Clinical Decision Analysis

BARUCH FISCHHOFF

Decision Research, A Branch of Perceptronics, Eugene, Oregon

(Received June 1978; accepted March 1979)

An analogy is drawn between decision analysis and the somewhat older profession of psychotherapy. Both offer a variety of techniques designed to help people function in a difficult and uncertain environment; both developed rapidly, sustained by a coherent underlying theory and anecdotal evidence of having helped some clients. Over the past half century, psychotherapy has faced a series of crises concerned with its transformation from an art to a clinical science. These include testing the effectiveness of various forms of therapy, validating elements of treatment programs and of the assumptions underlying therapy, improving the clinical skills of individual practitioners, and considering the broader political, social, ideological and ethical issues raised by psychotherapy. It is hoped that by considering the issues that a related profession has identified, the approaches it has developed to study those issues, and the (partial) conclusions it has reached, we can facilitate the development of decision analysis.

ENORMOUS PROGRESS has been made in developing formal models and computational aids for decision analysis. Much has been learned about how to represent both exotic and routine decision problems and how to compile judgments of value and likelihoods into composite recommendations with accompanying sensitivity analyses.

When it comes time to apply these tools, however, decision analysts must more or less rely on their own wits. There is no codified body of knowledge telling them when to use formal models and when to rely on intuitive judgments, how to approach decision makers and how to coax from them their true problems, which elicitation methods to use and when to trust their results, which parameters should be subjected to sensitivity analysis and what range of alternative values should be used, how to make certain that the assumptions and conclusions of an analysis are understood and heeded, or when decision analysis is likely to improve the understanding of a decision problem and when it is not likely to be cost effective. Such knowledge as does exist regarding these topics is largely anecdotal. It is acquired by trial and error in the field, perhaps aided by apprenticeship with a veteran practitioner [25].

In order for the application of decision analysis to progress as rapidly as its theoretical developments, a systematic basis is needed for these

Operations Research
Vol. 28, No. 1, January-February 1980

0030-364X/80/2801-0028 $01.25

practical skills. We need to know what works where and how well in order to evaluate the work of experienced analysts and to guide the professional training of aspiring ones. Creating such a fund of knowledge will require both empirical and theoretical work, the former to validate our techniques and the assumptions underlying them, the latter to understand how, in principle, these tools relate to particular settings. In essence, the application of decision analysis must be transformed from a clinical art to a clinical science.

How does one structure this complex task? The approach adopted here is to examine the patterns that have emerged from a related profession undergoing a similar transformation. The profession that I have chosen is psychotherapy, the broad collection of theories and procedures designed to help people live their lives better. Like decision analysis, these approaches attempt to help clients understand their world, their desires and their options. They acknowledge that indecision and bad decisions are due at least in part to the complexity and constraints of the world in which their clients live and that a precondition for effective action is explicitly facing difficult issues, like uncertainties and motives. Although the clients of decision analysts, if not the analysts themselves, might back off from the analogy with psychotherapy, the similarities between these two helping professions seem sufficiently strong to hope that psychotherapy research might provide a preliminary organization of the topics decision analysis must face, as well as some germane substantive results.

The transformation of psychotherapy began some 50 years ago with therapists' realization that they could not satisfy either critics or their own critical sense with evidence like "my clients say it helps them" and "the theory makes intuitive sense to me." Nor were they comfortable with sending their students out into the world with a bag of tricks and the admonition to use them wisely. The tale of their attempts to systematize their realm is not one of unremitting progress. Like other scientific endeavors, it has produced its share of dead ends, misconceptualized issues and misleading results. From the present perspective, one can identify a number of issues that have proved to be both critical and fruitful for psychotherapy and might serve the same role for decision analysis. These are: (a) Does it work? (b) How valid are its assumptions and assessment procedures? (c) How can the personal skills of practitioners be improved? (d) What are the bases of resistance to treatment? (e) How is the effectiveness and appropriateness of the approach limited by the social, political, psychological and ideological world in which the client lives?

DOES IT WORK?

The ideal way to evaluate a technique is through a controlled experimental design. Potential clients would be randomly divided between two

groups, one receiving the treatment of interest (decision analysis, psychotherapy), while the other receives no treatment at all or an alternative treatment. In a sophisticated design the alternative would be a placebo treatment, some form of advice that sounds useful but which should (from the perspective of decision analysis) have no systematic impact on decision-making effectiveness.

It is hard to imagine a situation in which such rigorous control would be possible. For example, both proprietary and ethical considerations might prohibit one from assigning clients to "decision analysis" and "other analysis" conditions. In such situations, evaluation might still be possible through the use of quasi-experimental designs in which statistical control substitutes for unobtainable experimental control [6, 28].

