REFERENCES

[1] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis.* New York: Wiley, 1958.
[2] P. F. Dienemann, "Estimating cost uncertainty using Monte Carlo techniques, Rand Memorandum RM-4854-PR, Santa Monica, CA, January 1966.
[3] W. Edwards, "Social utilities," *The Engineering Economist*, Summer Symposium Series, VI, 1971.
[4] ——, "How to use multiattribute utility measurement for social decisionmaking," this issue, pp. 326–340.
[5] P. C. Gardiner, "The application of decision technology and Monte Carlo simulation to multiple objective public policy decision making: A case study in California Coastal Zone management," Ph.D. dissertation, University of Southern California, 1974.
[6] P. C. Gardiner and W. Edwards, "Public values: Multiattribute utility measurement for social decision making," in Schwartz and Kaplan (Eds.), *Human Judgment and Decision Processes: Formal and Mathematical Approaches.* New York: Academic, 1975.
[7] R. L. Keeney, "Utility functions for multi-attributed consequences," *Management Science*, 18, 1972, 276–287.
[8] J. R. Newman, Unpublished SSRI working paper, 1976.
[9] H. Raiffa, "Preferences for multi-attribute alternatives," The RAND Corporation, RM-5868-DOT/RC, Santa Monica, CA, April, 1969.
[10] J. R. Miller, III, "A systematic procedure for assessing the worth of complex alternatives," Mitre Corporation, Bedford, MA, 1967.
[11] E. M. Scheuer and D. S. Stoller, "On the generation of normal random vectors," *Technometrics*, vol. 4, no. 2, May 1962.

A Framework for Processing Value Judgments

JUDEA PEARL, MEMBER, IEEE

***Abstract*—Traditional decision-analytic practice emphasizes the distinction between probability assessments and value (or utility) judgments. Whereas techniques for elicitation and integration of subjective probabilities often can be submitted to empirical tests of validity, the fidelity of encoding value judgments has so far defied measurements. A unified approach to the treatment of the two types of judgments is presented; value judgments are interpreted as conditional probability statements. Such formulation leads to rational methodologies and procedures for solving the following tasks: 1) empirical validation and refinement of value judgments and 2) aggregating value judgments obtained from a panel of experts.**

I. INTRODUCTION

PHILOSOPHERS have long separated between two kinds of human knowledge, informative (or positive) knowledge and normative (or value) knowledge. In the first category philosophers grouped the physical and socioeconomical sciences as well as our intuitive knowledge about objects, colors, geometrical shapes, etc. In the second category, which often enjoyed a more elevated status, we find the studies of ethics and religion.

Strangely, this philosophical tradition has retained its influence upon the thinking and practice governing the emerging technology of decision analysis (DA). From its early development in the 1940's until today, DA has prescribed separate interpretations and treatments to probabilistic and value judgments. Whereas the former is regarded as a formalization of experiential knowledge, the latter is treated as a subjective and capricious entity which aside from the requirement of satisfying transitivity may assume any structure whatsoever.

Manuscript received May 15, 1976; revised August 10, 1976. This work was supported in part by the National Science Foundation under Grant No. GJ-42732.

The author is with the School of Engineering and Applied Science, University of California, Los Angeles, CA.

This distinction has resulted in a technological disparity between the methodologies available for processing probabilistic and value judgments. Whereas the former enjoys the benefits of Bayesian analysis, no parallel formulation is available for processing value judgments. The disparity is most clearly reflected in the absence of rational procedures for refining, updating, calibrating, and aggregating value judgments. While judgmental probabilities can be modified by experts, updated by empirical observations, and rectified by calibration, none of those techniques seems to apply to value judgments. Value preferences seem to remain immune from analytical or empirical scrutiny, a thing to uphold and not question, satisfy but not alter.

