AD A047747

(12)

INTUITIVE PREDICTION:

BIASES AND CORRECTIVE PROCEDURES

DECISION RESEARCH • A BRANCH OF PERCEPTRONICS

Daniel Kahneman
Amos Tversky

DDC
RECEIVED
DEC 19 1977
B

ADVANCED DECISION TECHNOLOGY

CYBERNETICS TECHNOLOGY OFFICE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
Office of Naval Research • Engineering Psychology Programs

AD NO.
DDC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

The objective of the Advanced Decision Technology Program is to develop and transfer to users in the Department of Defense advanced management technologies for decision making. These technologies are based upon research in the areas of decision analysis, the behavioral sciences and interactive computer graphics. The program is sponsored by the Cybernetics Technology Office of the Defense Advanced Research Projects Agency and technical progress is monitored by the Office of Naval Research – Engineering Psychology Programs. Participants in the program are

Decisions and Designs, Incorporated
Harvard University
Perceptronics, Incorporated
Stanford University
The University of Southern California

Inquiries and comments with regard to this report should be addressed to

Dr. Martin A. Tolcott
Director, Engineering Psychology Programs
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

or

LT COL Roy M. Gulick, USMC
Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government. This document has been approved for public release with unlimited distribution.

TECHNICAL REPORT PTR-1042-77-6

INTUITIVE PREDICTION:

BIASES AND CORRECTIVE PROCEDURES

by

Daniel Kahneman and Amos Tversky

Sponsored by

Defense Advanced Research Projects Agency
Contract N00014-76-C-0074

June 1977

DECISION RESEARCH
A BRANCH OF PERCEPTRONICS
1201 Oak Street
Eugene, Oregon 97401
(503)485-2400

D D C
RECEIVED
DEC 19 1977
B

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

SUMMARY

Decisions vital to the accomplishment of military objectives are determined in large part by the intuitive judgments and educated guesses of decision makers or experts acting in their behalf. The critical role of intuitive judgments makes it important to study the factors that limit the accuracy of these judgments and to seek ways of improving them. Previous work in ARPA's Advanced Decision Technology Program has led to the discovery of major deficiencies in the unaided, intuitive judgments of probabilities for uncertain events. Of the many significant conclusions of this research, the following merit special mention:

(1) Errors of judgment are often systematic rather than random, manifesting bias rather than confusion. People suffer from mental astigmatism as well as from myopia, and any corrective prescription should deal appropriately with this diagnosis.

(2) There are no significant differences between the judgmental processes of experts, intelligence analysts, and physicians, to cite but a few, confirm the presence of common biases in the professional judgments of experts.

(3) Erroneous intuitions resemble visual illusions in a crucial respect: both types of error remain compellingly attractive even when the person is fully aware of their nature. In situations likely to produce illusions of sight or intuition,

ACCESSION for
NTIS White Section ☑
DDC Buff Section ☐
UNANNOUNCED ☐
JUSTIFICATION
BY
DISTRIBUTION/AVAILABILITY CODES
Dist. AVAIL. and/or SPECIAL
A

> we must let our beliefs and actions be guided by a critical and reflective assessment of reality, rather than by our immediate impressions, however compelling these may be.

This paper presents an approach to elicitation and correction of intuitive forecasts, which attempts to retain what is most valid in the intuitive process while correcting some errors to which it is prone. This approach is applied to two tasks that experts are often required to perform in the context of forecasting or in the service of decision-making: the prediction of uncertain quantities and the assessment of probability distributions. The analysis of these tasks reveals two common biases: non-regressiveness of predictions and overconfidence in the precision of estimates. In order to eliminate or reduce these biases, we propose specific procedures for the elicitation of expert judgments and for the assessment of corrected values. Our recommendations assume a dialogue between an expert and an analyst, whose role is to help the expert make most efficient use of his knowledge while avoiding some of the common pitfalls of intuition. Experts may, of course, act as their own analysts.

TABLE OF CONTENTS

	Page
SUMMARY	i
ACKNOWLEDGEMENT	iv
1. INTRODUCTION	1-1
2. SINGULAR AND DISTRIBUTIONAL DATA	2-1
3. REGRESSION AND INTUITIVE PREDICTION	3-1
3.1 A Corrective Procedure for Prediction	3-3
3.1.1 Step 1: Selection of a Reference Class	3-4
3.1.2 Step 2: Assessment of the Distribution For the Reference Class	3-5
3.1.3 Step 3: Intuitive Estimation	3-6
3.1.4 Step 4: Assessment of Predictability	3-6
3.1.5 Step 5: Correction of the Intuitive Estimate	3-8
4. THE OVERCONFIDENCE EFFECT	4-1
4.1 Debiasing Confidence Intervals	4-6
5. CONCLUDING REMARKS	5-1
DISTRIBUTION LIST	D-1
DD FORM 1473	

ACKNOWLEDGMENT

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under Contract N00014-76-C-0074 under subcontract from Decisions and Designs, Incorporated.

1. INTRODUCTION

Any significant activity of forecasting involves a large component of judgment, intuition and educated guesswork. Indeed, the opinions of experts are the source of many technological, political and social forecasts. Opinions and intuitions play an important part even where the forecasts are obtained by a mathematical model or a simulation. Intuitive judgments enter in the choice of the variables that are considered in such models, the impact factors that are assigned to them, and the initial values that are assumed to hold. The critical role of intuition in all varieties of forecasting calls for an analysis of the factors that limit the accuracy of expert judgments, and for the development of procedures designed to improve their quality.

The question of how people think under conditions of uncertainty has attracted increasing research interest in recent years. A comprehensive review of the findings and their implications has been assembled by Slovic, Fischhoff and Lichtenstein (1977) and some common biases have been described and analyzed by Tversky and Kahneman (1974). Several conclusions that emerge from this body of research are especially relevant to our present concern. First, errors of judgment are often systematic rather than random, manifesting bias rather than confusion. Thus, people suffer from mental astigmatism as well as from myopia, and any corrective prescription should fit this diagnosis. Second, many errors of judgment are shared by experts and laymen alike. Studies of stockbrokers (Stael von Holstein, 1972), electrical engineers (Kidd, 1970), intelligence analysts

(Brown, Kahr and Peterson, 1974) and physicians (Zieve, 1966), to cite but a few, confirm the presence of common biases in the professional judgments of experts. Third, erroneous intuitions resemble visual illusions in an important respect: the error remains compelling even when one is fully aware of its nature. Awareness of a perceptual or cognitive illusion does not by itself produce a more accurate perception of reality. Hopefully, it may enable us to identify situations in which our normal faith in our impressions must be suspended, and where judgment should be controlled by a more critical evaluation of the evidence.

