

COLLECTION

BF

1

.07

v. 13

no. 1

UNIVERSITY OF OREGON LIBRARY MAY 2 1974

Oregon Research Institute

Judgment Under Uncertainty: Heuristics and Biases

Amos Tversky
and
Daniel Kahneman

ORI RESEARCH BULLETIN
Vol. 13 No. 1

OREGON RESEARCH INSTITUTE

Judgment Under Uncertainty: Heuristics and Biases

Amos Tversky and Daniel Kahneman

The Hebrew University

Jerusalem

and

Oregon Research Institute

Research Bulletin
Vol. 13 No. 1

July 1973

JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES

Amos Tversky and Daniel Kahneman

The Hebrew University

Jerusalem

**Invited paper for the Fourth Conference on Subjective Probability, Utility,
Decision Making, Rome, September, 1973.**

JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES

Amos Tversky and Daniel Kahneman

The Hebrew University

Jerusalem

Most important decisions are based on beliefs concerning the likelihood of uncertain events such as the outcome of an election, the guilt of a defendant, or the future value of the dollar. These beliefs are usually expressed in statements such as "I think that...", "chances are...", "It is unlikely that...", etc. Occasionally, beliefs concerning uncertain events are expressed in a numerical form as odds or subjective probabilities. What determines such beliefs? How do people assess the likelihood of an uncertain event or the value of an uncertain quantity? The theme of the present paper is that people rely on a limited number of heuristic principles by which they reduce the complex tasks of assessing likelihoods and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they leads to severe and systematic errors.

The intuitive assessment of probability resembles the assessment of perceptual quantities such as distance or size. These judgments are all based on data of limited validity, which is processed according to heuristic rules. For example, the apparent distance of an object is determined in part by its clarity. The more sharply the object is seen, the closer it appears to be. This rule has some validity, because in any given scene the more

distant objects are seen less sharply than nearer objects. However, the reliance on this rule leads to systematic errors in the estimation of distance. Specifically, distances are often overestimated when visibility is poor because the contours of objects are blurred. On the other hand, distances are often underestimated when visibility is good because the objects are sharply seen. Three features of this example are worth noting. (i) People are not generally aware of the rules that govern their impressions: they are normally ignorant of the important role of blur in the perception of distance. (ii) People cannot deliberately control their perceptual impressions: a sharply seen hilltop looks near even if one has learned of the effect of clarity on the perception of distance. (iii) It is possible to learn to recognize the situations in which impressions are likely to be biased, and to deliberately make appropriate corrections. In making a decision to climb a hill, for example, one should consider the possibility that the summit is further than it looks if the day is particularly clear.

A similar analysis applies to the assessment of likelihoods and to the prediction of values. As in the perceptual example, people apply heuristic rules to their fallible impressions. Here too, people are rarely aware of the basis of their impressions, and they have little deliberate control over the processes by which these impressions are formed. However, they can learn to identify the heuristic processes that determine their impressions, and to make appropriate allowances for the biases to which they are liable. The following sections describe three heuristics that are commonly employed to

assess likelihoods and to predict values; enumerate systematic biases to which these heuristics lead; and discuss the applied and theoretical implications of this research.

REPRESENTATIVENESS

Many of the probabilistic questions with which people are concerned belong to one of the following types: What is the probability that an object A belongs to a class B? What is the probability that event A originates from process B? What is the probability that process A will generate an event B? In answering such questions people typically rely on the representativeness heuristic, in which probabilities are evaluated by the degree to which A is representative of B, i.e., by the degree of similarity between them. When A and B are very similar, e.g., when the outcome in question is highly representative of the process from which it originates, then its probability is judged to be high. If the outcome is not representative of the generating process, probability is judged to be low.

For an illustration of judgment by representativeness, consider an individual, Mr. X, who has been described as "meticulous, introverted, meek, solemn", and the following set of occupational roles: farmer, salesman, pilot, librarian, physician. How do people evaluate the likelihood that Mr. X is engaged in each of these occupations, and how do they order the occupations in terms of likelihood? In the representativeness heuristic, one assesses the similarity of Mr. X to the stereotype of each occupational role, and orders the occupations by the degree to which Mr. X is representative of these stereotypes. Research with problems of this type has shown that people

in fact order the occupations by likelihood and by similarity in exactly the same way (1). As will be shown below, this approach to the judgment of likelihood leads to serious biases, because several of the factors that should be considered in assessing likelihood play no role in judgments of similarity.

1. Insensitivity to prior probability of outcomes.

One of the factors that have a major effect on probability but has no effect on representativeness is the prior probability, or base-rate frequency, of the outcomes. In the case of Mr. X. for example, the fact that there are many more farmers than librarians in the population should enter into any reasonable estimate of the probability that Mr. X is a librarian rather than a farmer. Considerations of base-rate frequency, however, do not affect the similarity of Mr. X. to the stereotypes of librarians and farmers. If people evaluate probability by representativeness, therefore, prior probabilities will be neglected. This hypothesis was tested in an experiment where prior probabilities were explicitly manipulated (1). Subjects were shown brief personality descriptions of several individuals, allegedly sampled at random from a group of 100 professionals - engineers and lawyers. The subjects were asked to assess, for each description, the probability that it belonged to an engineer rather than to a lawyer. In one experimental condition, the subjects were told that the group from which the descriptions had been drawn consisted of 70 engineers and 30 lawyers. In another condition, subjects were told that the group consisted of 30 engineers and 70 lawyers. The odds that any particular description belongs to an engineer rather than to a lawyer should be higher in the first condition, where there is a majority of engineers, than in the

