

AD No. _____
DDC FILE COPY

AD A056667

LEVEL II

TECHNICAL REPORT PTR-1060-77-10

CAUSAL SCHEMATA IN JUDGMENTS UNDER UNCERTAINTY

DECISION RESEARCH • A BRANCH OF PERCEPTRONICS

Amos Tversky
Daniel Kahneman

12

October 1977

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DDC
JUL 24 1978
E

ADVANCED DECISION TECHNOLOGY PROGRAM

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

78

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 Research Institute
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

Dr. Stephen J. Andriole
Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

9 TECHNICAL REPORT, TR-1060-77-10

6 CAUSAL SCHEMATA IN JUDGMENTS UNDER UNCERTAINTY.

10 by
Amos Tversky Daniel Kahneman

Sponsored by

Defense Advanced Research Projects Agency
ARPA Order No. 3469

Under Subcontract from

Decisions and Designs, Incorporated

15 N00014-78-C-01005 ARPA Order-3469

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

11 October 1977

12 45 p.

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

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

78 07 21 040
390 664

slf

SUMMARY

Perhaps the major source of irrationality in risky decisions is the exaggerated response to the contrast between certainty and possibility, and the inadequate sensitivity to probability differences. We hypothesize that this tendency will have a major biasing effect on tradeoffs between attributes that involve both certain and probable changes. This hypothesis has important consequences to strategic decisions and negotiations. In these contexts, one often trades a certain gain or loss (e.g., gaining or losing a military advantage) in exchange for probabilistic gains or losses (e.g., an increase or a decrease in the probability of war). The preceding analysis suggests that probabilistic gains and losses will be undervalued in comparison to sure gains and sure losses. Methods for communicating this bias to decision makers and procedures for eliminating it will be explored.

Decision analysis views subjective probability as a degree of belief, i.e., as a summary of one's state of information about an uncertain event. This concept does not always coincide with the lay interpretation of probability. People sometimes think of the probability of an event as a measure of the propensity of some causal process to produce that event, rather than as a summary of their state of belief. The tendency to regard properties as belonging to the external world rather than to our own state of information characterizes much of our perception. We normally regard colors as properties of objects, not of our visual system, and we treat sounds as external rather than internal events. In a similar vein, people commonly interpret the assertion "the probability of heads on the next toss of this coin is $1/2$ " as a statement about the propensity of the coin to show heads, rather than as a statement about our ignorance regarding the outcome of the next toss.

The interpretation of probability as propensity leads people to base their judgments of likelihood primarily on causal considerations, and to ignore information that does not have causal interpretation produces characteristic errors in judgments of probability. Research to date has investigated this phenomenon in three different contexts. One experiment indicates that the assessment of a conditional probability, $P(A|B)$, is determined mainly by the perceived causal impact of B on A, even when this mode of judgment yields paradoxes and inconsistencies. A second study was concerned with the assessment of the posterior probability of an event, given the base-rate frequency of that event and some additional specific evidence. The results indicated that base-rate information is generally ignored unless it is given a causal interpretation. A third study examined the assessment of the probability of a hypothesis H given two items of information D_1 and D_2 , on the basis of the prior probability $P(H)$, and the conditional probabilities $P(H|D_1)$ and $P(H|D_2)$. It showed that people adopt different rules for combining evidence when the data (D_1 and D_2) are given a causal interpretation.

These results illustrate the discrepancy between intuitive judgments under uncertainty and the normative theory of subjective probability. In the latter, all probability assessments refer to one's degree of belief, and the same formal calculus applies to the combination of all probabilistic data. In contrast, people have several different ways of thinking about probability, and their approach to the combination of probabilistic data depends on the manner in which those data are interpreted. This discrepancy presents a major educational challenge: to help people achieve a synthesis between their natural modes of judgment and the logic of probability.

TABLE OF CONTENTS

	<u>Page</u>
SUMMARY	ii
ACKNOWLEDGMENT	v
1. INTRODUCTION	1-1
2. CAUSAL AND DIAGNOSTIC REASONING	2-1
2.1 Confidence in Causal and Diagnostic Inferences	2-1
2.2 Causal and Diagnostic Interpretations of Events	2-5
2.3 Prediction, Explanation and Revision	2-10
3. ON THE USE AND NEGLECT OF BASE-RATES	3-1
4. REFERENCES	4-1
DISTRIBUTION LIST	D-1
DD FORM 1473	

ACKNOWLEDGMENT

This research was sponsored by the Advanced Research Projects Agency of the Department of Defense and was monitored by Office of Naval Research under Contract N00014-78-C-0100 (ARPA Order No. 3469) under Subcontract from Decisions and Designs, Inc., to Oregon Research Institute and Subcontract 78-072-0722 to Perceptronics, Inc.

1. INTRODUCTION

Many of the decisions we make, in trivial as well as in crucial matters, depend on the apparent likelihood of events such as the keeping of a promise, the success of an enterprise, or the response to an action. Since we generally do not have adequate formal models to compute the probabilities of such events, their assessment is necessarily subjective and intuitive. The manner in which people evaluate evidence to assess probabilities has aroused much research interest in recent years, e.g., Edwards (1968), Slovic (1972), Tversky and Kahneman (1974), Slovic, Fischhoff and Lichtenstein (1977), Kahneman and Tversky (1978). This research has identified several judgmental mechanisms which are associated with characteristic errors and biases. The present paper is concerned with the role of causal thinking in judgments under uncertainty and with some biases that are associated with this mode of thinking.

We investigate judgments of the conditional probability $P(X/D)$ of some target event X , on the basis of some evidence or Data D . For a psychological analysis of the impact of evidence, it is useful to distinguish between different types of relations that the judge may perceive between D and X . If D is commonly perceived as a cause of the occurrence of non-occurrence of X , or as an indication of the presence of such a cause, we refer to D as a causal datum. On the other hand, if X is treated as a possible cause of D , we refer to D as a diagnostic datum. Finally, if the judge fails to see D either as a cause or as a consequence of X , we refer to D as an incidental datum. In a normative analysis, this distinction between causal, diagnostic and incidental data is immaterial, and the impact of data depends solely on their informativeness.

It is a psychological commonplace that people strive to achieve a coherent interpretation of the events that surround them, and that the cognitive organization of events into causal schemata is one of the primary means of achieving such coherence. The classical work of Michotte (1963) provided a compelling demonstration of the irresistible tendency to perceive sequences of events in terms of causal relations, even when the perceiver is fully aware that the relation between these events is incidental and that the imputed causality is illusory. The prevalence of causal schemata in the perception of elementary social relations was highlighted in Heider's (1958) seminal work and the study of causal attribution is one of the foci of contemporary social psychology (Jones et al. 1972; Ross, 1977).

