This article was downloaded by: [106.51.226.7] On: 09 August 2022, At: 21:43
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Exceptional Paper—Probability Encoding in Decision Analysis

Carl S. Spetzler, Carl-Axel S. Staël Von Holstein,

To cite this article:
Carl S. Spetzler, Carl-Axel S. Staël Von Holstein, (1975) Exceptional Paper—Probability Encoding in Decision Analysis. Management Science 22(3):340-358. https://doi.org/10.1287/mnsc.22.3.340

Full terms and conditions of use: **https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions**

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

© 1975 INFORMS

Please scroll down for article—it is on subsequent pages

informs®

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

MANAGEMENT SCIENCE
Vol. 22, No. 3, November 1975
Printed in U.S.A.

Exceptional Papers

PROBABILITY ENCODING IN DECISION ANALYSIS†‡*

CARL S. SPETZLER AND CARL-AXEL S. STAËL VON HOLSTEIN§

Stanford Research Institute

This paper presents the present philosophy and practice used in probability encoding by the Decision Analysis Group at Stanford Research Institute. Probability encoding, the process of extracting and quantifying individual judgment about uncertain quantities, is one of the major functions required in the performance of decision analysis. The paper discusses the setting of the encoding process, including the use of sensitivity analyses to identify crucial state variables for which extensive encoding procedures are appropriate. The importance of balancing modeling and encoding techniques is emphasized and examples of biases and unconscious modes of judgment are reviewed. A variety of encoding methods are presented and their applicability is discussed. The authors recommend and describe a structured interview process that utilizes a trained interviewer and a number of techniques designed to reduce biases and aid in the quantification of judgment.

1. Introduction

Probability encoding, the process of extracting and quantifying individual judgment about uncertain quantities, plays an important role in the application of decision analysis. This paper 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 informal experiments.

There is a vast literature that relates to probability encoding. Winkler [21] was the first to define a procedure that used more than one encoding technique. The books by Brown et al. [2] and Peterson et al. [10] each include an extensive section on probability encoding; however, their main concern is with the actual quantification step rather than the whole process of encoding. Hampton et al. [3] and Staël von Holstein [16] both offer comprehensive reviews of the literature on probability assessment and encoding techniques.

*This paper has been refereed under guidelines for exceptional contributions.
†Processed by Professor Ronald A. Howard, Departmental Editor for Decision Analysis; received July 1973, revised October 1974. This paper has been with the authors 11 months for revisions.
‡Presented at the ORSA-TIMS-AIEE 1972 Joint National Meeting, Atlantic City, New Jersey, November 8–10, 1972.
§During our research we have collaborated with Daniel Kahneman and Amos Tversky of Hebrew University, Jerusalem; the material in §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 James E. Matheson, and we are grateful for their valuable comments. The paper has further benefitted from careful reviews by Robert L. Winkler and Michael M. Menke.

Copyright © 1975, The Institute of Management Sciences

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

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 which provide insights relevant to probability encoding. Two recent overviews are given by Peterson and Beach [9] and Rapoport and Wallsten [13]. However, these psychological studies have restricted usefulness for probability encoding in practical situations for three reasons. Most studies deal with binary probability distributions (an event occurs or does not occur) rather than with continuous distributions. Moreover, they are based on laboratory experiments rather than actual decision situations. (The problems involved in extrapolating results from the laboratory to the real world have recently been discussed by Winkler and Murphy [22].) Finally, while the studies show how well or poorly subjects perform various tasks, they do not suggest procedures for improving performance.

This paper begins by discussing the role of probability encoding in the framework of decision analysis and the interaction that occurs between modeling and encoding. Since modeling and encoding represent two different but complementary ways of quantifying judgment, a balance of their use must be found in any decision problem. §4 discusses in detail the unconscious procedures or modes of judgment that individuals use in assessing uncertainties, and the biases that may follow from their use. §§5 and 6 describe a variety of probability encoding techniques and present a set of guidelines and a proposed structure for the probability encoding interview process. The paper concludes with a brief summary of our approach to probability encoding and with suggestions for areas of further research.

2. The Decision Analysis Framework

Probability encoding is a major element of decision analysis. It is usually performed in the context of a specific decision problem. However, the general benefit of probability encoding extends beyond the analysis of specific decisions. It improves a subject's awareness of his state of information and provides a clear means for communication and inference about uncertainty.

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 [4], [5] and Staël von Holstein [18]. Applications of decision analysis have been presented by Brown et al. [2], Howard et al. [6], and Spetzler and Zamora [15].

A decision analysis usually involves three phases—the deterministic, probabilistic, and informational phases. In the *deterministic phase* the decision problem is structured by defining relavant variables, characterizing their relationships in formal models, and assigning values to possible outcomes. The importance of the different variables is measured through sensitivity analysis.

During the *probabilistic phase*, uncertainty is explicitly incorporated into the analysis by assigning probability distributions to the important variables. These distributions are obtained by encoding the judgment of individuals knowledgeable about the problem. These judgments are processed using the models developed in the deterministic phase and are transformed into a probability distribution that expresses the uncertainty about the final outcome. After the decision maker's attitude toward risk has been evaluated and taken into account, the best alternative in the face of uncertainty is established.

In the *informational phase* the economic value of information is determined by calculating the worth of reducing uncertainty about each of the important variables in the problem. The value of additional information can then be compared with the cost

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

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 course of any decision analysis, the focus is on the decision and the decision maker. Expanding an analysis is considered valuable only if it helps the decision maker choose between available alternatives. Consequently, the time and effort expended on probability encoding depends on the importance of the decision problem, the importance of the variable under consideration, and so on. Sometimes it is useful to spend a great deal of time and effort on encoding probabilities. For such important variables, the judgments of more than one person are often encoded. In other situations only a brief encoding effort may be necessary.