This paper presents an approach to elicitation and correction of intuitive forecasts, which attempts to retain what is most valid in the intuitive process while correcting some errors to which it is prone. This approach is applied to two tasks that experts are often required to perform in the context of forecasting or in the service of decision-making: the prediction of uncertain quantities and the assessment of probability distributions. The analysis of these tasks reveals two common biases: non-regressiveness of predictions and overconfidence in the precision of estimates. In order to eliminate or reduce these biases, we propose specific procedures for the elicitation of expert judgments and for the assessment of corrected values. Our recommendations assume a dialogue between an expert and an analyst, whose role is to help the expert make most efficient use of his knowledge while avoiding some of the common pitfalls of intuition. Experts may, of course, act as their own analysts.

The rationale for our recommendations derives from a psychological analysis of judgmental biases. We have had some experience with the implementation of the proposed methods, which indicates that they are feasible. It should be emphasized, however, that the recommended procedures have not been subjected to systematic evaluation. They should be regarded as suggestions for improved practice, and as an illustration of a general approach to debiasing, rather than as a well established methodology of elicitation.

2. SINGULAR AND DISTRIBUTIONAL DATA

Experts are often required to provide a best guess, estimate or prediction concerning an uncertain quantity such as the value of the Dow-Jones index on a particular day, the future sales of a product, or the outcome of an election. We shall distinguish two types of information that are available to the forecaster: singular and distributional. Singular information, or case data, consists of evidence about the particular case under consideration. Distributional information, or base-rate data, consists of knowledge about the distribution of outcomes in similar situations. In predicting the sales of a new novel, for example, what one knows about the author, the style, and the plot is singular information, whereas what one knows about the sales of novels is distributional information. Similarly, in predicting the longevity of a patient, the singular information includes his age, state of health and past medical history, whereas the distributional information consists of the relevant population statistics. The singular information describes the specific features of the problem that distinguish it from others, while the distributional information characterizes the outcomes that have been observed in cases of the same general class. The present concept of distributional data does not coincide with the Bayesian concept of a prior probability distribution. The former is defined by the nature of the data, whereas the latter is defined in terms of the sequence of information acquisition.

Many prediction problems are essentially unique, in the sense that little, if any, relevant distributional information is available. Examples are the forecast of demand

for nuclear energy in the year 2000, or of the date by which an effective cure for leukemia will be found. In such problems, the expert must rely exclusively on singular information. However, the evidence suggests that people are insufficiently sensitive to distributional data even when such data are available. Indeed, recent research suggests that people rely primarily on singular information, even when it is scanty and unreliable, and give insufficient weight to distributional information (see, e.g., Kahneman and Tversky, 1973; Tversky and Kahneman, 1977).

The context of planning provides many examples in which the distribution of outcomes in past experience is ignored. Scientists and writers, for example, are notoriously prone to underestimate the time required to complete a project, even when they have considerable experience of past failures to live up to planned schedules. A similar bias has been documented in engineers' estimates of the completion time for repairs of power stations (Kidd, 1970). Although this 'planning fallacy' is sometimes attributable to motivational factors such as wishful thinking, it frequently occurs even when underestimation of duration or cost is actually penalized.

The planning fallacy is a consequence of the tendency to neglect distributional data, and to adopt what may be termed an 'internal approach' to prediction, where one focuses on the constituents of the specific problem rather than on the distribution of outcomes in similar cases. The internal approach to the evaluation of plans is likely to produce underestimation. A building can only be completed on time, for example, if there are no delays in the delivery of

materials, no strikes, no unusual weather conditions, etc. Although each of these disturbances is unlikely, the probability that at least one of them will occur may be substantial. This combinatorial consideration, however, is not adequately represented in people's intuitions (Bar-Hillel, 1973). Attempts to combat this error by adding a slippage factor are rarely adequate, since the adjusted value tends to remain too close to the initial value that acts as an anchor (Tversky and Kahneman, 1974). The adoption of an 'external approach', which treats the specific problem as one of many, could help overcome this bias. In this approach, one does not attempt to define the specific manner in which a plan might fail. Rather, one relates the problem at hand to the distribution of completion time for similar projects. We suggest that more reasonable estimates are likely to be obtained by asking the external question "*how long* do such projects usually last?", and not merely the internal question "what are the specific factors and difficulties that operate in the particular problem?"

The tendency to neglect distributional information and to rely mainly on singular information is enhanced by any factor which increases the perceived uniqueness of the problem. The relevance of distributional data can be masked by detailed acquaintance with the specific case, or by intense involvement in it. The perceived uniqueness of a problem is also influenced by the formulation of the question which the expert is required to answer. For example, the question of how much the development of a new product will cost may induce an internal approach where total costs are broken down into components. The equivalent question of the percentage by which costs will exceed the current budget is likely to call

to mind the distribution of cost-overruns for developments of the same general kind. Thus, a change of units, e.g. from costs to overruns, could alter the manner in which the problem is viewed.

The prevalent tendency to underweight, or ignore, distributional information is perhaps the major error of intuitive prediction. The consideration of distributional information, of course, does not guarantee the accuracy of forecasts. It does, however, provide some protection against completely unrealistic predictions. The analyst should therefore make every effort to frame the forecasting problem so as to facilitate the utilization of all the distributional information that is available to the expert.

3. REGRESSION AND INTUITIVE PREDICTION

In most problems of prediction, the expert has both singular information about the specific case and distributional information about the outcomes in similar cases. Examples are the counselor who predicts the likely achievements of a student, the banker who assesses the earning potential of a small business, the publisher who estimates the sales of a textbook, or the economist who forecasts some index of economic growth.

How do people predict in such situations? Psychological research (Kahneman and Tversky, 1973; Ross, 1977) suggests that the intuitive predictions are generated according to a simple matching rule: the predicted value is selected so that the standing of the case in the distribution of outcomes matches its standing in the distribution of impressions. The following example illustrates this rule. An editor reviewed the manuscript of a novel and was favorably impressed. He said: "This book reads like a best-seller. Among the books of this type that were published in recent years, I would say that only one in twenty impressed me more". If the editor were now asked to estimate the sales of this novel, he would probably predict that it will be in the top 5% of the distribution of sales.

