Poznań Studies in the Philosophy of the Sciences and the Humanities
1991, Vol. 21, pp. 225-249

Gerd Gigerenzer

ON COGNITIVE ILLUSIONS AND RATIONALITY

The controversy as to whether mathematical probability theory has anything to do with human rationality is as old as the theory itself (Daston 1988). The latest episode of this continuing controversy was sparked by developments in psychology, and has contributed a new, distinctly psychological wrinkle to the time-worn themes: the analogy between that stock-in-trade of perceptual psychology, the visual illusion, and what both sides of the current controversy have termed "cognitive illusions".[1] Amos Tversky and Daniel Kahneman (1974) introduced their research program on judgment under uncertainty with the analogy between biases in probabilistic reasoning and visual illusions, and others chose the generic term "cognitive illusions" for these biases in order "to suggest that these phenomena are quite similar to a variety of perceptual illusions extensively studied by psychologists" (Edwards & von Winterfeldt, 1986, p. 643). L. Jonathan Cohen has challenged the repeated claims that intuitive reasoning is biased and irrational, e.g., that experimental results have "bleak implications for human rationality" and that the human mind lacks "the correct programs for many important judgmental tasks" (Nisbett & Borgida, 1975, p. 935; Slovic, Fischhoff & Lichtenstein, 1976). Nevertheless, Cohen, too, described reasoning fallacies "as cognitive illusions ... to invoke the analogy with visual illusions" (Cohen, 1981, p. 324). Both sides, although diametrically opposed in almost all other respects, share the analogy between cognitive and visual illusions.

Ironically, the analogy originally pointed in the other direction, from visual illusions to statistical reasoning. The visual space is tridimensional whereas the retina, the perceiving organ, is bidimensional. How does the brain get three dimensions out of two?, asked George Berkeley in his *Essay Towards a New Theory of Vision* (1709). From Berkeley to the present, many theories have used reasoning, in particular, statistical reasoning, as an analogy for understanding what perception does. Hermann von Helmholtz spoke of "unconscious inferences", Egon Brunswik suggested parameter estimation by multiple regression as the possible mechanism of perception, and R. L. Gregory compared the inferential mechanism of perception to Fisherian statistical inference and Bayesian probability revision. For these theorists, perception is "betting against reality", and what we perceive is essentially illusory, where illusions are understood as deviations from physical reality. Thus, statistical reasoning is and has been a key analogy for understanding perception and perceptual illusions. Since the 1970s, the direction of the analogy has been reversed. In order to understand intuitive inductive reasoning, in particular so-called errors in probabilistic reasoning, visual illusions have now become the analogy of choice. A paradoxical, a circular, or a fruitful inversion of analogy?

In this paper, it will become clear that the analogy with visual illusions means different things to those who attack and to those who defend rationality in intuitive reasoning. I will first lay out three aspects of visual illusions that are relevant for the analogy. Second, I will argue that two of these aspects hold for neither version of the analogy. Third, I will suggest that the third aspect, which is largely ignored by both sides (though less so by Cohen), has real potential for an understanding of the nature of intuitive probabilistic reasoning. Developing this third point further, I will show how experimental studies can profit from philosophers' conceptual distinctions. This research vindicates Cohen's claim that even untutored intuition is capable of making conceptual distinctions of the sort philosophers make.

Visual illusions

One of the best-known visual illusions is the Müller-Lyer arrow figure (Figure 1). Visual illusions are a subset of perceptual illusions, to be distinguished from purely optical illusions such as the bent-stick-in-water

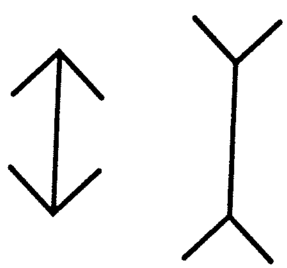

Figure 1: One variant of the Müller-Lyer arrow figure (originally devised by F. C. Müller-Lyer in 1889 in fifteen variants). The two vertical lines are of the same length, although the inward-going and the outward-going arrows make the lines appear smaller and larger, respectively.

effect, and sensory illusions, such as after-images occurring after intense stimulation of the retina. There is a great divide between two major accounts of perception with respect to the nature and the role of perceptual illusions. In the Helmholtzian view, as I mentioned above, only limited and uncertain sensory data is available, and to perceive means to go beyond the information given. Perception interprets incomplete information, and perceptual illusions are both the necessary consequence of this impossible task and the key to understanding the inference mechanism. In the Gibsonian view, in contrast, all the information needed for veridical perception is available in the "ambient array of light", and is simply "picked up" by the observer. Visual illusions are of no or little interest; in fact, there should be no uncertainty and consequently, no illusions. If visual illusions are dealt with, they are attributed to a restricted and unnatural stimulus environment, such as that created by experimenters in traditional laboratory experiments.