People normally think in terms of causes and effects. They can also invert the normal sequence and reason from consequences to causes. Incidental data which do not enter into causal schemata do not contribute to the coherent organization of events. These general considerations provide a background for the main theme of this paper, that the impact of evidence depends critically on whether it is perceived as causal, diagnostic or incidental. Specifically, we hypothesize (i) that causal data has a greater impact than diagnostic data of equal informativeness, and (ii) that incidental data are given little or no weight, in the presence of causal or diagnostic data.

In the first part of the paper, we compare the effects of causal and diagnostic data, and show that people assign greater impact to causal than to diagnostic data of equal informativeness. We also explore a class of problems where a particular datum has both causal and diagnostic significance, and demonstrate that intuitive assessments of $P(X/D)$ are dominated by the direct causal

impact of D on X, with insufficient regard for diagnostic considerations. The second part of the paper compares the impact of causal and incidental data of equal informativeness. Specifically, we investigate judgments of the posterior probability of an event, on the basis of the base-rate frequency of that event and some additional data. We show that base-rate data that are given a causal interpretation affect judgments, while base-rates that do not fit into a causal schema are given little or no weight.

2. CAUSAL AND DIAGNOSTIC REASONING

2.1 Confidence in Causal and Diagnostic Inferences

A causal schema represents an association between a cause and an effect, in which the cause precedes the effect both logically and temporally. Thus, it is hardly surprising that people find it easier and more natural to follow the normal sequence and reason from causes to consequences than to invert this sequence and reason from consequences to causes. If causal inferences are indeed easier and more natural than diagnostic inferences, then one would expect people to infer effects from causes with greater confidence than causes from effects - even when the effect and the cause actually provide the same amount of information about each other. We tested this hypothesis using two different measures: judgments of conditional probabilities and confidence in the accuracy of predictions.

In one set of questions we asked subjects to compare the two conditional probabilities $P(Y/X)$ and $P(X/Y)$ for a pair of events X and Y such that (i) X is naturally viewed as a cause of Y , and (ii) $P(X) = P(Y)$, i.e., the base-rate probabilities of the two events are equal. The latter condition implies that $P(Y/X) = P(X/Y)$. Our prediction was that most subjects would view the causal relation as stronger than the diagnostic relation, and would erroneously assert that $P(Y/X) > P(X/Y)$.

In another set of questions, we asked subjects to compare their confidence in predictions involving two continuous variables, depending on which of these variables was given and which was to be predicted. Here again, the problems are constructed so that one of the variables is naturally viewed as causal with respect to the other. If the two variables have

similar marginal distributions, there is no valid statistical reason to expect a difference in the accuracy with which one of the variables can be predicted from the other. We hypothesized that most subjects would state that a prediction from cause to effect can be made with greater confidence than a prediction from effect to cause.

The predicted asymmetry between causal and diagnostic inferences was strongly confirmed with both types of questions. The effect is illustrated by the following problems, where the numbers in parentheses indicate the number of subjects who chose each option.

Problem 1: Which of the following events is more probable?

- (a) That a girl has blue eyes if her mother has blue eyes (N = 106)
- (b) That the mother has blue eyes, if her daughter has blue eyes (N = 34)
- (-) The two events are equally probable (N = 35)

Problem 2: In a survey of high-school seniors in a city, the height of boys was compared to the height of their fathers. In which prediction would you have greater confidence?

- (a) The prediction of the father's height from the son's height (N = 15)
- (b) The prediction of the son's height from the father's height (N = 111)
- (-) Equal confidence (N = 24)

Since the distribution of height or eye color is essentially the same in successive generations, the correct answer to both problems is "Equal". Most subjects, however, expressed

greater confidence in the causal than in the diagnostic prediction.

Problems 1 and 2 illustrate a fairly direct causal relationship. Strictly speaking, of course, the color of the mother's eyes is not a cause of her daughter's eye color. It is merely a manifestation of a genetic mechanism which determines the color of the mother's eyes, and could also determine the eye color of her daughter. In common usage, however, it is quite acceptable to say that a boy is tall because his father is 6'4", while the statement that the father is 6'4" because his son is tall is clearly anomalous. There are many other situations in which the earlier manifestation of an underlying causal system is treated as a cause of subsequent manifestations of the same system.

A rather more subtle relation exists between manifestations of the same causal system which vary in the degree to which they directly represent that system. Here again, we expect the more direct manifestation of the underlying cause to serve as a proxy for it, in relation to less direct manifestations of the same system. Consequently, inferences from the more to the less direct expression of an underlying trait or disposition should be made with greater confidence than inferences in the inverse direction. Problems 3 and 4 illustrate the type of questions that were used to test this hypothesis.

Problem 3: Which of the following events is more probable?

- (a) That an athlete won the decathlon, if he won the first event in the decathlon (N = 19)
- (b) That an athlete won the first event in the decathlon, if he won the decathlon (N = 93)
- (-) The two events are equally probable (N = 34)

Problem 4: Two tests of intelligence were administered to a large group of students: a one-hour comprehensive test, and a 10-minute abbreviated version.

In which prediction would you have greater confidence?

- (a) The prediction of a student's score on the short test from his score on the comprehensive test (N = 143)
- (b) The prediction of a student's score on the comprehensive test from his score on the short test (N = 33)
- (-) Equal confidence (N = 17)

Here again, the correct answer is 'Equal' in both problems. In Problem 3, the prior probability that an (unspecified) athlete will win the decathlon is $1/N$, where N is the number of competitors. This is also the prior probability that an unspecified athlete will win the first event. Consequently, the two conditional probabilities must be equal. In Problem 4, the standard assumption of linear regressions entails equal accuracy in the prediction of one test from another. The responses to both problems, however, exhibit a marked preference for one direction of prediction over the other.

Problems 3 and 4 both involve two manifestations of the same underlying trait, which differ in reliability. Victory in the decathlon and victory in a single event are both manifestations of athletic excellence, but the former provides a stronger indication of excellence than the latter. Similarly, performance in intelligence tests reflects an underlying trait of intelligence, and the more comprehensive test provides a more reliable measure of this trait than does the abbreviated version. The results confirm the hypothesis that the prediction from the more reliable to the less reliable manifestation is associated with greater confidence than the inverse prediction.

2.2 Causal and Diagnostic Interpretations of Events

The previous section showed that the impact of causal data on the judged probability of a consequence is greater than the impact of diagnostic data on the judged probability of a cause. The present section investigates questions in which the evidence has both causal and diagnostic significance with respect to the target event. We study the hypothesis that people tend to focus on the causal impact of the data for the future, and tend to neglect their diagnostic implications about the past. We first discuss a class of problems in which the dominance of causal over diagnostic considerations produces inconsistent and paradoxical probability assessments. This type of problem was originally introduced by Turoff (1972) in a discussion of the cross-impact method of forecasting.