second condition, where there is a majority of lawyers. Specifically, it can be shown by applying Bayes' rule that the ratio of these odds should be $(.7/.3)^2 = 5.44$ for each description. In sharp contrast to Bayes' rule, the subjects in the two conditions produced essentially the same probability judgments. Apparently, subjects evaluated the likelihood that a particular description belonged to an engineer rather than to a lawyer by the degree to which this description was representative of the respective stereotypes, with little or no regard for the prior probabilities of the two outcomes.

The subjects correctly utilized prior probabilities when they had no other information. In the absence of a personality sketch they judged the probability that an unknown individual is an engineer to be .7 and .3 respectively, in the two base-rate conditions. However, prior probabilities were effectively ignored when a description was introduced, even when this description was totally uninformative. The responses to the following description illustrate this phenomenon:

Dick is a 30-year old man. He is married with no children. A man of high ability and high motivation, he promises to be quite successful in his field. He is well liked by his colleagues.

This description was intended to convey no information relevant to the question of whether Dick is an engineer or a lawyer. Consequently, the probability that Dick is an engineer should equal the proportion of engineers in the group, as if no description had been given. The subjects, however, judged the probability of Dick being an engineer to be .5 regardless of whether the stated proportion of engineers in the group was .7 or .3.

Evidently, people respond differently when given no evidence and when given worthless evidence (1). When no specific evidence is given - prior probabilities are properly utilized; when worthless evidence is given - prior probabilities are ignored.

2. Insensitivity to sample size.

To evaluate the probability of obtaining a particular result in a sample drawn from a specified population, people typically apply the representativeness heuristic. That is, they assess the likelihood of a sample result (e.g., that the average height in a random sample of ten men will be 6'0") by the similarity of this result to the corresponding parameter (i.e., to the average height in the population of men). The similarity of a sample statistic to a population parameter is unaffected by the size of the sample. Consequently, if probabilities are assessed by representativeness, then the judged probability of a sample statistic will be essentially independent of sample size. Indeed, when subjects assessed the distributions of average height for samples of various sizes, they produced identical distributions. For example, the probability of obtaining an average height greater than 6'0" was assigned the same value for samples of 1000, 100, and 10 men (2). Moreover, subjects failed to appreciate the role of sample size even when it was emphasized in the formulation of the problem. Consider the following question:

A certain town is served by two hospitals. In the larger hospital about 45 babies are born each day, and in the smaller hospital about 15 babies are born each day. As you know, about 50% of all babies are boys. The exact percentage of baby boys, however, varies from day to day. Sometimes it may be higher than 50%, sometimes lower.

For a period of one year, each hospital recorded the days on which more than 60% of the babies born were boys. Which hospital do you think recorded more such days?

- The larger hospital? (21)
- The smaller hospital? (21)
- About the same? (i.e., within 5% of each other) (53).

The values in parenthesis are the number of undergraduate students who chose each of the three answers.

Most subjects judged the probability of obtaining more than 60% boys to be the same in the small and in the large hospital, presumably because these events are described by the same statistic and are therefore equally representative of the general population. In contrast, sampling theory entails that the expected number of days on which more than 60% of the babies are boys is much greater in the small hospital than in the large one, because a large sample is less likely to stray from 50%. This fundamental notion of statistics is evidently not part of people's repertoire of intuitions.

A similar insensitivity to sample size has been reported in judgments of posterior probability, i.e., of the probability that a sample has been drawn from one population rather than from another. Consider the following example:

Imagine an urn filled with balls, of which $\frac{2}{3}$ are of one color and $\frac{1}{3}$ of another. One individual has drawn 5 balls from the urn, and found that 4 were red and 1 was white. Another individual has drawn 20 balls and found that 12 were red and 8 were white. Which of the two individuals should feel more confident that the urn contains $\frac{2}{3}$ red balls and $\frac{1}{3}$ white balls, rather than the opposite? What odds should each individual give?

In this problem, the correct posterior odds are 8 to 1 for the 4:1 sample and 16 to 1 for the 12:8 sample, assuming equal prior probabilities. However, most people feel that the first sample provides much stronger evidence for the hypothesis that the urn is predominantly red, because the proportion of red balls is larger in the first than in the second sample. Here again, intuitive judgments are dominated by the sample proportion and are essentially unaffected by the size of the sample, which plays a crucial role in the determination of the actual posterior odds (2). In addition, intuitive estimates of posterior odds are far less extreme than the correct values. The underestimation of the impact of evidence has been observed repeatedly in problems of this type (3,4). It has been labeled "conservatism."

3. Misconceptions of Chance.

People expect that a sequence of events generated by a random process will represent the essential characteristics of that process even when the sequence is short. In considering tosses of a coin, for example, people regard the sequence HTHTTH to be more likely than the sequence HHHHTT, which does not appear random, and also more likely than the sequence HHHHTH, which does represent the fairness of the coin (2). Thus, people expect that the essential characteristics of the process will be represented, not only globally in the entire sequence, but also locally in each of its parts. A locally representative sequence, however, deviates systematically from chance expectation: it contains too many alternations and too few runs. Another consequence of the same belief is the well-known gambler's fallacy.