For the purpose of this paper, I will list three aspects of the Helmholtzian view of visual illusions which I consider relevant to the analogy with cognitive illusions.

(1) *The "uncontroversial norm" aspect.* In visual illusions, there exists an uncontroversial norm, physical measurement, which defines systematically deviating judgments as illusions. For instance, in the Müller-Lyer arrow figure, this norm is the physical length of the lines.

(2) *The "stability" aspect.* Visual illusions are normally stable. Although visual illusions seem to be a function of early experience and age, the direction and extent of the illusion is relatively stable in humans and probably also in many animals (Gregory, 1974, p. 360). What can be done? Not much, even after having measured the lines, the illusion persists despite better knowledge.[2]

(3) *The "context and prior knowledge" aspect.* Since to perceive means to infer something that is not directly available in the proximal (sensory) information, these inferences have to rely on context information and prior knowledge. This inference mechanism produces both veridical perception (perceptual constancies) and perceptual illusions. According to Gregory (1974), for instance, the Müller-Lyer illusion originates from the way the context (the arrows at the ends of the lines) are taken into account. The arrows automatically generate a three-dimensional representation of the two-dimensional figure, inward-going arrows suggesting the line as the outside (front) corner of a three-dimensional object, and outward-going arrows suggest the line as the inside (back) corner.

Illusions are not due to contingent limitations of the brain, but rather to the nature of the brain's tasks; i.e., to read information from images which at best contain such information only implicitly. Illusions are, therefore, a *necessary consequence of the task*. Gregory (p. 376) conjectures that the brain possesses probably the best perceptual system available and that illusions will be a necessary part of all efficiently designed visual machines.

This is of course a simplified account of visual illusions, relying solely on the Helmholtzian tradition. Since, however, those who originally invoked the analogy have themselves referred to the importance of studying errors and to the betting metaphor of perception (Kahneman & Tversky, 1982b, p. 512), this may be a fair perspective to start with.

Cognitive illusions

Consider the *neglect of base rates* in one of the most well-known demonstrations by Kahneman and Tversky (1973). One group of subjects was given the following problem, known as the Engineer-Lawyer Problem:

> A panel of psychologists have interviewed and administered personality tests to 30 engineers and 70 lawyers, all successful in their respective fields. On the basis of this information, thumbnail descriptions of the 30 engineers and 70 lawyers have been written. You will find on your forms five descriptions, chosen at random from the 100 available descriptions. For each description, please indicate your probability that the person described is an engineer, on a scale from 0 to 100.
> The same task has been performed by a panel of experts, who were highly accurate in assigning probabilities to the various descriptions. You will be paid a bonus to the extent that your estimates come close to those of the expert panel (p. 241).

Subjects in a second group were given the same instructions but with inverted base rates — that is, they were told that there were 70 engineers and 30 lawyers. Both groups were given the same descriptions, for example:

"Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies, which include home carpentry, sailing, and mathematical puzzles" (p. 241).

The authors concluded that base rates were largely ignored by their subjects, since mean probability judgments were about the same in both groups. This base rate neglect was considered a systematic bias in reasoning and was explained by a general heuristic called representativeness: people judge the posterior probability (that a person is an engineer given

the above information) largely by the similarity between the descriptio and their stereotype of an engineer. The generality of the heuristic an the neglect of base rates seemed to be confirmed by the fact that th median probability judgments for an uninformative description, "Dick' were also the same (.50), although this description was constructed t contain only worthless information with respect to engineers and lawyer: Dick was described as, *inter alia*, a man of high ability, high motivatior and as well-liked by his colleagues. Only when no description was giver did the probability judgments follow the base rates.

According to the "systematic bias" view, this neglect of base rates satis fies the first two of the above aspects of visual illusions. First, there seem to be an uncontroversial norm that justifies calling a systematic deviatio from that norm properly a cognitive illusion, in the sense of a bias in rea soning. In the present case this norm is Bayes' theorem (which Kahnemar and Tversky used to calculate the "correct" differences between the prob ability judgments in both groups). More generally, the norm is referred tc as the rules of statistics or "the rules of probability in intuitive judg ments" (Tversky & Kahneman, 1977, p. 173). Second, this cognitive illu sion seems to be a stable one, at least in this task, it is "a highly robus effect", as we are told by several experimenters (e.g., Tversky & Kahneman, 1977, p. 176; von Winterfeldt & Edwards, 1986, p. 535). Even critics talk about an "overall *tendency* to neglect-base rates" (Adler, 1984, p. 173). Let us now have a closer look at the analogy.

Uncontroversial norms: a single yardstick for rationality?