There is considerable evidence that people often predict by matching prediction to impression. However, this rule of prediction is unsound because it fails to take uncertainty into account. The editor of our example would surely admit that sales of books are highly unpredictable. In such a situation of high uncertainty, the best prediction

of the sales of a book should fall somewhere between the value that matches one's impression and the average sales for books of its type.

One of the basic principles of statistical prediction, which is also one of the least intuitive, is that the extremeness of predictions must be moderated by considerations of predictability. Imagine, for example, that the publisher knows from past experience that the sales of books are quite unrelated to his initial impressions. Manuscripts that impressed him favorably and manuscripts that he disliked were equally likely to sell well or poorly. In such a case of zero predictability, the publisher's best guess about sales should be the same for all books (e.g., the average of the relevant category), regardless of his personal impression of the individual books. Predictions are allowed to match impressions only in the case of perfect predictability. In intermediate situations, which are of course the most common, the prediction should be regressive, i.e., it should fall between the class average and the value that best represents one's impression of the case at hand. The lower the predictability, the closer should the prediction be to the class average. Intuitive predictions are typically non-regressive: people often make extreme predictions on the basis of information whose reliability and predictive validity are known to be low.

The rationale for regressive prediction is most clearly seen in the prediction of the result of a repeated performance or a replication. The laws of chance entail that a very high score on the first observation is likely to be followed by a somewhat lower score on the second, while a poor score on

the first observation is likely to be followed by a higher score on the second. Thus, if we examine a group of firms that did exceptionally well last year we shall probably find that, on average, their current performance is somewhat disappointing. Conversely, if we select firms that did poorly last year we shall find that, on average, they are doing relatively better this year. This phenomenon, known as regression towards the mean, is a mathematical consequence of the presence of uncertainty. The best prediction for a repeated performance of an individual, a product, or a company is therefore less extreme (i.e., closer to the average) than the initial score. As was pointed out earlier, intuitive predictions violate this principle. People often make predictions as if measures of performance were equally likely to change toward the average and away from it.

The error of non-regressive prediction is common among experts as well as among laymen. Furthermore, familiarity with the statistics of prediction does not eliminate the erroneous strategy of matching predictions to impressions (Kahneman and Tversky, 1973). Thus, when an expert makes an intuitive prediction that is based on impression-matching, the analyst has grounds to suspect that the estimate is non-regressive, and therefore non-optimal.

3.1 A Corrective Procedure for Prediction

How can the expert be guided to produce properly regressive predictions? How can he be led to use the singular and the distributional information that is available to him, in accordance with the principles of statistical

prediction? In this section we propose a five-step procedure that is designed to achieve these objectives.

3.1.1 Step 1: Selection of a Reference Class. The goal of this stage is to identify a class to which the case at hand can be referred meaningfully, and for which the distribution of outcomes is known, or can be assessed with reasonable confidence.

In the prediction of the sales of a book, or of the gross earnings of a film, for example, the selection of a reference class is straightforward. It is relatively easy, in these cases, to define an appropriate class of books or films for which the distribution of sales or revenue is known.

There are prediction problems (e.g., forecasting the cost of developing a novel product, or the time by which it will reach the market), for which a reference class is difficult to identify, because the various instances appear to be so different from each other that they cannot be compared meaningfully. As was noted earlier, however, this problem can sometimes be overcome by redefining the quantity that is to be predicted. Development projects in different technologies, for example, may be easier to compare in terms of percentage of cost-overruns than in terms of absolute costs. The prediction of costs calls the expert's attention to the unique characteristics of each project. The prediction of cost-overruns, in contrast, highlights the determinants of realism in planning, which are common to many different projects. Consequently, it may be easier to define a reference class in the latter formulation than in the former.

More often than not, the expert will think of several classes to which the problem could be referred, and a choice among these alternatives will be necessary. For example, the reference class for the prediction of the sales of a book could consist of other books by the same author, of books on the same topic, or of books of the same general type, e.g., hard-cover novels. The choice of a reference class often involves a trade-off between conflicting criteria. Thus, the most inclusive class may allow for the best estimate of the distribution of outcomes but it may be too heterogeneous to permit a meaningful comparison to the book at hand. The class of books by the same author, on the other hand, may provide the most natural basis for comparison, but the book in question could well fall outside the range of previously observed outcomes. In this example, the class of books on the same topic could be the most appropriate.

3.1.2 Step 2: Assessment of the Distribution For The Reference Class. There are problems, e.g., sales of books, where statistics regarding the distribution of outcomes are available. In other problems, the relevant distribution must be estimated on the basis of various sources of information. In particular, the expert should provide an estimate of the class average, and some additional estimates that reflect the range of variability of outcomes. Sample questions are: "How many copies are sold, on the average, for books in this category?" "What proportion of the books in that class sell more than 15,000 copies?"

Many forecasting problems are characterized by the absence of directly relevant distributional data. This is always the case in long-term forecasting, where the relevant

distribution pertains to outcomes in the distant future. Consider, for example, an attempt to predict England's share of the world market in personalized urban transportation systems, in the year 2000. It may be useful to recast this problem as follows: "What is the likely distribution, over various domains of advanced technology, of England's share of the world market in the year 2000? How do you expect the particular case of transportation systems to compare to other technologies?" Note that the distribution of outcomes is not known in this problem. However, the required distribution could probably be estimated, on the basis of the distribution of values for England's present share of the world market in different technologies, adjusted by an assessment of the long-term trend of England's changing position in world trade.

3.1.3 Step 3: Intuitive Estimation. One part of the information which the expert has about a problem is summarized by the distribution of outcomes in the reference class. In addition, the expert usually has a considerable amount of singular information about the particular case, which distinguishes it from other members of the class. The expert should now be asked to make an intuitive estimate on the basis of this singular information. As was noted above, this intuitive estimate is likely to be non-regressive. The objective of the next two steps of the procedure is to correct this bias and obtain a more adequate estimate.

3.1.4 Step 4: Assessment of Predictability. The expert should now assess the degree to which the type of information which is available in this case permits accurate prediction of outcomes. In the context of linear prediction, the appropriate

measure of predictability is ρ, the product-moment correlation between predictions and outcomes. Where there exist records of past predictions and outcomes, the required value could be estimated from these records. In the absence of such data, one must rely on subjective assessments of predictability. A statistically sophisticated expert may be able to provide a direct estimate of ρ on the basis of his experience. When statistical sophistication is lacking, the analyst should resort to less direct procedures.