Problem 5: (Turoff). Let C be the event that within the next 5 years Congress will have passed a law to curb mercury pollution, and let D be the event that within the next 5 years, the number of deaths attributed to mercury poisoning will exceed 500. Let $\bar{C}$ and $\bar{D}$ denote the negations of C and D , respectively.²

Question: Which of the two conditional probabilities, $P(C/D)$ or $P(C/\bar{D})$, is higher?

Question: Which of the two conditional probabilities, $P(D/C)$ or $P(D/\bar{C})$, is higher?

The overwhelming majority of respondents state that Congress is more likely to pass a law restricting mercury pollution if the death toll exceeds 500 than if it does not, i.e., $P(C/D) > P(C/\bar{D})$. Most people also state that the death toll is likely to reach 500 if a law is enacted within the next five years than if it is not, i.e., $P(D/C) < P(D/\bar{C})$.

² The symbols C, D , etc. are introduced to facilitate the present exposition. They were not presented to the respondents who were given complete verbal descriptions of the events.

These judgments reflect the causal beliefs that a high death toll would increase the pressure to pass an anti-pollution measure, and that a measure would be effective in the prevention of mercury poisoning. In one sample of 68 students, 58 chose the modal answer to both questions. This seemingly plausible pattern of judgments violates the most elementary rules of conditional probability.

Clearly, $P(C/D) > P(C/\bar{D})$ implies $P(C/C) > P(C)$. Furthermore, the inequality

$$P(C/D) = \frac{P(C \& D)}{P(D)} > P(C)$$

holds if and only if $P(C \& D) > P(C) P(D)$ which holds if and only if

$$P(D) < \frac{P(C \& D)}{P(C)} = P(D/C)$$

which in turn implies $P(D/C) > P(D/\bar{C})$, provided $P(C)$ and $P(D)$ are non-zero. Hence, $P(C/D) > P(C/\bar{D})$ implies $P(D/C) > P(D/\bar{C})$, contrary to the prevailing pattern of judgment.

It is easy to construct additional examples of the same type in which people's intuitions violate the probability calculus. Such examples consist of a pair of events, A and B, such that the occurrence of B increases the likelihood of the subsequent occurrence of A, while the occurrence of A decreases the likelihood of the subsequent occurrence of B. For example, consider the following problem.

Problem 6: Let A be the event that before the end of next year, Peter will have installed a burglar alarm system in his home. Let B denote the event that Peter's home will be

burglarized before the end of next year.

Question: Which of the two conditional probabilities,
 $P(A/B)$ or $P(A/\bar{B})$, is higher?

Question: Which of the two conditional probabilities,
 $P(B/A)$ or $P(B/\bar{A})$, is higher?

A large majority of our subjects (56 of 68) stated that $P(A/B) > P(A/\bar{B})$ and that $P(B/A) < P(B/\bar{A})$, contrary to the laws of probability. We interpret this pattern of judgments as another indication of the dominance of causal over diagnostic considerations. To appreciate the nature of the effect, let us analyze the structure of Problem 6.

First consider $P(A/B)$, the conditional probability that Peter will install an alarm system in his home before the end of next year, assuming that his home will be burglarized sometime during this period. The alarm system could be installed either before or after the burglary. The information conveyed by the condition, i.e., the assumption of a burglary, has causal significance with respect to the future and diagnostic significance with respect to the past. Specifically, the occurrence of a burglary provides a cause for the subsequent installation of an alarm system, and it provides a diagnostic indication that the house had not been equipped with an alarm system at the time of the burglary. Thus, the causal impact of the burglary increases the likelihood of the alarm system while the diagnostic impact of the burglary decreases this likelihood. The nearly unanimous judgments that $P(A/B) > P(A/\bar{B})$ indicates that the causal impact of B dominates its diagnostic impact.

Precisely the same analysis applies to $P(B/A)$ - the probability that Peter's house will be burglarized before the end of next year, given that he will have installed an alarm system

sometime during this period. The presence of an alarm system is causally effective in reducing the likelihood of a subsequent burglary; it also provides a diagnostic indication that the occurrence of a burglary could have prompted Peter to install the alarm system. The causal impact of the alarm system reduces the likelihood of a burglary; the diagnostic impact of the alarm system increases this likelihood. Here again, the prevalence of the judgment that $P(B/A) < P(B/\bar{A})$ indicates that the causal impact of A dominates its diagnostic impact. Instead of weighting the causal and the diagnostic impacts of the evidence, people apparently assess the conditional probabilities $P(A/B)$ and $P(B/A)$ primarily in terms of the direct causal effect of the condition, which leads to contradictions in problems of this type.

A salient feature of Turoff's problems is the ambiguity of the temporal relation between the conditioning event and the target event. Even in the absence of temporal ambiguity, however, it is often the case that the conditioning event has both causal and diagnostic significance. The present analysis leads to the hypothesis that assessments of conditional probabilities are dominated by causal considerations, even when the temporal relation between the events is fully specified.

Problem 7: Which of the following two probabilities is higher?

$P(R/H)$ - The probability that there will be rationing of fuel for individual consumers in the US during the 1990's, if you assume that a marked increase in the use of solar energy for home heating will occur during the 1980's.

$P(R/\bar{H})$ - The probability that there will be rationing of fuel for individual consumers in the US during the 1990's, if you assume that no marked increase in the use of solar energy for home heating will occur during the 1980's.

It is perhaps instructive to consider the normative (Bayesian) approach to this problem, in the light of the distinction we have drawn between causal and diagnostic considerations. The event H that there will be a marked increase in the use of solar energy for home heating during the 1980's has both causal and diagnostic significance. The direct causal impact of H on R is clearly negative. Other things being equal, a marked increase in the use of solar energy can only alleviate a fuel crisis in later years. However, a marked increase in the use of solar energy during the 80's also provides a strong indication of an impending energy crisis. In particular, it suggests that fuel prices in the 80's are sufficiently high to make the investment in solar energy for home heating economical for a large number of consumers. High fuel prices in the 80's, in turn, suggest a state of shortage of fossil fuel, which increases the likelihood of fuel rationing in the subsequent decade. Thus, the direct causal impact of H on R reduces the likelihood of R, whereas the diagnostic implications of H indirectly increase the likelihood of R.

Although the question of the relative strength of these factors cannot be settled formally, we contend that the diagnostic implications of H outweigh its causal impact. The amount of fuel that may be saved by the increased use of solar energy for home heating is unlikely to be large enough to avert an impending crisis. On the other hand, the scarcity of fuel which is implied by H is highly indicative of a forthcoming energy crisis. According to this line of reasoning $P(R/H) > P(R/\bar{H})$, where $\bar{H}$ is the negation of H.