This first aspect of the perceptual/cognitive analogy maintains: For the reasoning problems posed, there exists a single correct answer just as there does for perceptual problems. A deviation between judgment and normative answer, if systematic, defines an illusion, visual or cognitive. Kahneman and Tversky (1982a, p. 493) explicitly draw this analogy between perceptual and cognitive illusions: "The presence of an error of judgment is demonstrated by comparing people's responses either with an established fact (e.g., that the two lines are equal in length) or with an accepted rule of arithmetic, logic, or statistics". In this analogy, standard probability theory is to inductive reasoning what a ruler is to perceptual judgment. This first aspect of the analogy is at best highly problematic, and, at worst, misleading at three levels.

Level of formal theories: There exist alternatives to standard probability theory (which is based on the Kolmogoroff axioms) as the norm of reasoning under uncertainty. Among others, such non-standard theories have been proposed by Jakob Bernoulli, Johann Heinrich Lambert,

L. Jonathan Cohen, Henry E. Kyburg, and Glenn Shafer. There are good reasons to prefer several alternative theories to a single one, if we want to model what we commonly call "rationality". Among these is the evolutionary reason that it could well be preferable to have a whole toolbox of probability approaches to varied environments, rather than a single specialized tool that works very well only for a few environments. A different kind of reason is that the meaning of "rationality" is notoriously various; it may for instance refer to prediction or to explanation. Standard probability theory does well at defining what it means to minimize errors of prediction in the long run, whereas non-additive probabilities can do well at modeling how the evidence of particular observations bears on several hypotheses.

Level of interpretation: Even within standard probability theory, the question of what belongs to its domain of application is open to interpretation. Is the theory about degrees of belief, relative frequencies, propensities, or about something else? I will argue later that these distinctions are relevant to understanding intuitive reasoning.

Level of representation: Even within standard probability theory and a particular domain of application, the question whether the structure of a particular reasoning problem can be represented or not by a specific statistical model must be carefully examined.

All three levels are important in answering the question of whether there exists exactly one answer to a given problem, or if not, what alternative solutions exist. None of these three complications exist for visual illusions. For measuring the length of two lines, for instance, alternative physical theories, interpretations of domain, and the question of representation play no role. The point Cohen made, and, which I want to follow up here, is that for many reasoning problems, in contrast, more than one resonable answer can be found, even within the narrow confines of a given formal theory and a given interpretation. This holds especially for real-world problems, but also for textbook problems of the real-world kind, and to a lesser degree for problems of the schematic urn-and-balls problems type. In general, the more the urn-and-balls problems are filled with content and context, the more alternative norms can be developed. To make this general statement more precise, I will attempt to state conditions under which we may be justified in claiming the contrary, that is, that a given reasoning problem does indeed have only one correct answer. To make my point stronger, I will consider only the level of representation, assuming consensus about formal theory and interpretation.

I propose that it is justified to claim that a reasoning problem has one and only one correct answer (as is regularly claimed in the literature on biases and heuristics without justification) if and only if two conditions hold, which I call "concept isomorphism" and "structural isomorphism" (Gigerenzer & Murray 1987, chapter 5). I will argue that in order to check these conditions, not only formal structure, but also context and content-related prior knowledge are relevant. I will use Bayes' theorem and textbook problems for illustration.

Concept isomorphism holds if and only if there is a one-to-one mapping between the concepts in a statistical model (e.g., the prior probabilities and the likelihoods in Bayes' theorem) and the *relevant* concepts of a reasoning problem (e.g., base rates mentioned for engineers and lawyers). This means that each statistical concept can be identified with one and only one concept in the problem, and vice versa, and no concept on either side remains. Figure 2 illustrates a situation in which concept isomorphism is violated. For instance, there may be two or more candidates for the prior probability; or some relevant context information in the problem may have no representation in the statistical model.

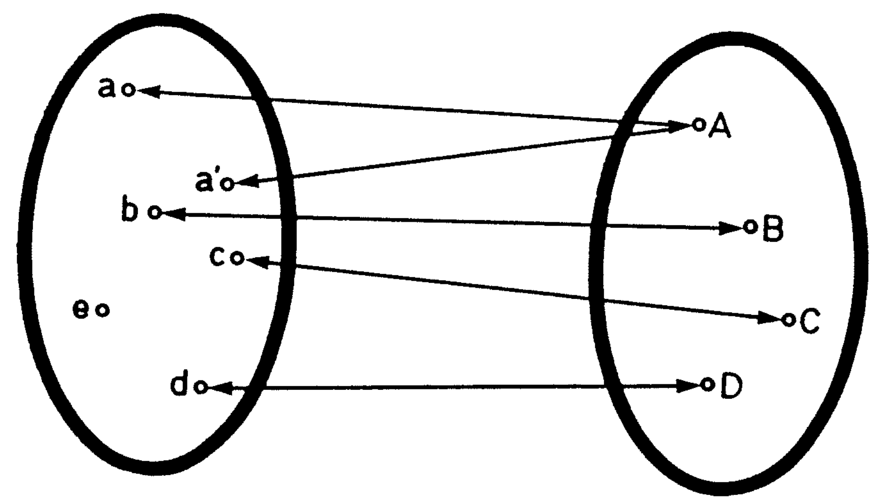

Figure 2: Points in the left circle represent relevant concepts in a reasoning problem, those in the right circle represent concepts of a statistical model. Concept isomorphism is violated in two respects, there exist two different candidates *a* and *a'* for *A*, and the concept *e* is not represented in the statistical model.

