

When Do Base Rates Affect Predictions?

Maya Bar-Hillel and Baruch Fischhoff
Decision Research, Eugene, Oregon

Recent studies have shown that when people make predictions, they often neglect base-rate considerations. Instead of considering what typically happens in situations like the one being judged, they rely on the extent to which the judged case is representative of the possible prediction categories—although if representativeness fails to provide a clear guide to prediction, people will resort to base-rate considerations. Manis, Dovalina, Avis, and Cardoze recently argued against this conclusion, presenting as evidence a series of experiments in which subjects predicted the category membership of individuals depicted in each of a set of photos. They found that base rates had a clear effect on discrete predictions (i.e., a majority of the photos were predicted to belong to the larger category) and a smaller effect on the confidence subjects attached to those predictions. However, only a minority of the photos could be readily classified by representativeness. As a result, Manis et al.'s findings can be reinterpreted in a way that makes them compatible with previous findings. In this light, their study emerges as a constructive replication of earlier results demonstrating judgment by representativeness.

A burgeoning area of research is the study of the role of background (or base-rate) information in people's predictions about specific target individuals (Borgida & Brekke, in press; Kassir, 1979). The impetus for much of this work was a study by Kahneman and Tversky (1973) demonstrating that base rates may be largely ignored in the presence of even flimsy specific evidence. They argued that predictions are often governed by judgments of "representativeness." The user of this heuristic predicts that people belong to groups whose prototypes they most resemble. Such predictions are made even when the target individuals are characterized by only a brief, unreliable description and even when they resemble low base-rate groups, that is, groups with few members in the population from which the individuals are drawn.

In the course of this base-rate research, Kahneman and Tversky's original design

and theory have been adapted to and tested in a variety of new situations. A study by Manis, Dovalina, Avis, and Cardoze (1980) affords an opportunity to reassess what Kahneman and Tversky said, how well their conclusions have borne the test of time, and what we now know about the conditions under which people attend to base rates.

The Representativeness Thesis

Kahneman and Tversky's (1973) original studies required subjects to predict the profession of an individual described in a thumbnail personality sketch. Subject's predictions were influenced heavily by the extent to which the description represented the dominant features of the stereotype associated with each profession, with little attention being given to the base rates of the various professions. Subjects who were given no individuating evidence and, hence, could not rely on representativeness utilized the base rate appropriately.

Subsequent studies have attempted to identify additional conditions under which base rates affect predictions, in order to refine our understanding of the representativeness heuristic's use and limits. The cumulative evidence suggests that base rates

This research was supported by the Office of Naval Research under Contract N00012-80-C-0150 to Perceptronics, Inc. Our thanks go to Sarah Lichtenstein, Don MacGregor, Paul Slovic, several anonymous reviewers, and the two editors for their comments on earlier drafts.

Requests for reprints should be sent to Baruch Fischhoff, Decision Research, a branch of Perceptronics, 1201 Oak Street, Eugene, Oregon 97401.

will affect predictions most in the following three situations:

1. *Judgment by representativeness fails to provide a clear guide to prediction.* This can happen in several ways. The prediction categories may evoke no stereotypes, the individuating information may not suggest any of the category stereotypes, or it may suggest all stereotypes equally. For example, Ginosar and Trope (1980) gave subjects a description whose details represented both of the two alternative categories. Unable to make a prediction on the basis of representativeness, subjects relied on base rates. The same pattern emerged when the two alternative categories evoked very similar stereotypes.

2. *The base rates have a causal relationship to the target outcome.* It appears that causally linked base rates are not ignored, since they can be incorporated directly into the prediction scheme, along with representativeness considerations (Ajzen, 1977; Bar-Hillel, 1980; Tversky & Kahneman, 1979). In Kahneman and Tversky's (1973) early studies, the base-rate information was non-causal whereas the specific information lent itself to judgment by representativeness.

3. *More than one value of the base-rate information is considered.* Drawing attention to the base rates apparently encourages subjects to utilize them, which they do with moderate success for the present kind of problem (Fischhoff, Slovic, & Lichtenstein, 1979).