3. Modeling and Encoding

The personal interpretation of probability is a cornerstone of the decision analysis philosophy. An assessment of probability reflects an individual's state of information about a given quantity or event. Since various people are likely to have different information, two persons can arrive at different probability assignments for the same uncertain quantity.

The decision maker is the person (or group of persons) who has 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.

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 meaningful to discuss encoding only with respect to state variables. At first it may seem difficult to classify some variables, such as price, as either decision or state variables. This difficulty, however, can be resolved by further structuring of the problem. For example, price may be separated into controllable price strategy and the uncertain market response. A similar problem of definition can arise when variables interact. For example, in new product decisions the variables of development time, program cost, and product performance are closely related. To solve this problem, one or two variables can be selected as decision variables and the others can then be treated as state variables. Before any probability encoding can begin, every decision problem must be carefully structured so that it is clear which variables are best considered decision variables and which state variables.

The Balance Between Modeling and Encoding

During the modeling stage of a decision problem, there is always the question of whether to encode the uncertainty in an important variable directly or to model the problem further. At one extreme, it is conceivable that the final worth or profit of a project could be encoded directly, thus bypassing the need for examining the underlying variables. Generally, however, a probability distribution for final worth is

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

more easily reached and engenders greater confidence if a model is constructed that relates the final worth to underlying variables. Modeling efforts tend to be most effective and most economical if they begin with a gross model that is successively refined. A model should be refined only as long as the cost of each additional refinement provides at least comparable improvement in information. The criterion for how much information is needed depends on how significantly the information bears on the decision at hand.

A decision about whether or not to launch a new-product development illustrates how the degree of modeling must be adapted to suit the problem at hand. Naturally, one of the most important factors in decisions regarding new-product introductions is the size of the market for the product. The simplest model might consider the market as a whole and define it by total market potential, company market share, and average growth rate. This is obviously a crude description, but in many cases it might be sufficient for the decision at hand. For a product with many potential markets with different characteristics, it might be necessary to expand the model to describe some of the markets separately. However, sensitivity analysis often shows that even though the markets are defined explicitly in the model, not all of them need be considered for probability encoding. In a recent decision analysis, probability distributions were encoded for the total international market and for three major domestic applications. At the same time, the remaining domestic applications were lumped together into one variable which, although uncertain, did not need to be considered probabilistically. In some situations, though, the complexity and importance of the problem necessitates a complicated model structure. For example, a recent study of a plant decision in the oil industry required a thorough model of the whole industry in order to assess the company's own market and price picture.

The choice between additional modeling and encoding may also need to be reconsidered during the encoding process. He may also reveal biases during the interview that can be counteracted by further structuring of the problem. The subject may find it easier to think about the problem in terms of a different structure.

Guidelines for Preparing to Encode Uncertain Quantities

We offer the following list of principles as an aid in defining and structuring any variable for which the uncertainty is to be encoded. From our experience, violating these principles leads to problems in the probability encoding process.

• Choose only uncertain quantities that are important to the decision, as determined by a sensitivity analysis. Be prepared to explain to the subject why the quantity is important to the decision. This demonstrates the relevance of the encoding process and is essential in gaining the subject's full cooperation.

• Define the quantity 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.

• Structure the quantity carefully. The subject may think of the quantity as conditional on other quantities. If so, conditionalities should be considered consciously and incorporated into the model structure because it is difficult for human minds to deal effectively with combinations of uncertain quantities. For example, a major consideration for someone forecasting the sales of a new product might be whether the main competitor will develop a similar product. The encoding might then be facilitated by making two separate probability assignments—one for the case where the competitor exists and one where he does not. Mental acrobatics should be minimized.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

• Clearly define the quantity. A good test is to ask whether a clairvoyant could reveal the value of the quantity by specifying a single number without requesting clarification. For example, it is not meaningful to ask for "the price of wheat in 1975," because the clairvoyant would need to know the quantity, kind of wheat, the date, the exchange, and whether you wanted to know the buying or selling price. However, "the closing price of 10,000 bushels of durum wheat on June 30, 1975, at the Chicago Commodity Exchange" is a well-defined quantity.

• Describe the quantity using a scale that is meaningful to the subject. For example, if the unknown quantity refers to a quantity of oil the subject may think in terms of gallons, barrels, or tank cars, depending on his occupation. 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 [1]

People seem to assess uncertainty in a manner similar to the way they assess distance. They use intuitive assessment procedures that are often based on cues of limited reliability and validity. Generally, these procedures or modes of judgment produce reasonable answers. For example, an automobile driver can generally estimate distance accurately enough to avoid accidents, and a business executive can generally evaluate uncertainties well enough to make his enterprise profitable. On the other hand, overreliance on certain modes of judgment may lead to answers that are systematically biased, sometimes with severe consequences.

To pursue the example with estimation of distance, it is known that people consistently overestimate the distance of a remote object when visibility is poor and underestimate the distance when the sky is clear. In other words, they exhibit a regular systematic bias. This is because they rely on the haziness of an object as a cue to its 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) Generally people are not aware of the cues on which their judgments are based. Few people know that they use haze to judge distances, although research shows that this applies to virtually everybody. (2) It is difficult to control the cues people use; the object seen through haze still appears more distant, even when we know why. (3) People can be made aware of the bias, and then can make a conscious attempt to control its effects, as does a pilot when flying on a hazy day.