To illustrate the latter: In the instructions to the Engineer-Lawyer Problem, subjects were told that a panel of experts "were highly accurate in assigning probabilities to the various descriptions". What can this mean? A subject who considers this information as relevant might infer since being "highly accurate" could only mean to assign probabilities close to zero or to one, this is a hint that the descriptions are highly informative and reliable, and that the solution lies in the descriptions, if only you could decipher them. Thus, this information can be read to imply, in essence, "use probabilities close to zero or one only, if you want to get a large bonus". There is, however, no way to account for this information in the standard application of Bayes' theorem to the problem. The general point is: What qualifies as "relevant" information in the reasoning problem depends on prior knowledge, and is not always exhausted by the concepts in the statistical model.

Structural isomorphism postulates that the structural assumptions underlying a statistical model hold in the reasoning problem (structural homomorphism), and, vice versa, that the structural properties of the reasoning problem are represented by a similar structure in the statistical model. For instance, structural assumptions of the formal model that must hold in the Engineer-Lawyer Problem are that the hypotheses considered are mutually exclusive and exhaustive, and that the descriptions have been obtained by random sampling from the population to which the base rates refer to. Among structural properties of a reasoning problem that may not be represented in Bayes' theorem are temporal and spatial relations. To see the latter point consider the notorious Cab Problem (Tversky & Kahneman, 1980), which describes a hit-and-run accident at night involving a cab and gives the following information:

(i) 85% of the cabs in the city are Green and 15% are Blue.

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

Question: What is the probability that the cab involved in the accident was Blue rather than Green? (p. 62).

Tversky and Kahneman look at this problem with the formal structure of Bayes' theorem firmly in mind, and identify the prior probabilities with the numbers given in (i) and the likelihoods p("Blue"/Blue) and p("Blue"/Green) with those given in (ii), .80 and .20, respectively (p("Blue"/Blue) is the probability that the witness says "Blue" if the cab is actually Blue).[3] Inserting these values into Bayes' theorem, they calculate p(Blue/"Blue") = .41, which they consider as the single normative answer. Therefore their subjects' median judgment of .80 appears to be an error of reasoning. Again, the error is identified as neglect of base rates. However, if one looks at Bayes' theorem (as applied) with the structure of the *problem* in mind, one realizes that one relevant structural aspect of the problem is not represented in the Bayesian modeling: the temporal sequence of the witness' two judgments, the first one at time t_1, the night of the accident, and the second one at time t_2, at the court's test. The likelihoods specified in (ii) characterize the witness' performance at t_2, but not at t_1. At issue are, however, the likelihoods at t_1. Several arguments can be made that and how the likelihoods have changed from t_1 to t_2. For instance, assume that both at t_1 and t_2 the witness knew the base rates and judged in a way that minimized judgmental errors, that is, the sum p("Blue"/Green)p(Green) + p("Green"/Blue) p(Blue). Since we are informed in (ii) that the base rates changed from t_1 to t_2, from 15% to 50%, we are then forced to conclude that the likelihoods also changed from t_1 to t_2. Calculating the likelihoods for t_1 and inserting them into Bayes' theorem gives p(Blue/"Blue") = .82, an alternative solution to Tversky and Kahneman's, who ignored the difference between t_1 and t_2 and calculated .41. Incidentally, .82 is close to the median judgment of untutored people (.80), from which latter figure base rate neglect was concluded. Paradoxically, .82 is calculated on the assumption that base rates are taken into account (see Birnbaum, 1983, Gigerenzer & Murray, 1987, pp. 168-173). Of course, there are other reasonable assumptions for a "criterion shift" from t_1 to t_2, which lead to alternative reasonable answers. This one, however, may suffice for making the point.

The example shows that structure that is in the problem but not in the formula matters for the claim that there is a single normative answer. Note that this holds despite the fact that we agreed on a particular formal theory and interpretation. Reasoning problems that are structurally equivalent from the point of view of a statistical model can have different surplus structures.