The fact that psychotherapy as a profession worries about evaluating itself is certainly to its credit. Failure to develop an evaluation methodology would have suggested that it had something to hide. However, the existence of a methodology and a commitment to its use does not guarantee the steady accumulation of wisdom. In articles reviewing research on the effectiveness of such diverse treatments as marathon encounter groups, sensitivity training, drug abuse reduction, marriage therapy, and behavior modification for juvenile delinquents, one finds a litany of methodological criticisms: lack of a control group, inappropriate control groups, impressionistic statistical analysis, biased data collection, lack of follow-up observations, failure to check observer reliability, unrepresentative samples, inappropriate outcome measures. Poor methodology often tends to produce results prejudiced against therapies whose efficacy is being tested. Sloppy research increases error variance (noise) and makes it hard to detect differences between groups [23]. While decision analysts may have little interest in the results of studies on marathon encounter groups, these *methodological* pitfalls are relevant to anyone interested in evaluating decision analysis.

Psychotherapy researchers have found specific effects that may mask the actual degree of success or failure encountered by a treatment:

(a) The fact that practitioners have been trained in a method and claim to be carrying it out is no guarantee that they are. Assessing the fidelity of implementation is crucial for knowing what is being evaluated.

(b) A well-designed therapeutic program may fail because of the tenacity of the client's problem or unanticipated and uncontrollable changes in the client's world. Thus "good therapy" does not necessarily imply "good outcome."

(c) Many people who apparently benefit from treatment would have improved anyway, due to changes in their life circumstances or outlook. Thus "good outcome" does not necessarily imply "good therapy."

(d) The success of some treatments may be less due to their substantive, theory-based message and manipulations than to the atmosphere they create. These "nonspecific treatment effects" include suggestion,

reduced apprehension, increased self-confidence and heightened attention to the problem.

(e) Unsubstantiated evaluations by practitioners are not to be trusted. Even dispassionate clinicians of high integrity may see treatment effects where statistical analysis shows random fluctuations [16], a record of past success which is exaggerated [13], or proven treatment programs where there are but folklore and bandwagon effects [29].

(f) Results can be biased by looking only for the positive effects a treatment produces and ignoring possible detrimental effects or by looking only for the negative effects [27].

(g) In some cases, defining a "good outcome" is far from trivial, for example, when one must weigh short-term and long-term well-being.

To draw a few of the possible parallels with decision analysis, some products labeled "decision analysis" really are not, and the craft should not be judged by their performance. The vicissitudes of life may "reward" well-analyzed decisions with unfortunate outcomes. Nor can it be presumed that everyone who seems to have done well after decision analysis would have floundered without it; good habits, luck and situational pressures would have "spontaneously" produced some good decisions. Decision analysis may help a decision maker simply because the analyst's deskside manner helps the decision maker focus attention and resources on the problem, and not because of the specific techniques used and their axiomatic justification. Although it is reassuring to hear clients say that our efforts help them, such claims are insufficient evidence. It may be obscurant to invoke unmeasurable benefits, like enhanced peace of mind or self-confidence, when we lack concrete proof of efficacy.

The possibility to treatments not being implemented as their designers intended raises a thorny problem for the evaluator. Obviously, it would be unfair to detract from decision analysis on the basis of crude, ineffectual analyses done by poorly trained individuals or under severe time constraints. Or would it? If the treatment "package" cannot be employed regularly by most practitioners, there is little point to it. If only a selected few can master the craft and the masters do little to monitor those acting in the craft's name, then its usefulness is limited. Its role is further limited if the experience is so unpleasant or expensive that few clients ever get the full treatment. A program with a relatively high drop-out rate but great success with those who complete it will not be highly regarded, particularly when one considers that people who stay in treatment are those most susceptible to persuasive messages of any kind [3]. Since all the resources (computer time, analyst fees, decision-makers' attention) needed for a full, proper decision analysis will seldom be available, a critical evaluation question becomes: Does decision analysis degrade gracefully? A partial analysis is obviously not as good as a full-blown one, but is it better than none at all?

No clear overview of the current state of decision analysis now exists.

Such an overview could be achieved by reviewing a random sample of recent reports of decision analysis and subjecting them to questions like those in Table I.

If the reviewer has opinions about the quality of the analyses or the competence of the analysts, such judgments can be related to these criteria to see what good reports contain and what good analysts do.

