

# NEW DIRECTIONS IN ATTRIBUTION RESEARCH

Volume 3

*Edited by*

**JOHN H. HARVEY**  
*Texas Tech University*  
**WILLIAM ICKES**  
*University of Missouri—St. Louis*  
**ROBERT F. KIDD**  
*Boston University*



1981

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS  
Hillsdale, New Jersey

# 3

## The Base Rate Fallacy in Attribution and Prediction

Eugene Borgida and Nancy Brekke  
*University of Minnesota*

At the heart of attribution theory, as formulated by Harold Kelley (1967, 1972a, 1972b, 1973), is the basic analogy-between the inferential skills of the layperson or intuitive scientist and those of the professional scientist (Kruglanski, Hamel, Maides, & Schwartz, 1978; Ross, 1977, 1978). This layperson-scientist analogy generally suggests that the intuitive scientist processes a variety of inferential tasks in a manner analogous to the scientist's normative approach to the solution of inferential problems. In an important and comprehensive treatment of the lay scientist's inferential abilities, however, Nisbett and Ross (1980) discuss a number of inferential "shortcomings" that characterize all stages of intuitive scientific judgment. In particular, Nisbett and Ross detail the extent to which the intuitive scientist's reliance on various cognitive representations or knowledge structures and intuitive inferential strategies may lead to inaccurate, as well as accurate, social inference.

The misapplication or the overapplication of, for example, cognitive schemata (Bartlett, 1932; Markus, 1977; Neisser, 1976; Taylor & Crocker, 1980), prototypes (Cantor & Mischel, 1979), implicit personality theories (Rosenberg & Sedlak, 1972; Wegner & Vallacher, 1977), or scripts (Abelson, 1976; Schank & Abelson, 1977) may affect the accuracy of inferences and predictions made by the intuitive scientist. Similarly, the lay scientist's use of intuitive judgmental heuristics such as representativeness and availability, often to the exclusion of more normatively appropriate strategies, results in persistent inferential errors on probabilistic judgment tasks (Kahneman & Tversky, 1972, 1973; Tversky & Kahneman, 1971, 1973, 1974, 1980). In fact, a reliance on these "generally valid, intuitive, inferential strategies," as Nisbett and Ross have suggested, may have pervasive ef-

fects on the coding of data, assessments of covariation, causal analysis, prediction, and theory testing.

This chapter focuses on the intuitive scientist's ability to use base rate or distributional information in two of these inferential tasks: causal analysis and prediction. It will be argued that essentially the same characteristics of human inference are implicated in these tasks. Kelley's covariation principle, for example, suggests that the intuitive scientist considers three sources of information in order to attribute causal explanations for social behavior: *consensus* (do other people behave in this manner in this situation, or is the response relatively rare?), *distinctiveness* (does the person behave in this manner in similar situations as well, or only in this particular situation?), and *consistency* (does the person usually or only occasionally behave in this manner?). Attribution of cause to either situational or dispositional factors, in this view, depends on the nature of the information pattern under consideration (Garland, Hardy, & Stephenson, 1975; McArthur, 1972, 1976; Orvis, Cunningham, & Kelley, 1975).

But in a series of experiments, Nisbett and Borgida (1975; Nisbett, Borgida, Crandall, & Reed, 1976) proposed that people often do not make normatively appropriate use of consensus information in explaining or predicting behavior. Nisbett and Borgida (1975), for example, described two previously conducted experiments and then asked their subjects to predict the behavior of a particular participant in the original experiment and to explain causally the behavior of another participant who had behaved in an extreme way. Prior to making predictions and attributions, some subjects were told that in one experiment, where participants had been asked to take a great deal of electric shock, the modal behavior had been to take the maximum amount of shock the apparatus could deliver. In the other experiment, where a confederate asked for help during what appeared to be a seizure, the modal participant never offered help. These behavioral base rates were surprising to subjects, who had believed that a moderate amount of shock-taking would have been typical in the shock experiment and that a fairly quick helping response would have been typical in the helping study. Nevertheless, Nisbett and Borgida argued that consensus information (i.e., knowledge of the behavioral base rate data), according to attribution theory, should have affected subjects' predictions about the behavior of any particular original participant and should have led to increased situational attributions as explanations for the behavior of any particular participant exhibiting modal behavior. They found, however, that their subjects' predictions, trait inferences, and causal attributions were unaffected by sample-based consensus information.

Nisbett and Borgida suggested that these failures to follow normative inferential rules were similar to Kahneman and Tversky's (1973) demonstration that people ignore population base rates when making predictions about the category membership of a target individual. Kahneman and Tver-

sky instructed some subjects that a panel of psychologists had administered personality tests to 70 engineers and 30 lawyers; for other subjects, the proportion of engineers and lawyers was reversed. Subjects were asked to estimate the probability that a specific respondent was one of the 70 (30) engineers. Kahneman and Tversky also presented subjects with target case or individuated information in the form of five paragraph-length personality sketches. Four of the five descriptions presented individuating information that was uninformative or *non-diagnostic* for predicting the target case's profession. In contrast, the fifth description provided subjects with individuating information that was informative or *diagnostic* for predicting the target's profession. Finally, in a *null* condition, subjects were told only that the target had been randomly selected from the sample of respondents. Kahneman and Tversky found that subjects who had no information about the target appropriately used the base rate when judging the probability that the target was an engineer or lawyer. When non-diagnostic and diagnostic individuating information was available, however, subjects underutilized the base rate distribution of engineers and lawyers from which the target had been sampled, and overweighted the individuating information.

The present chapter addresses several important questions with respect to this "base rate fallacy" (Bar-Hillel, 1980), or *the tendency to underutilize base rates in favor of singular or individuating information on various attribution and prediction tasks*. The first question, which is the focus of Part II, is: What limiting conditions for the base rate fallacy have emerged thus far in attribution and judgment research? The early experiments on consensus and base rate utilization by Kahneman and Tversky (1973) and Nisbett and Borgida (1975) have generated a sufficient body of research and controversy (Borgida, 1978; Wells & Harvey, 1978) such that several parameters of the phenomenon are now specifiable. As several investigators have suggested, then, the question is no longer whether people use consensus and base rate information, but rather, what factors seem to promote base rate utilization (Ajzen & Fishbein, in press; Borgida, 1978; Kassir, 1979b; Kulik & Taylor, in press).

A second question addressed by this chapter in Part III concerns the possible theoretical explanations for the base rate fallacy: *Why* are people susceptible to the base rate fallacy? Several researchers in both the attribution and judgment domains have offered such explanations. Part III, then, reviews these explanations and discusses a convergent framework for understanding the base rate fallacy in attribution.

Finally, Part IV illustrates the pervasiveness of the base rate fallacy in everyday reasoning about social behavior. Distributional beliefs, for example, are manifest in clinical diagnosis (Goldberg, 1968; Meehl, 1954; Meehl & Rosen, 1955), in hypothesis-testing (Snyder & Swann, 1978), predictions of recidivism risk (Carroll, 1977), attributions for task performance (Frieze & Weiner, 1971), in predicting dangerousness (Shah, 1978), and in the re-



quirements for establishing character proof in the law of evidence (Borgida, 1979). Part IV, however, focuses on the extent to which the dynamics of the base rate fallacy contribute to an understanding of the conditions under which social stereotypes may influence social judgments. Social stereotypes have the logical status of prior probabilities for social judgments about particular members of stereotyped social groups and Part IV discusses some recent research that suggests that stereotypic beliefs "behave" like distributional beliefs in the context of prediction or judgment tasks.

### SOME LIMITING CONDITIONS OF THE BASE RATE FALLACY

What limiting conditions of the base rate fallacy have emerged in attribution and judgment research since the initial work by Kahneman and Tversky (1973) and by Nisbett and Borgida (1975)? Recent attribution and judgment research suggests several conditions that seem to promote use of base rate information and an overview of these circumstances is provided in Table 3.1, which is referenced throughout this section of the chapter.

Before reviewing these conditions, however, several other considerations should be mentioned. First, and most important, our review and analysis presumes substantial overlap between the attribution and judgment research domains. Both inferential tasks are regularly performed by the intuitive scientist and both entail the same set of reasoning processes (Nisbett & Ross, 1980; Ross 1977). Although the two domains have a great deal in common (Fischhoff, 1976), there is an important conceptual difference. The very notion of a base rate *fallacy* suggests the existence of a normative standard with which to assess the normative appropriateness of certain attributions and predictions. The normative standard for causal explanations is somewhat ambiguous in that it rests on the logical dictates of attributional principles (Schneider, Hastorf, & Ellsworth, 1979) or, perhaps more suitably, on the informed opinions of "experts" when they agree on the normative rules governing a particular inference task (Nisbett & Ross, 1980; Stich & Nisbett, in press). In either case, however, there is no quantifiable standard of *optimality* for such rational base line models as attribution theory (Jones & McGillis, 1976). In attribution research one cannot assess the exact magnitude of judgmental deviations from optimality. Therefore most attributional studies of the base rate fallacy rely on null hypothesis testing (e.g., do people make statistically significant use of consensus information?).