These same characteristics pertain to the assessment of uncertain quantities. Here too, one relies on certain modes of judgment that may introduce systematic biases. Likewise, although modifying impressions and intuitions is exceedingly difficult, it is possible to learn to recognize the conditions under which such impressions are likely to lead us astray.

We will now briefly categorize some biases that may be encountered in probability

[1]Much of the material in this section is based on communications with Daniel Kahneman and Amos Tversky and is supported by their research]7], [8], [19], [20]. The last of the four references gives an overview of their research; each of the first three deals with a specific issue. The analogy between judgment of distance and judgment of uncertainty is due to them.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

encoding due to modes of judgment 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 by receiving new information. The task of the analyst is to elicit from the subject a probability distribution that describes his 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.* Some discrepancies in distributions 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.

The sources of bias can be classified as motivational or cognitive. *Motivational biases* are either conscious or subconscious adjustments in the subject's responses motivated by his perceived system of personal rewards for various responses. In other words, he may want to influence the decision in his favor by giving a particular set of responses. Or he may want to bias his response because he believes that 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 with a high degree of certainty 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 that are systematically introduced by the way the subject's responses that are systematically introduced by the way the subject intellectually processes his perceptions. For example, a response may be biased toward the most recent piece of information simply because that information is the easiest to recall. Cognitive biases, therefore, depend on subject's modes of judgment.

Basic Modes of Judgment

Since cognitive biases may result from the use of particular modes of judgment, an important responsibility of the interviewer is to discover what modes of judgment the subject might be using and then try to adapt the interview to minimize biases. In this section, we define five different modes of judgment and give examples of how each might operate in producing bias.

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 [19]. It is easy to recall information that made a strong impression at the time it was first presented. Recent results and present plans are also easily available. Because recent information is more available than old information, it is often given too much weight. For example, a piece of recent news

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

regarding a long-standing competitor may influence a sales forecast much more than it should be allowed to in the face of past experience with the competitor.

Availability appears to be an important mode of judgment in most probability encoding sessions. It can also be introduced deliberately by the interviewer to help compensate for a subject's bias. 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 [20]. 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 judgments about the extreme values for this year's sales. The initial response in an interview often 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. If this happens, we say that the response is "anchored" on the basis and the result is likely to be a central bias. Thus, anchoring can occur when some information has become overly available at the beginning of the interview. Anchoring results from a failure to process information about other points on the distribution independently from the point under consideration.

An experimental result, first reported by Alpert and Raiffa [1] and since replicated in many different settings, is that subjects seem to produce a central bias when they are asked first for the median for an uncertain quantity and then for the quartiles. Subsequent responses seem to be anchored on the first response, the median, which the subject usually views as the best single-number estimate.

Representativeness. Representativeness means that the probability of an event or a sample is evaluated according to the degree to which it is considered representative of, or similar to, some specific major characteristics of the process or population from which it originated [7], [8]. The effect of this mode of judgment is that probability judgments are reduced to judgments of similarity. For example, people tend to assign roughly the same distribution to the average of a group of uncertain quantities (e.g., the average production volume for a group of machines) as they do to each individual quantity forming the average when they should assign a much tighter distribution to the average. In this example an individual's judgment is based on the same characteristics, namely those of the population from which the group came.

There is a strong tendency to place more confidence in a single piece of information that is considered representative than in a larger body of more generalized information. For example, a company had to decide whether or not to introduce a new product that was considered to have a high demand potential. The product was test-marketed and there was a slightly unfavorable outcome; the revised assessment of the market said there was 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.

The biases resulting from representativeness can often be reduced or eliminated by

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

further structuring of the problem. In the marketing example, it would be easy to encode the prior probabilities for various levels of demand and then encode the probability distribution for the test result conditional on the demand. A simple application of probability calculus then will provide the posterior probabilities of demand level given the outcome of the market test.

Unstated Assumptions. Typically, a subject's responses are conditional on various unstated assumptions. Consequently, the resulting probability distribution does not properly reflect his total uncertainty. For example, the subject may not have considered such possibilities as future price controls, major strikes, currency devaluation, war, and so on, when expressing his judgment because he assumes that he is not responsible for considering such events. One result of operating on this kind of assumption 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 he could incorporate into his judgment.

While the subject cannot be held responsible for taking into account all possible eventualities that may affect the quantity he is assessing, it is his (and the interviewer's) responsibility to state the assumptions he is making about his own limits of responsibility. Once identified, they can be built into the model and an appropriate expert (who may or may not be the current subject) can assign their probabilities.

Coherence. People sometimes appear to 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; these could be scenarios on what markets might be penetrated and what the penetration rate might be with a reasonable marketing effort. Courtroom arguments are another case where evaluation of credibility is often based on the coherence of the sequence of evidence as presented by the prosecution or the defense. It is thus important that the discussion of scenarios leading to possible outcomes for an uncertain quantity be well balanced, since the relative coherence of various arguments can have a strong effect on 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. The different encoding methods used vary according to whether they ask a subject to assign probabilities (P), values (V), or both. The three basic types of encoding methods are listed below:

- P-methods require the subject to respond by specifying points on the probability scale while the values remain fixed.
- V-methods require the subject to respond by specifying points on the value scale while the probabilities remain fixed.
- PV-methods ask questions that must be answered on both scales jointly; the subject essentially describes points on the cumulative distribution.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

Any encoding procedure consists of a set of questions that the subject responds to either directly by providing numbers or indirectly by choosing between simple alternatives or bets. In the *direct response mode,* the subject is asked questions that require numbers as answers. Depending on the method being used, the answers will be given in the form of either values or probabilities.