Watson and Brown [37] have pioneered an alternative, reflexive evaluation strategy using decision analysis to analyze past decision analyses. Perhaps foretelling the difficulties awaiting such efforts, in two of the

TABLE I

CRITERIA FOR ANALYZING ANALYSES

Are the assumptions of the analysts listed?
Are the assumptions of the clients listed (e.g., those implicit in the way the problem was formulated)?
Are any of these assumptions tested, or is supporting evidence from other sources cited?
Are probabilities used? If so, is any justification given for the particular procedure by which they are elicited?
Are probabilities or utilities measured in more than one way?
Are values elicited from more than one person?
Are sensitivity analyses conducted, for probabilities, for utilities, with more than one factor varying at once?
Are interactions between impacts considered?
Is more than one problem structure used as a cross check?
Are possible alternatives given by the client or created with the client?
Are gaps in scientific knowledge noted?
Is a bottom line figure given, and if so, how is it hedged?
Is the public involved, and if so, at what stage?
Is there consideration of political feasibility or legal constraints?
Is there external criticism of the report, and if so, has the analysis been redone in its light?
Is there indication of when the analysis should be redone to consider possible changes of circumstance?
Is an attempt made to evaluate the analysis or to indicate how interested parties might do so on their own?
How much did the analysis cost?

three case studies chosen by Watson and Brown [36], the greatest benefits of the analyses seemed to come not from the decisions they recommended, but from their contribution to organizational processes (reduction of controversy and improvement of communication), considerations left out of Watson and Brown's formal model for the sake of simplicity.

HOW VALID ARE ITS ASSUMPTIONS AND ASSESSMENT PROCEDURES?

When technical difficulties preclude validating entire treatment programs, one may still be able to assess the validity of the theoretical assumptions upon which the programs are based and the effectiveness of

their component techniques. Such research can also point to what the treatment's strengths are and how they can be improved. In the context of psychotherapy, the most valuable results have emerged from attempts to test previously unquestioned theoretical assumptions, e.g., stable personality traits exist; feedback facilitates learning; psychopathology is related to unconscious libidinal and aggressive wishes; self-awareness is necessary for improvement.

As might be expected, the divide-and-conquer approach to evaluation has appealed to students of decision analysis. Tests of whether people accept the normative assumptions upon which decision analysis is based have shown (with varying degrees of definitiveness) that people often do not wish to accept Savage's independence axiom [32], that they occasionally want their judgments to be intransitive [34], that there do not seem to be consistent individual differences in risk aversion [42], and that verbally expressed preferences are not always consistent with those revealed in people's behavior [30]. Much less is known about the appropriateness of other assumptions: Are probability and utility judgments independent? Can we acceptably resolve inconsistencies in people's preferences due to theoretically irrelevant differences in elicitation procedure? Will people reply honestly to our questions about their values and, if not, can we spot their lies or "strategic responses?" Is it possible for the decision analyst to act as a neutral agent when eliciting judgments [14]?

The development of assessment procedures has long been a growth industry in psychotherapy. With the possibility of measuring every feasible personality and behavioral trait, psychologists produce over 3000 books, chapters and journal articles on assessment per year [15]. Unfortunately, there is no generally accepted characterization of the universe of traits and the relationship between seemingly similar traits (e.g., honesty and straightforwardness). As a result, it is difficult to know what conclusions to make from comparisons between studies.

By contrast, decision analysis is primarily interested in the assessment of two well-defined quantities, probabilities and utilities. Although the quantity of research here is perhaps 1/100 of that in personality assessment, the cumulative progress is probably greater. We know quite a lot about probabilities (e.g., they tend to reflect overconfidence although their validity depends heavily on context [19]). Somewhat less is known about eliciting utilities, although studies [11, 35] that compare a variety of methods using evaluative criteria drawn from the sophisticated methodology of psychometricians show great promise. Almost nothing is known about another topic which could be considered an assessment problem: determining the structure of a decision-maker's problem.

Psychometricians have discovered two threats to the generality of assessment procedures which should concern decision analysts. One is that the people's feelings about a particular object and the numbers they

assign to those feelings can vary greatly with arbitrary features of the elicitation procedure, such as the order in which alternatives are presented, the heterogeneity of the set of alternatives, the contrast established between the first two alternatives, whether the response scale is bounded, and the respondents' preconceptions about how the numbers are supposed to be used [14, 26, 27]. The second threat is that it is not tests but responses which have validities and reliabilities. Thus, the adequacy of an elicitation procedure in one context with one particular set of individuals is not a guarantee of universal applicability.