Although there is debate about the appropriateness of the Bayesian approach to various aspects of inductive inference (Ajzen & Fishbein, 1978; Fischhoff & Lichtenstein, 1978; Shafer, 1976; Tribe, 1971; Williams, 1978), research in the judgment domain traditionally has relied on Bayes' theorem for a normative description of the appropriate use of data for predictions

and judgments (Slovic, Fischhoff, & Lichtenstein, 1977; Slovic & Lichtenstein, 1971). Given two events, *A* and *B*, Bayes' theorem states that  $p(A/B) = p(A) \cdot p(B/A)/p(B)$ . In other words, beliefs and predictions should be affected both by the prior probability or base rate distribution of the criterion [ $p(A)$ ] and by the strength of the relationship between the datum and the criterion [ $p(B/A)/p(B)$ ].

Normatively, when the datum is non-diagnostic or uncorrelated with the criterion [ $p(B/A)/p(B) = 1$ ], people should rely solely on the prior probability of the criterion. When the datum is diagnostic [ $p(B/A)/p(B) \neq 1$ ], however, people should revise the prior probability of the criterion according to the informativeness or diagnosticity of the datum. Unlike attributional studies of the base rate fallacy, therefore, there is a standard of optimality that determines the normative appropriateness of judgments on Bayesian inference tasks. Thus, the reader should keep in mind the *type* of problem under investigation (attribution versus judgment). This information is indicated in Column 1 of Table 3.1. Conclusions about research findings should be tempered with the appropriate normative standard in mind. One can obtain significant effects for base rate utilization even when such utilization is highly sub-optimal from a Bayesian perspective.

Second, several investigators have stressed the importance of appropriate dependent variable measurement in research on the base rate fallacy (Wells & Harvey, 1977). Although it is, of course, important to choose theoretically appropriate dependent measures in any area of research, consideration of this dimension does not constitute a limiting condition of the phenomenon. The base rate fallacy has been demonstrated across a wide variety of measures and does not appear to be a function of the researcher's particular choice of dependent measures.

Similar concerns have been expressed with respect to the impact of the response format used to study the base rate fallacy. Manis, Dovalina, Avis, and Cardoze (1980), for example, have shown that discrete predictions about category membership may be more normatively appropriate than subjective probability estimates. In general, the impact of this factor is difficult to evaluate because differences between response formats to a certain extent are a function of the type of problem investigated. Judgment researchers tend to elicit probability estimates for Bayesian inference tasks whereas attribution researchers usually use Likert scales or a multiple-choice format. Even so, the variability of results within each type of problem, between studies using the same response format, suggests that response format per se does not represent a limiting condition of the base rate fallacy. The most direct evidence available in fact supports this contention. Lyon and Slovic (1976) compared the probability estimate and multiple-choice response formats for the same inference problem and found little difference between the two. Regardless of the response format used, subjects exhibited the base rate fallacy.

TABLE 3.1  
Review of Consensus and Base Rate Studies

Last	Type of Study		Order of Information					
	Judgment	Attribution	Between Consensus, Distinctiveness, & Consistency:		Between Base Rates and Target Info:		Sequential Base Rate Presentation	Base Rate Vividness
			Consensus 1st	Consensus	BR 1st	Target 1st		
Ajzen (1977)	X					X		
Bar-Hillel (1980)	X					X		
Borgida & Nisbett (1977)		X					X	X
Carroll (1977)	X					X		
Carroll (1980)	X					X		
Carroll & Siegler (1977)	X					X		
Cooper, Jones & Tuller (1972)		X						
DeJong & Amabile (1979)		X				X	X	
DiVitto & McArthur (1978)		X					X	
Dyck & Rule (1978)		X				X		
Feldman, Higgins, Karlovac, & Ruble (1976)		X					X	X
Fischhoff, Slovic, & Lichtenstein (1979)	X					X		X
Frieze & Weiner (1971)		X	X	X			X	
Garland, Hardy, & Stephenson (1975)		X				X		
Ginosar & Trope (1980)	X					X		
Hammerton (1973)	X					X	X	
Hansen (1979)		X	X	X				
Hansen & Donoghue (1977)		X						
Hansen & Lowe (1976)		X		X				
Hansen & Stonner (1978)		X						
Harvey, Cacioppo, & Yasuna (1977)		X				X		X
Kahneman & Tversky (1973)	X					X		
Karlovac, Feldman, Higgins, & Ruble (1976)		X					X	X
Kassin (1979b)		X				X		
Kassin & Lowe (1976)		X						
Kulik & Taylor (in press)		X					X	
Kulik & Taylor (1980)		X						
Lay, Burron, & Jackson (1973)		X						
Lowe & Kassin (1977)		X				X		
Lyon & Slovic (1976)	X					X	X	X
Manis, Dovalina, Avis, & Cardoze (1980)	X					X		X
Martin & Powers (1979)	X						X	
McArthur (1972)		X	X				X	
McArthur (1976)		X	X				X	
Nisbett & Borgida (1975)		X				X		
Nisbett, Borgida, Crandall, & Reed (1976)		X						
Orvis, Cunningham, & Kelley (1975)		X	X				X	
Plikonis (1977)		X				X		
Ruble & Feldman (1976)		X	X	X			X	
Wells & Harvey (1977)		X				X		
Wells & Harvey (1978)	X					X		
Zuckerman (1978a)		X					X	
Zuckerman (1978b)		X				X		

Continued

TABLE 3.1 (continued)

	<i>Extremity of Base Rate</i>	<i>Sampling Know- ledge</i>	<i>Base Rate Transla- tability</i>	<i>Individual Differ- ences</i>	<i>Implicit Con- sensus</i>	<i>Actions Versus Occur- rences</i>	<i>Causal Relevance</i>
Ajzen (1977)							X
Bar-Hillel (1980)	X						X
Borgida & Nisbett (1977)							
Carroll (1977)	X		X				
Carroll (1980)			X				
Carroll & Siegler (1977)	X	X	X				
Cooper, Jones & Tuller (1972)							
DeJong & Amabile (1979)							
DiVitto & McArthur (1978)							
Dyck & Rule (1978)							
Feldman, Higgins, Karlovac, & Ruble (1976)							
Fischhoff, Slovic, & Lichtenstein (1979)	X						
Frieze & Weiner (1971)							
Garland, Hardy, & Stephenson (1975)							
Ginosar & Trope (1980)				X			
Hammerton (1973)							
Hansen (1979)							
Hansen & Donoghue (1977)					X		
Hansen & Lowe (1976)					X		
Hansen & Stonner (1978)		X			X		
<b>Harvey, Cacioppo, &amp; Yasuna (1977)</b>							
Kahneman & Tversky (1973)							
Karlovac, Feldman, Higgins, & Ruble (1976)							
Kassin (1979b)	X						
Kassin & Lowe (1976)							
Kulik & Taylor (in press)					X		
Kulik & Taylor (1980)				X	X		
Lay, Burron, & Jackson (1973)					X		
Lowe & Kassin (1977)	X				X		
Lyon & Slovic (1976)	X				X		
Mannis, Dovalina, Avis, & Cardoze (1980)							
Martin & Powers (1979)							
McArthur (1972)							
McArthur (1976)							
Nisbett & Borgida (1975)							
Nisbett, Borgida, Crandall, & Reed (1976)							
Orvis, Cunsingham, & Kelley (1975)							
Pilkonis (1977)					X		
Ruble & Feldman (1976)							
Wells & Harvey (1977)	X	X					
Wells & Harvey (1978)							
Zuckerman (1978a)							X
Zuckerman (1978b)	X				X		



Finally, it should be noted that the base rate fallacy does not appear to be a function of faulty memory for base rate information. Unfortunately, many researchers, especially those in the judgment domain, failed to collect base rate recall data. Recall measures have been included more frequently in attributional studies of the base rate fallacy (Nisbett & Borgida, 1975; Wells & Harvey, 1977; Zuckerman, 1978b). These studies have yielded very consistent results suggesting that the base rate fallacy is not a function of subjects' limited memory capacities. It would appear that reasonably accurate base rate information is available to subjects but that they simply do not make appropriate use of it. Moreover, it is apparent from a number of studies that base rate information is used appropriately when base rates are presented in the absence of any individuating information (Bar-Hillel, 1980; Ginosar & Trope, 1980; Kahneman & Tversky, 1973; Kulik & Taylor, 1980; Locksley, Borgida, Brekke, & Hepburn, 1980; Locksley, Hepburn, & Ortiz, in press; Ruble, Feldman, Higgins, & Karlovac, 1979 [Experiment 1]; Zuckerman, 1978b). The consistency of these latter findings certainly implies that subjects understand the relevance of base rates and how to apply them in the absence of individuating information.

In the remainder of this part of the chapter, then, we review attribution and judgment studies that deal with base rate utilization in the absence of individuating information as well as studies that directly address the base rate fallacy, that is, examine base rate utilization in the presence of individuating information. Sections A-D discuss various factors that seem to promote the availability in memory of base rate information; sections E-G discuss several factors that seem to reduce sample size or sample bias problems associated with base rate utilization; and sections H-J review studies that examine the role of intuitive causal theories and base rates with causal implications.