After observing a long run of red on the roulette wheel, for example, most people erroneously believe that black is now due, presumably because the occurrence of black will result in a more representative sequence than the occurrence of an additional red. In general, chance is commonly viewed as a self-correcting process where a deviation in one direction induces a deviation in the opposite direction to restore the equilibrium. In fact, deviations are not "corrected" as a chance process unfolds, they are merely diluted.

Misconceptions of chance are not limited to naive subjects. A study of the statistical intuitions of experienced research psychologists (5) revealed a lingering belief in what may be called the "law of small numbers" according to which even small samples are highly representative of the populations from which they are drawn. The responses of these investigators reflected the expectation that a valid hypothesis about a population will be represented by a statistically significant result in a sample - with little regard for its size. As a consequence, the researchers put too much faith in the results of small samples, and grossly overestimated the replicability of such results. This bias has pernicious consequences for the conduct of research: it leads to over-interpretation of findings and to the choice of inadequate sample sizes.

4. Insensitivity to predictive accuracy.

People are sometimes called upon to make numerical predictions, e.g., of the future value of a stock, the demand for a commodity, or the outcome of a football game. Such predictions are often made by representativeness. For example, suppose one is given a description of a company, and is asked

to predict its future profit. If the description of the company is very favorable, a very high profit will appear most representative of that description; if the description is mediocre, a mediocre performance will appear most representative, etc. The degree of favorableness of the description, of course, is unaffected by the reliability of that description or by the degree to which it permits accurate prediction. Hence, if people predict solely in terms of the favorableness of the description, their predictions will be insensitive to the reliability of the evidence and to the expected accuracy of the prediction.

This mode of judgment violates the normative statistical theory according to which the extremity and range of predictions are controlled by considerations of expected accuracy. If expected accuracy is minimal, the same predictions should be made in all cases. Thus, if the descriptions of the various companies, for example, are unrelated to their profits, the same value (e.g., average profit) should be predicted for all companies. If expected accuracy is perfect, the range of predicted values should equal the range of actual values. In general, the greater the expected accuracy, the wider the range of predicted values.

Several studies of numerical predictions have demonstrated that intuitive predictions do not conform to this rule, and that subjects show little or no regard for considerations of expected accuracy (1). In one of these studies, subjects were presented with several paragraphs, each describing the performance of a student-teacher during a particular practice lesson. Some subjects were asked to evaluate the quality of the lesson described in the paragraph in percentile scores, relative to a specified population. Other subjects were asked to predict, also in percentile scores, the standing of each of the

student-teachers five years after the practice lesson. The judgments made under the two conditions were identical. That is, the prediction of a remote criterion, (success of a teacher after five years) was identical to the evaluations of the information on which the prediction was based (the quality of the practice lesson). The students who made these predictions undoubtedly knew that the prediction of teaching competence, on the basis of a single trial lesson five years earlier, can hardly be accurate. Nevertheless, their predictions were as extreme as their evaluations.

5. The illusion of validity.

As we have seen, people often predict by selecting the outcome (e.g., an occupation) that is most representative of the input (e.g., the description of a person). The confidence they have in their prediction depends primarily on the degree of representativeness attained in the prediction (i.e., on the quality of the match between the selected outcome and the input) with little or no regard for the factors that limit predictive accuracy. Thus, people express great confidence in the prediction that a person is a librarian when given a description of his personality which matches the stereotype of librarians, even if the description is scanty, unreliable or outdated. The unwarranted confidence which is produced by a good fit between the predicted outcome and the input information may be called the illusion of validity. This illusion persists even when the judge is aware of the factors that limit the accuracy of his predictions. It is a common observation that psychologists who conduct selection interviews often experience considerable confidence in their predictions, even when they know of the vast

literature that shows selection interviews to be notoriously fallible. The continued reliance on the clinical interview for selection despite repeated demonstrations of its inadequacy amply attests to the strength of this effect.

Given input variables of stated validity, a prediction based on several such variables can achieve higher accuracy when the input variables are independent of each other than when they are redundant or correlated. Redundant input variables generally yield input patterns that appear internally consistent, whereas uncorrelated input variables often yield input patterns that appear inconsistent. The internal consistency of the pattern of inputs (e.g., a profile of scores) is one of the major determinants of representativeness, and hence of confidence in a prediction. Consequently, people tend to have greater confidence in predictions based on redundant input variables than in predictions based on uncorrelated variables (1). Because redundancy among inputs usually increases accuracy and increases confidence, people tend to have most confidence in predictions that are very likely to be off the mark.

6. Misconceptions of Regression.

Suppose a large group of children have been examined on two equivalent versions of an aptitude test. If one selects ten children from among those who did best on one of the two versions, he will find their performance on the second version to be somewhat disappointing, on the average. Conversely, if one selects ten children from among those who did worst on one version, they will be found, on the average, to do somewhat better on the other version. More generally, consider two variables X and Y which have the same distribution. By and large, if one selects individuals whose average score deviates from the mean of X by k units, then their average deviation from the mean of Y will be less than k . These observations illustrate a general phenomenon known as regression toward the mean, which was first documented by Galton over one hundred years ago.