Once we understand the flaws in our assumptions and procedures, we need an error theory to tell us what their cumulative impact is. As Fischer [11] notes, without an error theory, we cannot know to what extent violations of assumptions and lack of robustness in responses threaten the results of a decision analysis. Important steps toward developing such a theory (or theories) are:

(a) Fischer's [11] work with multidimensional utility models;

(b) von Winterfeldt and Edwards' [40] finding that with continuous decision options (e.g., invest X dollars) some inaccuracy in individual probability and utility assessments will not produce terribly suboptimal decisions;

(c) Lichtenstein et al.'s [19] demonstration of how moderate miscalibration in probability assessment can substantially reduce expected utility with discrete decision options (e.g., operate/do not operate);

(d) von Winterfeldt and Edwards' [41] identification of the ease with which dominated alternatives can be selected through improper problem modeling;

(e) Aschenbrenner and Kasubek's [2] finding that two different, only partially overlapping, sets of attributes produced similar results in a multiattribute utility analysis;

(f) Kastenberg, McKone and Okrent's [17] discovery of the extreme sensitivity of risk assessments to the treatment of outliers; and

(g) Tihansky's [33] finding that errors in different estimates were positively correlated and, therefore, would not tend to cancel one another out.

These are but pieces of an error theory. Particularly useful additions would be guidelines to the way in which uncertainty from varying sources (people not knowing what they want, people being affected by choice of questioning procedure, people being confused by instructions, random error, etc.) is compounded. Until an adequate theory is developed, we will have to be very generous in performing sensitivity analyses for errors arising from judgmental sources.

HOW CAN CLINICAL SKILLS BE IMPROVED?

However powerful their measures and theories may be, clinicians realize that in the last analysis, they, themselves, are their own major

tool. They must instill confidence in clients, choose the appropriate questioning procedures to elicit sensitive information, handle crises, understand what is not being said, avoid imposing their own values and perceptions, and cooperate in creating solutions. To this end, clinical psychologists undergo 3–4 years of supervised practice, psychiatrists spend 1–2 years in internship and psychoanalysts undergo psychoanalysis to be fully aware of how they see and interact with others.

Such training assumes that the finer points of the craft can be learned only in the clinic of a master. To expedite this training, many researchers are attempting to discover just what it is that makes masters. Since these studies consider the interaction between therapists and clients with serious personal problems, one should use caution in drawing inferences regarding the interactions between decision analysts and corporate executives or government officials. One result that seems likely to generalize from the psychotherapeutic context is the extent to which one individual can shape another's responses by such subtle measures as appreciative grunts and nonverbal communication (posture, facial expressions, etc.). One can readily imagine an analyst subtly pressuring a client to change a probability assessment to a value the analyst believes to be more acceptable (analyst seems displeased; client thinks, "well, you're the expert on probabilities. Maybe what I meant was . . .") or an analyst and client "agreeing" that the latter's preferences on different attributes are really independent, making the elicitation procedure considerably less arduous. Slovic and Tversky [32] showed how direct pressure can induce clients to accept axiomatic principles. Further possibilities for influencing judgments emerge when the analyst works with groups. For example, the fact that group discussions tend to polarize opinions [21] suggests that the analyst can exert some control over the group's decision by deciding if and when the group should meet. Plott and Levine [24] demonstrated the extent to which group decisions can be manipulated by varying the order in which issues are considered. These effects must be understood if the analyst is to restrain, control or exploit them.

In general, however, we have little concrete evidence regarding clinical skills in decision analysis and their improvement. One place to start would be a taxonomy of decision situations indicating which techniques to use where. Several such guides have been derived from formal properties of the decision situation (e.g., [9, 18, 22]). Additional efforts might look at more subjective aspects: the public visibility of the issue at hand; how articulated people's values are; how much freedom the analyst and decision maker have to construct alternatives; and whether any evaluation of the analysis is planned. Such a guide should tell us, among other things: When, in order to avoid misplaced precision, should all resources be invested in problem structuring and none in attaching numbers? Can high-priced analysts be replaced by paraprofessionals? When is it advisable to acknowledge the poorly developed nature of people's preferences

and the limits on their information-processing abilities and to sacrifice axiomatic rigor for less demanding procedures [8]? Psychological theories have been likened to box cameras which take pretty good pictures because they require subjects to be at a great distance, in the sun and immobile [43]; is the same true of decision analysis?

WHAT CAUSES RESISTANCE TO TREATMENT?

Resistance to treatment takes many forms, all threatening its success. The client may reject the approach beause it is not expected to work, because its procedures (e.g., talking openly about sensitive matters) are threatening, because it is too expensive, because of objections to its underlying philosophy, or because of reluctance to admit that there is a problem. The client who accepts the approach may resist its recommendations because they require assuming too much responsibility for a situation, because it seems easier to stumble along than to undertake the needed action, or because the analysis mandates acknowledging one's own fallibility, desires and uncertainties.