#### A. Order of Information Presentation

*Between Consensus, Distinctiveness, and Consistency Information.* In a number of studies (McArthur, 1972; Orvis, Cunningham, & Kelley, 1975), consensus information may have been underutilized because it was presented prior to distinctiveness and consistency. Accordingly, Ruble and Feldman (1976) systematically varied the order of these three cues and found a recency effect for consensus utilization. When order was so counterbalanced, consensus, distinctiveness, and consistency information accounted for approximately equal amounts of variance, a finding that has since been replicated by Zuckerman (1978a).

It may be seen in Table 3.1 that there are actually very few studies in which consensus information was always presented first. The order in which consensus, distinctiveness, and consistency information are presented is

noted in Column 2. First, where all three cues are presented, the order is often varied (e.g., Frieze & Weiner, 1971; Hansen, 1979). Second, it may be seen in Column 2 that in most studies subjects are never even given consistency or distinctiveness information. And, finally, even when subjects are given consensus information, they are nevertheless insufficiently responsive to its implications.

*Between Base Rate and Individuating Information.* To what extent does base rate utilization differ as a function of the order in which base rate and individuating information are presented? Given that consensus utilization is subject to a recency effect (Ruble & Feldman, 1976), the impact of base rate information might be attenuated by individuating information presented after the base rate, but prior to the dependent measures. One might expect, then, that base rates would be more underutilized when followed by individuating information than when preceded by individuating information.

Column 2 also indicates the order in which base rate and individuating information have been presented. There appears to be no direct relationship between base rate utilization and order of presentation. For example, Nisbett and Borgida (1975) presented base rates first, whereas Borgida and Nisbett (1977) presented individuating information first. In both instances, subjects underutilized base rates. Lyon and Slovic (1976) also found that presenting the base rate information after the individuating information had no significant effects on subjects' predictions; subjects underutilized the base rate regardless of where it appeared in the problem. Overall, approximately half of the studies reviewed presented base rates first, and approximately half presented them second, yet there does not appear to be a corresponding pattern of base rate utilization.

#### B. Sequential Base Rate Presentation

To what extent is the underutilization of base rate information a function of the manner in which such information is presented? In everyday life, for example, people often obtain consensus information sequentially rather than simultaneously (Feldman, Higgins, Karlovac, & Ruble, 1976). Feldman et al. therefore manipulated the temporal presentation of consensus information. Subjects were shown videotapes in which four actors agreed or disagreed with the target actor either all at once (simultaneous condition) or one at a time (sequential condition). Subjects in the sequential conditions made significantly better use of the consensus information than subjects in the simultaneous conditions. Feldman et al. suggested that subjects may process simultaneous consensus as one unit of information, but process sequential consensus as many bits of information. Thus, subjects in sequen-



tial conditions are more likely to utilize consensus because they in effect process "more" consensus information.

A similar approach was taken by Manis, et al. (1980), who presented base rates in a case-by-case format. They showed subjects photographs of people from the target population along with information about the category membership of each case. When asked to make predictions about individual targets in a discrete-response format, subjects made significant use of the base rates.

Studies that have used this sequential format are noted in Column 3. It may be seen that there are only a few studies that directly address this issue. In fact, there are only two studies that use sequential base rate presentation besides the two already mentioned. Harvey, Cacioppo, and Yasuna (1977) obtained significant effects for sequential consensus, but on self-attributions of ability. Karlovac, Feldman, Higgins, and Ruble (1976) used virtually identical procedures and stimulus materials as those employed by Feldman et al. (1976), but found no significant differences in base rate utilization for subjects in simultaneous versus sequential presentation conditions.

More importantly, the results obtained by Feldman et al. and Manis et al. must be qualified somewhat. First, the significant effect for sequential presentation in the Feldman et al. study was obtained only for subjects who had no direct information about the stimulus array in question. When subjects had such information, there was no effect for temporal presentation. Second, in the Manis et al. studies, all base rate information was conveyed in a case-by-case manner; there were no simultaneous presentation conditions. Moreover, although base rate utilization was statistically significant, utilization was nevertheless sub-optimal.

### C. Vividness of the Base Rate

Nisbett and Borgida (1975; Nisbett et al., 1976) have proposed that base rates may be underutilized because they are abstract and pallid. Individuating or target case information, in contrast, is more concrete and vivid and it has been suggested that vivid information may be more salient (Taylor & Fiske, 1978) and therefore more cognitively available in memory than abstract information (Borgida & Nisbett, 1977; Nisbett & Ross, 1980). Studies that have operationalized base rate information in vivid and abstract forms are noted in Column 4. A number of studies not included in Column 4 have presented base rate information in a vivid mode but have not presented informationally equivalent base rates in an abstract form. Some investigators, for example, have used videotapes (Feldman et al., 1976), pictures (DiVitto & McArthur, 1978), signatures on a petition (DeJong & Amabile, 1979), audiotapes (Haney, 1979), and false physiological

feedback (Hansen & Lowe, 1976) to present base rate information. In general, these studies have yielded no better utilization than studies in which base rate information was presented in an abstract, summary form.

A more direct test of the hypothesis was provided by Borgida and Nisbett (1977). Subjects were given information about several psychology courses in the form of either summary tables containing mean course evaluations (comparable to the student course evaluation booklets distributed at many colleges and universities) or brief, face-to-face comments by students who had previously taken the courses. Subjects were asked to indicate those courses that they planned to take in the future. Thus, the vivid base rates were clearly concrete, emotionally interesting, and experienced in a proximate, sensory mode (Nisbett & Ross, 1980). As predicted, vivid, face-to-face information had a greater impact on subjects' course choices and confidence in those choices than did the statistical summary information, especially for those subjects who planned on majoring in psychology. In contrast, Manis et al. (1980) showed subjects photographs of people in each of two population subgroups in an attempt to present base rate information in a more concrete, vivid manner. This form of presentation, although not comparable to the Borgida and Nisbett operationalization of vividness, proved to be no more influential than base rates presented in a summary (abstract) form.

Fischhoff, Slovic, and Lichtenstein (1979), it should be noted, have recently taken a different approach to increasing the salience and therefore the availability of base rate information. By asking subjects to make a series of judgments across a variety of base rate values, that is, by using a within-subjects design instead of a between-subjects design, they were able to facilitate base rate utilization, though the latter was still sub-optimal.

### D. Base Rate Extremity

Kahneman and Tversky (1973) have suggested that base rates may be utilized appropriately if they are rather extreme. Column 5 indicates those studies in which base rate extremity (or magnitude of consensus) was manipulated. Although most studies include high and low base rate conditions (e.g., 70%/30%) only those in which a variety of base rates are investigated (e.g., 70%/30% versus 90%/10%) are noted in Column 5. Studies in which extremity was directly manipulated have produced equivocal results. Wells and Harvey (1977), for instance, used the stimulus materials from Nisbett and Borgida's (1975) study but added more extreme high consensus information. As predicted, they obtained significant effects for high consensus. In contrast, Lyon and Slovic (1976) obtained no significant effects when the base rate was made as extreme as 99%/1%.



### E. Sampling Knowledge

Several investigators have examined the extent to which sampling bias and sample size considerations influence base rate utilization, and these studies are noted in Column 6. For subjects to use base rate information appropriately, they must be confident that the behavior sample is valid and unbiased; if they believe that the sample is somehow unrepresentative, they may discount the relevance of the base rate when making predictions about entities or populations.

Wells and Harvey (1977), for example, explained to subjects the importance of using random samples in psychological research and assured them that the sample in question represented a randomly selected cross-section of the population. Their results, in fact, showed more consensus utilization in knowledge of random sampling conditions than in no-knowledge conditions. But sensitivity to sample bias may in fact require rather strong "sampling information" manipulations. There is considerable evidence suggesting that the lay scientist more often than not fails to recognize biased samples and to make appropriate allowances for predictions about entities and populations (Nisbett & Ross, 1980). Nisbett and Borgida (1975), for example, found that people were not sensitive to assurances about the randomness of a sample. Ross, Amabile, and Steinmetz (1977) have demonstrated that subjects do not make adequate inferential allowances even when the behavior sample is blatantly biased. And a recent experiment by Hamill, Wilson, and Nisbett (1980) has shown that subjects' inferences about a target person are dramatically insensitive to pertinent sampling information.

Sample size also has been proposed as a factor that may influence base rate utilization. Kassin (1979a), for example, reasoned that if subjects believed a sample was very small, they would assume that it was unrepresentative. Or, if the sample information contradicted their expectations, they might infer that the information was based on a small, unrepresentative sample. In either case, subjects are likely to reject the sample base rate as invalid and therefore useless for prediction purposes. To test this view, Kassin presented subjects with conflicting sets of base rates derived from samples of different sizes. He found that subjects relied more on larger sample base rates, although the pattern of results was sub-optimal.

Other studies, however, suggest that this sensitivity to inadequate sample size may be somewhat misleading. Borgida and Nisbett (1977), for example, found that subjects were not sensitive to the stability of population estimates based on large samples or the unreliability of estimates based on small samples. Moreover, Bar-Hillel (1979) demonstrated that although it is possible to elicit judgments indicating that perceived accuracy increases with sample size, "these judgments seem to reflect sensitivity to sample-to-population ratio rather than absolute sample size. In fact, people may trade

sample size for sample-to-population ratio, even when this actually decreases expected sample accuracy [pg. 245]."