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 to choosing between them. This indifference can then be translated into a probability or value assignment. When an external *reference process* is used, one bet is defined with respect to the uncertain quantity and the other with respect to a familiar reference event. Another procedure is to ask the subject to choose between events defined on the value scale for the uncertain quantity, where each event represents a set of possible outcomes for the uncertain quantity (for example, the event of sales being less than or equal to 2,000 units or that of sales being greater than 2,000 units). This type of response mode uses *internal events* for comparison.

Specific Encoding Techniques

Each probability encoding technique can be classified according to the encoding method and response mode used. The techniques we have found most useful are listed in Table 1 and described in the sections that follow.

TABLE 1. *Classification of Probability Encoding Techniques.*

	Response Mode		
	Indirect		
Encoding Method	External Reference Events	Internal Events	Direct
Probability (value fixed)	Probability Wheel	Relative likelihoods	Cumulative probability
Value (probability fixed)	Probability Wheel Fixed Probability Events	Interval technique	Fractiles
Probability-Value (neither fixed)	—	—	Drawing graph: Verbal encoding

Indirect Response Techniques

Probability Wheel. The probability wheel is one of the most useful tools we have discovered for encoding indirect responses from subjects. The wheel is a disk with two adjustable sectors, one blue and the other orange, with a fixed pointer in the center of the disk. (See Figure 1.) When spun, the disk will finally stop with the pointer either in the blue or the orange sector. 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 which of the two events he considers more likely—the event relating to the uncertain quantity (for example, the event that next year's production will not exceed x units), or the event that the pointer ends up in the orange sector. The amount of orange in the wheel is then varied until the subject finds

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

the two events equally likely. The relative amount of orange is then assigned as the probability of the event.

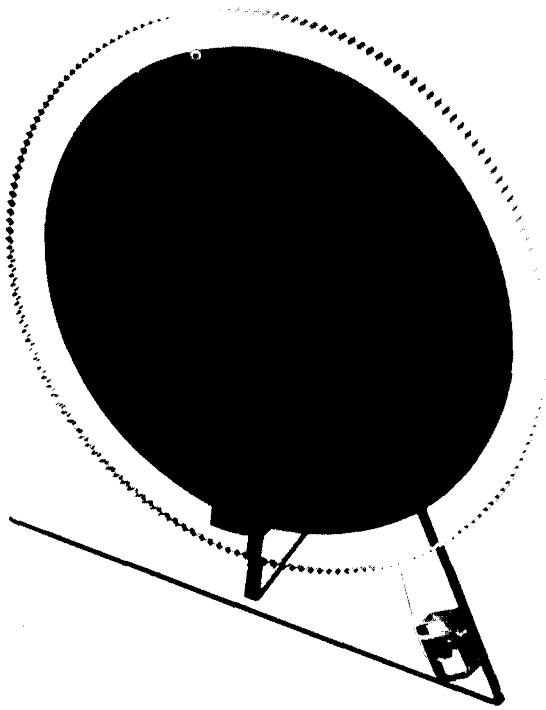

FIGURE 1. Probability Wheel.

Sometimes it helps to rephrase the questions in terms of choices among bets. For example, the subject is offered a choice between two propositions: he can win some amount, say $100, if the production does not exceed x units or he can win the same amount if the pointer ends up in the orange sector. This is an example of a P-method since the event (the value) is fixed and the subject is asked to assess the probability. The probability wheel can also be used as a V-method but this is generally less effective. The reason is that the P-method asks the subject to evaluate the uncertainty for a single event, whereas the V-method requires that the subject evaluate a sequence of events until he finds one that he considers equally likely as the reference event.

One advantage of the probability wheel is that the probability can be varied continuously from zero to one. However, because it is difficult for a subject to discriminate between the sizes of very small sectors, the wheel is most useful for evaluating probabilities in the range from 0.1 to 0.9.

Another tool that can be used instead of the probability wheel is a horizontal bar with a moveable marker than can be set to define two events—one to the left and one to the right of the marker. Another alternative is to ask the subject to visualize an urn with, say, 1000 balls of two colors. A ball is supposed to be drawn at random from the urn and the reference event in this case is that the ball will be orange. The composition of the urn can then be varied until it reflects the probability of the event in question [12]. We

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

prefer to use the probability wheel because subjects find it easier to visualize the chance process by looking at the wheel than by using the bar or visualizing the urn.

Fixed Probability Events. Other less tangible reference events may be useful in the encoding process, particularly when reference has to be made to low-probability events. For example, the event of tossing ten "heads" in a row with an unbiased coin has a probability close to 1/1000. Some subjects can identify easily with events relative to poker hands. For example, a royal flush has a probability of roughly 1/65,000. Typically reference processes such as these concern events with fixed probabilities and therefore basically work as V-methods, wherein the subject is asked to assess values that correspond to fixed probabilities. They could be used as reference events for P-methods if there were a large enough set of them; one would then go through the list of fixed probability events until indifference was reached.

Interval Technique. The interval technique is an example of an internal events response mode and is a V-method. To begin this technique, an interval is split into two parts. The subject is then asked to choose which part he would prefer to bet on, or which of the two parts he considers most likely. The dividing point is changed to reduce the size of the part considered most likely by the subject (thereby increasing the size of the other part), and the subject is asked to choose between the two new parts. The position of the dividing point is adjusted until the subject is indifferent between the two parts. These subintervals are then assigned equal probabilities. Beginning with an interval covering all possible outcomes and then splitting into two subintervals first gives the median, then the quartiles,[2] and so on. It is usually not meaningful to continue the interval technique after the quartiles have been obtained, because each question depends on earlier responses, and the errors are thus compounded. As noted earlier, subjects tend to arrive at their quartile estimates by adjustments from the median, which thus serves as an anchor. These adjustments are generally insufficient and result in distributions that are too tight. The interval technique can also be applied by splitting an interval into three or more parts at a time.