Even if the client is willing and able to adopt the approach and its directives, treatment may fail when the time comes to implement it in a hostile, unaccepting world. Classic failures of this type have been encountered by the *T*-group (or organizational development) movement, which tries to improve communication in a work setting by involving some workers and managers in intensive group experiences stressing openness and sensitivity. All too often, however, the behavioral changes induced by the pressure of the group situation and the manipulation of the group leader vanish when group members return to their hierarchic work settings. Although one might argue that it is not the client but the client's world which is "sick" and in need of help, the result is still a frustrating failure likely to reinforce old, bad habits.

Variants of all these problems seem possible with decision analysis, particularly when it is first introduced into organizations accustomed to less rigorous methods. In such contexts, its greatest potential advantages may prove to be stumbling blocks. Decision analysis requires explicit statement of problems. This, however, may produce great discomfort. Its computational procedures greatly relieve the decision-maker's mental load; however, for those unfamiliar with its logic, its recommendations may appear to be the output of numerical mumbo-jumbo with no intuitive appeal. Unlike other procedures, its logic is axiomatically grounded; however, for uninitiated superiors, subordinates, and constituents, abandoning the comfortable old maxims (e.g., "This is the way we've always done it") may come quite hard.

Resistance within the organization may come from people who feel that they have not been involved early enough and adequately enough in

the analysis. Like staunch believers in due process by law, they may believe that the decision-making process is more important than its product. Others may resist because they do not like the resultant recommendation. To achieve their ends, they may fight hard and dirty, questioning every fact and assumption in the analysis and casting aspersions on the integrity of its analysts, however well the analysis is done and however much its conclusions are qualified [4, 7]. Analysts who believe in their work may face an uncomfortable choice between orphaning their analyses and adopting an advocacy role for the analysis and, thereby, for the recommended alternative.

Some of these problems are due to the fact that decision analysis, some of the problems to which it is applied, and the very idea of analysis are new. As a result, the social forms needed to incorporate it are either missing or in a state of flux [39]. Westman [38], for example, complains that the legal mandate given regulators entrusted with improving U.S. water quality precludes their adopting the most cost-effective methods. Often projects are held up so long and altered so extensively in legal and administrative proceedings that their accompanying analyses become antiquated. Majone [20] has argued persuasively that alternatives are almost never adopted as proposed; rather, they are subject to continuous negotiation and alteration by the parties concerned.

Acknowledgment of these difficulties might lead to redirection of decision analysis. Brown [5] proposes that analysts treat action options as events and directly assess the uncertainty surrounding the form in which they will be realized. The preferred alternative might turn out to be one with dominated consequences but a better chance of being implemented. Another response would be for analysts to decide that feasibility is both a relative and mutable thing and append to each alternative a discussion of how it is likely to be waylaid en route to implementation and what needs to be done to keep it maximally intact.

In the long run, though, the adaptation should be mutual, with society and its component entities realizing the need to accommodate formal procedures. Toward this end, the educational potential of each analysis should be exploited. Broad participation should be viewed as an opportunity, not a burden. In some ways, it may be more important to build the analytic capacity of a society or organization than to guarantee the adoption of particular, desirable alternatives.

HOW DO POLITICS, IDEOLOGY AND ETHICS IMPINGE UPON ANALYSIS?

Attempts to shape and direct others' lives cannot be value neutral. The practitioner who is "only trying to help" has at the least made the evaluation that there is a situation needing help. The practitioner who is

"only trying to do what is best for the client" cannot avoid at least some subtle hints at what that "best" is. Even client-centered therapists, whose goal is to reflect and clarify their clients' own thoughts, are still promulgating the world view that people are responsible for their own predicaments and can extricate themselves if they understand themselves sufficiently well. Indeed, the very search for lasting solutions to problems implies that the client's universe has more orderliness than may be the case.

The ideological biases of many therapeutic interventions are familiar intellectual topics: the mechanistic image of people projected by behaviorism and its potential for control, the ethnocentrism of psychoanalysis, the narcissism of many contemporary therapies, the general tendency to treat clients as objects rather than colleagues in therapy, and the fatalism induced by approaches that teach people to accept their own life crises as inevitable.

Even when a therapy's philosophical basis is acceptable, it may be resisted because of ethical problems or political bias in the way it is used. Much opposition to behavior modification arose from its use in institutional settings (prisons, asylums) in which free, informed consent for treatment by the patient is impossible. A frequent problem for practitioners is determining who the true client is, the patient or someone else (e.g., a hospital administrator) interested in maintaining order. Other therapies have lost their credibility because therapists have become so dependent upon government and the politically powerful for their livelihood that they have lost the ability to make independent criticism, others because they can be afforded only by the rich, still others because they seem to be applied mainly to coerce the poor.