### F. Base Rate Translatability

Carroll and Siegler (1977) have suggested that subjects may use only base rates that are directly "translatable" to a given prediction problem. That is, subjects indeed may consider base rates relevant, but not know how to translate and apply them appropriately. For example, if they attempt to probability match (Weir, 1964) there would be no direct correspondence between a 70%/30% split in the population and a 70%/30% division in a sample of five people. Unable to solve the problem using the base rate information, subjects may simply decide to ignore it.

As indicated in Column 7, there are only a few studies in which translatable base rates have been used (Carroll, 1977, 1980; Carroll & Siegler, 1977), and they yield mixed empirical support. Carroll and Siegler (1977 [Experiment 2]) found that translatable base rates were used significantly more than non-translatable base rates, even when the latter were more extreme. In a third experiment, however, Carroll and Siegler (1977) found that the enhanced use of translatable base rates was limited to situations in which uninformative target case descriptions were provided. Furthermore, in subsequent studies, Carroll (1977, 1980) obtained no significant effects for base rate translatability on similar though not identical problems.

### G. Individual Differences

Some people may be more susceptible to the base rate fallacy than others, and research that adopts this individual differences approach to base rate utilization is noted in Column 8.

Using variations on the Kahneman and Tversky lawyer-engineer problem, for example, Ginosar and Trope (1980) found that a number of their subjects (28%) applied base rates incorrectly even in the absence of individuating information. These subjects presumably lacked a translation or "sampling rule" for converting relative frequencies into probabilities, and were susceptible to the base rate fallacy regardless of the amount and type of individuating information presented. Subjects who possessed such a sampling rule, however, consistently showed more appropriate use of base rates, even when individuating information was available. It would appear, then, that utilization of base rates may depend in part upon the possession of a relevant sampling rule.

The role of individual differences has been examined also by Kulik and



Taylor (in press) who investigated the relationship between self-monitoring (Snyder, 1979) and the use of consensus information. High self-monitoring individuals are more attuned to situational cues and information than low self-monitoring individuals; high self-monitoring individuals are generally more concerned with behaving in a situationally appropriate manner and are more sensitive to the expressive behavior of others in situations. Kulik and Taylor therefore hypothesized that high self-monitoring individuals would be more sensitive to the situational implications of consensus information than low self-monitoring individuals. Because the latter people tend to be less influenced by situational contingencies and exhibit more cross-situationally consistent behavior, they were expected to make much less use of consensus information.

Kulik and Taylor further suggested that persons high and low in self-monitoring might rely on different types of consensus information. They predicted that high self-monitoring individuals would be influenced by information about the behavior of others, or sample-based consensus. Low self-monitoring individuals, in contrast, should be more sensitive to *self-based consensus* (Hansen & Donoghue, 1977; Heider, 1958; Kulik & Taylor, 1980) or more prone to "false consensus" judgments (Ross, Greene, & House, 1977; Zuckerman & Mann, 1979); that is, they would assume that their own behavior in a situation was normative and thus would use it to infer population performance.

To test these notions, Kulik and Taylor conducted a study patterned after Nisbett and Borgida (1975). Although the prediction that low self-monitoring subjects would use self-based consensus to a greater extent than high self-monitoring subjects received only marginal support, there was evidence that individual differences in self-monitoring mediated the use of consensus information. High self-monitoring subjects made significantly greater use of consensus information than did low self-monitoring subjects, especially when predicting their own behavior. Moreover, the predictions of high self-monitoring subjects were affected by consensus information even when the modal behavior was socially undesirable.

There are a number of other questions that might be addressed by such an individual differences approach. Kulik and Taylor (1980), for example, suggest that an individual's self-concept might determine the relative weight that they assign to self-based versus sample-based consensus information. People who consider themselves to be typical or average in a certain behavioral domain may be more likely to ignore sample-based consensus that conflicts with their own self-based or self-generated consensus. In order to test such a hypothesis, therefore, one would have to measure the amount of experience various subjects have had in a given behavior domain, as well as whether or not they have a relevant self-schema (Markus, 1977).

#### H. *Implicit Consensus*

Implicit consensus, in contrast to Kelley's (1967) notion of sample-based consensus, generally refers to perceivers' expectations or intuitive causal theories for how others will behave in a particular situation. Such expectancies or a priori beliefs about behavior, which affect a variety of inferential tasks (Nisbett & Ross, 1980), may be based on self-generated or self-based consensus (i.e., the assumption that one's own behavior in a situation is normative) or they may be derived from knowledge of normative behavior or knowledge of the target's category membership (e.g., Jones & McGillis, 1976).

Regardless of the source, however, it has been suggested that implicit consensus may have more influence on judgments and attributions than sample-based consensus. Not surprisingly, implicit consensus or intuitive theories seem to affect attributions in the *absence* of sample-based consensus (Kassin, 1979b), when implicit consensus is congruent with normative expectancies (Lay, Burron, & Jackson, 1973), or when the normative expectancies in a given situation have been effectively "neutralized" (Hansen & Lowe, 1976). Furthermore, when sample-based consensus conflicts with normative expectancies or intuitive theories for how others should behave, subjects are more likely to rely on their own causal theories and underutilize information about the behavior of others (Lowe & Kassin, 1977; Zuckerman, 1978a). Studies that have investigated these issues are presented in Column 9.

As Kassin (1979b) has noted, the more important question concerning implicit consensus is the extent to which these intuitive causal theories interact with sample-based consensus when the latter is readily available. Hansen and Donoghue (1977), for example, used a sample-based consensus manipulation patterned after Nisbett et al. (1976) and found that implicit consensus had more influence than sample-based consensus on population estimates and causal attributions. Regardless of whether the consensus information was congruent or incongruent with implicit beliefs, subjects relied on their intuitive causal theories. Research by Zuckerman (1978b), however, suggests that sample-based consensus will be used only when it is congruent with implicit consensus. Zuckerman used the stimulus materials from Nisbett and Borgida (1975) but varied the social desirability of the consensus information given to subjects. He found that subjects used sample-based consensus only when it was socially desirable (i.e., when it was congruent with subjects' normative expectancies). When the sample-based consensus was socially undesirable (or incongruent with these expectancies), as was the case in the Nisbett and Borgida (1975) study, subjects underutilized consensus in their predictions about a target person's behavior. But the social desirability of consensus information does not seem



to explain why subjects are so willing to generate socially *undesirable* base rates based on their knowledge of the socially undesirable behavior of a few target individuals (Nisbett & Borgida, 1975).

Other studies, however, have yielded results suggesting that implicit consensus has a less substantial impact on behavioral predictions. Lowe and Kassin (1977), for example, manipulated consensus within the context of a helping experiment. They found that implicit consensus had no impact on population estimates or predictions about target individuals unless explicit consensus information was unavailable. Implicit consensus had more impact than explicit consensus only when subjects were evaluating a target whose behavior was already known. Kulik and Taylor (1980) found that implicit and explicit consensus had fairly equal impact. Either one alone affected population estimates and trait attributions, and when both were available, they were integrated more appropriately. That is, for subjects with implicit consensus, additionally congruent sample-based consensus enhanced the effects produced by implicit consensus in a normatively appropriate direction. When the sample-based consensus information was incongruent, however, these effects were attenuated.

### I. Actions versus Occurrences

Kruglanski (1975) has distinguished between behaviors that are completely voluntary (actions) and behaviors that are not completely voluntary (occurrences). Actions, because of their voluntary nature, almost always will be attributed to internal causes, whereas occurrences may be attributed to either internal or external causes.

Following this line of reasoning, Zuckerman (1978a) predicted that consensus information should have less impact on the attribution of actions than occurrences. Consensus information is informative only with respect to the stimulus because it provides information about the reactions of many actors to that stimulus. If the stimulus or entity has already been ruled out as a possible cause (because the behavior was completely voluntary) consensus information should be largely irrelevant. Zuckerman (1978a) investigated this hypothesis using procedures and measures similar to that of McArthur (1972). Consistent with this argument, he found that consensus had more impact on attributions for occurrences than attributions for actions. As noted in Column 10, however, Zuckerman is the only researcher thus far who has investigated the effects of this factor on consensus utilization.

### J. Causal Relevance

Ajzen (1977) proposed that when making predictions and explaining events,

people rely on a "causality heuristic," that is;

When asked to make a prediction, people look for factors that would cause the behavior or event under consideration. Information that provides evidence concerning the presence or absence of causal factors is therefore likely to influence predictions. Other items of information, even though important by the normative principles of statistical prediction, will tend to be neglected if they have no apparent causal significance [p. 304].