Consider, for example, the problem of calibration. Calibration procedures are used to bring probabilistic codings given by an individual closer to the truth frequencies of previous judgments [1], [2]. These procedures are used to remove systematic biases from both the individual's own perception of uncertainty and biases inherent in the elicitation method. Thus, if a probability assessor claims that his bread has fallen on the buttered side 90 percent of the time, we can use this overstatement to attempt to calibrate his prediction that his next car stands a 90 percent chance of being a lemon. However, if that same "expert" claims that he *prefers* arsenic to hot chocolate, we have no decision-analytic procedure of tempering this preference. It must be taken at face value and fed into an optimization procedure, which would produce an action conforming to almost any declaration depicting the expert's desires, however careless.

Next consider the problem of aggregation. When a panel of experts provides conflicting probability assessments, Bayesian techniques are available [3] for combining these assessments into a single probability structure reflecting the group overall knowledge. These techniques treat the experts as noisy information sources and incorporate the decision-maker's judgment of their reliability into the aggregation rule. No parallel techniques are available for aggregating value statements. It is strange that Bayesian methodology can resolve experts' disagreements on the failure probability of a nuclear reactor but could not handle disagreements concerning the social value of that reactor.

The purpose of this paper is to explore a paradigm of value semantics which may facilitate extensions of Bayesian decision-analytic techniques to the processing of value judgments and help rectify the imbalance now prevailing in our conception of the informative and normative modes of human knowledge.

II. A Paradigm for the Semantics of Value Judgments

The paradigm advanced in this paper asserts that *value judgments and probability statements are one and the same thing.* Both are numerical codes of experiential data, both are constructed by the same mental procedures, and both are utilized to guide future actions.

If probability quantifies a person's belief in the eventuality of a certain situation, then utility quantifies one's estimate of the benefit he might eventually draw from a certain situation. It is reasonable, therefore, to regard value judgments as conditional probability (or conditional expectation) statements. For example, the statement "I prefer outcome A to outcome B" may be interpreted to mean "I estimate the probability of reaching a certain state of satisfaction conditioned upon A to be higher than that conditioned upon B," or "I estimate the total benefit achievable as a result of realizing A to be higher than that achievable after realizing B." In either interpretation the entity estimated by the value statement is the strength of causal connection between two eventualities; one being the outcome A or B, the second being a person's achieving a certain level of satisfaction. In the framework of Bayesian analysis, though, the quantification of causal relations among events is usually formulated through conditional probability (or conditional expectation) assertions. To remain consistent with this formulation it behooves us to regard each value statement as a parsimonious form of conditional probability assertions. It should therefore be treated as an empirical statement about "matters of fact" and should acquire a status equal to that of any probabilistic assertion which summarizes observational data.

To complete the unification of value judgments with probability statements we need to clarify the concept of "personal satisfaction" or "total benefit." In most people's eyes satisfaction may appear in a variety of shades and colors not all of which seem commensurable. For example, the satisfaction one receives from dining in a certain restaurant seems noncommensurable with that induced upon hearing encouraging political news, or with the fear of adding so many extra calories to one's diet. The unlike character of these kinds of satisfaction, however, seems to be derived from people's tendency to attribute to satisfaction the character of the anticipatory procedures invoked for estimating its future likelihood or future magnitude. The knowledge that added weight stands in causal connection to overeating, for example, serves an anticipatory function of estimating the degree of future discomfort resulting from a certain meal. The nature of the discomfort itself, though, is fairly commensurable with the nature of the gratification one receives during the meal. Both involve biological sensations of pleasure and pain.