Conceptual and structural isomorphism are not easily attained for textbook reasoning problems, and this holds *a fortiori* for their real-world versions, such as guessing someone's profession in the cafeteria, or a legal judgment as to whether someone committed a hit-and-run accident. Personal judgment about what are the relevant concepts in a problem situation and about whether isomorphism holds is indispensable, and context and prior knowledge are necessary and legitimate cues in making this judgment. Thus, to speak of a single correct answer in reasoning in analogy to a single measure of length for judging visual accuracy, is to presuppose consensus between experimenters and subjects at all three levels. This shows why the analogy between cognitive and visual illusions is highly misleading with respect to the existence of an uncontroversial norm. We may not legitimately speak of a single norm unless we have first made all our premisses, such as random sampling, clear to the untutored layman. The results of the Engineer-Lawyer Problem reported in the next section indicate that such clarification does indeed matter for untutored judgments. The practical consequence of this lack of analogy is: Try to explain people's judgments as they are, not their deviation from some preconceived norm.

Cohen's use of the analogy with visual illusions is different. He challenges standard probability theory, which he calls Pascalian probability, as the only norm for rational reasoning. The alternative norm he proposes is called Baconian probability. Thus, for Cohen at least two alternative theories of probability exist, Pascalian and Baconian, and several interpretations of their domains are possible, but again only one norm for rationality exists. This norm is the intuition of our fellow adults, untutored in probability theory. Given this norm, systematic errors in probabilistic reasoning (i. e., competence rather performance errors) are by definition ruled out – except for the immature, the senile, etc. If, nevertheless, errors do occur, Cohen attributes them either to "cognitive illusions" or to lack of intelligence or mathematical education. Cohen's cognitive illusions are created by the experimenters, not committed by the subjects. How do experimenters do it? Like the Gibsonians, Cohen (1981, p. 324) proposes that such illusions are created by using (i) experimental tasks not representative of the normal conditions of life (e.g., check-ups are excluded), and (ii) unfamiliar or abstract material. For Cohen, the analogy with visual illusions means deception. The experimenter is compared to a conjurer who relies "on the visual inattentiveness of those who are watching, to hold the latter back from obtaining an appropriate additional input to their visual information-processing operation" (p. 325). Note that Cohen's use of the term "cognitive illusion" differs significantly from the notion of systematic biases or fallacies in reasoning. So does his notion of "visual illusion". Cohen is a Gibsonian in so far as illusions are attributed to the experimenter or some special conditions rather than to the very nature of the brain's task. Just as in Gibson's realism the organism simply "picks up" the appropriate information from the ambient array of light, so in Cohen's rationalism the organism "picks up" appropriate rules of reasoning from its untutored competence. Both premisses are in principle immune to falsification, it seems to me. But here the analogy ends. A realist is obliged to postulate one and only one physical world, but Cohen does not postulate *one* kind of probability theory as the standard for rationality.

For the Engineer-Lawyer Problem, Cohen (1979) proposes Baconian probability as the norm. The Baconian probability that the person described is an engineer is defined as the inductive reliability of the generalization that all persons described this way are engineers. The inductive reliability is a function of the outcome of tests of the generalization (hypothesis) with an appropriate list of inductively relevant variables. For testing the hypothesis that someone is an engineer, for example, such relevant variables may be the kind of hobbies and political attitudes; in general these are factors with potential causal impact. No measurement function is available for inductive reliability (Cohen, 1979, p. 390), although it may be ranked depending on the tests a hypothesis succeeded in passing. The Baconian probability that the person described is an engineer depends solely on the causal evidence in the description, and not on the base rates of engineers. Therefore, if a description is simply irrelevant to the traits of an engineer or of a lawyer, then the two generalizations "all descriptions of type D are engineers" and "all descritions of type D are lawyers" will both have zero-grade inductive reliability. That is, both of the two Baconian probabilities p_I(engineer/description) and p_I(lawyer/description) may be zero. This reveals a major contrast with conventional or Pascalian probabilities: Baconian probabilities run from nonproof to proof, whereas Pascalian probabilities run from disproof to proof.

To summarize: the "uncontroversial norm" aspect of the analogy with visual illusions is indispensable for the argument that so-called biases in probabilistic reasoning reveal that the untutored mind is running on shoddy software, that is, on programs that work only with a handful heuristics. I have argued, however, that this aspect presupposes consensus on formal theory, interpretation, and isomorphic representation. Therefore, this aspect of the analogy is not valid, unless consensus is established (for a particular situation). Consequently, the frequent joint classification of errors in probabilistic reasoning with errors in memory or in intuitive physics under the heading "cognitive illusions" (e. g., Edwards & von Winterfeldt 1986) is not valid, too, for the same reasons. For Cohen, who does not believe in standard probability theory as the norm of rationality in the first place, this important aspect of visual illusions is either irrelevant, if probability theory is considered as norm, or circular, if intuitive reasoning competence is considered the norm of rationality.

