The participants were informed that the gambles would actually be played by tossing a pair of fair coins, that one participant in ten would be selected at random to play the gambles of his or her choice. To ensure a positive return for the entire set, a third decision, yielding only positive outcomes, was added. These payoff conditions did not alter the pattern of preferences observed in the hypothetical problem: 67 percent of respondents chose prospect A and 86 percent chose prospect D. The dominated combination of A and D was chosen by 60 percent of respondents, and only 6 percent favored the dominant combination of B and C.
12. S. Lichtenstein and P. Slovic, *J. Exp. Psychol.* 101, 16 (1973); D. M. Grether and C. R. Plott, *Am. Econ. Rev.* 69, 623 (1979); I. Lieblich and A. Lieblich, *Percept. Mot. Skills* 29, 467 (1969); D. M. Grether, *Social Science Working Paper No. 245* (California Institute of Technology, Pasadena, 1979).
13. Other demonstrations of a reluctance to integrate concurrent options have been reported: P. Slovic and S. Lichtenstein, *J. Exp. Psychol.* 78, 646 (1968); J. W. Payne and M. L. Braustein, *ibid.* 87, 13 (1971).
14. M. Allais, *Econometrica* 21, 503 (1953); K. McCrimmon and S. Larsson, in *Expected Utility Hypotheses and the Allais Paradox*, M. Allais and O. Hagan, Eds. (Reidel, Dordrecht, 1979).
15. Another group of respondents ($N = 205$) was presented with all three problems, in different orders, without monetary payoffs. The joint frequency distribution of choices in problems 5, 6, and 7 was as follows: ACE, 22; ACF, 65; ADE, 4; ADF, 20; BCE, 7; BCF, 18; BDE, 17; BDF, 52. These data confirm in a within-subject design the analysis of conditional evaluation proposed in the text. More than 75 percent of respondents made compatible choices (AC or BD) in problems 5 and 6, and less than half made compatible choices in problems 6 and 7 (CE or DF) or 5 and 7 (AE or BF). The elimination of payoffs in these questions reduced risk aversion but did not substantially alter the effects of certainty and pseudocertainty.
16. For further discussion of rationality in protective action see H. Kunreuther, *Disaster Insurance Protection: Public Policy Lessons* (Wiley, New York, 1978).
17. W. H. McGlothlin, *Am. J. Psychol.* 69, 604 (1956).
18. R. Thaler, *J. Econ. Behav. Organ.* 1, 39 (1980).
19. B. Fischhoff, P. Slovic, S. Lichtenstein, in *Cognitive Processes in Choice and Decision Behavior*, T. Wallsten, Ed. (Erlbaum, Hillsdale, N.J., 1980).
20. J. C. Hershey and P. J. H. Schoemaker, *J. Risk Insur.*, in press.
21. J. Pratt, A. Wise, R. Zeckhauser, *Q. J. Econ.* 93, 189 (1979).
22. R. H. Strotz, *Rev. Econ. Stud.* 23, 165 (1955); G. Ainslie, *Psychol. Bull.* 82, 463 (1975); J. Elster, *Ulysses and the Sirens: Studies in Rationality and Irrationality* (Cambridge Univ. Press, London, 1979); R. Thaler and H. M. Shifrin, *J. Polit. Econ.*, in press.
23. P. Slovic and A. Tversky, *Behav. Sci.* 19, 368 (1974).
24. A. Tversky and D. Kahneman, *Science* 185, 1124 (1974); P. Slovic, B. Fischhoff, S. Lichtenstein, *Annu. Rev. Psychol.* 28, 1 (1977); R. Nisbett and L. Ross, *Human Inference: Strategies and Shortcomings of Social Judgment* (Prentice-Hall, Englewood Cliffs, N.J., 1980); H. Einhorn and R. Hogarth, *Annu. Rev. Psychol.* 32, 53 (1981).
25. H. A. Simon, *Q. J. Econ.* 69, 99 (1955); *Psychol. Rev.* 63, 129 (1956).
26. J. March, *Bell J. Econ.* 9, 587 (1978).
27. This work was supported by the Office of Naval Research under contract N00014-79-C-0077 to Stanford University.

LIST OF THESES IN DECISION ANALYSIS
AT STANFORD UNIVERSITY

LIST OF THESES SUPERVISED IN THE DECISION ANALYSIS PROGRAM

Stanford University
Stanford, California 94305

<u>Name of Student</u>	<u>Title of Dissertation and Date of Degree</u>
Allen S. Ginsberg	Decision Analysis in Clinical Patient Management with An Application to the Pleural Effusion Problem, July 1969
Arnold B. Pollard	A Normative Model for Joint Time/Risk Preference Decision Problems, August 1969
D. Warner North	The Invariance Approach to the Probabilistic Encoding of Information, March 1970
Otto R. Bekman	A Decision Analysis Approach to Two-Person Games, March 1970
Dean W. Boyd	A Methodology for Analyzing Decision Problems Involving Complex Preference Assessments, May 1970
Edward G. Cazalet	Decomposition of Complex Decision Problems with Applications to Electrical Power Systems Planning, May 1970
Bruce S. Plasch	Competitive Economic Theory under Uncertainty with Costly Information, December 1970
Peter A. Morris	Bayesian Expert Resolution, May 1971
Jose E. Guisard Ferraz	A Methodology for Evaluating Corporate Investment Decisions, May 1973