In a series of clever experiments, Ajzen demonstrated the influence of this causality heuristic on the utilization of base rate information. In one such experiment, subjects were given base rates concerning success (failure) on a final exam, and were then asked to judge the likelihood that a particular target individual passed (failed) the exam. In *causal* base rate conditions, subjects were told that 70% of all students who take that particular class pass (fail) that exam. This base rate allowed them to infer exam difficulty, one of the plausible causes of success or failure on the exam. In *non-causal* base rate conditions, subjects were told that the target came from a sample of students interviewed by an educational psychologist doing research. Because the psychologist was specifically interested in reactions to success (failure), he chose a sample in which 70% succeeded (failed). This description conveyed the base rate, and was even slightly more elaborate than the causal base rate description, yet it contained no information about a possible causal factor such as exam difficulty. As expected, Ajzen found that base rate information strongly influenced predictions only when base rates seemed causally relevant to the criterion. When base rates were not causal, predictions were based primarily on individuating information.

Research in which the causal relevance of the base rate was manipulated is noted in Column 11. Although there are relatively few studies of this sort, the results have been very consistent. In every experiment, causal base rates were used significantly more than non-causal base rates.

Tversky and Kahneman (1980), for instance, have also demonstrated the importance of causal relevance in base rate utilization. They presented subjects with the following problems:

*Problem 1:* A cab was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. Eighty-five percent of the cabs in the city are Green and 15% are Blue. A witness identified the cab as a Blue cab. The court tested his ability to identify cabs under appropriate visibility conditions. When presented with a sample of cabs (half of which were Blue and half of which were Green) the witness made correct identifications in 80% of the cases and erred in 20% of the cases.

*Problem 2:* A cab was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. Although the two companies are roughly equal in size, 85% of the cab accidents in the city in-



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

Subjects were asked to estimate the probability that the cab involved in the accident was Blue rather than Green. The two problems are formally identical and have identical solutions. They differ subjectively, however, inasmuch as the second problem enables the judge to infer that the drivers of Green cabs are more reckless than the drivers of Blue cabs. Thus, the base rate information appears to be causally relevant in Problem 2 but causally irrelevant in Problem 1. Consistent with the hypothesis that people are more likely to use base rate information when it appears to be causally relevant, Tversky and Kahneman (1980) found that the median response of subjects to the first problem was 80%, whereas the median response of subjects to the second problem was 55%, which is much closer to the correct solution of 41%.

### Summary

What, then, are the circumstances or factors that seem to promote base rate utilization and influence the base rate fallacy? First, there are several availability-mediated manipulations that seem to improve the memorability of base rates and lead to more appropriate utilization (Ajzen & Fishbein, in press). The order of presentation, sequential rather than simultaneous presentation of base rates, as well as base rate vividness and extremity represent manipulations that essentially enhance the immediate memorability of base rate information and increase the likelihood of obtaining more appropriate, though not necessarily optimal, use of the base rates.

Second, there have been several attempts to improve base rate utilization by drawing attention to sample bias and sample size considerations. Although there is some evidence that people are less prone to the base rate fallacy when they are assured that the sample in question is representative or large, there is also considerable evidence suggesting that in many domains people lack the sophistication and sampling rules to make adequate inferential allowances for sample size or sample bias. In any case, the base rate fallacy does not appear to be a function of a lack of understanding of the relevance and appropriate application of base rates. Aside from some people who simply lack the necessary sampling rule, people use base rates in a normatively appropriate manner in the absence of individuating information.

Third, it is clear that preconceptions in the form of various knowledge structures and intuitive theories that people hold about how others behave

in situations can bias predictions and lead to normatively appropriate and inappropriate use of base rate information. In the absence of individuating information, normative expectancies and intuitive, self-generated theories about events influence attributions and prediction estimates. This influence is in the direction of normatively appropriate utilization when these expectancies, whatever their form, are congruent with base rate information. When these intuitive theories about behavior in a given situation are incongruent with base rate information, however, the latter is consistently underutilized. In addition, when base rates are made to appear causally relevant to a prediction problem, that is, when the base rate is cognitively available and implies a direct causal connection to the criterion, people are much less susceptible to the base rate fallacy.

Yet a more compelling question with respect to these three sets of factors is why one must even consider such circumstances in order to achieve appropriate utilization of base rate information. In the next section of this chapter, then, several theoretical frameworks for understanding the base rate fallacy will be discussed.

### THE BASE RATE FALLACY—A MATTER OF PERCEIVED RELEVANCE

Several sets of factors that may promote base rate utilization were discussed in the previous section. The focus of Part III is on the various explanations that have been proposed to account for the base rate fallacy. That is, why do people underutilize information about population base rates in favor of individuating information when making predictions and attributions about the category membership of a target individual?

Kahneman and Tversky (1973) originally proposed that the base rate fallacy reflects the tendency for people to rely on the representativeness heuristic when making probabilistic judgments. In the lawyer-engineer problem, for example, people judge probability by assessing the degree to which the individuating information (i.e., in the form of brief personality descriptions) is representative of or similar to their stereotype or preconceived theory about lawyers or engineers (Tversky, 1977). Because perceived representativeness is unaffected by the base rate, subjects ignored prior probability in favor of diagnostic as well as non-diagnostic individuating information.

An alternative explanation of the base rate fallacy was proposed by Nisbett and Borgida (1975; Nisbett, et al., 1976). They argued that base rate information, almost by its very nature, is abstract and pallid, and simply may lack the force to trigger cognitive work of any kind. In contrast, individuating information about a target individual is more concrete, vivid, and salient. Thus, base rate information may be underutilized in favor of



individuating information because the latter is more vivid and therefore more likely to be remembered and utilized than abstract, statistical information (Nisbett & Ross, 1980).

Subsequent research on base rate utilization and the base rate fallacy, however, suggests that neither of these explanations is entirely sufficient to account for the base rate fallacy phenomenon. Although people may engage in a feature-matching process in accordance with the representativeness heuristic for the lawyer-engineer problem, representativeness thinking alone does not account for *why* base rate information was not integrated into their probability judgments (Bar-Hillel, 1980). And although the vividness or salience explanation considers characteristics of both the base rate and individuating information, it is unable to account for the base rate fallacy when the base rate and individuating information are presented in a comparably abstract form (Tversky & Kahneman, 1980), or, for that matter, when both are presented in a comparably concrete manner (DiVitto & McArthur, 1978). In the latter context, it may be that consensus information is underutilized because it pertains to others' behavior and does not relate specifically or causally to an actor's behavior. Distinctiveness and consistency information, in contrast, are utilized because they are more specific and relevant to judgments about the particular actor's behavior (Nisbett & Ross, 1980).

A third and more compelling explanation of the base rate fallacy is based on the notion of causal thinking. As discussed in section II-J, Ajzen (1977) and Tversky and Kahneman (1980) have shown that the underutilization of base rates is a more general phenomenon than representativeness thinking. They have demonstrated that when base rate information is represented in a form that links the base rate to an intuitive causal theory and enables people to infer a probable cause of the criterion, base rates are more likely to affect social judgments. In contrast, base rate information that is not given this causal interpretation will be dominated by the presence of more causally relevant individuating information. In the absence of the latter, base rates will be used more appropriately. Thus, the causal relevance interpretation suggests that base rate information "which is not incorporated into a causal schema, either because it is not interpretable as an indication of propensity or because it conflicts with an established schema, is given little or no weight (Tversky & Kahneman, 1980, p. 71)."

But the causal relevance interpretation of the base rate fallacy does not explain, for example, certain research findings that suggest that non-causal base rates significantly affect judgments (Ajzen, 1977; Ginosar & Trope, 1980; Locksley et al., 1980). In order to account for such findings, Bar-Hillel (1980) has proposed an explanation of the base rate fallacy that is compatible with but encompasses the causal relevance interpretation. When making probability judgments, according to this view, people order the available information in terms of its relative *relevance* to the judgment

problem in question. If base rate information and individuating information appear equally relevant to a judgment problem, then the two will be combined in some manner to determine the final judgment. Information that appears to be more relevant to the problem, however, will tend to dominate information of lower relevance. Furthermore, according to Bar-Hillel (1980): "since less relevant items are discarded prior to any considerations of diagnosticity, an item of no diagnostic value, if judged more relevant, may dominate an item of high diagnosticity [p. 9]."

What determines informational relevance? Bar-Hillel suggests that one important factor might be *specificity*. That is, information about a population base rate may be regarded as less relevant than information about a specific subset of that population. Individuating information and causal base rates alike relate more specifically to the individual case in this manner. When both the base and individuating information, for example, were of low perceived relevance, Bar-Hillel found in one experiment that the median response was very close to the normative solution. When the base rate and individuating information were of equally high relevance, Bar-Hillel found that subjects tended to average the two types of information, although there was no consensual strategy for integrating the two types of information. When Bar-Hillel presented a judgment problem that juxtaposed two causal base rates, however, subjects' judgments were based entirely on the base rate that seemed to relate more specifically to a sub-population rather than to the overall population. Similarly, when subjects were given two non-causal base rates, her subjects used the more specific (sample-based) base rate instead of the more general (population-related) base rate. In this view, making a base rate seem causally relevant about specific members of a given population is just one means of increasing the perceived relevance of base rate information. It is not causality per se that is important, according to this view, but rather the fact that causal base rates are more specifically related (and therefore more relevant) to the target outcome than non-causal base rates. Thus, when base rate information pertains to a more specific subset of the population, it is perceived as more relevant and has more impact on judgments. People fall prey to the base rate fallacy, therefore, when they regard population base rates as low-relevance information, and individuating information as more specific and causally related to the target outcome.