The Probability-Learning Paradigm

Manis et al. (1980) added several new dimensions to the contexts within which base-rate usage has been studied using a modification of the probability-learning paradigm. In that paradigm, subjects must predict the identity of each item in a series drawn from a binomial population. Typically, subjects learn the base rates of the item categories from their trial-by-trial feedback and then "match" them in their predictions (Estes, 1976). That is, prediction proportions come to match sample proportions. Although this response strategy fails to maximize the expected rate of correct predictions (as would happen if subjects al-

ways predicted the most common category), it does show that "the observed base rates . . . have a clear and replicable influence, in contrast to the weak, inconsistent effects reported by Kahneman and Tversky" (Manis et al., 1980, p. 232).

Manis et al. observed that the probability-learning paradigm differs from that used by Kahneman and Tversky in two ways. One is that it presents the base-rate information sequentially rather than simultaneously. The second is that its response mode requires a discrete prediction of category membership rather than a continuous evaluation of category membership probabilities. Manis et al. merged the two paradigms by presenting the base-rate information case by case and eliciting discrete responses.

In each trial of Manis et al.'s Experiment 1, subjects viewed a photograph of a young male, guessed his attitude (pro or con) on an issue such as legalization of marijuana, and then were informed what his "true" position was. This "true position" feedback was assigned at random according to experimenter-determined base rates. Subjects showed little initial preference for predicting either a "pro" or a "con" position. But over the course of 50 trials, the proportion of pro predictions came close to matching the proportion in the sample, thereby replicating the pattern of responses typically found in probability-learning studies.

One important difference between the present study and probability-learning studies makes it unlikely that Manis et al.'s subjects were responding like probability-matching subjects. Probability-learning experiments provide no individuating information. Each trial is generated randomly according to the appropriate base rate. In practice, probability-learning subjects seem to individuate each trial on the basis of the pattern they perceive in preceding trials. In particular, they predict events that will create random-looking sequences. Most frequently, the result is a negative recency effect, which leads subjects to predict events that they have not seen recently (Jarvick, 1951; Tune, 1964).

In Manis et al.'s paradigm, on the other hand, the trials themselves are naturally distinctive. Each presents a different face, one

that may be seen as more or less representative of the possible categories. Manis et al. argue that their subjects relied somewhat on this differential representativeness in allocating category labels. We would go a step further and argue that their subjects relied entirely on differential representativeness whenever that was possible. That is, subjects who encountered a face that looked like a marijuana advocate or opponent predicted as much. It is only when subjects encountered "neutral" photographs (for which representativeness provides no guide) that they were swayed by the base rates. If correct, this claim would cast a rather different light on Manis et al.'s results. A closer look at Manis et al.'s Experiment 1 is needed to see how far its results can support this reinterpretation.

Reinterpreting Experiment 1

In order to assess the differential representativeness of the photographs they used, Manis et al. had a separate group of (no-feedback) subjects categorize each photo as pro or con. When most subjects agreed that a face looked like a supporter (or opponent) of legalizing marijuana, it was considered to be stereotypically pro (or con). For over half of the pictures, however, there was no such consensus among subjects. These could have been "neutral" photos, which seemed neither pro nor con, or "controversial" photos, which seemed pro to some subjects and con to others (in roughly equal proportions). Since subjects could not label photos "neutral," there is no way of knowing how prevalent such judgments would have been. In the discussion that follows, we shall assume that some photos in all three consensus categories seemed neutral to some subjects.

Figure 1 is in part a redrafting of Manis et al.'s Figure 2, relabeled to fit the distinctions made in the preceding paragraph.¹ The solid lines represent the percentages of cases in which subjects predicted that a photo showed a pro person. For all three categories, there was a higher percentage of pro predictions with the high pro base rate. Manis et al. (1980) interpreted this trend to mean that "the base-rate data proved to be equally influential, whether the individuating information (the photograph in question) seemed