Neurophysiology has found that many mammals' brains (including humans) contain a so-called "pleasure center," stimulation of which encourages that individual to seek further stimulation there. In rats, a conditioned reflex can be erected on this basis; in humans, the overt effect was that the patient asked the doctor to "Do that again, please." From this low level of neural sensation, therefore, it is reasonable to assume that only two parameters suffice to characterize the force which drives an organism's behavior: intensity and time duration. Formally we can express this assumption by stating that at any time an organism is attempting to maximize the expected value of a *satisfaction integral*:

$$S = \int_{t=t_0}^{\infty} s(t)\, dt \tag{1}$$

where $s(t)$ measures the intensity of the pleasure sensation at time t (a negative s should reflect painful sensations). Toda [4] used an identical formulation which he named an *adaptance* integral. S is a concept comparable to utility in the sense that, like utility, S is unidimensional and unique up to a linear transformation. However, unlike utility, S is not an attribute of a decision outcome, neither is it subjective in nature. It measures an actual neural activity of an organism undergoing a specific life history.

To connect S with the usual concept of utility, we assume that the organism configures its experiential knowledge in a form which permits him to calculate (or estimate) the function $E(S \mid \theta)$, where E stands for the expected value operation and θ is a description of a state of the world. We then say that the utility of the state θ is given by

$$V(\theta) = E(S \mid \theta). \tag{2}$$

Clearly, the mental calculation of $E(S \mid \theta)$ is subject to many biases and inconsistencies inherent in humans' computing machinery. The main deficiency lies in man's ability to organize past experience in a data structure that would permit him to calculate the probabilistic connection between S and θ (for all levels of S and all states θ). A second limitation lies in man's capacity to compute $E(S \mid \theta)$ on the basis of such probabilistic information. With these deficiencies in mind, our main assumptions at this point are 1) that an individual would attempt to harness his computational resources, however limited, to facilitate an approximate calculation of $E(S \mid \theta)$, and 2) that once a crude estimate of $E(S \mid \theta)$ is obtained, it is this estimate which governs individual choice behavior.

In our normal speech, the variable S does not appear explicitly as an argument of utility. We talk about utility of lotteries, outcomes or commodity bundles, tacitly implying that the latter are only vehicles serving to achieve or estimate the total future benefit S to be experienced by our pleasure and pain centers.[1]

The reason that the role of the senses in affecting value judgments has almost completely disappeared from our everyday language is several fold. First, the argument S remains constant in all value judgments issued by an individual at time t_0, and so, it is redundant and may be dropped from a normal discourse. Second, the neural activity represented by S is so logically remote and perceptually uncertain that it rarely participates explicitly in the calculation of $E(S \mid \theta)$. Chess masters, for example, concern themselves primarily with heuristic concepts such as mobility, tempo, center-control, etc., rarely mentioning the ultimate goal of the game: checkmate. In much the same way our cognition system is equipped with heuristic concepts such as the pursuits of social virtues, economical wealth, or political power which serve to guide our choices toward an intermediate milestone. These concepts do not in themselves guarantee a realization of high S but where chosen by our culture as useful computational tools in the estimation of $E(S \mid \theta)$, by virtue of their tight correlation with future pleasure/pain balance.

Formal models for the evolution and utilization of such value concepts were developed in the literature on artificial intelligence [6], [7], and will be used later in Section III. The main point to notice here is that the usefulness of these heuristic concepts lies not in an extraphysical domain but rather depends on the degree to which they constitute a faithful code of experiential data. For example, a utility curve for money is only useful so long as it reflects one's experience regarding one's ability to convert certain sums of money to streams of neural gratification. The rationale for using this utility curve to govern future actions is based on inductive projection of this past experience onto an estimate of the individual's conversion abilities in the future (including sums of money not yet experienced), and should also be based on factual knowledge codified in probabilistic forms. Adherence to this paradigm of value should strip value judgments of their metaphysical, extrafactual status and reinstate them on equal footings with conditional probability statements, vulnerable to empirical scrutiny and manageable by Bayesian technology.

III. Potential Applications

In this section we shall explore some of the operational implications of the value paradigm developed in Section II. Two tasks are studied: validation and aggregation of value judgments.