Relative Likelihoods. A P-method that uses the internal events response mode asks the subject to assign relative likelihoods (or odds) to two well-defined events. For example, the subject 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 many times more likely is it? This method is used primarily for uncertain quantities that have only a few possible outcomes.

Direct Response Techniques

Cumulative Probability and Fractiles. In the direct response mode the subject can be asked either to assign the 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 to assign the value corresponding to a probability (e.g., What is the level of sales that corresponds to a 10 percent probability?). The probability response can be expressed either as an absolute number (0.20), as a percentage (20 percent), or as a fraction ("one in five" or "two in ten"). This last way is particularly useful for small probabilities because subjects usually 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 fractional form is closely related to expressing it in terms of odds or ratios, particularly

[2]Examples of applications of this technique are presented by Pratt et al. [11] and Raiffa [12].

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

for probabilities close to zero. The form chosen to express probability should be the one that is most familiar to the subject.

Graphs. Graphing is a direct response mode that requires a subject to provide joint probability and value assignments, thereby making it a *PV*-method. It requires that the subject either draw a density function or a cumulative distribution or state a series of pairs of numbers (value and probability). Another approach is to show the subject a series of density functions and then ask him to choose the one that corresponds most closely to his judgment. The density functions can be generated easily by taking a family of distributions (e.g., beta distributions) and varying the parameters. This has been done in some psychological experiments with the help of CRT displays, where the subject has had two levers to change the parameters and thereby vary the displayed density function.

Verbal Encoding. Verbal encoding uses verbal descriptors to characterize events in the first phase of the encoding procedure. The descriptors used are those to which the subject is accustomed, such as "high," "medium," and "low" production cost. Quantitative interpretation of the descriptors is then encoded in a second phase. This method could be of particular use in dealing with quantities that have no ordinal value scale. Like the graphing technique, verbal encoding is a *PV*-method that requires the subject to provide joint probability and value assignments.

Applicability of the Various Techniques

Subjects seem to fall into two categories: those who feel capable of (and often prefer) giving direct numerical probability assignments and those who experience difficulty in making such judgments. Most people seem to fall into the second category. Furthermore, many individuals who prefer direct numerical responses are later found to have little confidence in their initial numerical responses. For this reason, the indirect response mode is generally the better way to begin encoding. Of these methods, most subjects find it easiest to use the probability wheel. Later, when the subject no longer finds the device helpful, the interviewer can shift to the direct response mode.

The interval technique, in which the subject is asked to generate the median, quartiles, and sometimes tertiles, is especially useful for arriving at a meaningful assessment of the median. However, it is generally unwise to begin the encoding process by eliciting the median, since that value tends to serve as an anchor for subsequent responses. Instead, we prefer to use the interval technique as a consistency check after other techniques have been used. Similarly, relative likelihood questions are usually used for verification of earlier responses.

Subjects are seldom able to express their uncertainty in terms of a density function, a cumulative distribution, or moments of a distribution, Therefore, it is usually not meaningful to try eliciting a distribution or its moments directly. There are, for example, procedures that ask the subject for the parameters of a special distribution—for example, the mean and standard deviation of a normal distribution or a beta distribution. Our experience indicates that subjects will give such parameters, but that usually they do not understand the full implications. We consider the choice of special distributions to be a modeling consideration and believe that it should not normally be made part of the encoding process. However, we do find that graphical displays of distributions drawn from indirect responses can provide useful feedback to the subject.

Encoding for Rare Events

Rare events present special problems in probability encoding since the standard

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

techniques do not work well for small probabilities. Subjects often find that it is difficult to discriminate between sizes of small sectors on a probability wheel. Similarly, the interval technique is more effective with the central part of the distribution even though theoretically it can also be used to generate the tails of a distribution. For example, continuing to split the lowest interval into equal parts generates an event with a probability of roughly 0.001 after ten steps. However, the final response is the composite of ten different responses and even slight biases in the responses lead to substantial error when compounded ten times.

As mentioned previously, fixed probability events such as poker hands or coin-tossing sequences can be used as external reference events for low-probability events. One can also develop reference processes that can serve as *P*-methods, at least when it comes to discriminating between orders of magnitude for the size of a probability. An example of such a technique is to show the subject a chart divided into squares one thousand by one thousand, that is, one million squares in all. (It is easy to make such a chart from standard graph paper.) The subject is asked to imagine that each of the one million squares has equal probability of being selected by some random mechanism. The event that a particular square will be chosen then has a probability of 10^{-6}, which is small enough for almost any rare event that might practically be encoded. Reference events with other probabilities are defined by selecting the relevant number of squares.

However, our experience with probability encoding for rare events indicates that probabilistic modeling is generally more effective than direct encoding. For example, in order for an event to occur it may be necessary that a sequence of other events occur. These intermediate events may not be low-probability events and standard encoding procedures can then be used. The problem of encoding the probability of a single rare event is thus transformed into the task of modeling the probabilistic structure of a sequence of events and then encoding the probabilities of these events. Some aspects of this modeling approach are discussed by Selvidge in her presentation of results from experiments on encoding probabilities for rare events [14].