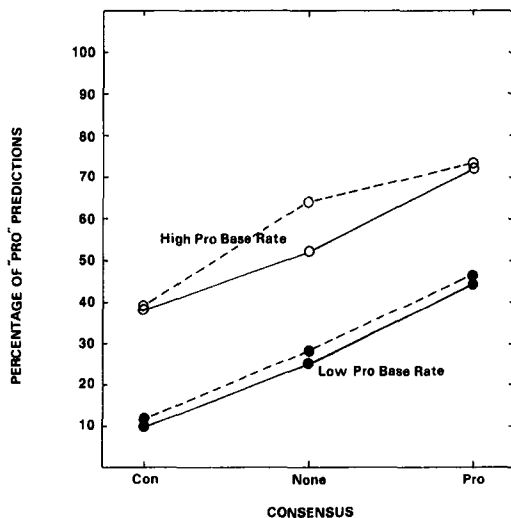


Figure 1. Predictions (solid line = actual, broken line = reconstructed) concerning attitudes toward legalization of marijuana as a function of the observed base rate and the consensual classification of the photos. (From "Base Rates Can Affect Individual Predictions," by Melvin Manis, Ismael Dovalina, Nancy E. Avis, and Steven Cardoze, *Journal of Personality and Social Psychology*, 1980, 38, 231-248, Figure 2. Copyright 1980 by the American Psychological Association. Reprinted by permission. The figure has been relabeled for the sake of consistent usage with other parts of the text.)

to be informative [i.e., differentially representative of the two categories] or not" (p. 235).

The pattern of Manis et al.'s data can, however, be accounted for without assuming any violations of prediction by representativeness. Table 1 substantiates this claim with a hypothetical reconstruction of the judgments underlying Manis et al.'s data. It assumes that the base rate that subjects learn affects predictions only for neutral photos, whereas photos that seem pro or con are judged according to representativeness. Regarding the prevalence of neutral figures, we make the following auxiliary assumptions: (a) 45% of the time photos in the pro consensus and con consensus groups were viewed as neutral (and would have been judged as such had there been a "neutral"

¹ Manis et al. labeled the abscissa "phototype," with the three categories "positive," "negative," and "neutral." They labeled the ordinate "percentage of 'for' predictions."

Table 1
A Hypothetical Distribution of Photo Evaluations Reconstructing the Results of Manis et al.'s (1980) Experiment 1

Variable	Consensual classification of photos		
	Con	None	Pro
% of time photos are perceived to be:			
Con	52	24	18
Neutral	45	60	45
Pro	3	16	37
Total	100	100	100
% of time photos are predicted to be pro:			
Low pro base rate			
Judged by representativeness	3	16	37
Judged by probability matching	9	12	9
Total	12	28	46
High pro base rate			
Judged by representativeness	3	16	37
Judged by probability matching	36	48	36
Total	39	64	73

category), and (b) in the no-consensus group, photos were viewed as neutral 60% of the time.

If subjects probability match when making predictions for these neutral stimuli, they would classify 80% of them in the majority category and 20% in the minority category. In particular, 20% of the neutral photos would be labeled "pro" in the high-pro base-rate condition. This represents 9% of all predictions for photos consensually classified as con ($.20 \times .45$). Since another 3% of the time photos in the con consensus group were judged as pro by representativeness, we would have 12% (3% + 9%) pro predictions for con consensus photos in the low-pro base-rate condition. This 12% corresponds roughly to the 10% observed by Manis et al. The bottom sections of Table 1 carry through analogous computations for the other five photo-classification base-rate combinations. The results that would be produced by this hypothesized process are represented by the broken lines in Figure 1. The similarity between these simulated results and those actually observed indicates that Manis et al.'s results are not inconsistent with consistent reliance on representativeness.²

In this light, the base-rate effect might be

² Two other comments on the details of the simulated results in Table 1 may be pertinent: (a) Manis et al. (1980) report that when a group of subjects who received no feedback performed the prediction task, "fully 70% of the faces were classified as nonpolarized (or neutral), since they elicited inconsistent expectations from different respondents. By contrast, only 53% of the same faces were classified as neutral when considered in the context of the marijuana issue" (p. 234, Footnote). "Neutral" in this sense means not consistently categorized as pro or con. That inconsistency could reflect arbitrary responses by subjects unable to predict by representativeness, or it could reflect deliberate responses of subjects who disagreed about which prediction category was best represented by a particular photo. If we take the percentage of "neutral" photos (in Manis et al.'s sense) as equal to the percentage of the time that no judgment could be made by representativeness ("neutral" in the sense of Table 1), then the percentages in the second row of Table 1 correspond roughly to those observed by Manis et al. (b) The asymmetry between the *perception* of those who were consensually classified as con or pro (i.e., the fact that there was a 52:3 split in one case and 18:37 in the other) was introduced to maintain consistency with Manis et al.'s data. As can be seen from the solid lines in Figure 1, Manis et al.'s con photos were more strongly con than their pro photos were pro. For example, for the cases in which the consensus judgment of the photo matched the base rate, 90% of their con photos were judged con in the high-con (low-pro) base-rate condition, whereas only 72% of pro photos were judged pro in the high-pro base-rate condition. Similar asymmetries can be seen where the consensual photo classification disagreed with the base rate and where there was no consensus.