The hypothesis of this section, however, was that people generally overweight the direct causal contribution of the conditioning event in assessments of conditional probabilities,

and do not give sufficient weight to its diagnostic significance. This hypothesis entails, in the example of Problem 7, that the stipulation of an increase in the use of solar energy for heating in the 1980's will reduce the judged probability of fuel rationing in the 1900's. Indeed, 68 of 83 respondents stated that $P(R/H) < P(R/\bar{H})$. The same pattern of judgments is observed in other problems of this type, where the indirect diagnostic implications of the condition are in conflict with its direct causal implications. Although this pattern of judgments does not violate the rules of probability, as was the case in Turoff's problems, it reflects, we believe, a common tendency to neglect the diagnostic significance of the conditioning event in judgments of conditional probability.

2.3 Prediction, Explanation and Revision

In the preceding sections we presented some evidence in support of the hypothesis that causal inferences have greater efficacy than diagnostic inferences. First we showed that inferences from causes to consequences are made with greater confidence than inferences from consequences to causes. Second, we showed that when the same data have both causal and diagnostic significance, the former is generally given more weight than the latter in judgments of conditional probability.

We turn now to the more general question of the relation between an image, model or schema of a system, e.g., the energy situation or the personality of an individual, and some outcome or manifestation of that system, e.g., an increased use of solar energy or a display of hostility. The relation between models and outcomes is relevant to several types of judgments,

including prediction, explanation and model-revision. Thus, a person may apply the model to predict the outcome or to assess its probability. He may use the model to explain the occurrence of a particular event or consequence. Finally, he may employ the information provided by the occurrence of a particular event to correct or revise his model.

Prediction and explanation represent two different types of causal inference, while model-revision is a prototype of diagnostic inference. In prediction, the judge selects that outcome which is most congruent with his model of the system. In explanation, the judge identifies those features of the model that are most likely to give rise to the specified outcome. In revision, on the other hand, the judge corrects or completes the elements of the model that are least congruent with the data.

Most inferences in everyday life rely on models or images which are imprecise, incomplete and occasionally incorrect. Moreover, people are commonly willing to acknowledge that their models of systems such as the intentions of a person or the energy situation could be in error. The presence of uncertainty regarding the accuracy of a model has implications for the proper conduct of prediction, explanation and revision. If a model is subject to error, predictions from that model should be moderate or regressive, i.e., they should not greatly depart from base-rate predictions. For instance, one should be more reluctant to predict that a person will engage in a rare or unusual behavior when one's information about the person comes from an unreliable source than when the same information comes from a more reliable source.

Explanations that are based on uncertain models should also be tempered with caution, since the causal factors that are used in the explanation may not exist in reality. Furthermore, explanation in the presence of uncertainty should always be combined with model-revision. For example, if a person engages in an activity that appears incompatible with our impression of his personality, we should seriously consider the possibility that our impression was incorrect, and that it should be revised in the direction suggested by the new data. The greater the uncertainty about the model and the more surprising the behavior, the greater should the revision be. An adequate explanation should take into account the changes in the model that are implied or suggested by the event that is to be explained. From a normative point of view, therefore, explanation in the presence of uncertainty about the model involves both diagnostic and causal inferences.

Previous research has shown that people commonly overpredict from highly uncertain models. For example, subjects confidently predict the professional choice or academic performance of an individual on the basis of a brief personality sketch, even when this sketch is attributed to an unreliable source (Kahneman and Tversky, 1973). The intentions and traits that are inferred from a personality sketch are naturally viewed as causes of such outcomes as professional choice or success in school. The overprediction that is observed in such problems is therefore compatible with the high efficacy of causal data that was discussed in the preceding sections.

In the context of explanation and revision, the strength of causal reasoning and the weakness of diagnostic reasoning are manifest in the great ease with which people construct causal accounts for outcomes which they could not predict, and in the difficulty that they have in revising uncertain models to accomodate new data. It appears easier to assimilate a new fact within an existing causal model than to revise the model in the light of this fact. Moreover, the revisions that are made to accommodate new facts are often minimal in scope and local in character.

To illustrate this notion, we turn to previously unreported observations from an earlier study of intuitive prediction (Kahneman and Tversky, 1973). In that study 114 graduate students in psychology were presented with a paragraph-length description of a graduate student, Tom W., which had alledgedly been written during his senior year in high school by a clinical psychologist, on the basis of projective tests. The following description was given:

"Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense."

The subjects were first asked to predict Tom W.'s field of graduate specialization by ranking nine possibilities in terms of their likelihood. There was a strong consensus among the respondents that Tom W. is most likely to be in Computer Science or Engineering, and least likely to be in Social Sciences and Social Work or in the Humanities and Education. Responses to an additional question also exhibited general agreement that projective tests do not provide a valid source of information for the prediction of professional choice. After completing the prediction task, the subjects were asked the following question.

"In fact, Tom W. is a graduate student in the School of Education and he is enrolled in a special program for training for the education of handicapped children. Please outline very briefly the theory which you consider most likely to explain the relation between Tom W.'s personality and his choice of career."

What is the proper approach to this question? The respondents were faced with an apparent conflict between a hard fact, Tom W.'s choice of career, and a detailed but unreliable description of his personality. The high confidence with which people predict professional choice from personality descriptions implies a belief in a high correlation between personality and vocational choice. This belief, in turn, entails that professional choice is highly diagnostic with respect to personality. In the above example, Tom W.'s vocational choice is unlikely in view of his personality description, and that description is attributed to a source of low credibility. A reasonable diagnostic inference should

therefore lead to a substantial revision of one's image of Tom W.'s character, to make it more compatible with the stereotype of his chosen profession. If one believes that students of special education are generally compassionate, then Tom W.'s professional choice should raise doubts about his having "little feel and little sympathy for other people," as stated in the psychologist's report. An adequate response to the problem should at least raise the possibility that Tom W.'s personality is not as described, and that he is in fact kinder and more humane than his description suggests.

Our subjects did not follow this approach. Only a small minority (21%) even mentioned any reservations about the validity of the description. The overwhelming majority of respondents, including the sceptics, resolved the conflict either by reference to suitably chosen aspects of Tom W.'s description (e.g., his deep moral sense) or by a reinterpretation of the psychological significance of his choice (e.g., as an expression of a need for dominance).

It could be argued that our subjects' failure to revise their image of Tom W. merely reflects the demand characteristics of the task which they had been assigned, namely to "explain the relation between Tom W.'s personality and his choice of career." According to this account, the task naturally interpreted as calling for an attempt to relate Tom W.'s professional choice to the description of his personality without questioning its validity. We believe, however, that the prevalent tendency to treat the image of Tom W. as if it were perfectly valid, in spite of severe doubts, exemplifies a much broader phenomenon: the tendency to explain without revising, even when the model that is used in the explanation is highly uncertain.

In our view, the subjects' responses illustrate the reluctance to revise a rich and coherent model, however uncertain, and the ease with which such a model can be used to explain new facts, however unexpected. We were impressed by the fluency which our respondents displayed in developing causal accounts of Tom W.'s unexpected choice of vocation, and have no reason to believe that they would have been less facile in explaining other unexpected behaviors on his part.