Under certain circumstances, however, individuating information may not be regarded as more specific and causally relevant. Individuating information is more specific and causally relevant to the extent that it is congruent with the intuitive causal theories and implicit personality theories that people hold about the relationship between various personality characteristics and category membership. Any information that might affect the degree of congruence between these *a priori* theories and the individuating information, therefore, should reduce the extent to which the



individuating information is regarded as specific and causally relevant. In fact, Ginosar and Trope (1980) recently have demonstrated that the introduction of such information leads to more appropriate utilization of non-specific, abstract, non-causal base rate information.

By introducing inconsistent elements into the personality descriptions used by Kahneman and Tversky (1973), Ginosar and Trope were able to reduce the usefulness and relevance of individuating information for diagnosing category membership and thereby increase the likelihood of base rate utilization. In contrast, when the individuating information was internally consistent and congruent with intuitive causal theories, Ginosar and Trope's subjects, as expected, exhibited the base rate fallacy. By introducing inconsistent elements into the individuating information, Ginosar and Trope in effect disrupted subjects' well-developed reliance on individuating information and induced a more controlled mode of thinking (Langer, 1978). For example, Ginosar and Trope combined contradictory attributes from Kahneman and Tversky's (1973) engineer and lawyer descriptions into a single description (e.g., solves mathematical puzzles, interested in political and social issues, and argumentative). As a result, the salience and availability of the appropriate (but less preferred) sampling rule was increased and subjects recognized its relevance to the required judgment. Interestingly, although Ginosar and Trope were able to induce this shift to use of the base rate, they also found that subjects were less confident of their probability judgments when they used the more appropriate but less preferred sampling rule to make probability judgments.

The relevance explanation of the base rate fallacy, grounded in the judgment research literature, also clarifies the relationship between implicit consensus and sample-based consensus in attribution research. Simply put, when sample-based consensus is congruent with people's well-developed causal theories (based on either self-generated consensus or normative expectancies), consensus information will be regarded as more specific and causally relevant to attributional judgments and behavioral predictions about target individuals. In the Kulik and Taylor (1980) study, for example, this congruence enhanced population estimates and trait attributions in a normative direction. Such theory-guided reasoning, therefore, may result in what appears to be more normatively appropriate reasoning (Nisbett & Ross, 1980). When sample-based consensus is incongruent with intuitive causal theories, however, sample-based consensus will be regarded as less relevant to judgments about a specific target person. But the *degree* of incongruence must be taken into consideration. If people's ability to rely on intuitive theories is drastically diminished, as in the Ginosar and Trope (1980) study, then the availability of base rates and the likelihood of appropriate use is enhanced. But when the two types of information are more comparably relevant, as was the case in the Kulik and Taylor (1980) study, subjects clearly made an attempt to integrate the two in their judgments. In

such cases, regardless of the combinatorial rule that subjects use, sample-based information will be underutilized.

In sum, a consideration of various theoretical explanations of the base rate fallacy suggests a convergent framework for understanding the base rate fallacy in attributional and judgment contexts. Further research on other determinants of perceived relevance, various factors that may affect the assessment of specificity, on a critical role of diagnosticity, as well as the combinatorial rules that govern the relationship between implicit consensus and sample-based consensus, is clearly in order. The ramifications of an understanding of the base rate fallacy for other reasoning domains such as social stereotyping are discussed next.

### SOCIAL STEREOTYPES AND THE BASE RATE FALLACY

For investigators in the area of social cognition, one of the most important implications of the relationship between base rates and individuating information, as discussed in previous sections, is with respect to the domain of social stereotypes. The traditional approach to social stereotype research, for example, has focused on stereotyped assessment (Brigham, 1971). For over 40 years researchers have employed variants on the adjective checklist measure introduced by Katz and Braly (1933) and in general have obtained very consistent results. In the case of sex stereotypes, for example, the typical man has been described as more assertive, active, objective, rational, and competent than the typical woman, and the typical woman has been described as more passive, emotional, submissive, compassionate, and socially sensitive (Broverman, Vogel, Broverman, Clarkson, & Rosenkrantz, 1972; Ruble & Ruble, 1980; Spence & Helmreich, 1978). Sex stereotypic beliefs are thus well documented and fairly stable over time.

Given the stability and conceptual consistency of these results, one might assume that sex stereotypes affect social judgments; that is, knowledge of an individual's sex as well as information about his or her behavior should affect sex-associated trait attributions to that person. Attempts to demonstrate these effects of sex stereotypes on social judgments, however, have produced variable and inconsistent results. Locksley, et al. (1980), for example, recently conducted an experiment designed to test this hypothesis. Subjects read a purportedly real, first-person account of a college student's behavior in three ordinary problem situations. The trait characteristic of the target's behavior was systematically varied along a dimension that is a central and stable component of sex stereotypes: assertiveness/passivity. Information about the target's sex was fully crossed with information about the target's behavior, so that in two versions of the account the target's behavior was stereotypically consistent (i.e., a male who behaved in a consistently passive manner or a female who behaved in a con-



sistently assertive manner). There were also two versions in which the target's sex was unidentified (and the target was assertive or passive). Thus, one could assess not only the extent to which information about sex leads subjects to infer personality characteristics, but also the extent to which information about behavior leads subjects to infer sex and sex-associated personality characteristics. The dependent measures used in the experiment included predictions of the target's behavior in four novel problem situations and ratings of the target's personality on a set of sex stereotypic trait dimensions.

Surprisingly, information about the target's sex had no impact on predictions about the target's behavior or ratings of the target's personality characteristics. Regardless of whether the target's behavior was consistent or inconsistent with sex stereotypic expectations, subjects' judgments were based entirely on behavioral information. Moreover, this effect was not merely due to the consistency of the target's behavior; subjects relied just as much on behavioral information when the cross-situational consistency of the target's behavior was reduced. It thus appeared that subjects viewed sex as rather uninformative with respect to the target's personality and unknown behavior.

The contrast between the consistent results of sex stereotype assessment research and the failure of experiments such as these to obtain effects of sex stereotypes on social judgments is striking. If certain traits are strongly associated with sex, why doesn't information about sex increase the probability of attributing sex-associated traits to the target?

Locksley, Borgida, Brekke, and Hepburn (1980; also see Borgida, Locksley, and Brekke, in press) have suggested that research in the psychology of prediction can account for many of the apparent inconsistencies in stereotype research. One way to view stereotypes is to regard them as popular beliefs about base rates. As McCauley and Stitt (1978; McCauley, Stitt, & Segal, 1980) have suggested, stereotypes reflect the extent to which category membership can be regarded as a source of *probabilistic*, rather than perfect information about an individual's personality traits. For example, sex stereotypes may be seen as equivalent to statements like "On the average, women are less assertive than men" or "More men are assertive than women."

Once the distributional nature of stereotypes is recognized, research on the base rate fallacy becomes pertinent to an understanding of the circumstances under which stereotypes are likely to affect social judgments about a target individual. Stereotypes should affect judgments of individuals about whom little else is known but their social category, just as base rates affect judgments in the absence of individuating information (Nisbett & Ross, 1980). But as soon as specific, individuating information about an individual is known, and perceived as relevant, stereotypes may have minimal, if any, impact on judgments of that person (Kahneman &

Tversky, 1973; Nisbett & Ross, 1980); that is, individuating information should be regarded as more relevant to judgments about a particular individual than distributional beliefs.

Results of the Locksley et al. (1980) experiment are consistent with this explanation. Because subjects had a great deal of both diagnostic and non-diagnostic individuating information available, it is not surprising that sex stereotypes had no effects on judgments. Locksley et al. conducted a second experiment to provide a more direct test of this theoretical position. In the second experiment, subjects' sex stereotypic beliefs about the trait, assertiveness, were assessed in a distributional format. In addition, subjects were asked to indicate in terms of a percentage the likelihood that each one of six individuals was assertive. Two of these targets were identified by name only, in one case a female name and in the other case a male name. Thus for these targets, subjects had only *social category or base rate information* on which to base their judgments. Two more targets were identified by a male or female name in the context of a brief description of a single behavioral event, which was written so as to be subjectively *non-diagnostic* of the trait, assertiveness. The last two targets were identified by a female and a male name in the context of a brief description of a single behavioral event written to be subjectively *diagnostic* of the trait, assertiveness.

Results from this second experiment indicated, first, that subjects did have stereotypic beliefs about the distribution of assertiveness. The mean estimate of the percentage of males who are assertive was 56.1%, which differed reliably from the mean estimate of the percentage of females who are assertive (43.5%). Table 3.2 presents the mean judgments of the six targets, as a function of their sex and as a function of whether they were identified by social-category information only; by name in the context of a brief, non-diagnostic description of a single behavioral episode (non-diagnostic individuating information); or by name in the context of a brief, diagnostic description of a single behavioral episode (diagnostic individuating information).