At first blush, the image of people and society fostered by decision analysis seems to be a highly flattering one. With proper coaching, people are capable of understanding and expressing what they know and what they want. Acknowledging their own information-processing limitations, people will prefer to have their values and beliefs combined mechanically and then will accept the indicated course of action.

There may, however, be problems with this seemingly innocuous perspective. One is that it may create an illusion of analyzability for problems that are insoluble, contributing to the mystique of science and "technical fixes." Because it asks us about everything important, it may lead us to believe that we have and should have beliefs and opinions about everything. We may be forced, for the sake of answering the analyst, to create preferences that are only superficially understood. Forcing people to have (necessarily shallow) opinions about many things may be an excellent way to guarantee that they have articulated views about nothing. Persistent questioning about poorly formulated beliefs may lead to responses designed to make the elicitor happy and to

overreliance on easily measured and justified standards like monetary values.

In the public domain, the very reasonableness of decision analysis is based upon a political-ideological assumption: that society is sufficiently cohesive that its problems can be resolved by reason and without struggle. Although this "get on with business" orientation will be pleasing to many, it will not satisfy all. For those who do not believe that society is in a fine-tuning stage, any technique that fails to mobilize public consciousness and involvement has little to recommend it.

Like therapy, if decision analysis is not biased at its core it can be biased in its application. For example, most applications to societal problems seem to foster the transfer of decision-making power to a technical elite by offering little opportunity for effective citizen participation [31]. Although this trend seems inevitable due to the highly technical nature of the issues studied, in principle it might be countered by a concerted effort by analysts to go beyond the narrow dictates of their analytical mandate. The theoretical problems of aggregating group opinions need not forestall efforts to elicit them. To take another example, most analyses ignore the issue of equitable distribution of good and bad consequences. Although this is not a necessary feature of decision analysis, repeated omission of equity considerations will suggest lack of interest in such issues, or even evasiveness, on the part of analysts and those who hire them.

When analytic resources are limited, the analyst must take cues from someone about how to restrict the alternatives and consequences considered. That someone is likely to be the one who commissioned the study. If commissioners come consistently from one sector of society and consistently prefer (or reject out of hand) particular kinds of solutions or consequences, a persistent bias may be produced. Such bias would also determine what issues are never analyzed and how results are presented. If the commissioners are public officials, there may be a predisposition toward reports that bury uncertainties and delicate assumptions in sophisticated technical machinations and masses of undigested data [12].

Psychotherapy's response to charges of ideological bias has been fairly minimal, with the most dramatic proposals within the profession being to encourage truth-in-packaging (e.g., providing potential clients with a description of the assumptions and procedures of an approach). Its response to charges of improprieties in therapeutic interactions has been more extreme. Clinical psychologists, for example, have organized as a guild with rigorous standards for entry, state and national licensing, censure mechanisms (albeit not often used), external review of research proposals and papers, and a strict code of ethics. That code addresses issues like recognizing and acknowledging the limits to one's competence, protecting clients' confidentiality and policing one's colleagues [1].

Whether a guild structure is needed or appropriate for decision analysis is a moot point. Certainly, all that calls itself decision analysis does not glitter. However, the costs of controlling incompetent analysts might be substantial, draining the efforts of qualified analysts, discrediting the profession by unrepresentative public quibbling and raising prices through restraint of trade. Perhaps more modest steps might be appropriate, if any are needed at all:

(a) Setting up a "public interest decision analysis group" similar to that set up by the largest accounting firms in the United States in order to "give accounting away,"

(b) Insisting that some fixed amount of funds (say, 10%) in all analysis contracts be allocated to independent external review,

(c) Establishing a professional norm of participating in voluntary review networks,

(d) Teaching students to conduct and document enough sensitivity analyses to satisfy a report's most skeptical critics, or

(e) Adopting informal guidelines like those proposed by Fairly [10] for experts called upon to assess probabilities of rare accidents.

Because it functions in the public domain, as well as in the private sector, decision analysis faces ethical dilemmas at least as challenging as those faced by psychotherapy. For example, the American Psychological Association's ethics committtee [1] was unable to agree on how to revise their standards regarding confidentiality (leaving them unchanged from 1964) even without having to consider (as the decision analyst might) the additional problems of what to do with proprietary information or information that could cause public panic if released. Therapists may find themselves forced to treat a delinquent when they should be treating the family. Similarly, analysts may get well into a problem before realizing that the wrong problem has been attacked, or that the wrong information has been provided, or that they are being set up to produce an advocacy rather than an honest analysis. Therapists often face the problem of how to assure informed consent by psychologically incompetent clients, whereas analysts are often asked to pursue their craft on behalf of clients, perhaps a whole society, judged by someone to be technically incompetent.