Highly developed explanatory skills probably contribute to the proverbial robustness and stability of impressions, models, conceptions and paradigms in the face of incompatible evidence (Abelson, 1959; Hovland, 1959; Kuhn, 1962; Janis, 1972; Jervis, 1975). The impetus for revising a model can only come from the recognition of an incongruency between that model and some new evidence. If people can explain most occurrences to their own satisfaction with minimal and local changes in their existing conceptions, they will rarely feel the need for drastic revision of these conceptions. In this manner, the fluency of causal thinking inhibits the process of diagnostic revision.

3. ON THE USE AND NEGLECT OF BASE-RATES

When people judge the probability of an event, they usually have access to information of two types: 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 relative frequency of that event in a relevant population. For example, the presenting symptoms of a patient or the results of a laboratory test provide singular information about the presence of a particular disease, while the base-rate frequency of that disease provides distributional information. Although the distinction between singular and distributional data is not always so clear, it serves a useful role in the analysis of intuitive inferences. Note that the present concept of distributional information does not coincide with the Bayesian concept of prior probability. The former is defined by the nature of the data, whereas the latter is defined by the temporal sequence of information acquisition.

Previous research has shown that judgments of probability are often dominated by singular evidence, with little or no regard for distributional data. To illustrate, consider the following problem (Kahneman and Tversky, 1973).

Problem 8: A panel of psychologists interviewed a sample of 30 engineers and 70 lawyers, and summarized their impressions in thumbnail descriptions of those individuals. The following description has been drawn at random from the sample of 30 engineers and 70 lawyers.

"John is a 39-year-old man. He is married and has two children. He is active in local politics. The hobby that he most enjoys is rare book collection. He is competitive, argumentative, and articulate."

Question: What is the probability that John is a lawyer rather than an engineer?

Problem 8 was given to a group 85 respondents. Another group of 86 respondents answered another version of the same problem in which the sample from which John was allegedly drawn was said to consist of 70 engineers and 30 lawyers. The median answer to Problem 8 was .95 in both groups. The base-rate frequency of engineers and lawyers had no effect whatsoever in this problem. In other problems, the manipulation of base-rate had a slight, albeit significant effect. It follows readily from Bayes' rule that probability of John's being a lawyer should be considerably higher when the base-rate of lawyers is high than when it is low. Specifically, the ratio of the posterior odds of the two groups should be $(7/3)^2$ (Kahneman and Tversky, 1973).

In that paper, we related the neglect of base-rate data to the hypothesis that people evaluate the probability that John is a lawyer rather than an engineer by the degree to which he is more representative of the stereotype of lawyers than of the stereotype of engineers. Since the similarity of a thumbnail description to the stereotype of a category is unaffected by the base-rate frequency of the category, a judgment of probability that is based exclusively on similarity or representativeness should be essentially independent of base-rate frequency. While judgments of probability by representativeness are inherently insensitive to base-rate

frequencies, the neglect of base-rate information appears to be a more general phenomenon, which occurs even when considerations of representativeness or similarity are not involved in the assessment of probability.

We now believe that, regardless of whether or not probability is judged by representativeness, base-rate information will be dominated by case data, except when it is given a causal interpretation. We now turn to some elementary demonstrations of this hypothesis in problems where both base-rate and case data are provided in numerical form.

Problem 9: A cab was involved in a hit-and-run accident at night: Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

- (i) 85% of the cabs in the city are Green and 15% are Blue.
- (ii) A witness identified the cab as a Blue cab. The court tested his ability to identify cabs under the appropriate visibility conditions. When presented with a sample of cabs (half of which were Blue and half of which were Green) the witness made correct identifications in 80% of the cases and erred in 20% of the cases.

Question: What is the probability that the cab involved in the accident was Blue rather than Green?

Several hundred subjects have been given slightly different versions of this question. For all versions, the modal and median response was 80%. Thus, the intuitive judgment of probability coincides with the credibility of the witness and ignores the relevant base-rate, i.e., the relative frequency of Green and Blue cabs. It is instructive to contrast people's answers with the formal solution of the problem.

To obtain the correct answer, let B and G denote respectively the hypotheses that the cab involved in the accident was Blue or Green, and let W be the witness' report. By Bayes' rule in odds form:

$$\frac{P(B/W)}{P(G/W)} = \frac{P(W/B)P(B)}{P(W/G)P(G)} = \frac{(.8)(.15)}{(.2)(.85)} = \frac{12}{17}$$

and hence $P(B/W) = \frac{12}{12+17} = .41$

In spite of the witness' report, therefore, the hit-and-run cab is less likely to be Blue than Green, because the base-rate is more extreme than the witness is credible. The overwhelming majority of subjects fail altogether to take the base-rate into account.

Base-rate information, however, is properly used in the absence of case data. When item (ii) is omitted from Problem 9, almost all subjects give the correct answer that the probability of the cab being Blue is .15. In the absence of case data the question is naturally viewed as a sampling problem, and the accident serves to define a sampling trial in which a single cab is drawn from the population of cabs in the city. As soon as pertinent case data about the hit-and-run cab is introduced, the base-rate no longer seems relevant. Apparently, people are unable to relate the color of the hit-and-run cab simultaneously to two different types of processes: random sampling of cabs, and imperfect identification of color by a witness.

The neglect of base-rate data is a highly robust effect, which has been confirmed in a variety of contexts ranging from simple questions such as Problem 9 (Ajzen, 1977; Hammerton, 1973; Kahneman and Tversky, 1973; Lyon and Slovic, 1976) to complex realistic problems (Nisbett and Borgida, 1975; Nisbett, Borgida, Crandall and Reed, 1976). In her doctoral dissertation, Maya Bar-Hillel (1975) investigated the neglect of base-rate using a wide variety of questions. Her results, and those of Lyon and Slovic (1976), show that base-rate data have little or no impact regardless of whether they are presented before or after the case data, and regardless of whether they agree or disagree with the case data. There is a small effect of extreme base-rate (e.g., 1:99), although the discrepancy between estimates and correct values is actually most pronounced in these situations. We have found that base-rate data are also neglected when the evidence is presented in verbal rather than numerical form, e.g., "the great majority of cabs are Green," "the witness was correct in most of the cases."

Base-rate data, however, are not always ignored, as demonstrated in the following modification of the accident problem.

Problem 10: A cab was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

(i) Although the two companies are roughly equal in size, 85% of cab accidents in the city involve Green cabs, and 15% involve Blue cabs.

(ii) As in Problem 9 above.

Question: What is the probability that the cab involved in the accident was Blue rather than Green?