6. The Interview Process

Structure of 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:

1. Motivating. Rapport with the subject is established and possible motivational biases are explored.
2. Structuring. The structure of the uncertain quantity is defined.
3. Conditioning. The subject is conditioned to think fundamentally about his judgment and to avoid cognitive biases.
4. Encoding. The subject's judgment is quantified in probabilistic terms.
5. 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 and probabilistic predictions.

The second purpose is to explore whether any motivational biases might operate. The interviewer and the subject should discuss openly any payoffs that might be associated with the probability assignment as well as possible misuses of the information. The subject may be aware of misuses of single-number predictions—for

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

example, the fact that they are often interpreted as firm projections or commitments. It should be pointed out that no commitment is inherent in a probability distribution and that the only aim of the encoding process is to elicit a probability distribution that represents as clearly as possible the complete judgment of the subject.

Structuring. The next phase in the interview process is to define and clearly structure the uncertain quantity. It should be defined as an unambiguous state variable and the definition should pass the clairvoyant test described in §3. The structure should be expanded as necessary 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 asked 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. The interviewer should probe any areas that seem unclear. Of course, this phase will vary greatly depending on the decision problem, the uncertain quantity, and the subject.

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, this 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 unstated assumptions are being made, and so on. The interviewer should be on the alert for biases symptomatic of the modes of judgment discussed in §4. We have found the following suggestions useful in many applications.

- Ask the subject to specify the most important bases for his judgment. Often these will be values from the current business plan or results from previous years. Such values could then be expected to act as anchors and possibly lead to a central bias.
- Ask the subject what information he is taking into account in making his estimates. This will indicate what information is most easily available. If the interviewer suspects that the subject has a central bias, he can use availability to correct this tendency, by asking the subject to compose scenarios that would produce extreme outcomes. This may have the effect of increasing the probabilities of the extreme outcomes.

If the uncertain quantity represents an average, such as average productivity or average reliability, try to determine whether the distribution assigned by the subject really is a distribution for the average or a distribution for an individual unit. Our experience has shown that subjects often have difficulty discriminating between the two situations. If the subject is assigning a distribution for individual units, 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. Often the best way to handle such a situation is to ask for the probability distribution of the quantity without the new information and for the probability of the information conditional on the revealed outcome of the quantity. By applying probability calculus to these two distributions, one can obtain the distribution for the quantity given the information.

Specify all assumptions on which the subject should condition his response, as well as those factors that the subject is expected to integrate into his judgment. The structure may be changed as some conditions are made explicit.

When a subject is assigning a probability to the occurrence or nonoccurrence of some event (for example, the probability that a product will or will not be successful in the market), he may base his assignment on whether he can generate plausible

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

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 influencing his judgment. 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 an equally coherent scenario that implies the opposite outcome might considerably change the subject's final probability assignment.

Encoding. The procedures outlined for this phase of the interview process are provided as a guideline. They rest primarily on use of the probability wheel as the encoding technique. Often a subject's responses will indicate a need to return to the tasks in the previous three phases. In particular, there may be a need for further structuring when the subject's responses and arguments indicate that they are based on different underlying assumptions.

- Begin by asking the subject for what he considers to be extreme values for the uncertain quantity. Then ask for scenarios that might lead to outcomes outside of those extremes. Also ask for the probabilities of outcomes outside of the extremes. This deliberate use of availability is designed to counteract the central bias that is otherwise likely to occur. Eliciting the scenarios for extreme outcomes makes them available to the subject, and he is thus more likely to assign higher probabilities to extreme outcomes. This has the overall effect of increasing the variability of his assigned distribution for the variable.
- Next take a set of values and use the probability wheel to encode the corresponding probability levels. Do not choose the first value in a way that may seem significant to the subject, otherwise you may 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. This means, for example, that you should begin by making the orange sector much smaller than what might actually correspond to the subject's probability. It is then easy for the subject to state which event is most likely and he becomes more comfortable with the procedure. Next, choose a sector that is much too large. After two easy choices, there is generally no problem in homing in on the indifference point.
- As you question the subject, plot each response as a point on a cumulative distribution and number them sequentially. An example is shown in Figure 2. This will point out any inconsistencies and will also indicate gaps in the distribution that need one or more additional points. Do not show the plotted points to the subject at this point in time. Otherwise, he may try to make subsequent responses consistent with the plotted points.
- Next, use the interval technique 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 toward responses that are consistent with each other. Be alert to shifts in the subject's attention (for example, shifting attention from the encoding process to the actual problem), changes in the subject's modeling of the situation, and the appearance of new information. After enough points have been encoded, a curve should be visually fitted to the points. An example is shown in Figure 3.

Verifying. In the last phase of the interview, the judgments are tested to see if the subject really believes them. If the subject is not comfortable with the final distribution, it may be necessary to repeat some of the earlier steps in the interview process.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