In the normal course of life, we encounter many instances of regression towards the mean, e.g., in the comparison of the height of fathers and sons, of the intelligence of husbands and wives, or of the performance of individuals on consecutive examinations. Nevertheless, people do not develop correct intuitions about this phenomenon. First, they do not expect regression in many contexts where it is bound to occur. Second, when they recognize the occurrence of regression, they typically invent spurious causal explanations for it (1). We suggest that the phenomenon of regression remains elusive because it is incompatible with the belief that the predicted outcome should be maximally representative of the input, and hence that the value of the outcome variable

should be as extreme as the value of the input variable.

The failure to recognize the import of regression can have pernicious consequences, as illustrated by the following observation. In the training of pilots, the successful execution of a complex flight maneuver is likely to be followed by a deterioration on the next attempt, while a poor performance is likely to be followed by an improvement - a standard manifestation of regression toward the mean. This effect will occur even when the instructor does not respond to the trainee's performance. However, people do not recognize regression effects for what they are, and invent unwarranted causal explanations for them. Since flight instructors typically praise the trainee after a good performance, and admonish him after a poor performance, they tend to come to the erroneous and potentially harmful conclusion that verbal rewards are detrimental to learning whereas punishments are beneficial.

In social interaction as well as in intentional training, rewards are typically administered when performance is good and punishments are typically administered when performance is poor. By regression alone, therefore, behavior is most likely to improve after punishment and most likely to deteriorate after reward. Consequently, the human condition is such that, by chance alone, one is most often rewarded for punishing others and most often punished for rewarding them. People are generally not aware of this contingency. In fact, the elusive role of regression in determining the apparent consequences of reward and punishment seems to have escaped the notice of students of this area.

AVAILABILITY

There are situations in which people assess the frequency of a class or the probability of an event by the ease with which instances or occurrences could be brought to mind. For example, one may assess the risk of heart attack among middle aged people by recalling such occurrences among one's acquaintances. Similarly, one may evaluate the probability that a given business venture will fail by imagining various difficulties which it could encounter. This judgmental heuristic is called availability. In general, availability is a useful clue for assessing frequency or probability, because instances of large classes are recalled better and faster than instances of less frequent classes. However, availability is also affected by other factors besides frequency and probability. Consequently, the reliance on availability leads to predictable biases, some of which are illustrated below.

7. Biases due to the retrievability of instances.

When the frequency of a class is judged by the availability of its instances, a class whose instances are easily retrieved will appear more numerous than a class of equal frequency whose instances are less retrievable. In an elementary demonstration of this effect, subjects heard a list of well-known personalities of both sexes and were subsequently asked to judge whether the list contained more names of men than of women. Different lists were presented to different groups of subjects. In some of the lists the men were relatively more famous than the women, and in others the women were relatively more famous than the men. In all lists, the subjects erroneously judged the classes consisting of the more famous personalities to be

the more numerous (6).

In addition to familiarity, there are other factors (e.g., salience) which affect the retrievability of instances. For example, seeing a house burned down will have a greater impact on the subjective probability of such accidents than merely reading about a fire in the local paper. Furthermore, recent occurrences are likely to be relatively more available than earlier occurrences. It is a common experience that the subjective probability of an accident rises temporarily when one sees a car overturned by the side of the road.

8. Biases due to the effectiveness of a search set.

Suppose you sample a word (of three letters or more) at random from an English text. Is it more likely that the word starts with r or that r is its third letter? People approach this problem by recalling words that begin with r (e.g., road) and words that have r in the third position (e.g., car) and assess relative frequency by the ease with which words of the two types come to mind. Because it is much easier to search for words by their first than by their third letter, most people judge words that begin with a given consonant to be more numerous than words in which the same consonant appears in the third position. They do so even for consonants (e.g., r or k) that are actually more frequent in the third position than in the first (6).

Different tasks elicit different search sets. For example, suppose you are asked to rate the frequency with which abstract words (e.g., thought, love) and concrete words (e.g., door, water) appear in written English. A natural way to answer this question is to search for contexts in which the word could

appear. It seems easier to think of contexts in which an abstract concept is mentioned (e.g., 'love' in love stories) than to think of contexts in which a concrete word (e.g., 'door') is mentioned. If the frequency of words is judged by the availability of the contexts in which they appear, abstract words will be judged as relatively more numerous than concrete words. This bias has been observed in a recent study (7) which showed that the judged frequency of occurrence of abstract words was much higher than that of concrete words of the same objective frequency. Abstract words were also judged to appear in a much greater variety of contexts than concrete words.

9. Biases of imaginability.

Sometimes, one has to assess the frequency of a class whose instances are not stored in memory but can be generated according to a given rule. In such situations, one typically generates several instances, and evaluates frequency or probability by the ease with which the relevant instances can be constructed. However, the ease of constructing instances does not always reflect their actual frequency, and this mode of evaluation is prone to biases. To illustrate, consider a group of 10 people who form committees of k members, $2 \leq k \leq 8$. How many different committees of k members can be formed? The correct answer to this problem is given by the binomial coefficient $\binom{10}{k}$ which reaches a maximum of 252 for $k = 5$. Clearly, the number of committees of k members equals the number of committees of $(10 - k)$ members because any elected group of, say, two members defines a

unique non-elected group of 8 members.