Although Problem 10 is formally identical to Problem 9, the responses to the two problems were radically different. To highlight the difference, we compare the responses of a group of 69 subjects who answered Problem 9 to those of another group of 72 subjects who answered Problem 10. The proportion of subjects whose answers coincided with the witness' hit-rate (80%) was .45 in Problem 9 and .18 in Problem 10. The proportions of subjects whose responses coincided with the base-rate (15%) were .07 and .18, respectively for the two problems. The proportion of intermediate responses, which may be taken as evidence for the use of both items, was .35 in Problem 9 and .60 in Problem 10. The median answer to Problem 10 was .55, which is not too far from the correct answer of .41. It is worth noting that the variance of responses to this problem was large, indicating a lack of consensus in the weighting of the base-rate and the singular data.

Problem 10 differs from Problem 9 only in that the base-rate refers to the frequency of accidents rather than to the frequency of cabs. According to our analysis, however, this is a crucial change. The difference between the base-rates of accidents for the Blue and Green companies readily elicits a causal explanation, i.e., that the drivers of the Green cabs are more reckless and less competent than the drivers of the Blue cabs. This property applies to any specific Green cab and increases the perceived likelihood that a Green rather than a Blue cab was involved in the accident. In contrast, the difference between the proportions of Blue and Green cabs in the city does not elicit a causal explanation which makes any specific Green cab more likely to be involved in an accident than any specific Blue cab.

In other words, we suggest that base-rate data will be used when the difference in the base-rate of outcomes is interpreted as a difference in the propensities to produce these outcomes. In Problem 9, the difference in the proportions of Green and Blue cabs is ignored because it cannot be related to the propensity of any particular cab to be involved in an accident. In Problem 10, on the other hand, the difference in the frequency of accidents has impact on judgments because it is interpreted as a difference in the propensity to cause accidents.

Bar-Hillel (1975) has constructed several families of problems to explore the conditions under which base-rate data are utilized or neglected. Her results support the hypothesis that base-rate information which is interpreted as a propensity dominates other base-rate information which cannot be interpreted in this manner. The following question is adopted from one of her studies.

Problem 11: (Bar-Hillel) Consider the following assumptions regarding suicide. In a population of young adults, 80% of the individuals are married and 20% are single. The percentage of deaths by suicide is three times higher among single individuals than among married individuals.

Question: What is the probability that an individual, selected at random from those that had committed suicide, was single?

The correct answer to this problem follows readily from Bayes' rule. Let D denote death by suicide, and let M and S denote, respectively, married and single.

Hence

$$\frac{P(S/D)}{P(M/D)} = \frac{P(D/S)}{P(D/M)} = \frac{(3)(.20)}{(1)(.80)} = \frac{3}{4}$$

and $P(S/D) = 3/7 = .43$. The modal and median answer of a group of 65 subjects was .75. This value corresponds exactly to the ratio of suicide rates among single and married people, and reflects total neglect of the base-rate frequencies of these categories.

Two items of information are supplied in this question: the proportions of married and single individuals, and the ratio of suicide-rates in these categories. The second of these items is readily related to a causal schema of suicide, while the former is not. The differential suicide-rates imply that a single person has a stronger propensity to commit suicide than a married person. In contrast, the proportions of single and married individuals have no bearing on the causation of suicide. As hypothesized, base-rate information which is not linked to a causal schema is ignored in the presence of causally relevant evidence. This result does not depend on prior beliefs regarding suicidal tendencies in the population. When respondents were told that the suicide-rate is higher among married than among single people, the base-rate frequencies of these categories were again ignored (Bar-Hillel, 1975).

If the above account is correct, the base-rate should not be ignored when it is causally related to the target outcome, as in the following problem, which was given to a group of 68 subjects.

Problem 12: Consider the following assumptions regarding suicide. In a population of adolescents, 80% of suicide attempts are made by girls, and 20% by boys. The percentage of suicide attempts that result in death is three times higher among boys than among girls.

Question: What is the probability that an adolescent, selected at random from those who had died by suicide, was a boy?

Problem 12 is formally identical to Problem 11, but the answers to the two problems are different. The proportion of subjects whose answers correspond to the ratio of fatalities was .54 in Problem 11 and .28 in Problem 12. The median answer to Problem 12 was .50, which differs significantly from the median of .75 observed for Problem 11.

Why is the base-rate frequency ignored in Problem 11 but not in Problem 12? According to the present analysis, the base-rate has impact when it can be interpreted as a propensity that is causally related to the target outcome. The proportions of single and married individuals do not affect the propensity of any particular individual to commit suicide, much as the proportions of Blue and Green cabs do not affect the propensity of any particular cab to be involved in an accident. In contrast, the different rates of suicide attempts among boys and girls indicate that girls are more prone than boys to attempt suicide. In addition, the subject is informed in Problem 12 that girls are less likely than boys to die from a suicide attempt. Since both items of information are causally linked to death by suicide, neither item dominates the other. In Problem 11, on the other hand, one item

(the differential suicide rate) has causal character, while the other (the proportions of single and married individuals) does not, and the former dominates the latter.

Another manner in which base-rate data can be incorporated into a causal schema is illustrated in the following problem that was presented to two separate groups of 53 subjects.

Problem 13: Among high school seniors in a given school, 5% (or 40%) were awarded scholarships.

David is a senior in that school. He was described by his counselor as "industrious, intelligent and responsible, does well above average in structured tasks, but lacks intellectual curiosity"

Question: What is the probability that David obtained a scholarship?

The responses to Problem 13 were sensitive to the base-rate frequency of recipients of the scholarship. The median estimate was .25 when the base-rate was 5%, and .50 when the base-rate was changed to 40%. The change of base-rate in this problem alters the perception of the scholarship from one that is hard to get to one that is easy to get, and the estimates of David's chances vary accordingly.

Similar problems were used by Ajzen (1977) in a compelling demonstration of the contrast between causal and non-causal base-rates. In one experiment, subjects were required to assess the probability that a student, for whom a brief description was provided, had passed a particular examination.

Base-rate information was given in different forms to two groups of subjects.

Causal base-rate information: "Two years ago, a final exam was given in a course at Yale University. About 75% of the students passed (failed) the exam."

Non-causal base-rate information: "Two years ago, a final exam was given in a course at Yale University. An educational psychologist interested in scholastic achievement interviewed a large number of students who had taken the course. Since he was primarily concerned with reactions to success (failure), he selected mostly students who had passed (failed) the exam. Specifically, about 75% of the students in his sample had passed (failed) the exam."

Ajzen (1977) notes: "It can be seen that whereas the causal base-rate implies that the examination was easy (75% passed) or difficult (75% failed), the non-causal base-rate had no such implication."

The results indicated that the causal base-rate was much more potent than the non-causal, although both types of base-rates produced significant effects. For the causal base-rate, the judged probability of success (averaged across descriptions) was higher by .34 when the base-rate of success was high than when it was low. For the non-causal base-rate, the corresponding difference was only .12. In the terms of the present analysis, the ease or difficulty of an examination is one of the factors that affect a student's propensity to pass that examination, and it is therefore integrated with other determinants of propensity, such as the student's intelligence and motivation.