TABLE 3.2  
Mean Attributions of Assertiveness as a Function of  
Target Case Information  
(From Locksley, Borgida, Brekke, & Hepburn, 1980)

Target Case Information	Sex of Target		t(df)
	Male	Female	
Social category information only	49.44	46.59	1.95(125) <sup>a</sup>
Nondiagnostic information	48.73	44.05	2.56(127) <sup>b</sup>
Diagnostic information	67.62	67.30	0.47(128)

<sup>a</sup>.06 > p > .05

<sup>b</sup>.05 > p > .01



As may be seen in Table 3.2, when subjects only had information about the target's sex, their judgments of the percentage of time in which the target was assertive were higher for the male target than for the female target. There was a similar difference between judgments for male and female targets when subjects were given low-relevance, non-diagnostic individuating information in addition to social category information. In contrast, when subjects had diagnostic individuating information in addition to social-category information, their judgments about male and female targets were virtually identical, even though these subjects had indicated that they believed that males in general were more likely to be assertive than females. These results are generally consistent with research by Nisbett, Zukier, and Lemley (in press) and Locksley et al. (in press) using non-controversial social stereotypes. These investigators, however, also found that even non-diagnostic individuating information is sufficient to weaken the effects of stereotypic beliefs on social judgments about a target individual.

In conclusion, the results of the experiments by Locksley et al. (1980) provide preliminary support for a general argument about the effects of social stereotypes that reflects the dynamics of the base rate fallacy as discussed in this chapter. Perhaps the single most important implication of this approach to social stereotypes that warrants further attention by researchers in social cognition is that social stereotypes may not exert as pervasive or as powerful an effect on social judgments of individuals as social psychologists traditionally have assumed (Hamilton, 1979). Social stereotypes and other knowledge structures (Cantor & Mischel, 1979; Nisbett & Ross, 1980) may affect recall of information and judgments of individuals about whom little else is known besides their social category. But as soon as individuating, subjectively diagnostic attributes of a person are known, stereotypes may have minimal impact on judgments on that person. Thus, people may sustain general prejudices while simultaneously viewing specific individuals with whom they frequently interact in a non-prejudicial manner.

#### ACKNOWLEDGMENTS

The preparation of this chapter was supported in part by NIMH Grant No. R01 MH30724-01 to Eugene Borgida. Nancy Brekke was supported by a University of Minnesota Graduate School Fellowship. The authors would like to acknowledge the helpful comments and suggestions of Icek Ajzen, Baruch Fischhoff, Steve Gangestad, Anne Locksley, Richard Nisbett, Mark Snyder, and Yaacov Trope on an earlier version of this manuscript.

#### REFERENCES

- Abelson, R. P. Script processing in attitude formation and decision making. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- 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.
- Ajzen, I., & Fishbein, M. Use and misuse of Bayes' theorem in causal attribution: Don't attribute it to Ajzen and Fishbein either. *Psychological Bulletin*, 1978, 85(2), 244-246.
- Ajzen, I., & Fishbein, M. Relevance and availability in the attribution process. In J. Jaspers, F. Fincham, & M. Hewstone (Eds.), *Attribution theory: Essays and experiments*. London: Academic Press, in press.
- Bar-Hillel, M. The role of sample size in sample evaluation. *Organizational Behavior and Human Performance*, 1979, 24, 245-257.
- Bar-Hillel, M. The base-rate fallacy in probability judgments. *Acta Psychologica*, 1980, 44, 211-233.
- Bartlett, F. C. *Remembering*. Cambridge, Mass.: Cambridge University Press, 1932.
- Borgida, E. Scientific deduction—Evidence is not necessarily informative: A reply to Wells and Harvey. *Journal of Personality and Social Psychology*, 1978, 36, 477-482.
- Borgida, E. Character proof and the fireside induction. *Law and Human Behavior*, 1979, 3, 189-202.
- Borgida, E., Locksley, A., & Brekke, N. Social stereotypes and social judgment. In N. Cantor & J. F. Kihlstrom (Eds.), *Cognition, social interaction, and personality*. Hillsdale, N.J.: Lawrence Erlbaum Associates, in press.
- Borgida, E., & Nisbett, R. E. The differential impact of abstract vs. concrete information on decisions. *Journal of Applied Social Psychology*, 1977, 7(3), 258-271.
- Brigham, J. C. Ethnic stereotypes. *Psychological Bulletin*, 1971, 76, 15-38.
- Broverman, I., Vogel, S., Broverman, D., Clarkson, F., & Rosenkrantz, P. Sex role stereotypes: A current appraisal. *Journal of Social Issues*, 1972, 28, 59-78.
- Cantor, N., & Mischel, W. Prototypes in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12). New York: Academic Press, 1979.
- Carroll, J. S. Judgments of recidivism risk: Conflicts between clinical strategies and base-rate information. *Law and Human Behavior*, 1977, 1(2), 191-198.
- Carroll, J. S. Judgments of recidivism risk: The use of base-rate information in parole decisions. In P. D. Lipsett & B. D. Sales (Eds.), *New directions in psycholegal research*. New York: Van Nostrand Reinhold, 1980.
- Carroll, J. S., & Siegler, R. S. Strategies for the use of base-rate information. *Organizational Behavior and Human Performance*, 1977, 19, 392-402.
- Cooper, J., Jones, E. E., & Tuller, M. S. Attribution, dissonance, and the illusion of uniqueness. *Journal of Experimental Social Psychology*, 1972, 8, 45-57.
- DeJong, W., & Amabile, A. M. *Consensus information and the foot-in-the-door effect*. Paper presented at the meeting of the American Psychological Association, New York, 1979.
- DiVitto, B., & McArthur, L. Z. Developmental differences in the use of distinctiveness, consensus, and consistency information for making causal attributions. *Developmental Psychology*, 1978, 14(5), 474-482.
- Dyck, R. J., & Rule, B. G. Effect on retaliation of causal attributions concerning attack. *Journal of Personality and Social Psychology*, 1978, 36(5), 521-529.
- Feldman, N. S., Higgins, E. T., Karlovac, M., & Ruble, D. N. Use of consensus information in causal attribution as a function of temporal presentation and availability of direct information. *Journal of Personality and Social Psychology*, 1976, 34, 694-698.
- Fischhoff, B. Attribution theory and judgment under uncertainty. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 1). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Fischhoff, B., & Lichtenstein, S. Don't attribute this to Reverend Bayes. *Psychological Bulletin*, 1978, 85(2), 239-243.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. Subjective sensitivity analysis. *Organizational Behavior and Human Performance*, 1979, 23, 339-359.