seen as due primarily to the preponderance of neutral stimuli. If, as this account suggests, neutral photos are the only ones that are responsive to base rates, one would expect stronger base-rate effects to be associated with higher proportions of neutral stimuli. This additional prediction seems to be borne out by Manis et al.'s data. The base-rate effect with the marijuana legalization task was somewhat smaller than that with an analogous task asking subjects to predict whether people in photographs supported or opposed legislation mandating seat belt usage.³ Manis et al. (1980) argued that "there are no clear stereotypes regarding the physical characteristics of those who favor such legislation" (p. 233) and found a higher proportion of no-consensus pictures with that task.

Reinterpreting Experiment 2

The feedback given to Manis et al.'s subjects clearly taught them the base rates. Since that feedback was random, it may also have taught them something about the validity of representativeness as a guide to action, namely, that the differential representativeness of the photos is not as diagnostic a predictor as it might appear. Although that bitter experience does not point to an alternative prediction strategy, it might still reduce subjects' confidence in representativeness-derived predictions. This would be particularly true for predictions of the low base-rate category which tend to encounter negative feedback more often. The effect of feedback on confidence can be tested if confidence is elicited in addition to discrete predictions of category membership. This was done in Manis et al.'s Experiment 2.

In this experiment, Manis et al. presented 40 labeled photographs one by one, ostensibly as a memory task, without requesting predictions. Subjects then guessed the post-graduate plans of each of 10 new cases and indicated confidence in their predictions on a 3-point scale anchored at "guess" and "very certain." Manis et al. translated these responses into "subjective probabilities" that predictions were correct as follows: "Guess" was translated to .51, medium confidence to .755, and "certain" to 1.

Although intended to eliminate the "ex-

PLICIT feedback (reinforcement) following each prediction" (Manis et al., 1980, p. 238) in Experiment 1, this design does not preclude the possibility of self-generated feedback, a possibility that may be almost unavoidable in such tasks (Fischhoff & Slovic, 1980). That is, subjects in the initial phase of Experiment 2, having been told they would later be asked to remember both the faces and their labels, may have been trying to "learn" something about the face-label relationship, both as a mnemonic device and as a means of making the task more meaningful and interesting. If that was the case, similar results concerning discrete predictions would have been expected in the two experiments. This was found.

The novel aspect of this experiment is its confidence assessments. Figure 2 shows these assessments after they were translated into mean "subjective probabilities." For each face, mean confidence in the high base-rate prediction is contrasted with the mean assigned in the low base-rate condition, following Kahneman and Tversky's (1973) Figure 1. Of the 10 points plotted this way (one for each face), 2 are on the identity line, reflecting complete insensitivity to base rates, and all the others are closer to the identity line than to the curve appropriate for judges ideally sensitive to base rates. Thus, Figure 2 is not unlike Kahneman and Tversky's comparable plot in that both show a slight, though significant, effect of base rates. The fact that none of the faces in Figure 1 received a very high mean probability of be-

³ It is possible to compute the expected *pro* predictions for each of the issues and each of the consensus groups. Regarding the seat belt issue, suppose 70% of the time photos are viewed as neutral and the rest are split evenly between positive and negative. Then, by the same reconstruction as is offered in Table 1, we would expect 71% *pro* predictions in the positive consensus group and 29% in the negative consensus group. Seventy-one percent represents 80% of the neutrals plus the positives, that is, $56\% + 15\% = 71\%$. Likewise, regarding the marijuana issue, if only 53% of the time the photos are viewed as neutral, we would expect 66% *pro* predictions in the positive consensus group and 34% in the negative consensus group. The actual numbers obtained by Manis et al. were 70% vs. 30% in the marijuana issue and 77% vs. 30% in the seat belt issue. Although the predictions do not match the observations exactly, they do account for the somewhat larger difference between the two proportions in the issue with the higher proportion of neutrals.