CONCLUSION

The analyst's job is extremely difficult. Confronting the issues raised above will make it even more difficult. However, the fact that they can be explicitly identified is in some sense a tribute to the clarity and comprehensiveness of decision analysis and its potential for development. As a result, I believe that efforts to implement a research program exploring these problems in the context of decision analysis would be well rewarded. Some of these issues have obvious pecuniary importance for

the long-term prosperity of the field and its practitioners (e.g., proving its effectiveness and buttressing its foundations). Others, like examining ideological and ethical questions, will be intellectually stimulating. Still others, though, will seem like exercises in validating what common sense knows to be true (e.g., that there is more to decision analysis than putting on a good act). However, study of even these issues may have merit, for common sense may be superficial or wrong, and may vary across individuals, as psychotherapy's concerted effort to test and refine common sense has shown. Furthermore, examining the obvious can help convince others that we are right, improve our confidence in (and willingness to act upon) our knowledge, and help us learn why we were right all along.

ACKNOWLEDGMENTS

Preparation of this report was supported by Advanced Research Projects Agency of the Department of Defense and was monitored by Office of Naval Research under Contract N00014-79-C-0029 (ARPA Order No. 3668) to Perceptronics, Inc. I would like to thank Craig Kirkwood, Sarah Lichtenstein, Leon Rappoport, Paul Slovic, Mark Weinrott and anonymous reviewers for comments on earlier drafts.

REFERENCES

1. American Psychological Association, "Revised Ethical Standards for Psychologists," *APA Monitor* **8,** 22-23 (1977).
2. K. M. Aschenbrenner and W. Kasubek, "Convergence of Multiattribute Evaluations when Different Sets of Attributes Are Used," in *Decision Making and Change in Human Affairs,* H. Jungermann and G. deZeeuw (eds.), D. Reidel, Amsterdam, 1977.
3. A. Bandura, *Principles of Behavior Modification,* Holt, Rinehart & Winston, New York, 1969.
4. S. M. Barrager, B. R. Judd and D. W. North, "Decision Analysis of Energy Alternatives: A Comprehensive Framework for Decision Making," Stanford Research Institute, Palo Alto, Calif., 1976.
5. R. V. Brown, "Modeling Subsequent Acts for Decision Analysis," Decisions & Designs, Inc. (McLean, Va.) *DDI Technical Report 75-1,* 1975.
6. D. T. Campbell and J. C. Stanley, "Experimental and Quasi-Experimental Designs for Research in Teaching," in *Handbook of Research on Teaching,* N. L. Gage (ed.), Rand McNally, Chicago, 1966.
7. J. L. Creighton, "The Limitations and Constraints on Effective Citizen Participation," Address to the Interagency Council on Citizen Participation, Washington, D. C., Dec. 8, 1976.
8. W. Edwards, "How to Use Multi-attribute Utility Measurement for Social Decision Making, *IEEE Trans. Systems Man Cybernet.* **7,** 326-340 (1977).
9. S. V. Emelyanov and V. M. Ozernoi, "Decision Making in Multi-objective Problems: A Survey," *Problems of Control and Information Theory,* pp. 51-64 (1975).