A possible way to answer this question without computation is to imagine several committees of k members, and to evaluate the number of such committees by the ease with which they come to mind. Committees of few members, say 2, are more available than committees of many members, say 8. The simplest scheme for the construction of committees is a partition of the group into disjoint sets. One readily sees that it is easy to construct five disjoint committees of 2 members, while it is impossible to generate even two disjoint committees of 8 members. Consequently, if frequency is assessed by imaginability, or by availability for construction, the small committees will appear more numerous than larger committees, in contrast to the correct symmetric bell-shaped function. Indeed, when naive subjects were asked to estimate the number of distinct committees of various sizes, their estimates were a decreasing monotonic function of the committee size (6). For example, the median estimate of the number of committees of 2 members was 70, while the estimate for committees of 8 members was 20 (the correct answer is 45 in both cases).

Imaginability plays an important role in the evaluation of probabilities in real-life situations. The risk involved in an adventurous expedition, for example, is evaluated by imagining contingencies with which the expedition is not equipped to cope. If many such difficulties are vividly portrayed, the expedition can be made to appear exceedingly dangerous, although the ease with which disasters are imagined need not reflect their actual likelihood. Conversely, the risk involved in an undertaking may be grossly underestimated if some possible dangers are either difficult to conceive, or simply do not come to mind.

10. Illusory correlation

Chapman and Chapman (8) have described an interesting bias in the judgment of the frequency with which two events co-occur. They presented naive judges with clinical diagnoses and with test material for several hypothetical patients. Later the subjects estimated the frequency with which each diagnosis (e.g., paranoia or suspiciousness) had been accompanied by various symptoms (e.g., peculiarities in the drawing of the eyes). The subjects markedly overestimated the frequency of co-occurrence of natural associates, such as suspiciousness and peculiar eyes. This effect was labeled illusory correlation. In their erroneous judgments of the data to which they had been exposed, naive subjects "rediscovered" much of the common but unfounded clinical lore concerning the interpretation of the draw-a-person test. The illusory correlation effect was extremely resistant to contradictory data. It persisted even when the correlation between symptom and diagnosis was actually negative, and it prevented the judges from detecting relationships that were in fact present.

Availability provides a natural account for the illusory-correlation effect. The judgment of how frequently two events co-occur could be based on the strength of the associative bond between them. When the association is strong, one is likely to conclude that the events have been frequently paired. Consequently, strong associates will be judged to have occurred frequently together. According to this view, the illusory correlation between suspiciousness and peculiar drawing of the eyes, for example, is due to the fact that suspiciousness is more readily associated with the eyes than with any other part of the body.

Life-long experience has taught us that, in general, instances of

large classes are recalled better and faster than instances of less frequent classes; that likely occurrences are easier to imagine than unlikely ones; and that the associative connections between events are strengthened when they frequently co-occur. As a consequence, man has at his disposal a procedure (i.e., the availability heuristic) for estimating the numerosity of a class, the likelihood of an event or the frequency of co-occurrences, by the ease with which the relevant mental operations of retrieval, construction, or association can be performed. However, as the preceding examples have demonstrated, this valuable estimation procedure is subject to systematic errors.

ADJUSTMENT AND ANCHORING

In many situations, people make estimates by starting from an initial value which is adjusted to yield the final answer. The initial value, or starting point, may be suggested by the formulation of the problem, or else it may be the result of a partial computation. Whatever the source of the initial value, adjustments are typically insufficient (4). That is, different starting points yield different estimates, which are biased towards the initial values. We call this phenomenon anchoring.

11. Insufficient adjustment.

In a demonstration of the anchoring effect people were asked to estimate various quantities, stated in percentages (e.g., the percentage of African countries in the U.N.). For each question a starting value between 0 and 100 was determined by spinning a wheel of fortune in the

subjects' presence. The subjects were instructed to indicate whether the given (arbitrary) starting value was too high or too low, and then to reach their estimate by moving upward or downward from that value. Different groups were given different starting values for each problem. These arbitrary values had a marked effect on the estimates. For example, the median estimates of the percentage of African countries in the U.N. were 25% and 45%, respectively, for groups which received 10% and 65% as starting points. Payoff for accuracy did not reduce the anchoring effect.

Anchoring occurs not only when the starting point is given to the subject but also when the subject bases his estimate on the result of some incomplete computation. A study of intuitive numerical estimation illustrates this effect. Two groups of high-school students estimated, within 5 seconds, a numerical expression that was written on the blackboard. One group estimated the product $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$, while another group estimated the product $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$. To rapidly answer such questions people perform a few steps of computation and estimate the product by extrapolation or adjustment. Because adjustments are typically insufficient, this procedure should lead to underestimation. Furthermore, because the result of the first few steps of multiplication (performed from left to right) is higher in the descending sequence than in the ascending sequence, the former expression should be judged larger than the latter. Both predictions were confirmed. The median estimate for the ascending sequence was 512, while the median estimate for the descending sequence was 2,250. The correct answer is 40,320.