[1] This idea, that human conception of value derives its structure from sensory information, goes back to John Locke. In his words, "And thus we see how Moral Being and Notions, are founded on, and terminated in these simple ideas, we have received from Sensation or Reflection, besides which, we have nothing at all in our Understanding, to employ our Thoughts about" [5, p. 195]. He further states, "So that whencesoever, we take the Rule of Moral Actions; or what standard soever we frame in our Minds the *ideas* of Virtues or Vices, they consist only, and are made up of Collections of simple *ideas* which we originally received from Sense or Reflection" [5, p. 196].

Validation

The problem of establishing an external means for demonstrating the effectiveness of decision-analytic methods is one of the most controversial issues in decision analysis [8]. The difficulty stems from the fact that by its very nature the idea of externally validating a value-processing procedure seems to imply a paradox: the proclaimed aim of such procedures is to remain as faithful as possible to the subjective preference structure of the value assessor. That structure, though, is only observable via the procedure under test, and so, value processing and decision-analytic procedures appears to be self-validated on *a priori* grounds.

Our interpretation of value judgments, however, offers a sound resolution of this dilemma. No longer is fidelity with internal value structure an end in itself. One's value structure is merely a code of empirical data, and so, fidelity with that data should provide the necessary validation test.

If one regards the individual value assessor as a transducer for empirical data, then the objectives of value processing procedure are two-fold: first, gaining a faithful access to the code of experience and second, rectifying distortions between that code and the original data. Identical objectives are assigned to probability-processing methods, coding distortions are handled by calibration, and elicitation distortions are minimized by the celebrated "divide and conquer" approach [9].

Let us begin by examining this last approach more closely. It is devised to rectify a mismatch between the format in which a problem (or inquiry) is presented to an individual and the format in which that individual coded his experiential data. A model of the latter contains a combination of specific facts or scenarios (e.g., it rained on my last birthday), generic relations (e.g., when it rains I am normally moody) and procedures for assessing and integrating the former to provide answers to queries [10]. By and large, the formats of the queries confronting an individual do not match the structure of his knowledge organization. For example, the query "Was I moody on my last birthday?" can be readily answered from the assumed stored knowledge. However, the query "Did it rain the last time I was moody?" requires further processing. This would be the case if my memory contains no explicit recollection of that last moody occasion, and if the memory file was not organized to be accessed by referring to my moods. An answer to such a query, though, can be reconstructed using knowledge fragments such as "When I have to write a paper I am almost always moody" and "Last season it hardly rained, and I had quite a few papers to write," etc. The reconstruction procedure may be formalized in terms of the Bayes rule which provides the basis for the methodology of probabilistic information processing [11]. The rationale for this methodology lies with the assumption that judgments *copied* from memory are more reliable than those *reconstructed* from memory, and so, the methodology provides mechanical aids to perform the reconstruction on the basis of probabilistic judgments explicitly stored in one's knowledge base. Note that not every division guarantees conquest, but only that which decomposes a query into a set of components which are coded in explicit forms. Were I, for example, to decompose the query "Was I moody last time it rained?" (assumed to be coded explicitly) into a set of queries concerning the weather characteristics conditioned upon my state of mood (not coded explicitly), the reliability of the reconstructed answer would only be diminished.

Let us examine now if similar methodology could be applied to value elicitation. According to the paradigm proposed in Section II, value judgments require the mental calculation of $E(S \mid \theta)$, where θ is a description of a state of nature. It is clearly the case that no individual can expand the amount of memory necessary to store an explicit value of $E(S \mid \theta)$ for all conceivable θ. Rather, people seem to configure their experience in a form of probabilistic relations between S and various combinations of state attributes or aspects. For example, θ may represent the possession of a *specific* house, while our experience with the benefits of house dwelling is coded in terms of the benefit derived from certain combinations of aspects such as privacy, locality, number of rooms, etc. These entities which serve as primitives for our value code are called value *aspects*, (the name "attributes" or "dimensions" are also used) and will be designated by the vector $\boldsymbol{a} = (a_1,a_2,\cdots)$. In accordance with the "divide and conquer" principle, we should express the value $U(\theta)$ in terms of $\boldsymbol{a}$, which serves as a reference key to our code of knowledge. Thus we obtain