10. W. B. Fairley, "Evaluating the "Small" Probability of a Catastrophic Accident from the Marine Transportation of Liquefied Natural Gas, in *Statistics and Public Policy*, W. B. Fairley and F. Mosteller (eds.), Addison-Wesley, Reading, Mass., 1977.
11. G. W. Fischer, "Multidimensional Utility Models for Risky and Riskless Choice," *Organizational Behavior Human Performance* **17,** 127-146 (1976).
12. B. Fischhoff, "Cost-Benefit Analysis and the Art of Motorcycle Maintenance, *Policy Sci.* **8,** 177-202 (1977).
13. B. Fischhoff and R. Beyth, "I Knew It Would Happen": Remembered Probabilities of Once-Future Things, *Organizational Behavior Human Performance* **13,** 1-16 (1975).
14. B. Fischhoff, P. Slovic and S. Lichtenstein "Knowing What You Want: Measuring Labile Values," in *Cognitive Processes in Choice and Decision Behavior*, T. Wallsten (ed.), Erlbaum, Hilsdale, N. J., in press.
15. L. R. Goldberg, "Objective Diagnostic Tests and Measures," *Ann. Rev. Psychol.* **25,** 343-366 (1974).
16. R. R. Jones, M. Weinrott and R. S. Vaught, Effects of Serial Dependency on the Agreement between Visual and Statistical Inference, *J. Appl. Behavior Anal.* **11,** 277-283 (1978).
17. W. E. Kastenberg, T. E. McKone and D. Okrent, "On Risk Assessment in the Absence of Complete Data," *UCLA Report No. UCLA-ENG-7677*, July 1976.
18. R. Keeney and H. Raiffa, "A Critique of Formal Analysis in Public Decision Making, in *Analysis of Public Systems*, A. W. Drake, R. L. Keeney and P. M. Morse (eds.), MIT Press, Cambridge, Mass., 1972.
19. S. Lichtenstein, B. Fischhoff and L. D. Phillips, "Calibration of Probabilities: The State of the Art," in *Decision Making and Change in Human Affairs*, H. Jungermann and G. deZeeuw (eds.), D. Reidel, Amsterdam, 1977.
20. G. Majone, "Choice Among Policy Instruments for Pollution Control, *Policy Anal.* **7,** 589-613 (1976).
21. D. G. Myers and H. Lamm, "The Polarizing Effect of Group Discussions," *Am. Sci.* **63,** 297-303 (1975).
22. D. Pearce, "The Limits of Cost-Benefit Analysis as a Guide to Environmental Policy," *Kyklos* **29,** 97-112 (1976).
23. R. Perloff, E. Perloff and E. Sussna, "Program Evaluation," *Ann. Rev. Psychol.* **27,** 569-594 (1976).
24. C. R. Plott and M. E. Levine, "A Model of Agenda Influence on Committee Decisions," *Am. Econ. Rev.* **68,** 146-160 (1978).
25. M. Polanyi, *Personal Knowledge*, Routledge & Kegan Paul, London, 1962.
26. E. C. Poulton, "The New Psychophysics: Six Models for Magnitude Estimation," *Psychol. Bull.* **69,** 1-19 (1968).
27. E. C. Poulton, "Quantitative Subjective Judgments Are Almost Always Biased, Sometimes Completely Misleading," *Br. J. Psychol.* **68,** 409-425 (1977).
28. H. W. Riecken and R. F. Boruch, *Social Experimentation: A Method for Planning and Evaluating Social Intervention*, Academic Press, New York, 1974.

29. F. SCHECTMAN, "Convention and Contemporary Approaches to Psychotherapy," *Am. Psychol.* **32,** 197–204 (1977).
30. H. SCHUMAN AND M. P. JOHNSON, "Attitudes and Behavior," *Ann. Rev. Sociol.* **40,** 161–207 (1976).
31. W. R. D. SEWELL AND T. O'RIORDAN, "The Culture of Participation in Environmental Decision Making," *Natural Resources J.* **16,** 1–21 (1976).
32. P. SLOVIC AND A. TVERSKY, "Who Accepts Savage's Axiom?" *Behav. Sci.* **19,** 368–373 (1974).
33. D. TIHANSKY, "Confidence Assessment of Military Air Frame Cost Predictions," *Opns. Res.* **24,** 26–43 (1976).
34. A. TVERSKY, "Intransitivity of Preferences," *Psychol. Rev.* **76,** 31–48 (1969).
35. I. VERTINSKY AND E. WONG, "Eliciting Preferences and the Construction of Indifference Maps: A Comparative Empirical Evaluation of Two Measurement Methodologies," *Socio-Econ. Planning Sci.* **9,** 15–24 (1975).
36. S. R. WATSON AND R. V. BROWN. "Issues in the Value of Decision Analysis. Decisions and Designs (McLean, Va.), *DDI Tech. Report 75-10,* 1975.
37. S. R. WATSON, "The Valuation of Decision Analysis," *J. Roy. Statist. Soc. Ser. A* **141,** 69–78 (1978).
38. W. E. WESTMAN, "Problems in Implementing U.S. Water Quality Goals," *Am. Sci.* **65,** 197–203 (1977).
39. A. F. WICHELMAN, "Administrative Agency Implementation of the NEPA of 1969: A Conceptual Framework for Explaining Differential Response," *Natural Resources J.* **16,** 263–300 (1976).
40. D. VON WINTERFELDT AND W. EDWARDS, "Evaluation of Complex Stimuli Using Multi-attribute Utility Procedures," University of Michigan, Engineering Psychology Laboratory (Ann Arbor, Mich.), *Technical Report 011313-2-T,* 1973.
41. D. VON WINTERFELDT, "Error in Decision Analysis: How to Create the Possibility of Large Losses by Using Dominated Strategies," University of Southern California, Social Science Research Institute (Los Angeles, Ca.). *SSRI Research Report 75-4,* 1975.
42. G. N. WRIGHT AND L. D. PHILLIPS, "Personality and Probabilistic Thinking," *Br. J. Psychol.,* in press.
43. R. B. ZUNIGA, "The Experimenting Society and Radical Social Reform," *Am. Psychol.* **30,** 99–115 (1975).