One such procedure requires the expert to compare the predictability of the variable with which he is concerned to the predictability of other variables. For example, the expert could be fairly confident that his ability to predict the sales of books exceeds the ability of sportcasters to predict point-spread in football games, but is not as good as the ability of weather forecasters to predict temperature two days ahead of time. A skillful and diligent analyst could construct a rough scale of predictability based on computed correlations between predictions and outcomes for a set of phenomena that range from highly predictable (e.g., temperature) to highly unpredictable (e.g., stock prices). He would then be in a position to ask the expert to locate the predictability of the target quantity on this scale, thereby providing a numerical estimate of ρ.

An alternative method for assessing predictability involves questions such as "if you were to consider two novels that you are about to publish, how often would you be right in predicting which of the two will sell more copies?" An estimate of the ordinal correlation between predictions and outcomes can now be obtained as follows: If ρ is the estimated proportion of pairs in which the order of outcomes

was correctly predicted, then $\tau = 2p - 1$ provides an index of predictive accuracy, which ranges from zero when predictions are at chance level to unity when predictions are perfectly accurate. In many situations τ can be used as a crude approximation for ρ.

Estimates of predictability are not easy to make, and they should be examined carefully. The expert could be subject to the hindsight fallacy (Fischhoff, 1975), which leads to an overestimate of the predictability of outcomes. He could also be subject to an availability bias (Tversky and Kahneman, 1973), and mostly recall surprises or memorable cases where strong initial impressions were later confirmed.

3.1.5 Step 5: Correction of the Intuitive Estimate. To correct for non-regressiveness, the intuitive estimate should be adjusted toward the average of the reference class. If the intuitive estimate was non-regressive, then under fairly general conditions the distance between the intuitive estimate and the average of the class should be reduced by a factor of ρ, where ρ is the correlation coefficient. This procedure provides an estimate of the quantity, which is hopefully free of the non-regression error.

For example, suppose that the expert's intuitive prediction of the sales of a given book is 12,000 and that, on the average, books in that category sell 4,000 copies. Suppose further that the expert believes that he would correctly order pairs of manuscripts by their future sales on 80% of comparisons. In this case, $\tau = 1.6 - 1 = .6$, and the regressed estimate of sales would be $4,000 + .6(12,000 - 4,000) = 8,800$.

The effect of this correction will be substantial when the intuitive estimate is relatively extreme and predictability is moderate or low. The rationale for the computation should be carefully explained to the expert, who will then decide whether to stand by his original prediction, adopt the computed estimate, or correct his assessment to some intermediate value.

The procedure that we have outlined is open to several objections, which are likely to arise in the interaction between analyst and expert. First, the expert could question the assumption that his initial intuitive estimate was non-regressive. Fortunately, this assumption can be verified by asking the expert to estimate (i) the proportion of cases in the reference class (e.g., of manuscripts) which would have made a stronger impression on him; (ii) the proportion of cases in the reference class for which the outcome exceeds his intuitive prediction (e.g., the proportion of books that sold more than 12,000 copies). If the two proportions are approximately the same, then the prediction was surely non-regressive.

A more general objection may question the basic idea that predictions should be regressive. The expert could point out, correctly, that the present procedure will usually yield conservative predictions that are not far from the average of the class, and is very unlikely to predict an exceptional outcome, which lies beyond all previously observed values. The answer to this objection is that a fallible predictor can retain a chance to correctly predict a few exceptional outcomes only at the cost of erroneously identifying many other cases as exceptional. Non-regressive predictions over-predict: they are associated with a

substantial probability that any high prediction is an overestimate and any low prediction is an underestimate. In most situations, this bias is costly, and should be eliminated.

4. THE OVERCONFIDENCE EFFECT

A forecaster is often required to provide, in addition to his best estimate of a quantity, some indication of confidence in his estimate or, equivalently, some expression of his uncertainty about the value of the quantity, (see, e.g., Spetzler and Staël von Holstein, 1975). These judgments can take the form of confidence intervals or probability distributions. To construct confidence intervals, the expert selects a value X_π of the uncertain quantity $\underset{\sim}{X}$, such that he has a probability π that the outcome will fall below X_π, i.e., $P(\underset{\sim}{X} < X_\pi) = \pi$. Values obtained in this manner are called fractiles. A probability distribution can be constructed by assessing fractiles, e.g., X_{01}, X_{25}, X_{50}, X_{99}. The range between symmetric fractiles is called a (symmetric) confidence interval. For example, the interval between X_{01} and X_{99} is the 98% confidence interval: the expert's probability is .98 that the true value will be contained within the interval, and only .02 that it will be below X_{01} or above X_{99}.

Consider, for example, a publisher who attempts to forecast the sales of a new textbook. Suppose he thinks that there is only one chance in 100 that the book will sell less than 3,000 copies (i.e., X_{01} = 3,000), and that there is a probability of .99 that the book will sell less than 25,000 copies (i.e., X_{99} = 25,000). The range between 3,000 and 25,000 is the 98% confidence interval for the number of copies that will be sold. Another expert may select X_{01} = 5,000 and X_{99} = 15,000 for the sales of the same textbook.

The narrower confidence interval of the second expert expresses greater confidence in his ability to predict the sales of the book in question.

Stimulated by the widely cited unpublished work of Alpert and Raiffa (1969), a considerable amount of research has established the existence of a highly consistent bias in the setting of confidence intervals and probability distributions. The bias can be demonstrated by noting the proportion of cases in which the actual value of the uncertain quantity falls outside the confidence interval, in a large number of problems. These cases are called surprises. If the expert's confidence adequately reflects his knowledge, the true value should fall outside the 98% confidence interval (i.e., below X_{01} or above X_{99}) on approximately 2% of problems. If the percentage of surprises is much higher, the judge is said to be overconfident: his confidence intervals are narrower than his knowledge justifies. Conversely, a proportion of surprises that is much lower than the designated value exhibits underconfidence.

A large number of studies recently reviewed by Lichtenstein, Fischhoff and Phillips (1977) have reported considerable overconfidence in the estimation of uncertain quantities. For 98% confidence intervals, where the rate of surprises should be 2%, the actual proportion of surprises is typically above 25%! All one need do to verify this effect is to select a few quantities from a standard almanac (e.g., population of countries, air distance between cities, yearly consumption of various foods), ask a few friends to assess X_{01} and X_{99} for each of these quantities, and record the percentage of surprises.