Stability

Stability of cognitive illusions is essential for the view that Cohen attacks. If untutored intuition lacks the correct programs for probabilistic reasoning, and relies on a few general heuristics such as representativeness and availability, then cognitive illusions should be highly stable. Let us consider the Engineer-Lawyer Problem, which seems to me a fair example, since here both sides of the debate seem to accept the stability of the base rate neglect. "Regardless of what kind of information is presented, subjects pay virtually no attention to the base rate in guessing the profession of the target." (Holland et al. 1986, p. 217). Cohen (1979, p. 401), too, refers in his discussion of that problem to an "overall human tendency to prefer reasoning from causal rather than statistical data". How stable in fact is the neglect of base rates? My own research suggests, not very.

First, let us consider the uninformative description "Dick". The puzzling fact reported as evidence for stability is that even if a description is worthless, base rates are ignored. Recall that in addition to "Dick" each subject judged four informative descriptions, which function as context. In a replication of the Engineer-Lawyer study (Gigerenzer, Hell & Blank, 1988) we systematically varied the position of the uninformative description within a set of six descriptions, and separated those cases where untutored intuition encountered the uninformative description first (no context of informative description available) from those where an informative description was given first and "Dick" was in a later position. Base rate neglect largely disappeared in the former case, whereas it was maintained in the latter. Neglecting position and averaging across positions, however, as in the original study, results in an average difference (between base rate groups) which is zero or close to zero, since with six descriptions there are five times as many second-and-later positions than there are first positions.

This tells us that base rates are not automatically neglected when an uninformative description is judged. Quite the contrary. Only if one or more informative descriptions were presented first, did the subject continue to use a representativeness strategy for Dick, too. The fact that a strategy, once adopted, tends to be maintained in subsequent similar problems is well-known as "set" effect or "functional fixedness" from the Würzburg and Gestalt schools of thinking. Furthermore, this result explains the apparently contradictory findings in the litarature on the probability that Dick is an engineer. Studies reporting base rate neglect used *many* descriptions in addition to Dick, those reporting a difference between base rate groups that is about as large as the difference between base rates (Ginossar & Trope, 1980) used only Dick (i.e., always in the first position), and those reporting differences in-between (Wells & Harvey, 1978; Zukier & Pepitone, 1984) used one informative description in addition to Dick, and switched orders randomly. Note that none of these studies analyzed the position effect.

Second, let us see how to eliminate the neglect of base rates for both uninformative and informative descriptions. The interesting question is, "Do untutored people represent the problem as a probability-revision problem?" There is a crucial assumption that must hold in order to use the specified base rates as prior probabilities. The descriptions have to be *randomly* sampled from the 100 descriptions available. If not, the base rates of engineers and lawyers would be irrelevant. Kahneman and Tversky asserted random sampling of descriptions to their subjects, although this was not true. (Descriptions were deliberately constructed, e.g., to fit the American stereotype of an engineer, such as "Jack".). What about the subjects' prior experience? It is not safe to assume that in the subjects' previous experience with guessing a person's profession, these persons were randomly drawn from a population with known base rates or, at least, could be considered to be. For instance, people in several countries watch TV programs in which a panel of experts guesses the profession of a candidate, who answers only yes or no to their questions. Here, the main heuristic strategy is to ask for new information in order to increase what Cohen calls the inductive reliability of a hypothesis, whereas base rates are not important since the candidates were selected and not randomly drawn. In fact, the experts would perform badly if they started with the known base rates of professions and revised them according to Bayes' theorem. I doubt that previous experience with profession guessing suggests to the subjects the same mental representation of the problem as the experimenters had in mind. A mere verbal assertion, the word "randomly" in the instruction, may not be a contextual cue strong enough to revise an already established representation. Note that this issue is not limited to this particular experiment: In all studies on profession guessing I know of which claim a base rate fallacy, either random sampling was merely claimed and never true, or it was not even mentioned (e.g., in Kahneman & Tversky's (1973) "Tom W." problem).

How can we make sure that random sampling gets into the subjects' problem representation? There is an easy way: let the untutored people themselves *do* the random sampling of descriptions. We performed an experiment (Gigerenzer et al. 1988) in which one group of subjects did the random sampling and a second group was tested in a straightforward replication of Kahneman and Tversky's study as a control. With random sampling being explicit by doing it, base rate neglect disappeared for both informative and uninformative descriptions. Average judgments were closer to Bayesian predictions than to what the representativeness heuristic (and Baconian probability) predicts. The control condition, in contrast, largely replicated Kahneman and Tversky's original findings.

This experiment illustrates two points. First, stability is not a general aspect of cognitive illusions. What is considered a general bias of reasoning can be *easily* eliminated, one does not need thousands of repetitions or an upbringing in an entire different environment, as in the case of the Müller-Lyer illusion. Second, the easy elimination indicates the powerful role of contextual information and prior knowledge in cueing a particular mental representation of the problem.