12. Biases in the evaluation of conjunctive and disjunctive events.

In a recent study (9), subjects were given the opportunity to bet on one of two events. Three types of events were used: (i) simple events, e.g., drawing a red marble from a bag containing 50% red marbles and 50% white marbles; (ii) conjunctive events, e.g., drawing a red marble 7 times in succession, with replacement, from a bag containing 90% red marbles and 10% white marbles; (iii) disjunctive events, e.g., drawing a red marble at least once in 7 successive tries, with replacement, from a bag containing 10% red marbles and 90% white marbles. In this problem, a significant majority of subjects preferred to bet on the conjunctive event (the probability of which is .48) rather than on the simple event, the probability of which is .50. Subjects also preferred to bet on the simple event rather than on the disjunctive event which has a probability of .52. Thus, most subjects bet on the less likely event in both comparisons. This pattern of choices illustrates a general finding. Studies of choice among gambles and of judgments of probability indicate that people tend to overestimate the probability of conjunctive events (10) and to underestimate the probability of disjunctive events. These biases are readily explained as effects of anchoring. The stated probability of the elementary event (e.g., of success at any one stage) provides a natural starting point for the estimation of the probabilities of both conjunctive and disjunctive events. Since adjustment from the starting point is typically insufficient, the final estimates remain too close to the probabilities of the elementary events in both cases. Note that the overall probability of a conjunctive event is lower than the probability of each elementary event, whereas the overall probability of a disjunctive event is higher than the probability of each elementary event. As a consequence of anchoring, the overall probability will be overestimated

in conjunctive problems and underestimated in disjunctive problems.

Biases in the evaluation of compound events are particularly significant in the context of planning. The successful completion of an undertaking (e.g., the development of a new product) typically has a conjunctive character: for the undertaking to succeed each of a series of events must occur. Even when each of these events is very likely, the overall probability of success can be quite low when the number of events is large. The general tendency to overestimate the probability of conjunctive events leads to unwarranted optimism in the evaluation of the likelihood that a plan will succeed, or that a project will be completed on time. Conversely, disjunctive structures are typically encountered in the evaluation of risks. A complex system (e.g., a nuclear reactor or a human body) will malfunction if any of its essential components fails. Even when the likelihood of failure in each component is slight, the probability of an overall failure can be high if many components are involved. Because of anchoring, people will tend to underestimate the probabilities of failure in complex systems. Thus, the direction of the anchoring bias can sometimes be inferred from the structure of the event. The chain-like structure of conjunctions leads to overestimation, the funnel-like structure of disjunctions leads to underestimation.

13. Anchoring in the assessment of subjective probability distributions.

For many purposes (e.g., the calculation of posterior probabilities, decision-theoretical analyses) a person is required to express his beliefs about a quantity (e.g., the value of the Dow-Jones on a particular day) in the form of a probability distribution. Such a distribution is usually con-

structed by asking the person to select values of the quantity that correspond to specified percentiles of his subjective probability distribution. For example, the judge may be asked to select a number X_{90} such that his subjective probability that this number will be higher than the value of the Dow-Jones is .90. That is, he should select X_{90} so that he is just willing to accept 9 to 1 odds that the Dow-Jones will not exceed X_{90} . A subjective probability distribution for the value of the Dow-Jones can be constructed from several such judgments corresponding to different percentiles (e.g., X_{10} , X_{25} , X_{75} , X_{99} , etc.)

By collecting subjective probability distributions for many different quantities, it is possible to test the judge for proper calibration. A judge is properly (or externally) calibrated in a set of problems if exactly $\Pi\%$ of the true values of the assessed quantities fall below his stated values of X_{Π} . For example, the true values should fall below X_{01} for 1% of the quantities and above X_{99} for 1% of the quantities. Thus, the true values should fall in the confidence interval between X_{01} and X_{99} on 98% of the problems.

Several investigators (e.g., 11, 12, 13) have obtained probability distributions for many quantities from a large number of judges. These distributions indicated large and systematic departures from proper calibration. In most studies, the actual values of the assessed quantities are either smaller than X_{01} or greater than X_{99} for about 30% of the problems. That is, the subjects state overly narrow confidence intervals which reflect more certainty than is justified by their knowledge about the assessed quantities. This bias is shared by naive as well as sophisticated subjects,

and it is not eliminated by introducing proper scoring rules which provide incentives for external calibration. This effect is readily interpreted as an instance of anchoring. To select X_{90} for the value of the Dow-Jones, for example, it is natural to begin by thinking about one's best estimate of the Dow-Jones and to adjust this value upward. If this adjustment - like most others - is insufficient, then X_{90} will not be sufficiently extreme. A similar anchoring effect will occur in the selection of X_{10} which is presumably obtained by adjusting one's best estimate downwards. Consequently, the confidence interval between X_{10} and X_{90} will be too narrow, and the assessed probability distribution will be too tight. In support of this interpretation it can be shown that the tightness of subjective probability distributions is eliminated by a procedure in which one's best estimate does not serve as an anchor.