There is some evidence that the degree of overconfidence increases with ignorance. For example, we found 28% of surprises in assessments of the air distance between New Delhi and Peking, which compares to 15% for assessments of the distance between London and Tel Aviv. The two distances, in fact, are approximately equal. Naturally, the confidence intervals were considerably wider in the former problem, about which our respondents knew little than in the latter problem, about which they knew more. Confidence intervals were too narrow in both problems, as indicated by the high rate of surprises, but overconfidence was much more pronounced in the more difficult question.

It appears that overconfidence does not occur when the expert has considerable information about the conditional distribution of the outcomes. In extensive studies of confidence intervals given by weather forecasters for the temperature on the next day, Murphy and Winkler (1974, 1977) found that the proportion of surprises corresponded quite precisely to the designated probabilities. We believe that this exception to the overconfidence effect is due to the repetitive nature of the situation with which these experts are concerned, and to the availability of feedback about the outcome following each forecast. The recurrence of an identifiable pattern of indicators, which is followed by different outcomes on different occasions, allows the expert to learn the distribution of outcomes which is associated with that pattern. In this case, the forecaster could judge the probability of different outcomes in terms of their relative frequency. Since people are fairly accurate in their perception of relative frequency (see, e.g., Vlek, 1970), the overconfidence effect is not expected to occur in essentially repetitive situations.

Few forecasting tasks are likely to offer the scope for frequency learning which is available to the meteorologist. In the absence of such distributional data, confidence intervals can only be assessed on the basis of singular information, and overconfidence prevails.

Psychological studies of judgment under uncertainty implicate several factors that contribute to the overconfidence effect. First, people are not sufficiently sensitive to some factors that determine the quality of evidence, e.g., the amount and the reliability of the available information, and often express high confidence in predictions that are based on small samples of unreliable data. Studies of naive and sophisticated respondents (Tversky and Kahneman, 1971; Kahneman and Tversky, 1972) showed that the confidence in conclusions based on sample data did not vary sufficiently with the size of the sample. Similarly, it has been shown that people predict a person's occupation with unwarranted confidence from a brief and unreliable description of his personality (Kahneman and Tversky, 1973). Apparently, sample size and reliability have little impact on judgments of confidence, contrary to the normative principles of statistics.

Insensitivity to the quality of evidence could help explain the overconfidence effect. In many problems of prediction and estimation, available information is limited, incomplete, and unreliable. If people derive almost as much confidence from poor data as from good data, they are likely to produce overly narrow confidence intervals when their information is of inferior quality. That is, they will have too much confidence in the statement that the actual value of the uncertain quantity is included in a narrow range around

the best estimate. This account is supported by the observation that overconfidence is reduced when one has more information about a particular problem, i.e., when the quality of the evidence is high. In fact, overconfidence could disappear in the presence of a large quantity of reliable data.

Oversensitivity to the consistency of available data is a second cause of overconfidence. People tend to draw more confidence from a small body of consistent data than from a much larger body of less consistent data. For example, we instructed subjects to predict students' class standing on the basis of grades obtained in the freshman year. Our subjects made essentially the same prediction on the basis of a single B in one course and on the basis of A in one course and C in another. However, they expressed much more confidence in predicting from a single grade than from an inconsistent pair of grades, a pattern which is not readily justified on statistical grounds. Similarly, we suspect that the public is likely to have more confidence in a conclusion that was unanimously supported by a panel of three experts than in a conclusion that was supported by ten experts in a panel of twelve. This pattern is also difficult to justify.

The effect of consistency indirectly contributes to overconfidence. In their search for coherence, people often see patterns where none exist, reinterpret data so as to increase their apparent consistency, and ignore evidence that does not fit their views. In this manner, people are likely to overestimate the consistency of data, and to derive too much confidence from them.

Two additional factors that contribute to overconfidence in the assessment of uncertain quantities are conditionality and anchoring. Conditionality refers to the adoption of unstated assumptions regarding the assessed quantity. An expert who attempts to estimate the future revenue of a firm, for example, typically assumes normal operating conditions, and may not take into account the possibility that these conditions could change because of war, depression or sabotage. Indeed, experts often claim that their expertise is limited to normal conditions, and that if these conditions are drastically altered "all bets are off". A probability distribution that is conditioned on restrictive assumptions reflects only part of the existing uncertainty regarding the quantity, and is therefore likely to yield too many surprises.

Anchoring refers to the biasing effect of an initial value on subsequent judgments. When constructing a probability distribution over a quantity, one normally considers a best guess before assessing extreme fractiles. The best guess therefore acts as an anchor, and the extreme fractiles, e.g., X_{01} and X_{99}, are pulled toward it. This common bias further contributes to the setting of confidence intervals that are overly narrow.

4.1 Debiasing Confidence Intervals

Because the choice of action is often sensitive to the possibility of extreme outcomes, the best estimate of an uncertain quantity may be less relevant to decision making than the 98% confidence interval. The presence of a large overconfidence bias in the setting of such intervals implies

that the element of uncertainty is typically underestimated in risky decisions. The elimination of overconfidence is therefore an important objective in an attempt to improve the quality of the intuitive judgments that serve decision making.

The preceding analysis of overconfidence suggests that this effect may be quite difficult to overcome. Merely acquainting people with the phenomenon and exhorting them to "spread those extreme fractiles!" does little to reduce the bias (Alpert and Raiffa, 1969). The attempt to do so may destroy the intuitive basis for the initial judgment, without substituting an alternative for it: how is one to know how far the extreme fractiles should be spread? Indeed, the overconfidence effect may be too large to yield to such blandishments. For most people, a change in the probability of an event from .02 to .30 is a qualitative shift, which alters the character of the event from very unlikely to fairly probable. Since this is the magnitude of the shift that is required to abolish overconfidence, the basic view of the problem must be modified for the corrected fractiles to be intuitively acceptable. This becomes vividly evident when one first constructs a probability distribution, then attempts to reallocate 30% of the total area of the distribution outside the original 98% confidence interval. The attempt could induce a sense of confusion, a loss of any confident intuition about the problem and a tendency to wild guessing.

What can be done, then, to eliminate the overconfidence bias in intuitive assessments of confidence intervals and probability distributions? The most radical suggestion is to replace such assessments by computations. This is sometimes possible, when appropriate information is available.