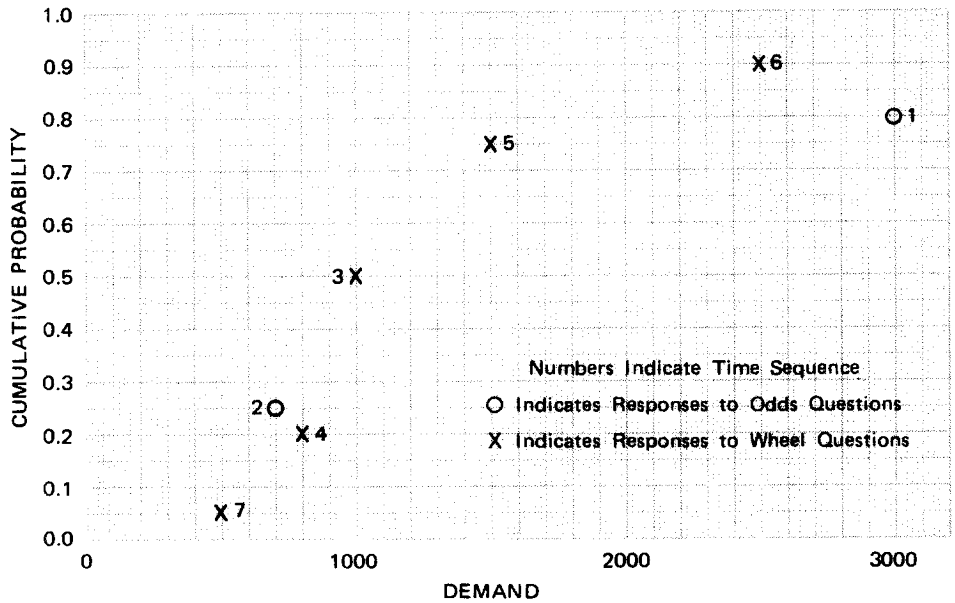

FIGURE 2. Graphical Representation of Responses.

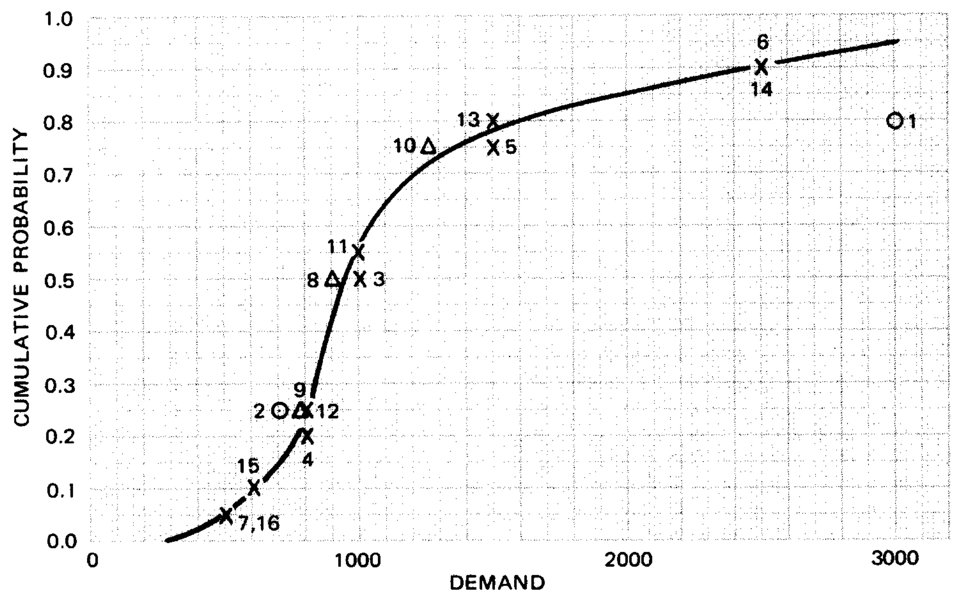

FIGURE 3. Example of a Curve Fitted to Responses.

Graphically representing the responses as points on a cumulative distribution and interpretating this distribution (perhaps in terms of a density function) provides an important test and feedback. The interviewer will naturally have to explain how the responses were plotted and how the fitted curve should be interpreted; this generally does not present problems. An examination of the distribution itself cannot show

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

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 to present the subject with a sequence of pairs of bets drawn from the distribution curve. Each pair is chosen so that the two bets should be equally attractive if the curve is consistent with the subject's judgment.

These checking procedures should be continued until several (perhaps three to five) indifference responses have been elicited. This provides the subject and the interviewer with confidence that the curve represents the subject's judgment.

Length of the Interview Process

A typical interview, in our experience, lasts anywhere from thirty to ninety minutes. The length of the interview depends on many factors, such as the importance and complexity of the uncertain quantity and the subject's previous experience in probability encoding. The pre-encoding steps are usually more time-consuming than the actual encoding step. This is particularly true when it is important to understand the structure underlying the subject's judgment and when dealing with subjects who are deeply involved in the project under analysis, especially if they have not had any exposure to the interviewer or to the decision analysis effort. The pre-encoding steps alone can sometimes take up to a couple of hours. The encoding step may take up to an hour if the quantity is very important or it may last only five or ten minutes if only a few points are needed on the distribution.

Comparison with Other Interview Processes

It should be clear that the encoding techniques discussed in this paper stress the interaction between interviewer and subject. We find that having the subject assign a probability distribution without the help of an analyst often leads to poor assignments. This is true even for subjects who are well trained in probability or statistics. The main reason for our emphasis on interaction is that it is difficult to avoid serious biases without having an analyst present.

The technique of having a subject fill out a questionnaire without an interviewer present also suffers from the lack of the interaction between interviewer and subject. Questionnaires can be used as a first approximation to the encoding process, but the subjects should preferably be experienced in probability encoding.

An interactive computer interview can make use of iterative techniques, such as the interval technique, and thereby avoid some of the pitfalls inherent in direct response modes.[3] However, the balancing effect of personal interaction is still missing. Again, we do not recommend using a computer program unless the subject has been through a number of actual interviews which dealt with similar uncertain quantities. Moreover, even when a subject has had long experience with the computer interview, it should not be used for encoding new quantities. However, in situations where the decision problem is not important enough to justify the cost of having an analyst perform the interview or when an organization uses probabilities regularly and extensively to communicate about uncertainty, interactive computer interviews and questionnaires might be quite valuable.