Subjective probability distributions for a given quantity (e.g., the Dow-Jones) can be obtained in two different ways. (a) By asking the subject to select values for the Dow-Jones that correspond to specified percentiles of his probability distribution. (b) By asking the subject to assess the probability that the true value of the Dow-Jones will exceed some specified values. The two procedures are formally equivalent and should yield identical distributions. However, they suggest different modes of adjustment from different anchors. In procedure (a), the judge states his answer in units of the assessed quantity, and the natural starting point is his best estimate. In procedure (b) the answers are stated in odds or probabilities and the natural starting point is even odds or a probability of one-half. Anchoring on the starting point in procedure (b) will yield conservative estimates of

odds, i.e., odds that are too close to 1:1 and probability distributions that are too flat.

To contrast the two procedures, a set of 24 quantities (e.g., the air distance New Delhi-Peking) was presented to one group of subjects who assessed either X_{10} or X_{90} for each problem. Another group of subjects received the median judgment of the first group for each of the 24 quantities. They were asked to assess the odds that each of the given values exceeded the true value of the relevant quantity. In the absence of any bias, the subjects in the second group should retrieve the odds specified to the first group, i.e., 9:1. If the subjects in the second group are anchored on even odds, however, their stated odds should be less extreme, i.e., closer to 1:1. Indeed, the median odds stated by this group, across all problems, were 3:1. When the judgments of the two groups were tested for external calibration, it was found that the judgments of the first group were indicative of overly tight probability distributions, in accord with earlier results, whereas the odds stated by the second group were indicative of overly flat probability distributions. This observation suggests the intriguing possibility that an appropriate combination of the two methods could yield properly calibrated probability distributions.

DISCUSSION

The preceding sections described some heuristics that are commonly employed in judgments about uncertain events, and demonstrated several biases to which these judgments are susceptible. In the present section we discuss the nature of these heuristics and biases and their place in the analysis of rational judgment.

The biases with which we are concerned, like perceptual errors and illusions, are characteristic of the cognitive operations by which impressions and judgments are formed. These cognitive biases are distinct from the better-known intrusions of emotional and motivational factors into judgment, such as wishful thinking and the intentional distortions of judgment induced by payoffs and penalties. The biases described in this paper are consequences of the reliance on heuristics such as representativeness and availability, and they are not attributable to motivational considerations. Indeed, several of the severe errors of judgments reported earlier were observed despite the fact that subjects were encouraged to be accurate and were rewarded for the correct answers. For example, the common erroneous belief that there are more words in an English text that begin with r than words in which r is the third letter was not shaken by monetary payoffs for accuracy (6). Similarly, offering the subjects a \$1 bonus for the correct answer did not increase the prevalence of the belief that a daily list of births in which more than 60% of the babies are boys is more likely in a small hospital than in a large one (2).

The reliance on heuristics and the presence of common biases are general characteristics of intuitive judgments under uncertainty. They apply not only to laymen untutored in the laws of probability, but also to experts - when they think intuitively. For example, the tendency to predict the outcome that best represents the individuating information, with insufficient regard for prior probability, has been observed in the intuitive judgments of individuals who had extensive training in statistics (1). Although some common errors (e.g., the gambler's fallacy) are easily avoided by the statistically sophisticated, there is evidence that

the intuitive judgments of experts are prone to similar errors in more intricate and less familiar questions (e.g., the birthday problem).

It is not surprising that heuristics such as representativeness and availability are not discarded, even though they occasionally lead to errors in prediction or estimation. What is perhaps surprising is the failure of people to infer from life-long experience such fundamental statistical rules as regression towards the mean, or the effect of sample size on sampling variability. Although everyone is exposed in the normal course of life to numerous examples from which these rules could have been induced, very few people discover the principles of sampling and regression on their own. The main cause for the failure to develop valid statistical intuitions is that events are normally not coded in terms of the features that are crucial to the learning of statistical rules. Although we encounter many samples of different sizes from the same population (e.g., lines, paragraphs and pages in texts) we rarely compare the statistical properties of such samples, e.g., their average word length. Consequently, we do not have an effective opportunity to discover that, in general, successive pages differ less in average word length than do successive lines. People just do not think about texts in this manner.

When events are coded into natural categories, the probabilities or relative frequencies of these categories are learned without difficulty. It is the lack of an appropriate code that explains why people usually do not detect the biases in their own judgments. A person could conceivably learn whether his probability judgments are externally calibrated by keeping a tally of the proportion of events that actually occur among those to which he assigns

the same probability. However, it is not natural to group events by their judged probability. In the absence of appropriate grouping of events, the only available feedback is whether individual events did or did not occur. This dichotomous feedback provides little information concerning the adequacy of one's judgments of probability. Thus, the failure to realize that judgmental operations are repetitive - even when they apply to unique events - is a major obstacle for effective learning.

Modern decision theory (14,15) regards subjective probability as the quantified opinion of an idealized person. Specifically, the subjective probability of a given event is defined by the set of bets about this event which such a person is willing to accept. An internally consistent, or coherent, subjective probability measure can be derived for an individual if his choices among bets satisfy certain principles (i.e., the axioms of the theory). The derived probability is subjective in the sense that different individuals are allowed to have different probabilities for the same event. Naive or intuitive judgments of probability typically fail to satisfy the necessary axioms. The theory of subjective probability provides a rationale for a procedure in which estimates are modified or corrected to achieve internal consistency. The inherently subjective nature of probability judgments has led many writers to the belief that internal consistency is the only criterion by which judged probabilities should be evaluated. From the standpoint of the formal theory of subjective probability, any set of internally consistent probability judgments is as good as any other.