Such results vindicate Cohen's critique of the "systematic biases" view. His concept of cognitive illusions does not imply stability. On the contrary, it means temporary deception by the experimenter. But do our results also vindicate Cohen's explanation of intuitive reasoning? Cohen (1979, p. 397) says "in the experiment about the engineers and lawyers the subjects naturally tended to go by the weight of evidence and use Baconian reasoning". And he asserts that he and Tversky and Kahneman explain intuitive reasoning in the same way: "Representativeness, as Tversky and Kahneman call it [i.e., Baconian probability], is thus not just a heuristic here, as they regard it, but rather the rationally appropriate criterion" (p. 396). I have three comments on this issue. First, although they look alike, representativeness is here not the same as Baconian probability. Consider the uninformative description, where subjects' median probabilities that the person described was an engineer were .50 in the original experiment. That is, the description is no more "representative" of an engineer than of a lawyer, and vice versa. If subjects reasoned in Baconian probabilities, the evidence that the person was either an engineer or a lawyer would be negligible, and they would have assigned probabilities of zero, or close to zero, instead.[4] Second, this shows that even in the original experiment, subjects did not reason consistently in a Baconian way. Probability judgments about informative descriptions were consistent with both Baconian and representativeness reasoning, those about the uninformative description with the latter but not the former, and those without descriptions (base rate information only) were consistent with neither. Our experiment replicates this pattern, and adds more: Judgments about an uniformative description presented first (or alone) are close to Bayesian reasoning (using the base rates specified as priors); and the same holds for all kinds of descriptions if random sampling was performed *by the subject*. Thus, Baconian reasoning is not stable, even with reference to the same task and the same person.[5] Third, Cohen (1979, p. 397) suggests that we "construe at least some such experiments as revealing how subjects do decide between Pascalian and Baconian responses". The one hypothesis, he offers is that non-experts tend to apply Baconian reasoning when causally relevant information is offered in the descriptions. On this hypothesis, real random sampling as opposed to its mere assertion should make no difference, since the amount of inductively relevant evidence remains the same. Our results contradict this particular hypothesis. Nevertheless, Cohen's general suggestion opens a highly valuable research perspective, which I will attempt to develop briefly.

Context and prior knowledge

In the Helmholtzian view, the brain must rely on context and content-dependent prior knowledge in order to make progress in its task, i.e., to infer from proximal cues (such as the bidimensional retinal picture) to the distal world (the tridimensional visual space). Proximal cues are in principle uncertain cues, even for such elementary judgments as the size of an object: The size of the retinal picture of an object is by itself only a poor indicator of its true size. As mentioned in the introduction, the Helmholtzians understand perception via the analogy of reasoning: To perceive means to select a hypothesis about the world, inferred from generalizations based on past experience and context. If content and context are indispensable for improving elementary perception, then this should hold *a fortiori* for probabilistic reasoning. Thus, if we invert the direction of the analogy, an explanation of this aspect should be quite revealing.

Let us look at the Engineer-Lawyer Problem for illustration. How do Kahneman and Tversky and how does Cohen deal respectively with the context and content of the problem? The former look at the problem with the formal structure of Bayes' theorem in mind. No context information and no content-relevant prior knowledge matters, except for that necessary to calculate the likelihoods for the formula (base rates are already numerically specified). Rational reasoning is reduced to analyzing a problem situation solely in terms of a preconceived formal theory; here, in terms of likelihoods and prior probabilities. The Helmholtzian analogy suggests a different kind of rationality. Of primary importance is the presentation of a problem, which mediates, together with prior knowledge, the subjects' representation of the problem. As we have seen above, base rate neglect disappears if subjects do the random sampling themselves, and I conjecture that this change in the presentation of a problem generates a new mental representation of the problem, one that is different from that suggested by prior experience with profession guessing. Kahneman and Tversky, however, did not analyze their subjects' representation of the Engineer-Lawyer Problem and this representation is built up from context and prior knowledge. Their focus is on deviations from formal structures.

In the present case, even their explanation of base rate neglect, representativeness, can be shown to be equivalent to the formal concept of likelihood in Bayes' theorem (Gigerenzer & Murray 1987, pp. 153-155: see also Note 4). Consequently, this "explanation" (intuition relies on likelihoods, not on base rates) is but a redescription of the phenomenon (intuition neglects base rates, but relies on likelihood).

Ironically, here Cohen seems to be the better psychologist. In various places (e.g., Cohen, 1979, 1982, 1983) he has pointed to the role of the subjects' representation of a problem. What he calls the "Norm Extraction Method" is basically a research program to generate interesting psychological explanations (of the after-the-fact kind) about how subjects could have understood a problem, and how this could explain their reasoning. He opposes his "Norm Extraction Method" to the "Preconceived Norm Method" of those he criticizes. The latter presupposes that a given reasoning problem has one single correct answer, based on the (usually implicit) assumption that there is consensus on the levels of formal theory, interpretation and isomorphic representation (see above). The Norm Extraction Method, in contrast, assumes "that, unless their judgment is clouded at the time by wishful thinking, forgetfulness, inattentiveness, low intelligence, immaturity, senility, or some other competence-inhibiting factor, all subjects reason correctly about probability: none are programmed to commit fallacies or indulge in illusions" (Cohen, 1982, p. 251). Cohen is explicit that his motivation for this claim is to protect the foundations of analytical philosophy. If we were programmed to use cognitive heuristics that sometimes generate correct answers and sometimes do not, this would, Cohen fears, "seriously discredit the claims of intuition to provide — other things being equal — dependable foundations for inductive reasoning in analytical philosophy" (Cohen, 1986, p. 150).