Consider, for example, a publisher who wishes to estimate the 90% confidence interval for the sales of a new textbook. Instead of making an intuitive estimate of the interval, which is likely to be too narrow, the publisher could proceed as follows. First, he should assess X_{05} and X_{95} for the distribution of sales of textbooks in the appropriate reference class. This provides a 90% confidence interval for the class. The width of the 90% confidence interval for the particular book can now be estimated from the width of the corresponding interval for the class.

The statistical theory of prediction entails a simple relation between a confidence interval for an individual case and the corresponding confidence interval for the reference class. This relation is mediated by predictability, i.e., by the correlation between predictions and outcomes (e.g., between predicted and actual sales). Under standard assumptions (e.g., linear regression, normal distributions) the width of the confidence interval for an individual case is $c\sqrt{1 - \rho^2}$ where c is the width of the interval for the class and ρ is the correlation between predicted and actual values. Thus, if one has assessed ρ to be .40, the interval between X_{05} and X_{95} for the particular book should be 92% of the interval between the corresponding fractiles in its class ($.92 = \sqrt{1 - .40^2}$). Many of our students find this statistical relation counter-intuitive: a gain of 8% in precision is smaller than would be expected on the basis of a correlation of .40.

The computational procedure that was illustrated for the prediction of the sales of a book is applicable, in

principle, whenever the statistical assumptions are met at least approximately, and when there are sufficient distributional data. If the relevant data are sparse, the assessment of extreme fractiles cannot be reliable. The main advantage of this procedure is that it relies on an assessment of the distribution in the reference class, which is likely to be more precise and less biased than intuitions about a particular case.

In a less radical vein, the computational approach can provide a check on subjective probability distributions obtained in the standard manner. When an expert who admits that predictability is low sets confidence intervals for a particular case that are much narrower than corresponding intervals for the reference class, there are strong grounds to suspect that he is overconfident. In such cases, the analyst would do well to suggest to the expert that his confidence interval should be bracketed between his initial assessment for the case and his estimate for the class.

The procedures of the debiasing of confidence intervals and for the correction of non-regressive predictions share the same rationale. The need for correction arises in both cases because of the inadequate sensitivity of intuition to considerations of predictability. The suggested procedures involve an assessment of predictability and the explicit use of distribution data. The corrections consist of regressing the expert's intuitive best guess toward the average of the reference class, and expanding his intuitive confidence interval toward the corresponding interval for the class.

5. CONCLUDING REMARKS

Our view of forecasting rests on the following notions. First, that most predictions and forecasts contain an irreducible intuitive component. Second, that the intuitive predictions of knowledgeable individuals contain much useful information. Third, that these intuitive judgments are often biased in a predictable manner. Hence, the problem is not whether to accept intuitive predictions at face value or to reject them, but rather how they can be debiased and improved.

The analysis of human judgment shows that many biases of intuition stem from the tendency to give little or no weight to certain types of information, e.g., the base-rate frequency of outcomes and their predictability. The strategy of debiasing which has been presented in this paper attempts to elicit from the expert relevant information which he would normally neglect, and to help him integrate this information with his intuitive impressions in a manner that respects basic principles of statistical prediction. This approach has been illustrated in an analysis of two tasks, the prediction of uncertain values and the assessment of confidence intervals. The basic approach of adapting procedures of forecasting and decision-making to the recognized limitations of human judgment could be extended to many other activities, such as the evaluation of evidence from multiple sources, the design of effective communication between expert and decision-maker, and the weighting of advantages and disadvantages of alternative policies.

6. REFERENCES

Alpert, M. and Raiffa, H. A Progress Report on the Training of Probability Assessors. Unpublished manuscript, 1969.

Bar-Hillel, M. On the Subjective Probability of Compound Events. Organizational Behavior and Human Performance, 1973, 9: 396-406.

Brown, R.V., Kahr, A.S. and Peterson, C. Decision Analysis for the Manager. New York: Holt, Rinehart and Winston, 1974.

Fischhoff, B. Hindsight and Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 1975, 1: 288-299.

Kahneman, D. and Tversky, A. Subjective Probability: A Judgment of Representativeness. Cognitive Psychology, 1972, 3: 430-454.

Kahneman, D. and Tversky, A. On the Psychology of Prediction. Psychological Review, 1973, 80: 237-251.

Kidd, J.B. The Utilization of Subjective Probabilities in Production Planning. Acta Psychologica, 1970, 34: 338-347.

Lichtenstein, S.C., Fischhoff, B. and Phillips, L. Calibration of Probabilities: The State of the Art. In Proceedings of the Fifth Research Conference on Subjective Probability, Utility, and Decision Making. Jungermann, H. and deZeeuw, G. (Eds.), 1977, in press.

Murphy, A.H. and Winkler, R.L. Subjective Probability Forecasting Experiments in Meteorology: Some Preliminary Results. Bulletin of the American Meteorological Society, 1974, 55: 1206-1216.

Murphy, A.H. and Winkler, R.L. The Use of Credible Intervals in Temperature Forecasting: Some Experimental Results. In Proceedings of the Fifth Research Conference on Subjective Probability, Utility, and Decision Making. Jungermann, H. and deZeeuw, G. (Eds.), 1977, in press.

Ross, N. The Intuitive Psychologist and His Shortcoming: Distortions in the Attribution Process. In Berkowitz, L. (Ed.) Advances in Social Psychology, 1977, in press.

Slovic, P., Fischhoff, B. and Lichtenstein, S.C. Behavioral Decision Theory. Annual Review of Psychology, 1977, in press.

Spetzler, C.S. and Staël von Holstein, C.A.S. Probability Encoding in Decision Analysis. Management Science, 1975, 22: 340-358.

Staël von Holstein, C.A.S. Probabilistic Forecasting: An Experiment Related to the Stock-Market. Organizational Behavior and Human Performance, 1972, 8: 139-158.

Tversky, A. and Kahneman, D. The Belief in the "Law of Small Numbers". Psychological Bulletin, 1971, 76: 105-110.

Tversky, A. and Kahneman, D. Availability: A Heuristic for Judging Frequency and Probability. Cognitive Psychology, 1973, 5: 207-232.

Tversky, A. and Kahneman, D. Judgment Under Uncertainty: Heuristics and Biases. Science, 1974, 185: 1124-1131.