Our position is that internal consistency alone does not guarantee the adequacy of a set of probability judgments because an internally consistent set of subjective probabilities can be incompatible with other beliefs held by the individual. Consider a person whose subjective probabilities for all possible outcomes of a coin-tossing game reflect the gambler's fallacy. That is, his estimate of the probability of tails on any toss increases with the number of consecutive heads that preceded that toss. The judgments of such a person could be internally consistent and therefore acceptable as adequate subjective probabilities according to the criterion of the formal theory. These probabilities, however, are incompatible with the generally-held belief that a coin has no memory and is therefore incapable of generating sequential dependencies.

For judged probabilities to be considered adequate, or rational, internal consistency is not enough. The judgments must be compatible with the entire web of beliefs held by the individual. Compatibility among beliefs is the essence of rational judgment. This criterion is more stringent than internal consistency but also more appropriate because it requires that a set of judgments be compatible with the judge's entire body of knowledge and not only consistent within itself. Unfortunately, there can be no simple formal procedure for assessing the compatibility of a set of probability judgments with the judge's total system of beliefs. Nevertheless, the rational judge will strive for compatibility, even though internal consistency is more easily achieved and assessed. In particular, he will attempt to make his probability judgments compatible with his knowledge about (i) the subject-matter; (ii) the laws of probability; (iii) his own judgmental heuristics and biases. The present view provides a rationale for a procedure in which judged probabilities are modified or corrected to achieve a higher degree of compatibility with all these types of knowledge.

We began this paper with the question of how people make intuitive judgments of probability. The answer appears to be that such judgments are based on the outcomes of some specified mental operations such as the assessment of representativeness or availability. In the formal theory of subjective probability, the mental operation performed by the idealized judge is a choice between bets. Although subjective probabilities can sometimes be inferred from choices between bets, people normally do not evaluate probabilities in this manner. In fact, judgments of likelihood usually determine preferences among bets and are not derived from them, as in the axiomatic theory of subjective probability (14, 15).

SUMMARY

This paper describes three heuristics, or mental operations, that are employed in judgment under uncertainty. (i) An assessment of representativeness or similarity, which is usually performed when people are asked to judge the likelihood that an object or event A belongs to a class or process B. (ii) An assessment of the availability of instances or scenarios, which is often employed when people are asked to assess the frequency of a class or the plausibility of a particular development. (iii) An adjustment from a starting point, which is usually employed in numerical prediction when a relevant value is available. These heuristics are highly economical and usually effective, but they lead to systematic and predictable errors. A better understanding of these heuristics and of the biases to which they lead could improve judgments and decisions in situations of uncertainty.

REFERENCES AND NOTES

1. D. Kahneman and A. Tversky, On the psychology of prediction, Psychological Review, 1973, in press.
2. D. Kahneman and A. Tversky, Subjective probability: A judgment of representativeness, Cognitive Psychology, 1972, 3, 430-454.
3. W. Edwards, Conservatism in human information processing. In B. Kleinmuntz (Ed.), Formal representation of human judgment. New York: Wiley, 1968. 17-52.
4. P. Slovic and S. Lichtenstein, Comparison of Bayesian and regression approaches to the study of information processing in judgment. Organizational Behavior and Human Performance, 1971, 6, 649-744.
5. A. Tversky and D. Kahneman, The belief in the law of small numbers. Psychological Bulletin, 1971, 76, 105-110.
6. A. Tversky and D. Kahneman, Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 1973, in press.
7. R.C. Galbraith and B.J. Underwood, Perceived frequency of concrete and abstract words. Memory & Cognition, 1973, 1, 56-60.
8. L.J. Chapman and J.P. Chapman, Genesis of popular but erroneous psychodiagnostic observations. Journal of Abnormal Psychology, 1967, 73, 193-204.
L.J. Chapman and J.P. Chapman, Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. Journal of Abnormal Psychology, 1969, 74, 271-280.
9. M. Bar-Hillel, Compounding subjective probabilities, Organizational Behavior and Human Performance, 1973, in press.
10. J. Cohen, E.I. Chesnick and D. Haran, A confirmation of the inertial- ψ effect in sequential choice and decision, British Journal of Psychology, 1972, 63, 41-46.

11. M. Alpert and H. Raiffa, A report on the training of probability assessors. Unpublished manuscript, Harvard University, 1969.
12. C. Staël von Holstein, Two techniques for assessment of subjective probability distributions - an experimental study. Acta Psychologica, 1971, 35, 478-494.
13. R.L. Winkler, The assessment of prior distributions in Bayesian analysis. Journal of the American Statistical Association, 1967, 62, 776-800.
14. L.J. Savage, The foundations of statistics. New York: Wiley, 1954.
15. B. De Finetti, Probability: interpretation, in D.L. Sills (Ed.) International Encyclopedia of the Social Sciences, 1968, 13, 496-504.
16. This study was supported by a grant from the Research and Development Authority of the Hebrew University and by the following grants to the Oregon Research Institute: Grants MH 12972 and MH 21216 from the National Institute of Mental Health, Grant RR 05612 from the National Institute of Health, U.S. Public Health Service; Grant GS 3250 from the National Science Foundation.