The base-rate of success was used in the preceding study to define an examination as easy or hard. In a second study, the base-rate of preferences was used to define options as more or less attractive (Ajzen, 1977). Subjects were required to assess the probability that students, for whom a personality sketch was provided, would choose either history or economics as an elective general-interest course. The causal base-rate, which served to define the relative attractiveness of the two options, consisted of the proportions of students enrolled in the two courses (.70 and .30). The non-causal base-rate was introduced as follows:

"To obtain student reaction, the history (economics) professor recently interviewed 70 students who had taken his general interest course in history (economics). In order to enable comparisons, he also interviewed 30 students who had taken the course in economics (history)."

Note that, unlike the causal base-rate, the non-causal version provides no information about the popularity of the two courses. The effect of the non-causal base-rate was not significant in this study, although there was a probability difference of .025 in the expected direction. In contrast, the causal base-rate had a strong effect: the mean judged probability of choice was .65 for a popular course (high base-rate), and .36 for an unpopular course (low base-rate). Evidently, the attractiveness of courses is inferred from the base-rate of choices and is integrated with personal characteristics in assessing the probability that a particular student will select one course rather than the other.

Ajzen's experiments, as well as Problem 13 demonstrate the effectiveness of base-rate data, in situations where these data serve to specify a causal schema which is otherwise incomplete in an essential respect. Scholarships can be more or less selective, examinations can be easy or hard, and general-interest courses vary in attractiveness. In the absence of pertinent data, people probably complete the appropriate schema of success or choice by a default value, e.g., by assuming a typical level of difficulty for the examination, or equal popularity for the two courses. If the base-rate provides information about difficulty or popularity, this information is used to complete the causal schema and it therefore affects the judged probability that a particular student will obtain the scholarship, pass the examination, or select economics over history. Base-rate data which do not provide information about difficulty or popularity are not incorporated in the causal schema and are therefore given little or no weight in judgments of probability.

Base-rate data which do provide information about the difficulty of an undertaking or the prevalence of an action may nevertheless be neglected, if they conflict with an established prior conception or causal schema. Impressive demonstrations of this phenomenon have been reported by Nisbett and Borgida (1975). They presented subjects with detailed information about the procedure and results of a well-known experimental investigation of people's willingness to help a stranger in distress (Darley and Latane, 1969). In that experiment, six men (of whom one was a confederate of the experimenter) participated in a discussion of personal problems. Under the pretext of maintaining privacy, the participants were seated in separate booths and the conversation was held through an intercom system.

The confederate, who was the first to speak, mentioned that he was prone to seizures. Then he began to stammer, indicated that one of his seizures was coming on, and asked for help. His last comments were "I'm gonna die-- er-- er. I'm, ..., die--er--er--seizure--er-- (choking sounds, then silence)."

The object of the helping experiment was to find out what proportion of participants would try to help, and how soon they would do so. The surprising result was that not one of the fifteen subjects rushed to the confederate when he first asked for help. Four subjects came out of their booths by the end of the speech when the man was choking, and six never came out of their booths at all. Nisbett and Borgida presented these data to their subjects in order to test their effect of predictions of the behavior of individual participants. The subjects were shown brief film clips of interviews, allegedly conducted with three of the participants in the helping experiment. The interviews were not directly pertinent to helping behavior but they provided a general impression of personality. The subjects were then asked to guess how each of the three individuals had behaved in the helping experiment.

Nisbett and Borgida found that knowledge of the distribution of helping responses in the original experiment had no effect. Subjects who knew the results of the helping experiment and subjects who were not given this information made identical predictions concerning the helping behavior of the men shown in the films.

Since these men appeared normal and friendly in the films, informed and uninformed subjects alike tended to guess that they had tried to help the stranger. The knowledge that helping had been infrequent did not affect predictions. Nisbett and Borgida comment "It is interesting to speculate just what kind of monstrous target case these subjects would have had to witness before they would guess that he would behave in a fashion that they knew to be modal" (p. 940).

Exposure to the results of the helping experiment should have led the subjects to realize that rushing out of one's booth to help a stricken stranger is more difficult than it appears to be. Recall that in Ajzen's (1977) experiment, information about the difficulty of an examination affected judgments of probability for individual cases, while Nisbett and Borgida (1975) found no such effect. The apparent conflict between the studies of Ajzen and Nisbett and Borgida suggests the hypothesis that base-rate data which describe the difficulty or attractiveness of an action are used when they complete a schema that is not fully specified, but not when they conflict with an existing schema. In the problems studied by Ajzen, there is no incompatibility between the image of the individual student and any plausible value of the base-rate that specifies the difficulty of an examination or the attractiveness of a course. In contrast, Nisbett and Borgida created a conflict between the callous behavior suggested by the base-rate and common expectations about the behavior of normal and friendly people. The low frequency of helping in the original experiment calls for a significant revision of accepted conceptions of helping behavior. Such revisions, as we have argued, are difficult to make.

Nisbett and Borgida (1975) explained the neglect of base-rate information by the observation that this "information is, almost by its very nature, abstract, pallid and remote. In contrast, target case information is generally concrete, vivid and salient." In support of this contention, Nisbett et al. (1976) subsequently reported a series of ingenious demonstrations in which exposure to a few concrete examples had greater impact on judgments than the presentation of a large body of statistical evidence. For example, they reported that hearing brief comments from a few students has greater impact on the choice of courses than reading a detailed survey of student opinions about these courses.

Vividness and concreteness are undoubtedly among the major determinants of the impact of information. We propose, however, that the conditions under which base-rate is used or neglected are best understood in terms of the role of this information in causal schemata. In the cab problem, for example, we argued that the differential base-rate of accidents for two companies of equal size affects judgments because it suggests a difference in accident-proneness between the drivers of the two companies. In contrast, the proportions of Blue and Green cabs does not induce a differential propensity to be involved in accidents and this information is therefore neglected. The critical difference between the two problems involves causality rather than vividness. Similarly, the statement that 70% of students preferred history over economics is hardly more vivid or concrete than the description of the instructor's decision to sample 70 of his students and 30 students who chose another course (Ajzen, 1977). Here again, the essential difference is that the base-rate of preferences has causal significance, while the base-rate of sampling does not.

The preceding discussion suggests the following generalization. Distributional data affect predictions when they induce a causal model which (i) explains the base-rate, and (ii) applies to the individual case. Thus, the notion that the drivers of the Green company are reckless explains the high rate of accidents in this company, and increases the propensity of any Green cab to be involved in an accident. Similarly, the assumption that the history course is attractive explains why the majority of students chose that course, and also increases the propensity of any individual student to prefer this course over another. Distributional information which is not incorporated into a causal schema, either because it is not interpretable as an indication of propensity or because it conflicts with an established schema, is given little or no weight in the presence of singular data.