[3]An example of an interactive computer program is the Probability Encoding Program (PEP) developed by the Decision Analysis Group at Stanford Research Institute.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

7. Summary

In this paper we have presented our general approach to probability encoding in the form of a set of general guidelines rather than a cookbook with well-structured recipes on many different procedures. The procedures actually used for a particular uncertain quantity will always depend on the quantity and its importance for a decision, on the subject, and on the interviewer. The interviewer is an important factor since it is his responsibility to perceive problems or biases that the subject may have and to adapt the encoding procedures accordingly.

The methodology we have described differs considerably from what is generally described in the literature. We feel the following points distinguish our approach. First, the pre-encoding steps are essential and usually take longer than the encoding step. The purpose of the entire pre-encoding stage is to establish rapport between the interviewer and the subject, to make sure the problem is well defined, to model the problem, and to ensure that possible sources of bias have been detected. Second, probability assignments should be inferred from questions that require only ordinal judgments rather than direct probability assignments. Third, reference processes, such as the probability wheel, appear to provide an effective encoding technique for most subjects. Fourth, more than one technique should be used as a consistency check in both the encoding and verifying steps. A technique should also be changed if the subject seems to have a difficulty in using it.

There are many special topics related to probability encoding that may be relevant in particular situations, including encoding of discrete distributions, encoding of probabilities for rare events, accuracy and calibration of probability assignments, and the use of multiple subjects. Several of these topics are presently being researched at Stanford Research Institute and will be discussed in subsequent papers.

References

1. Alpert, M. A. and Raiffa, H., "A Progressive Report on the Training of Probability Assessors," unpublished manuscript, Harvard University, 1969.
2. Brown, R. V., Kahr, A. S. and Peterson, C. R., *Decision Analysis for the Manager*, Holt, Rinehart and Winston, New York, 1974.
3. Hampton, J. M., Moore, P. G. and Thomas, H., "Subjective Probability and Its Measurement." *Journal of the Royal Statistical Society, Ser. A*, Vol. 136 (1973), pp. 21 – 42.
4. Howard, R. A., "Decision Analysis: Applied Decision Theory," in D. B. Hertz and J. Melese (eds.), *Proceedings of the Fourth International Conference on Operations Research*, Wiley, New York, 1966 (1968), pp. 55–71.
5. ———, "The Foundations of Decision Analysis," *IEEE Transactions on Systems Science and Cybernetics*, Vol. SSC–4 (1968), pp. 211–219.
6. ———, Matheson, J. E. and North, D. W., "The Decision to Seed Hurricanes," *Science*, Vol. 176 (1972), pp. 1191–1202.
7. Kahneman, D. and Tversky, A., "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology*, Vol. 3 (1972), pp. 430–454.
8. ——— and ———, "On the Psychology of Prediction," *Psychological Review*, Vol. 80 (1973), pp. 237–251.
9. Peterson, C. R. and Beach, L. R., "Man as an Intuitive Statistician," *Psychological Bulletin*, Vol. 68 (1967), pp. 29–46.
10. ———, Kelly, C. W., III, Barclay, S. and Hazard, T. H., "Probability Distribution for a Continuum," a preliminary chapter for *Handbook for Decision Analysis*, Engineering Psychology Programs, Office of Naval Research, 1972.
11. Pratt, J. W., Raiffa, H. and Schlaifer, R., *Introduction to Statistical Decision Theory* (preliminary edition), McGraw-Hill, New York, 1965.
12. Raiffa, H., *Decision Analysis*, Addison-Wesley, Reading, Massachusetts, 1968.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

13. Rapoport, A. and Wallsten, T. S., "Individual Decision Behavior," *Annual Review of Psychology,* Vol. 23 (1972), pp. 131–176.

14. Selvidge, J. E., "Assigning Probabilities to Rare Events," unpublished dissertation, Harvard University, 1972.

15. Spetzler, C. S. and Zamora, R. M., "Decision Analysis of a Facilities Investment and Expansion Problem," *Proceedings of the Sixth Triennial Symposium sponsored by the Engineering Economy Division of ASEE and AIIE,* June 1971, published by *Engineering Economist.*

16. 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.

17. ———, "A Bibliography on Encoding of Subjective Probability Distributions," unpublished manuscript, Stanford Research Institute, 1972.

18. ———, "A Tutorial on Decision Analysis," unpublished manuscript, Stanford Research Institute, 1972.

19. Tversky, A. and Kahneman, D., "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology,* Vol. 5 (1973), pp. 207–232.

20. ——— and ———, "Judgment under Uncertainty: Heuristics and Biases," *Science* (1974), 185, pp. 1124–1131.

21. Winkler, R. L., "The Assessment of Prior Distribution in Bayesian Analysis," *Journal of the American Statistical Association,* Vol. 62 (1967), pp. 776–800.

22. ——— and Murphy, A. H., "Experiments in the Laboratory and the Real World," *Organizational Behavior and Human Performance,* Vol. 10 (1973), pp. 252–270.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.

Published in Management Science on November 01, 1975 as DOI: 10.1287/mnsc.22.3.340.
This article has not been copyedited or formatted. The final version may differ from this version.

Copyright 1975, by INFORMS, all rights reserved. Copyright of Management Science is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.

Downloaded from informs.org by [106.51.226.7] on 09 August 2022, at 21:43 . For personal use only, all rights reserved.