- Frieze, I., & Weiner, B. Cue utilization and attributional judgments for success and failure. *Journal of Personality*, 1971, 39, 591-605.
- Garland, J. H., Hardy, A., & Stephenson, L. Information search as affected by attribution type and response category. *Personality and Social Psychology Bulletin*, 1975, 1, 612-615.
- 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.
- Goldberg, L. R. Simple models or simple processes? Some research on clinical judgments. *American Psychologist*, 1968, 23, 483-496.
- Hamill, R., Wilson, T. D., & Nisbett, R. E. Insensitivity to sample bias: Generalizing from atypical cases. *Journal of Personality and Social Psychology*, 1980, 39, 578-589.
- Hamilton, D. L. A cognitive attributional analysis of stereotyping. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12). New York: Academic Press, 1979.
- Hammerton, M. A case of radical probability estimation. *Journal of Experimental Psychology*, 1973, 101, 252-254.
- Haney, C. *Consensus information and attitude change: Modifying the effects of counter-attitudinal behavior with information about the behavior of others*. Unpublished manuscript, University of California, Santa Cruz, 1979.
- Hansen, R. D. *Describing attributors' use of causal information: Can the scientist and the naive perceiver agree?* Unpublished manuscript, Oakland University, 1979.
- Hansen, R. D., & Donoghue, J. M. The power of consensus: Information derived from one's own and others' behavior. *Journal of Personality and Social Psychology*, 1977, 35(5), 294-302.
- Hansen, R. D., & Lowe, C. A. Distinctiveness and consensus: The influence of behavioral information on actors' and observers' attributions. *Journal of Personality and Social Psychology*, 1976, 34, 425-433.
- Hansen, R. D., & Stonner, D. M. Attributes and attributions: Inferring stimulus properties, actors' dispositions, and causes. *Journal of Personality and Social Psychology*, 1978, 36, 657-667.
- Harvey, J. H., Cacioppo, J. T., & Yasuna, A. Temporal pattern of social information and self-attribution of ability. *Journal of Personality*, 1977, 45(2), 281-296.
- Heider, F. *The psychology of interpersonal relations*. New York: Wiley, 1958.
- Jones, E. E., & McGillis, D. Correspondent inferences and the attribution cube: A comparative reappraisal. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 1). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Kahneman, D., & Tversky, A. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 1972, 3, 430-454.
- Kahneman, D., & Tversky, A. On the psychology of prediction. *Psychological Review*, 1973, 80, 237-251.
- Karlovac, M., Feldman, N., Higgins, E. T., & Ruble, D. N. *The effect of consensus information on causal attributions*. Paper presented at the meeting of the American Psychological Association, Washington, D.C., 1976.
- Kassin, S. M. Base rates and prediction: The role of sample size. *Personality and Social Psychology Bulletin*, 1979, 5(2), 210-213. (a)
- Kassin, S. M. Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology*, 1979, 37(11), 1966-1981. (b)
- Kassin, S. M., & Lowe, C. A. *Content and structure of one-sentence descriptions of behavior as determinants of observers' attributions*. Paper presented at the meeting of the American Psychological Association, Washington, D.C., 1976.
- Katz, D., & Braly, K. W. Racial stereotypes of one hundred college students. *Journal of Abnormal and Social Psychology*, 1933, 28, 280-290.
- Kelley, H. H. Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska Symposium on Motivation* (Vol. 15). Lincoln, Neb.: University of Nebraska Press, 1967.
- Kelley, H. H. Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown, N.J.: General Learning Press, 1972. (a)
- Kelley, H. H. Causal schemata in the attribution process. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown, N.J.: General Learning Press, 1972. (b)
- Kelley, H. H. The process of causal attribution. *American Psychologist*, 1973, 28, 107-128.
- Kruglanski, A. W. The endogenous-exogenous partition in attribution theory. *Psychological Review*, 1975, 82, 387-406.
- Kruglanski, A. W., Hamel, J. Z., Maides, S. A., & Schwartz, J. M. Attribution theory as a special case of lay epistemology. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 2). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.
- Kulik, J. A., & Taylor, S. E. Self-monitoring and the use of consensus information. *Journal of Personality*, in press.
- Kulik, J. A., & Taylor, S. E. Premature consensus on consensus?: Effects of sample-based versus self-based consensus information. *Journal of Personality and Social Psychology*, 1980, 38, 871-878.
- Langer, E. Rethinking the role of thought in social interaction. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 2). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1978.
- Lay, C., Burrton, B., & Jackson, D. Base rate and information value in impression formation. *Journal of Personality and Social Psychology*, 1973, 28, 390-395.
- Locksley, A., Borgida, E., Brekke, N., & Hepburn, C. Sex stereotypes and social judgment. *Journal of Personality and Social Psychology*, 1980, 39, 821-831.
- Locksley, A., Hepburn, C., & Ortiz, V. Social stereotypes and judgments of individuals: An instance of the base-rate fallacy. *Journal of Experimental Social Psychology*, in press.
- Lowe, C. A., & Kassin, S. M. On the use of consensus information: Prediction, attribution, and evaluation. *Personality and Social Psychology Bulletin*, 1977, 3, 616-619.
- Lyon, D., & Slovic, P. Dominance of accuracy information and neglect of base rates in probability estimation. *Acta Psychologica*, 1976, 40, 287-298.
- 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.
- Markus, H. Self-schemata and processing information about the self. *Journal of Personality and Social Psychology*, 1977, 22, 171-193.
- Martin, J., & Powers, M. *If case examples provide no proof, why under-utilize statistical information?* Paper presented at the meeting of the American Psychological Association, New York, 1979.
- McArthur, L. A. The how and what of why: Some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology*, 1972, 22, 171-193.
- McArthur, L. Z. The lesser influence of consensus than distinctiveness information on causal attributions: A test of the person-thing hypothesis. *Journal of Personality and Social Psychology*, 1976, 33, 733-742.
- McCauley, C., & Stitt, C. L. An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology*, 1978, 36(3), 929-940.
- McCauley, C., Stitt, C. L., & Segal, M. Stereotyping: From prejudice to prediction. *Psychological Bulletin*, 1980, 87, 195-208.



- Meehl, P. E. *Clinical vs. statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis, Minn.: University of Minnesota Press, 1954.
- Meehl, P. E., & Rosen, A. Antecedent probability and the efficacy of psychometric signs, patterns, or cutting scores. *Psychological Bulletin*, 1955, 52, 194-216.
- Neisser, U. *Cognition and reality: Principles and implications of cognitive psychology*. San Francisco: Freeman, 1976.
- Nisbett, R. E., & Borgida, E. Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, 1975, 32, 932-943.
- Nisbett, R. E., Borgida, E., Crandall, R., & Reed, H. Popular induction: Information is not necessarily informative. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Nisbett, R. E., & Ross, L. *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, N.J.: Prentice-Hall, 1980.
- Nisbett, R. E., Zukier, H., & Lemley, R. The dilution effect: Nondiagnostic information weakens the implications of diagnostic information. *Cognitive Psychology*, in press.
- Orvis, B. R., Cunningham, J. D., & Kelley, H. A. A closer examination of causal inferences: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology*, 1975, 32, 605-616.
- Pilkonis, P. A. The influence of consensus information and prior expectations on dispositional and situational attributions. *Journal of Social Psychology*, 1977, 101, 267-279.
- Rosenberg, S., & Sedlak, A. Structural representations of implicit personality theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6). New York: Academic Press, 1972.
- Ross, L. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10). New York: Academic Press, 1977.
- Ross, L. Some afterthoughts on the intuitive psychologist. In L. Berkowitz (Ed.), *Cognitive theories in social psychology*. New York: Academic Press, 1978.
- Ross, L., Amabile, T. M., & Steinmetz, J. L. Social roles, social control, and biases in social perception processes. *Journal of Personality and Social Psychology*, 1977, 35, 485-494.
- Ross, L., Greene, D., & House, P. The false consensus phenomenon: An attributional bias in self-perception and social perception processes. *Journal of Experimental Social Psychology*, 1977, 13, 279-301.
- Ruble, D. N., & Feldman, N. S. Order of consensus, distinctiveness, and consistency information and causal attributions. *Journal of Personality and Social Psychology*, 1976, 34, 930-937.
- Ruble, D. N., Feldman, N. S., Higgins, E. T., & Karlovac, M. Locus of causality and the use of information in the development of causal attributions. *Journal of Personality*, 1979, 47, 595-614.
- Ruble, D. T., & Ruble, T. L. Sex stereotypes. In A. G. Miller (Ed.), *In the eye of the beholder: Contemporary issues in stereotyping*. New York: Holt, Reinhart & Winston, 1980.
- Schank, R., & Abelson, R. P. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Schneider, D. J., Hastorf, A. H., & Ellsworth, P. C. *Person perception* (2nd ed.). Reading, Mass.: Addison-Wesley, 1979.
- Shafer, G. *A mathematical theory of evidence*. Princeton, N.J.: Princeton University Press, 1976.
- Shah, S. A. Dangerousness: A paradigm for exploring some issues in law and psychology. *American Psychologist*, 1978, 33, 224-238.
- Slovic, P., Fischhoff, B., & Lichtenstein, S. Behavioral decision theory. *Annual Review of Psychology*, 1977, 28, 1-39.
- Slovic, P., & Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, 1971, 6, 649-744.
- Snyder, M. Self-monitoring processes. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12). New York: Academic Press, 1979.
- Snyder, M., & Swann, W. B., Jr. Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology*, 1978, 36, 1202-1212.
- Spence, J., & Helmreich, R. *Masculinity and femininity*. Austin, Texas: University of Texas Press, 1978.
- Stich, S. P., & Nisbett, R. E. *Philosophy of science*. Ann Arbor, Mich.: University of Michigan Press, in press.
- Taylor, S. E., & Crocker, J. C. Schematic bases of social information processing. In E. T. Higgins, P. Herman, & M. P. Zanna (Eds.), *The Ontario Symposium on Personality and Social Psychology* (Vol. 1). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1980.
- Taylor, S. E., & Fiske, S. T. Salience, attention, and attribution: Top of the head phenomena. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 11). New York: Academic Press, 1978.
- Tribe, L. H. Trial by mathematics: Precision and ritual in the legal process. *Harvard Law Review*, 1971, 84, 1329-1393.
- Tversky, A. Features of similarity. *Psychological Review*, 1977, 84, 327-352.
- Tversky, A., & Kahneman, D. Belief in the law of small numbers. *Psychological Bulletin*, 1971, 76, 105-110.
- Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, 5, 207-232.
- Tversky, A., & Kahneman, D. Judgment under uncertainty: Heuristics and biases. *Science*, 1974, 185, 1124-1131.
- Tversky, A., & Kahneman, D. Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1980.
- Wegner, D. M., & Vallacher, R. R. *Implicit psychology*. New York: Oxford University Press, 1977.
- Weir, M. W. Developmental changes in problem-solving strategies. *Psychological Review*, 1964, 71, 473-490.
- Wells, G. L., & Harvey, J. H. Do people use consensus information in making causal attributions? *Journal of Personality and Social Psychology*, 1978, 35(5), 279-293.
- Wells, G. L., & Harvey, J. H. Naive attributors' attributions and predictions: What is informative and when is an effect and effect? *Journal of Personality and Social Psychology*, 1978, 36, 483-490.
- Williams, P. M. On a new theory of epistemic probability. *The British Journal for the Philosophy of Science*, 1978, 29, 375-387.
- Zuckerman, M. Actions and occurrences in Kelley's cube. *Journal of Personality and Social Psychology*, 1978, 36(b), 647-656. (a)
- Zuckerman, M. Use of consensus information in prediction of behavior. *Journal of Experimental Social Psychology*, 1978, 14, 163-171. (b)
- Zuckerman, M., & Mann, R. W. The other way around: Effects of causal attributions on estimates of consensus, distinctiveness, and consistency. *Journal of Experimental Social Psychology*, 1979, 15, 582-597.