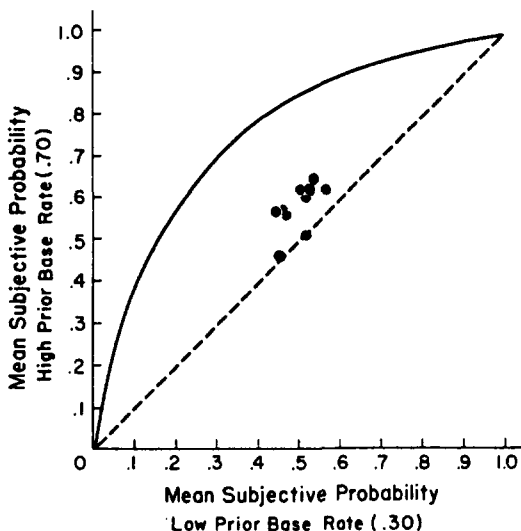


Figure 2. Mean derived subjective probabilities that the individuals depicted in 10 pictures have plans to attend graduate school, plotted separately for groups that had seen a sample of either 70% or 30% students with such plans. (From "Base Rates Can Affect Individual Predictions," by Melvin Manis, Ismael Dovalina, Nancy E. Avis, and Steven Cardoze, *Journal of Personality and Social Psychology*, 1980, 38, 231-248, Figure 4. Copyright 1980 by the American Psychological Association. Reprinted by permission.)

longing to either category in either condition suggests than none evoked a very strong consistent stereotype. Indeed, the means are so close to .50 that in every case relatively few subjects could have had confidence scores of .755 or more.

Response Mode Effects

Manis et al. noted the contrast between the modest base-rate effect with confidence judgments (Figure 2) and the pronounced base-rate effect with categorical judgments (Figure 1). Table 2 attempts to clarify the various patterns of results that are possible with these two response modes using several hypothetical sets of stimuli and prediction strategies. For the sake of simplicity, the table makes two restrictive assumptions: (a) Subjects have learned the base rate, and (b) subjects agree about the differential representativeness of each face. Each table entry represents a subject's confidence that a face is pro, using Manis et al.'s scheme for de-

termining probability assessments; hence, probabilities of .51 or greater mean that the subject predicted pro. In Example 1, both subjects agree that Faces 1, 2, and 3 are certainly pro and that 8, 9, and 10 are certainly con; they predict as much. They also agree that Faces 4, 5, 6, and 7 are neutral and classify them in the majority category, although with little confidence in those predictions. Subject 1, who has seen 70% con, guesses that they are con, whereas Subject 2, who has seen 70% pro, guesses that they are pro. The results is that the two subjects' category predictions reflect their respective base rates perfectly, whereas all 10 points fall very close to or on the identity line, appearing like disregard for the base rates.

In Example 1, the subjects respond to feedback simply by biasing their categorization of "neutral" cases. In Example 2, they also lower their confidence, particularly for minority-category predictions, to reflect the reduced credibility of the individuating evidence. The result is a pattern of discrete choices and "probability" judgments roughly similar to that observed by Manis et al. (i.e., a large difference in discrete choices accompanied by a modest difference in confidence judgments). This similarity supports the idea that feedback affects confidence but does not induce predictions that run contrary to representativeness.