This conclusion accords with the main thesis developed in this paper, that the impact of evidence on judgments depends critically on whether this evidence can be incorporated into causal schemata. Causal schemata are directional. Consequently, causal data have more impact than diagnostic data and explanation is easier than revision. Causal schemata are specific. Consequently, base-rate information affects judgments about a specific case only if this information is entered into the schema of that case. The directionality and specificity of causal schemata produce major errors in prediction and explanation, which reflect the deep contrast between intuitive reasoning and the normative theory of evidence.

4. REFERENCES

- Abelson, R.P. Modes of resolution of belief dilemmas. Journal of Conflict Resolution, 1959, 3, 343-352.
- Ajzen, I. Intuitive theories of events and the effects of base-rate information on prediction. Journal of Personality and Social Psychology, 1977, in press.
- Bar-Hillel, M. The base-rate fallacy in subjective judgments of probability. Unpublished doctoral dissertation, Hebrew University, 1975.
- Darley, J.M. and Latane, B. Bystander intervention in emergencies: Diffusion of responsibility. Journal of Personality and Social Psychology, 1968, 8, 377-383.
- Edwards, W. Conservatism in human information processing. In B. Kleinmuntz, (Ed.), Formal Representation of Human Judgment. New York: Wiley, 1968, 17-52.
- Hammerton, M. A case of radical probability estimation. Journal of Experimental Psychology, 1973, 101, 252-254.
- Heider, F. The Psychology of Interpersonal Relations. New York: Wiley, 1958.
- Hovland, C.I. Reconciling conflicting results derived from experimental and survey studies of attitude change. American Psychologist, 1959, 14, 8-17.
- Janis, I.L. Victims of Groupthink. Boston: Houghton-Mifflin, 1972.
- Jervis, R. Perception and Misperception in International Relations. Princeton: Princeton University Press, 1975.
- Jones, E.E., Kanouse, D.E., Kelly, H.H., Nisbett, R.E., Valins, S., and Weiner, B. Attribution: Perceiving the Causes of Behavior. Morristown: General Learning Press, 1972.
- Kahneman, D., and Tversky, A. On the psychology of prediction. Psychological Review, 1973, 80, 237-351.
- Kahneman, D., and Tversky, A. Intuitive prediction: Biases and corrective procedures. Management Science, 1978, in press.

Kuhn, T.S. The Structure of Scientific Revolutions. Chicago: University of Chicago Press, 1962.

Michotte, A. The Perception of Causality. New York: Basic Books, 1963.

Nisbett, R.E., and Borgida, E. Attribution and the psychology of prediction. Journal of Personality and Social Psychology, 1975, 32, 932-943.

Nisbett, R.E., Borgida, E., Crandall, R., and Reed, H. Popular induction: Information is not necessarily informative. In J.S. Carroll and J.W. Payne (Eds.) Cognition and Social Behavior. Hillsdale: Lawrence Erlbaum Associates, 1976.

Ross, L. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.) Advances in Experimental Social Psychology. New York: Academic Press, 1977.

Slovic, P., Fischhoff, B. and Lichtenstein, S. Behavioral decision theory. Annual Review of Psychology, 1977.

Turoff, M. An alternative approach to cross-impact analysis. Technological Forecasting and Social Change, 1972, 3, 309-339.

Tversky, A., and Kahneman, D. Judgment under uncertainty: Heuristics and biases, Science, 1974, 185, 1124-1131.

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
Decisions and Designs, Incorporated 8400 Westpark Drive, Suite 600 McLean, Virginia 22101	20 copies

Revised March 1978

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-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
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

Superintendent
Naval Postgraduate School
Attention: R. J. Roland, (Code 52R1)
C³ Curriculum
Monterey, CA 93940

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

Department of the Army

Alan H. Curry
Operations and Management Science Division
U.S. Army Institute for Research in Management Information and Computer Science
730 Peachtree St., N.E. (Suite 900)
Atlanta, Georgia 30308

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
5001 Eisenhower Avenue
Alexandria, VA 22333

Director, Organizations and Systems
Research Laboratory
U.S. Army Institute for the Behavioral
and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

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

Marine Corps Representative
U.S. Army War College
Carlisle Barracks, PA 17013

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

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

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

Deputy Chief of Staff, Plans, and
Operations
Directorate of Concepts (AR/XOCCC)
Attention: Major R. Linhard
The Pentagon, Room 4D 1047
Washington, DC 20330

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

Mr. Richard Heuer
Methods & Forecasting Division
Office of Regional and Political Analysis
Central Intelligence Agency
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 (20 copies)

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

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

Decision Research (20 copies)
1201 Oak Street
Eugene, OR 97401

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

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

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

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

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

Professor Howard Raiffa
Morgan 302
Harvard Business School
Harvard University
Cambridge, MA 02163

Department of Psychology
University of Oklahoma
Attention: Dr. Charles Gettys
455 West Lindsey
Dale Hall Tower
Norma, OK 73069

Institute of Behavioral Science #3
University of Colorado
Attention: Dr. Kenneth Hammond
Room 201
Boulder, Colorado 80309

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER PTR-1060-77-10 ✓	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) CAUSAL SCHEMATA IN JUDGMENTS UNDER UNCERTAINTY		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Amos Tversky Daniel Kahneman		8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0100 ^{2a} Subcontract 78-072-0722
9. PERFORMING ORGANIZATION NAME AND ADDRESS Decision Research A Branch of Perceptronics 1201 Oak Street, Eugene, Oregon 97401		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order #3469
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, Virginia 22217		12. REPORT DATE October 1977
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research, Code 455 800 No. Quincy Street Arlington, Virginia 22217		13. NUMBER OF PAGES
		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) probability assessment heuristics causality judgmental biases		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) In contrast to the normative theory of evidence, where the impact of data is determined solely by their informativeness, we ^{the paper} develop the thesis that the impact of evidence on intuitive judgments of probabilities depends critically on whether it is perceived as causal, diagnostic or incidental. The first part of the paper shows that people assign greater impact to causal data than to diagnostic data of equal informativeness. When the same datum has both causal and diagnostic implications, the former (over)		

20. (Continued)

dominate the latter. The ease with which people explain unexpected facts and the reluctance to revise old conceptions in the light of new facts are related to the dominance of causal over diagnostic reasoning. The second part of the paper analyzed the use and neglect of base-rate data in terms of the role of these data in causal schemata. It is shown that base-rate information which is given a causal interpretation affects judgments, while base-rate information which cannot be interpreted in this manner is given little or no weight.

To appear in: M. Fishbein (Ed.) Progress in Social Psychology. Hillsdale: Lawrence Erlbaum Associates, 1977.