$$U(\theta) = E(S \mid \theta) = \sum_{\boldsymbol{a}} E(S \mid \boldsymbol{a},\theta)P(\boldsymbol{a} \mid \theta) \tag{3}$$

where $P(\boldsymbol{a} \mid \theta)$ represents the probability that aspect combination $\boldsymbol{a}$ would be attained assuming state θ is reached. Since the quantities $E(S \mid \boldsymbol{a},\theta)$ are never found in one's memory, we must replace them with $E(S \mid \boldsymbol{a})$:

$$U(\theta) = \sum_{\boldsymbol{a}} E(S \mid \boldsymbol{a})P(\boldsymbol{a} \mid \theta) \tag{4}$$

reflecting the fact that the value of any state θ can only be evaluated through its aspects, which in turn are defined by the language used in the organization of one's knowledge base.

If $P(\boldsymbol{a} \mid \theta)$ peaks sharply in the neighborhood of some aspect combination $\boldsymbol{a}(\theta)$, then (4) can be approximated by

$$U(\theta) \simeq E[S \mid \boldsymbol{a}(\theta)]. \tag{5}$$

For example, if a_3 represents the number of rooms in the example above, then $P(a_3 \mid \theta)$ is sharp since the number of rooms in a prospective house is usually known. Advanced knowledge of the degree of privacy (a_1) provided by a certain house, on the other hand, is generally less accurate and may require retaining the summation over $P(\boldsymbol{a} \mid \theta)$ in (4).

Among all techniques used in multiattribute utility analysis, the one most widely practiced and yet most hotly debated is the linear *decomposition* model [12], [13]. Here the analyst elicits the value (utility) of each attribute separately and the relative importance of these attributes and then combines these in a linear fashion to arrive at an overall utility evaluation. In terms of our model, the issue of whether the linear model is adequate reduces to the question of whether the conditional expectation $E(S \mid \boldsymbol{a})$ can be decomposed in a linear fashion to give:

$$E(S \mid \boldsymbol{a}) = f_1(a_1) + f_2(a_2) + \cdots. \tag{6}$$

Formally speaking, such a decomposition can be guaranteed only when the a's are both marginally and conditionally independent, a rare case for an arbitrarily chosen set of attributes. One should recall though that $\boldsymbol{a}$ is not arbitrary but represents that selected collection of descriptors used by an individual to organize his experiential knowledge. The selection of these a's has probably undergone a long process of perceptual and linguistic evolution aimed at economizing computations. A linear decomposition as in (6) would save a substantial amount of storage space and time in the calculation of $E(S \mid \boldsymbol{a})$. It stands to reason, therefore, that if a linear decomposition exists (using readily computable a's), it would be chosen by the race and utilized in knowledge organization.

All indications point to the fact that the attributes used by people to represent knowledge have evolved in such a way as to render the decomposition in (6) feasible. Anderson [14] has found that in a wide variety of judgmental situations (ranging from perceptual to linguistic experiments), human judgment of a unified "whole" can be modeled by very simple arithmetic combination (e.g., addition or multiplication) rules over the information components. Einhorn [15] and Dawes [16] reported that linear composities of attribute values elicited from individuals outperformed intuitive judgments obtained from the same individuals in such tasks as predicting a student's college success or the longevity of patients with Hodgkins disease. Samuel [7] has written one of the most successful checker-playing programs which evaluates game positions by a linear combination of attributes (e.g., material advantage, mobility) defined by a human player.