In these examples, subjects' ability to probability match without violating representativeness hinged on the availability of enough neutral faces. Examples 3, 4, and 5 in Table 1 show what might happen if each face was clearly a pro or clearly a con and the split between these categories differed from that of the base rate. Now, subjects can either probability match or classify by representativeness, but not both. In Examples 3 and 4, subjects choose the latter course. Here, the "subjective probabilities" show a clear base-rate effect, whereas the discrete choices show none, thereby reversing the pattern of Example 1. Example 3 assumes that all five pros and all five cons represent their category equally. Consequently, if subjects reduce their confidence in one minority-label prediction, they should reduce it equally in all minority-label predictions. On the other hand, the presence of

Table 2
Continuous Versus Discrete Response Modes

Variable	Pro base rate	Face										% predicted as pro
		1	2	3	4	5	6	7	8	9	10	
Example 1												
Subject 1	30	100	100	100	49	49	49	49	0	0	0	30
Subject 2	70	100	100	100	51	51	51	51	0	0	0	70
Face seems: ^a		Pro	Pro	Pro	Neut ^b	Neut	Neut	Neut	Con	Con	Con	
Example 2												
Subject 3	30	100	75.5	51	49	49	49	49	24.5	0	0	30
Subject 4	70	100	100	75.5	51	51	51	51	49	24.5	0	70
Face seems:		Pro	Pro	Pro	Neut	Neut	Neut	Neut	Con	Con	Con	
Example 3												
Subject 5	30	75.5	75.5	75.5	75.5	75.5	0	0	0	0	0	50
Subject 6	70	100	100	100	100	100	24.5	24.5	24.5	24.5	24.5	50
Face seems:		Pro	Pro	Pro	Pro	Pro	Con	Con	Con	Con	Con	
Example 4												
Subject 7	30	100	100	100	51	51	0	0	0	0	0	50
Subject 8	70	100	100	100	100	100	49	49	0	0	0	50
Face seems:		Pro	Pro	Pro	Pro	Pro	Con	Con	Con	Con	Con	
Example 5												
Subject 9	30	100	100	100	49	49	0	0	0	0	0	30
Subject 10	70	100	100	100	100	100	51	51	0	0	0	70
Face seems:		Pro	Pro	Pro	Pro	Pro	Con	Con	Con	Con	Con	

Note. Faces on the left of the dividing line are judged to be pro in the discrete-response mode.

^a Indicates choice according to representativeness. ^b Neut. = neutral.

some less clear-cut cases might lead to a pattern like that in Example 4. Finally, in Example 5, representativeness is sacrificed for the sake of probability matching. We suspect that this pattern will not be commonly encountered, that is, we believe representativeness will typically dominate probability matching when the two are in conflict.

Reinterpreting Experiment 3

In Experiment 3 Manis et al. abandoned the sequential presentation of base-rate information in favor of a summary format: Subjects were told the composition of the population from which the judged cases were drawn. Their task was still to predict the category membership of individual cases.

The salience of the base-rate summary was manipulated as one independent vari-

able. According to our account, this variable would have an effect only if the least salient summary failed to attract enough attention for the base rate to be used in predictions for neutral targets. The fact that salience had no significant effect may mean that even a minimal summary is adequate.

A second independent variable was the position or "issue" being predicted. Manis et al. intended that the 10 issues they used would vary from highly stereotypic ones (e.g., "belongs to a fraternity") to unsterotypic ones (e.g., "likes to swim"). With respect to each issue, a stimulus was judged "positive" (e.g., likes swimming) or "negative" (e.g., does not like swimming). "Each subject responded to 20 items, 10 of which were fillers. . . . Each issue was presented in combination with all levels of the remaining variables" (Manis et al., 1980, p. 241).

The representativeness hypothesis predicts that those issues that evoke clearer stereotypes facilitate judgment by representativeness and reduce the impact on base rates. However, the issues variable failed to produce an effect. Whether this result constitutes evidence against the representativeness hypothesis depends on whether one accepts Manis et al.'s intuitive classification of issues as strongly or weakly stereotypic and whether one believes that the stimuli were differentially representative of the various categories. Since no relevant evidence is presented, the absence of an issues effect may be due to the weakness of the manipulation. Alternatively, even if the different issues do vary in the strength of the stereotypes they evoke, having subjects consider a variety of issues may have encouraged them to generate stereotypes for even lackluster issues. Fischhoff and Slovic (1980) found subjects able to create discriminatory schemes even with highly unfamiliar stimuli and "issues."