Tversky, A. and Kahneman, D. Causal Thinking in Judgment Under Uncertainty. In Logic, Methodology and Philosophy of Science, (Eds.) Butts, R. and Hintikka, J. Dordrecht, Holland: D. Reidel Publishing Co., 1977, in press.

Vlek, C.A.J. Multiple Probability Learning: Associating Events With Their Probabilities of Occurrence. Acta Psychologica, 1970, 33: 207-232.

Zieve, L. Misinterpretation and Abuse of Laboratory Tests By Clinicians. Annals of N.Y. Academy of Science, 1966, 134: 563-572.

CONTRACT DISTRIBUTION LIST
(Unclassified Technical Reports)

Director Advanced Research Projects Agency Attention: Program Management Office 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Office of Naval Research Attention: Code 455 800 North Quincy Street Arlington, Virginia 22217	3 copies
Defense Documentation Center Attention: DDC-TC Cameron Station Alexandria, Virginia 22314	12 copies
DCASMA Baltimore Office Attention: Mr. K. Gerasim 300 East Joppa Road Towson, Maryland 21204	1 copy
Director Naval Research Laboratory Attention: Code 2627 Washington, D.C. 20375	6 copies
Office of Naval Research Attention: Code 102IP 800 North Quincy Street Arlington, Virginia 22217	6 copies

Revised September 1977

SUPPLEMENTAL DISTRIBUTION LIST
(Unclassified Technical Reports)

Department of Defense

Director of Net Assessment
Office of the Secretary of Defense
Attention: MAJ Robert G. Gough, USAF
The Pentagon, Room 3A930
Washington, DC 20301

Assistant Director (Net Technical Assessment)
Office of the Deputy Director of Defense Research and Engineering (Test and Evaluation)
The Pentagon, Room 3C125
Washington, DC 20301

Assistant Director (Environmental and Life Sciences)
Office of the Deputy Director of Defense Research and Engineering (Research and Advanced Technology)
Attention: COL Henry L. Taylor
The Pentagon, Room 3D129
Washington, DC 20301

Director, Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Director, Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Director, ARPA Regional Office (Europe)
Headquarters, U.S. European Command
APO New York 09128

Director, ARPA Regional Office (Pacific)
Staff CINCPAC, Box 13
APO San Francisco 96610

Dr. Don Hirta
Defense Systems Management School
Building 202
Ft. Belvoir, VA 22060

Chairman, Department of Curriculum Development
National War College
Ft. McNair, 4th and P Streets, SW
Washington, DC 20319

Defense Intelligence School
Attention: Professor Douglas E. Hunter
Washington, DC 20374

Vice Director for Production
Management Office (Special Actions)
Defense Intelligence Agency
Room 1E863, The Pentagon
Washington, DC 20301

Command and Control Technical Center
Defense Communications Agency
Attention: Mr. John D. Hwang
Washington, DC 20301

Department of the Navy

Office of the Chief of Naval Operations (OP-987)
Attention: Dr. Robert G. Smith
Washington, DC 20350

Office of the Chief of Naval Operations (OP-951)
Washington, DC 20450

Office of Naval Research
Assistant Chief for Technology (Code 200)
800 N. Quincy Street
Arlington, VA 22217

Office of Naval Research (Code 230)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Naval Analysis Programs (Code 431)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Operations Research Programs (Code 434)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research (Code 436)
Attention: Dr. Bruce MacDonald
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Information Systems Program (Code 437)
800 North Quincy Street
Arlington, VA 22217

Director, ONR Branch Office
Attention: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Director, ONR Branch Office
Attention: Dr. J. Lester
495 Summer Street
Boston, MA 02210

Director, ONR Branch Office
Attention: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Director, ONR Branch Office
Attention: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Office of Naval Research
Scientific Liaison Group
Attention: Dr. M. Bertin
American Embassy - Room A-407
APO San Francisco 96503

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps (Code RD-1)
Washington, DC 20380

Headquarters, Naval Material Command
(Code 0331)
Attention: Dr. Heber G. Moore
Washington, DC 20360

Head, Human Factors Division
Naval Electronics Laboratory Center
Attention: Mr. Richard Coburn
San Diego, CA 92152

Dean of Research Administration
Naval Postgraduate School
Attention: Patrick C. Parker
Monterey, CA 93940

Naval Personnel Research and Development
Center
Acquisition and Initial Service (Code 310)
Attention: Dr. Leonard P. Kroeker
San Diego, CA 92152

Naval Personnel Research and Development
Center (Code 305)
Attention: LCDR O'Bar
San Diego, CA 92152

Navy Personnel Research and Development
Center
Manned Systems Design (Code 311)
Attention: Dr. Fred Muckler
San Diego, CA 92152

Naval Training Equipment Center
Human Factors Department (Code N215)
Orlando, FL 32813

Naval Training Equipment Center
Training Analysis and Evaluation Group
(Code N-00T)
Attention: Dr. Alfred F. Smode
Orlando, FL 32813

Director, Center for Advanced Research
Naval War College
Attention: Professor C. Lewis
Newport, RI 02840

Naval Research Laboratory
Communications Sciences Division (Code 5403)
Attention: Dr. John Shore
Washington, DC 20375

Dean of the Academic Departments
U.S. Naval Academy
Annapolis, MD 21402

Chief, Intelligence Division
Marine Corps Development Center
Quantico, VA 22134

Commander, U.S. Army Research Office (Durham
Box CM, Duke Station
Durham, NC 27706

Department of the Army

Deputy Under Secretary of the Army
(Operations Research)
The Pentagon, Room 2E621
Washington, DC 20310

Director, Army Library
Army Studies (ASDIRS)
The Pentagon, Room 1A534
Washington, DC 20310

U.S. Army Research Institute
Organizations and Systems Research Laboratory
Attention: Dr. Edgar M. Johnson (PERI-OK)
1300 Wilson Boulevard
Arlington, VA 22209

Director, Organizations and Systems
Research Laboratory
U.S. Army Institute for the Behavioral
and Social Sciences
1300 Wilson Boulevard
Arlington, VA 22209

Technical Director, U.S. Army Concepts
Analysis Agency
8120 Woodmont Avenue
Bethesda, MD 20014

Director, Strategic Studies Institute
U.S. Army Combat Developments Command
Carlisle Barracks, PA 17013

Commandant, Army Logistics Management Center
Attention: DRXMC-LS-SCAD (ORSA)
Ft. Lee, VA 23801