These experiments seem to support the notion that most human knowledge is organized in terms of linearly decomposable variables. Assuming now that value information is coded in a similar fashion, these experiments would constitute an external validation of the adequacy of the linear decomposition techniques. Indeed, it would be an extreme waste of computational resources for an organism to manage two different forms of knowledge representation, one for factual knowledge, the other for normative knowledge. (Would knowledge about shoes be organized differently than knowledge about socks?) The superior performance of linear decomposition techniques in predicting factual data [13] provides an external validation for the effectiveness of multiattribute utility elicitation, although the latter is not observable directly.

Value Aggregation

Bayesian techniques for aggregating probabilistic information obtained from several assessors are based on the concept that such assessments represent outcomes of a noisy experiment which are probabilistically related to the underlying state of the world [3]. Let $\boldsymbol{r} = (r_1,r_2,\cdots,r_n)$ constitute a vector of probabilistic reports obtained from experts regarding the probability of a specific variable θ. A straightforward application of the Bayes rule yields

$$P(\theta \mid \boldsymbol{r}) = kP(\boldsymbol{r} \mid \theta)P(\theta). \tag{7}$$

The term $P(\boldsymbol{r} \mid \theta)$ formalizes the decisionmaker's (or aggregator) knowledge regarding the manner in which the reports $\boldsymbol{r}$ are influenced by the underlying variable θ. It contains, for instance, his knowledge concerning the reputation of various experts insofar as issuing a faithful and accurate report. It also contains knowledge about the correlation expected among the experts (e.g., it protects against counting repetitive reports with equal weights).

The problem of integrating expert opinions is very similar to that of integrating knowledge received from various sections of one's memory. The internal clues or attributes used to predict an outcome serve the same function as external experts; the two represent a condensed code of past experience which is no longer available to the integrator. There is one difference, though, which makes the problem of expert aggregation harder than that of internal aggregation. A person is usually much more familiar with the nature of his internal procedures than with the nature of the experts he consults. For one thing, a person may have actively participated in the original design of his heuristic procedures (i.e., to render them linear decomposable) or may have observed their performance over a variety of inferential problems. In addition, tests on the operations of one's internal procedures can be readily conducted, as they involve only mental exercises. As a result, the articulation of the function $P(\boldsymbol{r} \mid \theta)$ may, in general, be a much harder task than the characterization of $P(\boldsymbol{a} \mid \theta)$ or $E(S \mid \boldsymbol{a})$ in (4).

Let us now imagine that a panel of experts provides a list of value judgments instead of a probabilistic report. For example, each of the n experts may issue a real number $V_i(\theta)$, $i = 1,\cdots,n$, reflecting his assessment of the social value of a certain project θ. Our guiding paradigm dictates that each of these assessments should be treated as a condensed code for the expert's personal experience with related events. It should be processed, therefore, as an outcome of a physical measurement on the underlying process relevant to the value of θ, in the same fashion as probabilistic reports.

Let us first examine the case of competitive experts. Assume that a leasing authority offers a certain parcel of undeveloped land for bid in the open market and that each bidder issues his personal judgment as to the value of the land. Such judgments are usually issued after a careful survey of the property and after assessing the manner in which the development project might fit into the bidder's frame of business. It represents, therefore, the quantity V_i which measures the expected benefit to the ith bidder, assuming he wins the contract. The leasing authority attempts to determine the real worth of the property $V(\theta)$ on the bases of its own study as well as the aggregate reports it receives from the bidders. That worth depends on both the total income stream the property can generate, and the ability of the decisionmaker to enjoy that income. (Unlike a neural activity stream $s(t)$, an income stream should be discounted to account for uncertainties in one's ability to consume the income [17] or convert it into neural satisfaction.) Assume, for the moment, that the decisionmaker's expected benefit $E(S \mid \boldsymbol{\Theta})$ could be calculated given a set of attributes $\boldsymbol{\Theta} \doteq (\theta_1,\cdots,\theta_K)$ of the said property, (e.g., $\boldsymbol{\Theta}$ could include the quantity of oil deposits, the depth of that oil, etc.). The value judgments $V_1,V_2,\cdots,V_n$ could be used to modify the decisionmaker's probability distribution on $\boldsymbol{\Theta}$, yielding an aggregate utility:

$$\begin{aligned} U &= E[S \mid V_1,V_2,\cdots,V_n] \\ &= \sum_{\boldsymbol{\Theta}} E(S \mid \boldsymbol{\Theta})P(\boldsymbol{\Theta} \mid V_1,V_2,\cdots,V_n) \\ &\quad \cdot \sum_{\boldsymbol{\Theta}} E(S \mid \boldsymbol{\Theta}) \frac{P(V_1,V_2,\cdots,V_n \mid \boldsymbol{\Theta})P(\boldsymbol{\Theta})}{P(V_1,\cdots,V_n)}. \end{aligned} \tag{8}$$

The term $P(V_1,\cdots,V_n \mid \boldsymbol{\Theta})$ represents, like the terms $P(\boldsymbol{r} \mid \boldsymbol{\Theta})$ in (7), the decisionmaker's model of the experts, especially the dependence of their reports $V_1,\cdots,V_n$ on the value-determining parameter $\boldsymbol{\Theta}$. It should reflect the professional reputation of their surveyors, a possible correlation among the latter, and even the possibility of a concerted attempt to underbid the agency. A Bayesian scheme identical to (8) was used by Kaplan [18] in analyzing off-shore oil bidding strategy. Its salient point is the exploitation of value judgments as means for sharpening the decisionmaker's estimation of a state parameter $\boldsymbol{\Theta}$, which affects his own personal utility.

Let us now return to a more cooperative situation and consider, for example, a public official polling his constituents regarding the value of a certain public policy. It is customary in problems of this sort to view the official as an echo to public wishes and to seek equitable entity called public utility, which should be obtained by some arithmetic aggregation rule applied to the individual values of the constituents [19]. Aside from the celebrated conceptual difficulties connected with the concept of group utility [20], [21], this model also suffers from a practical drawback: it is utterly unrealistic. Regardless of how benevolent the public official is, he is bound by the very nature of his neural system wiring to act in a way that would maximize his *own* expected gratification. Part of this gratification is derived from materially contributing to public welfare, part from enjoying the image and reputation of a good-doer, and part from power-derived personal benefits. In short, it contains all factors and forces which play a role in the normal political arena. The maximization of the gratification function S requires knowledge of certain parameters $\boldsymbol{\Theta}$ which underly the sociopolitical scene. $\boldsymbol{\Theta}$ could include, for example, the degree of public trust enjoyed by the official at any given time, his likelihood of retaining power, etc. A rational scheme of aggregating public values would therefore consist of utilizing the latter as information signals to update one's knowledge of $\boldsymbol{\Theta}$ in a manner similar to that of (8). If, for example, a certain value-poll $V_1,\cdots,V_n$ strengthens the official's hypothesis that the wishes of a certain vocal section of the public are erratic or transitory and therefore could be ignored, it makes both ethical and logical sense for him to incorporate this knowledge into his action scheme.

On this note we wish to argue that it is in the systemization of such knowledge-updating procedures, rather than the advocation of *ad hoc* value-aggregation rules that Bayes technology would provide a realistic assistance to the decision process.

IV. Summary

The basic paradigm expanded in this paper is that value judgments are to be interpreted as representations of factual information. The meaning of such judgments, therefore, lies solely in their connection to empirical data recorded by an individual or by a cultured society. Viewed in this light, the problem of processing value judgments becomes almost identical to that of processing other codes of experiential knowledge, i.e., probabilistic judgments. This view implies that, aside from the desire to satify one's value structure, other factors should also be considered, such as the faithfulness of the judgment and the appropriateness of the experience it represents to the facts surrounding the problems at hand. Bayes analysis, the formalism used for representing factual knowledge and probabilistic judgments, can also be used to capture such consideration of faithfulness and appropriateness of value judgments, and incorporate the latter in a unified representation of human knowledge.