Collapsing over all 10 test issues, Manis et al. found a significant, though small, effect for a third independent variable: the base rates of "positives" for the different issues. This result replicates the previous finding that base rates are utilized somewhat when subjects consider alternative base-rate values (Fischhoff et al., 1979). On a 1-10 likelihood scale, a typical description was judged about 1.2 points more likely to belong to a given category when that category was a majority one (70% of the population) than when it was a minority one (30%).

The final independent variable was the "type" of individuating information given about each judged case. The five types were: (a) last initial only; (b) name, age, hair color, and eye color; (c) "positive picture" (i.e., one strongly resembling the positive stereotype for the issue); (d) "negative picture"; and (e) "neutral picture." The representativeness hypothesis predicts that "when the target case resembles neither [category stereotype], we should find that base-rate information is more substantially involved in the judgmental process" (Manis et al., 1980, p. 238). From this perspective, Type a might be called "worthless" information, since representativeness provides no guide to prediction, whereas Types c and d might

constitute "diagnostic" information, since one can predict by representativeness. The diagnosticity of information of Type b would seem to depend on the issue being judged (e.g., hair color might be diagnostic for "is Jewish" but not for "is an only child"). It is unclear whether the neutrality of the Type e pictures is due either to their resembling neither category stereotype or to their resembling different stereotypes in the judgment of different subjects.

Consistent with prediction by representativeness, the modest overall base-rate effect noted above was somewhat greater for worthless information (Type a $M = 1.73$) than for diagnostic information (Types c and d $M = 1.17$), with Type b falling in between ($M = 1.42$). The relatively weak base-rate effect with the "neutral picture" (Type e $M = 1.13$) would conflict with judgment by representativeness if these pictures resembled neither stereotype for most subjects; it would be consistent, however, if the pictures were merely controversial, that is, if subjects were divided between those who saw each as a positive and those who saw each as a negative.

Reinterpreting Experiment 4

In Manis et al.'s Experiment 4, three formats were used to present base-rate information about a sample of 40 college students' plans to continue schooling. Group I sequentially viewed photographs with labels that "did not seriously violate the stereotypes [they] evoked" (p. 244). Group II received only a statistical summary of the sample base rate. Group III received Group II's statistical summary and Group I's pictures, minus the labels. All groups were tested on a set of 20 "relatively neutral" new faces reportedly drawn from the same population. Subjects classified each picture and then indicated their confidence on a scale of 50% (pure guess) to 100% (certainty).

Since the test faces were "relatively neutral," representativeness would provide a poor guide to prediction. Hence, we would expect a base-rate effect on the discrete predictions in all three conditions. Indeed, "the results . . . showed a strong main effect for base rate. . . . The base-rate format was not

significant as a main effect . . . nor did it interact with the base-rate variable" (Manis et al., 1980, p. 245). This effect was nonetheless somewhat weaker than that found in Experiments 1 and 2, suggesting that the set of "relatively neutral" faces included many stimuli evoking inconsistent stereotypes (across subjects).

When Do Base Rates Affect Predictions?

The tendency to ignore base-rate information has been most clearly demonstrated in experiments that present subjects with a single question requiring the integration of base-rate and individuating information. When no integration is called for (i.e., when no individuating information is presented), people use the base rate correctly. Even when worthless or nondiagnostic information is added, people still rely on base rates. It is only when seemingly diagnostic information is given that people make predictions on the basis of representativeness. The opportunity for a hybrid response strategy is provided by tasks offering a large number of cases that vary in the degree to which they are differentially representative of the possible prediction categories. In such cases, people can classify the diagnostic cases by representativeness and the neutral ones by base-rate considerations.

Many existing studies can be classified according to whether or not they allow subjects to express both the base rates and representativeness in their judgments. One cannot probability match in making a single category prediction; studies using such tasks have found that people judge by representativeness and ignore base rates (e.g., Ginosar & Trope, 1980; Kahneman & Tversky, 1973; Nisbett & Borgida, 1975). On the other hand, in multiple-case studies that include many neutral descriptions (which can be assigned to the majority category without violating representativeness), a base-rate effect is typically found (e.g., Carroll & Siegler, 1977; Manis et al., 1980; Basok, Note 1). Table 1 suggests alternative strategies for resolving the conflicting demands of probability matching and representativeness when neutral descriptions are absent.