Department of Engineering
United States Military Academy
Attention: COL A. F. Grum
West Point, NY 10996

Chief, Studies and Analysis Office
Headquarters, Army Training and Doctrine
Command
Ft. Monroe, VA 23351

Department of the Air Force

Assistant for Requirements Development
and Acquisition Programs
Office of the Deputy Chief of Staff for
Research and Development
The Pentagon, Room 4C331
Washington, DC 20330

Air Force Office of Scientific Research
Life Sciences Directorate
Building 410, Bolling AFB
Washington, DC 20332

Commandant, Air University
Maxwell AFB, AL 36112

Chief, Systems Effectiveness Branch
Human Engineering Division
Attention: Dr. Donald A. Topmiller
Wright-Patterson AFB, OH 45433

Director, Advanced Systems Division
(AFHRL/AS)
Attention: Dr. Gordon Eckstrand
Wright-Patterson AFB, OH 45433

Commander, Rome Air Development Center
Attention: Mr. John Atkinson
Griffis AFB
Rome, NY 13440

IRD, Rome Air Development Center
Attention: Mr. Frederic A. Dion
Griffis AFB
Rome, NY 13440

HQS Tactical Air Command
Attention: LTCOL David Dianich
Langley AFB, VA 23665

Other Government Agencies

Chief, Strategic Evaluation Center
Central Intelligence Agency
Headquarters, Room 2G24
Washington, DC 20505

Director, Center for the Study of
Intelligence
Central Intelligence Agency
Attention: Mr. Dean Moor
Washington, DC 20505

Office of Life Sciences
Headquarters, National Aeronautics and
Space Administration
Attention: Dr. Stanley Deutsch
600 Independence Avenue
Washington, DC 20546

Other Institutions

Department of Psychology
The Johns Hopkins University
Attention: Dr. Alphonse Chapanis
Charles and 34th Streets
Baltimore, MD 21218

Institute for Defense Analyses
Attention: Dr. Jesse Orlansky
400 Army Navy Drive
Arlington, VA 22202

Director, Social Science Research Institute
University of Southern California
Attention: Dr. Ward Edwards
Los Angeles, CA 90007

Perceptronics, Incorporated
Attention: Dr. Amos Freedy
6271 Variel Avenue
Woodland Hills, CA 91364

Stanford University
Attention: Dr. R. A. Howard
Stanford, CA 94305

Department of Psychology
Montgomery College
Attention: Dr. Victor Fields
Rockville, MD 20850

Director, Applied Psychology Unit
Medical Research Council
Attention: Dr. A. D. Baddeley
15 Chaucer Road
Cambridge, CB 2EF
England

Department of Psychology
Brunel University
Attention: Dr. Lawrence D. Phillips
Uxbridge, Middlesex UB8 3PH
England

Decision Analysis Group
Stanford Research Institute
Attention: Dr. Miley W. Merkhofer
Menlo Park, CA 94025

Human Factors Research, Incorporated
Santa Barbara Research Park
Attention: Dr. Robert R. Mackie
6780 Cortona Drive
Goleta, CA 93017

Wharton School
University of Pennsylvania
Attention: Dr. H. L. Morton
Philadelphia, PA 19174

Decision Research
1201 Oak Street
Eugene, OR 97401

Department of Psychology
University of Washington
Attention: Dr. Lee Roy Beach
Seattle, WA 98195

Analytics, Incorporated
Attention: Mr. William Martin
7405 Colshire Drive, Suite 200
McLean, VA 22101

Eclectech Associates, Incorporated
Attention: Mr. Alan J. Pesch
Post Office Box 179
North Stonington, CT 06359

Department of Electrical and Computer
Engineering
University of Michigan
Attention: Professor Kan Chen
Ann Arbor, MI 94135

Department of Government and Politics
University of Maryland
Attention: Dr. Davis B. Bobrow
College Park, MD 20747

CACI, Incorporated-Federal
Attention: Dr. Bertram Spector
1815 North Fort Myer Drive
Arlington, VA 22209

Department of Psychology
Hebrew University
Attention: Dr. Amos Tversky
Jerusalem, Israel

Dr. T. Owen Jacobs
P.O. Box 3122
Ft. Leavenworth, KA 66027

Dr. Gary Ireland
500 23rd Street, NW - Apt. B301
Washington, DC 20037

Dr. Andrew P. Sage
School of Engineering and Applied Science
University of Virginia
Charlottesville, VA 22901

Professor Raymond Tantor
Political Science Department
The University of Michigan
Ann Arbor, MI 48109

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: PTR-1042-77-6

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle): 6 INTUITIVE PREDICTION: BIASES AND CORRECTIVE PROCEDURES

5. TYPE OF REPORT & PERIOD COVERED: 9 TECHNICAL REP[redacted]T

6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s): 10 Daniel Kahneman, Amos Tversky

8. CONTRACT OR GRANT NUMBER(s): 15 N00014-76-C-0074

[redacted]

ARPA Order [redacted]-3052

9. PERFORMING ORGANIZATION NAME AND ADDRESS:
Decision Research
A Branch of Perceptronics
1201 Oak Street, Eugene, Oregon 97401

11. CONTROLLING OFFICE NAME AND ADDRESS:
Defense Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, Virginia 22217

12. REPORT DATE: 11 Jun[redacted]77

13. NUMBER OF PAGES: ~~38~~ 12 38p.

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office):
Office of Naval Research, Code 455
800 No. Quincy Street
Arlington, Virginia 22217

15. SECURITY CLASS. (of this report): Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report):
Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES
NONE

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Decision Making
Intuitive Prediction
Judgmental Bias
Judgment Aids

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This paper presents an approach to elicitation and correction of intuitive forecasts, which attempts to retain the valid component of intuitive judgments while correcting some biases to which they are prone. This approach is applied to two tasks that experts are often required to perform in the context of forecasting and in the service of decision making: the prediction of values and the assessment of confidence intervals. The

(continued)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

390664

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. (continued)

analysis of these judgments reveals two major biases: non-regressiveness of predictions and overconfidence. Both biases are traced to people's tendency to give insufficient weight to certain types of information, e.g., the base-rate frequency of outcomes and their predictability. The corrective procedures described in this paper are designed to elicit from experts relevant information which they would normally neglect, and to help them integrate this information with their intuitive impressions in a manner that respects basic principles of statistical prediction.