Acknowledgment

The author acknowledges with thanks his discussions with Dr. Ward Edwards of USC, Dr. Norman Dalkey of UCLA, and Dr. Donald O. Walter of the Brain Research Institute at UCLA.

References

[1] H. Raiffa, "Assessments of probabilities," unpublished manuscript, Harvard University, 1969.
[2] S. Lichtenstein, B. Fischoff, and L. D. Phillips, "Calibration of probabilities: The state of the art," *Proceedings of the 5th Conference on Subjective Probability, Utility and Decision Making*, Darmstadt, Germany, September 1975.
[3] P. A. Morris, "Decision analysis export use," *Management Science*, vol. 20, no. 9, May 1974.
[4] M. Toda, "Brain-computer approach to the theory of choice," in *Proc. 1960 Int. Congress Logic Methodology and Philosophy of Science*, Nagel, Suppes, and Tarski, Eds., Stanford Univ. Press, pp. 434–441.
[5] J. Locke, *An Essay Concerning Humane Understanding*, 3rd Edition. London: 1695.
[6] N. J. Nilsson, *Problem Solving Methods in Artificial-Intelligence*. New York: McGraw-Hill, 1971.
[7] A. Samuel, "Some studies in machine learning using the game of checkers," reprinted in E. Feigenbaum and J. Feldman (Eds.), *Computers and Thought*. New York: McGraw-Hill, 1963, pp. 71–105.
[8] J. Pearl, "On demonstrating effectiveness of decision analysis," UCLA-ENG-1272, September 1972, also: R. J. Newman, "Assessing the reliability and validity of multi-attribute utility procedures: An application of the theory of generalizability," University of Southern California, SSRI Research Report 75-7, July 1975.
[9] H. Raiffa, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Reading, MA: Addison-Wesley, 1968.
[10] D. E. Rumelhart, P. H. Lindsay, and D. A. Norman, "A process model for long-term memory," in E. Tulving and W. Donaldson (Eds.), *Organization of Memory*. New York: Academic Press, 1972, p. 198.
[11] W. Edwards, "Dynamic decision theory and probabilistic information processing," *Human Factors*, pp. 59–73, 1962.
[12] P. L. Gardiner and W. Edwards, "Public values: Multi-attribute utility measurement for social decision making," *Proceedings of the 1975 International Conference on Cybernetics and Society*, San Francisco, CA, pp. 309–314, September 23–25, 1975.
[13] R. M. Dawes, "Predictive models as a guide to preference," *Proceedings of the 1975 International Conference on Cybernetics and Society*, San Francisco, CA, pp. 320–322, September 23–25, 1975.
[14] N. H. Anderson, "Information integration theory: A brief survey," in R. C. Atkinson *et al.* (Eds.), *Contemporary Developments in Mathematical Psychology*, Volume 2. San Francisco: Freeman, 1974.
[15] H. J. Einhorn, "Expert judgment: Some necessary conditions and an example," *Journal of Applied Psychology*, *59*, pp. 562–571, 1974.
[16] R. M. Dawes and B. Corrigan, "Linear models in decision making," *Psychological Bulletin*, *81*, pp. 95–106, 1974.
[17] R. F. Meyer, "On the relationship among the utility of assets, the utility of consumption, and investment strategy in an uncertain, but time-variant world," *OR 69: Proceedings of the Fifth International Conference on Operation Research*, Venice, 1964, Tavistock Publications, 1970.
[18] S. Kaplan, "On the application of Bayes inference to accept-reject decisions on off-shore lease bids," *Journal of Petroleum Technology*, March 1976.
[19] J. F. Nash, "The bargaining problem," *Econometrica*, vol. 18, no. 2, April 1950.
[20] K. Arrow, *Social Choice and Individual Values*. New York: Wiley, 1951.
[21] N. C. Dalkey, "Group decision analysis," UCLA-ENG-7571, August 1975.