To rely on representativeness, judges must be able to establish a differential match between the salient features of the individuating information and those of the possible prediction categories. When unable to do so, they rely on base rates. As mentioned, representativeness may fail to provide a guide when the individuating information is either absent altogether or devoid of relevant (i.e., diagnostic) features. It may also fail when the prediction categories have either no salient features or equivalent salient features. Empirical identification of a category's salient features and of the representativeness of a target case may prove to be fairly difficult. For example, even when independent assessment is attempted, one cannot be certain that under the set induced by an experiment, subjects will not conjure up some image for a category that otherwise seems bland and unevocative.

The impact of the individuating information can be modified by experience. In particular, feedback regarding the validity of judging by representativeness can change the apparent diagnosticity of individuating information. However, unless respondents are convinced that the individuating information has no predictive validity, they will still predict according to representativeness. All that will change is their confidence in those predictions. Such a pattern can only be revealed with an experimental design that elicits both predictions and confidence. Manis et al.'s study showed such a pattern, as did a study by Kahneman and Tversky (1973) in which two groups of subjects made predictions based on the same information. One was told that only a small proportion of such predictions proved accurate, and the other was given a higher proportion. Although these subjects did not actually receive any feedback, they had grounds for anticipating what such feedback would be. Both groups made essentially the same discrete predictions, but the "low expected accuracy" group expressed lower confidence in their judgments than did the "high expected accuracy" group.

The effect of providing (implicit or explicit) feedback is logically independent of the effect of affording multiple prediction opportunities (without which probability

matching is impossible). The first affects the confidence ratings, whereas the second affects the pattern of discrete predictions. Given this independence, there is no need to reconcile the discrete-prediction and subjective-confidence results in a study like Manis et al.'s, which gives both feedback and multiple predictions. As a result of this compound manipulation, Manis et al.'s results do not contradict those of Kahneman and Tversky. Rather, they provide an instructive complement, clarifying judgmental processes and their assessment.

Reference Note

1. Basok, M. [*The effect of kind of individuating information on the impact of frequentistic information on subjective probabilities.*] Unpublished MA thesis, Hebrew University of Jerusalem, 1980.

References

- Ajzen, I. Intuitive theories of events and the effects of base-rate information on prediction. *Journal of Personality and Social Psychology*, 1977, 35, 303-314.
- Bar-Hillel, M. The base-rate fallacy in probability judgments. *Acta Psychologica*, 1980, 44, 211-233.
- Borgida, E., & Brekke, N. The base-rate fallacy in attribution and prediction. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 3). Hillsdale, N.J.: Erlbaum, in press.
- Carroll, J. S., & Siegler, R. S. Strategies for the use of base-rate information. *Organizational Behavior and Human Performance*, 1977, 19, 392-402.
- Estes, W. K. The cognitive side of probability learning. *Psychological Review*, 1976, 83, 37-64.
- Fischhoff, B., & Slovic, P. A little learning . . . : Confidence in multicue judgment. In R. Nickerson (Ed.), *Attention and performance VIII*. Hillsdale, N.J.: Erlbaum, 1980.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. Subjective sensitivity analysis. *Organizational Behavior and Human Performance*, 1979, 23, 339-359.
- Ginosar, Z., & Trope, Y. The effects of base rates and individuating information on judgments about another person. *Journal of Experimental Social Psychology*, 1980, 16, 228-242.
- Jarvick, M. E. Probability learning and negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology*, 1951, 41, 291-297.
- Kahneman, D., & Tversky, A. On the psychology of prediction. *Psychological Review*, 1973, 80, 237-251.
- Kassin, S. M. Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology*, 1979, 37, 1966-1981.
- Manis, M., Dovalina, I., Avis, N. E., & Cardoze, S. Base rates can affect individual predictions. *Journal of Personality and Social Psychology*, 1980, 38, 231-248.
- Nisbett, R. E., & Borgida, E. Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, 1975, 32, 932-945.
- Tune, G. S. Response preferences: A review of some relevant literature. *Psychological Bulletin*, 1964, 61, 286-302.
- Tversky, A., & Kahneman, D. Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology*. Hillsdale, N.J.: Erlbaum, 1979.

Received May 26, 1980

Revision received February 27, 1981 ■