Allen C. Miller	The Value of Sequential Information, January 1974
Thomas R. Rice	The Economics of Decision Analysis, June 1974
Miley W. Merkhofer	Flexibility and Decision Analysis, May 1975
Steven N. Tani	Modeling and Decision Analysis, May 1975
Stephen M. Barrager	Preferences for Dynamic Lotteries: Assessment and Sensitivity, October 1975
John Nafeh	Markov Decision Processes with Policy Constraints, April 1976
Bruce R. Judd	Decision Analysis of Auto Emission Control, May 1976
Thomas W. Keelin	A Protocol and Procedure for Assessing Multiattribute Preference Functions, September 1976
Dennis M. Buede	The Assessment of a Joint Probability Distribution on An Evolving Dynamic Uncertain Process, September 1977
Allen F. Grum	Information Perishing and Replenishment, October 1977
Kenneth R. Oppenheimer	A Proxy Approach to Multi-Attribute Decision Making, October 1977
Khaled H. Nahas	Preference Modeling of Utility Surfaces, December 1977
Daniel Owen	The Concept of Influence and Its Use in Structuring Complex Decision Problems, November 1978

Alimohammad Sharifnia	Optimal Information Acquisition for Randomly Occurring Decisions, April 1979
Patricia A. Owen	Decisions That Affect Outcomes in the Distant Future, May 1979
Burke E. Robinson	Crisis Decision Analysis, June 1979
Larry D. Brandt	Coordination of Decision Analysis within Decentralized Organizations, March 1980
Barry Karlin	A Decision Analytic Approach to Measuring Effectiveness in Decision-Making Organizations, May 1981
Ted Dintersmith	A Separating Choice Hyperplane Algorithm for Evaluating Multi-Attribute Decisions, April 1982
William E. Balson	The Foundations of the Design of Agreement: A Quantitative Theory of Arbitration, May 1982
Jeffrey D. Roughgarden	Rules, Choice and Cooperation, June 1982
Thomas W. Parkinson	Using Complex Profit Models in Decision Analysis, July 1982
Soung Hie Kim	Markovian Methodology for Encoding and Updating the Prior Probability Assessment on Dynamic Processes, January 1983



ABOUT SDG

Strategic Decisions Group (SDG) specializes in helping capital-intensive, risk-intensive, and research-intensive companies analyze their most critical decisions, develop strategies, and create business opportunities. Formed in 1981 and located in Menlo Park, SDG comprises a staff of consultants who have been pioneers in strategy development and decision analysis for more than fifteen years. The SDG combination of broad business experience, specific industry knowledge, and technical expertise in decision analysis has proven outstandingly effective to clients in developing strategies, creating business innovations, allocating resources, managing risk, and selecting business portfolios. SDG consultants work in joint teams with clients to transfer knowledge and to build the strong client commitment necessary for successful implementation. By blending this participative consulting style with an ongoing professional development program, SDG enriches its clients' own corporate capabilities. *Readings on the Principles and Applications of Decision Analysis* is a product of SDG's continued dedication to advancing the principles of good decision making.

ABOUT THE EDITORS

Dr. Ronald A. Howard, Professor of Engineering-Economic Systems in the School of Engineering of Stanford University since 1965, directs teaching and research in the Decision Analysis Program, and is the Director of the Decision and Ethics Center. Dr. Howard defined the profession of decision analysis in 1964 and has since supervised several doctoral theses in decision analysis every year. His experience includes dozens of decision analysis projects in virtually all fields of application, from investment planning to research strategy, and from hurricane seeding to nuclear waste isolation. He has been a consultant to several companies and is presently a founder and director of Strategic Decisions Group. He has published three books, written dozens of technical papers, and provided editorial service to seven technical journals. He has lectured on decision analysis at universities in several foreign countries, including the Soviet Union and the People's Republic of China. His current research interests are life-and-death decision making and the creation of a coercion-free society.

Dr. James E. Matheson, a founder and director of Strategic Decisions Group and a leading figure in developing professional decision analysis, has supervised hundreds of decision analysis applications in such areas as corporate strategy, capital investment, research and development, environmental safety, contract bidding, space exploration, and public investment. He created the SRI International Decision Analysis Group and directed it for fourteen years. He is responsible for innovations in methodology that have made decision analysis a powerful tool throughout many fields and industries. In 1967, Dr. Matheson was appointed to the consulting faculty of Stanford University where he is currently Consulting Professor of Engineering-Economic Systems. He continues to help businesses develop their own decision analysis capabilities through SDG's extensive executive and professional development program. His current research interests are managing risk in business portfolios and integrating research and development decisions into business strategy.

STRATEGIC DECISIONS GROUP
3000 Sand Hill Road • Menlo Park, California 94025