Is such a strong defense of human rationality necessary? The analogy with visual illusions may help us to see a third approach, besides the Preconceived Norm and Norm Extraction Methods. This is suggested by the Helmholtzian perspective, in particular, the third aspect of the visual illusion analogy. We have already seen that what Cohen terms the Preconceived Norm Method does not correspond to a Helmholtzian view in this respect. (It corresponds much more to the early Artificial Intelligence notion that perception works by following simple, general rules or algorithms.) But neither does Cohen's own program. Again, it recalls Gibson's program of determining the invariants in the ambient light that enables the realistic perception of the environment. Where the Gibsonians search for invariants, Cohen searches for norms, and both need these to prove convincingly that untutored perception or reasoning

has direct access to reality and rationality, respectively. The third view, the Helmholtzian, allows both for optimal cognitive functioning *and* for systematic illusions. Such a view would require altering the Norm Extraction Method in some important ways. I will list a few such alterations.

First, according to the "content and prior knowledge" aspect, systematic reasoning errors are to be expected, and it is not necessary to attribute them to inattentiveness or to "cognitive illusions" created by the experimenter.

Second, such errors are a byproduct of the difficult nature of the task in real-world situations. In order to improve reasoning, the mind must take advantage of contextual information and prior knowledge. For instance, consider the "functional fixedness" or "set effect" mentioned above. On the hypothesis that similar problems demand similar solution strategies, it is optimal to use only the first description for creating a mental representation of the problem and not to repeat this for the second, third, etc. An illusion occurs if this similarity does not in fact hold. In perception we continuously "fill in" by expectation: we watch the first few steps, but we do not have to watch the steps unblinkingly when we walk; we are seldom aware of blinking, and, if I close one eye, there is no blank in the visual world that corresponds to the blind spot in my eye, for the blind spot is filled in using the context and prior experience. Of course, we always can stop and check the steps, as we can check the structure of each new problem. However, it is precisely the fact that we do not have to check and recheck that makes such inference mechanisms so efficient.

Third, the dependence on context and content suggests that we will not get far by posing question of the following, formal kind: Do people intuitively understand the law of large numbers? Is intuitive reasoning Bayesian reasoning? Do people confuse $p(A/B)$ with $p(B/A)$? The Helmholtzian analogy predicts: By choosing a proper content and context, you then will be able to produce reasoning that either does or does not follow a given formal rule, or any point in-between. For instance, when the same subjects who neglected base rates in the above replication of the Engineer-Lawyer study were given further problems of the same Bayesian structure, but with soccer games as the content, their judgments immediately became indistinguishable from Bayesian reasoning. Prior knowledge connects soccer games with statistics, and when judging the probability that a team will win a particular game, subjects combine information about the particular game with base rates of the general performance of the team in the long run (Gigerenzer et al. 1988).

561

I prefer to interpret such results as indicating that untutored intuition does not consist of formal rules plus a recognition mechanism that abstracts corresponding formal structure from a problem. Even when intuitive reasoning gives the same answer as Bayesian reasoning, as when estimating the probability that a soccer team will win, this does not imply that the untutored mind has the formal rule in its competence. In contrast, I understand Cohen to mean that untutored intuition contains logical and statistical rules. He emphasizes the distinction between a subject's not having a rule and simply not applying it (Cohen 1983, p. 511). For instance, he interprets the subjects' reasoning in the Wason selection task as a failure to apply the law of contraposition, a cognitive illusion in his sense of the term, but rejects the conclusion that the competence of most untutored people does not embrace the law (Cohen 1981, pp. 323-324).

However, as I argued, the Helmholtzian analogy shows that there is at least one alternative to Cohen's premisse of untutored competence in statistical formulae which both supports intuition as basically rational and at the same time allows for systematic illusions.

Can Cohen's premisse of untutored rationality be falsified by observation or experiment? Cohen (1981, p. 330) claims flatly "nothing in the existing literature ... or in any possible future results of human experimental enquiry, could ... establish a faulty competence". But he seems less definite in other places (e. g., Cohen, 1982, p. 252). I think that he can safely use his competence/performance distinction to protect his premisse against empirical falsification. In fact, Cohen takes those contents and contexts where reasoning conforms to probability theory as a proof of lay competence in statistical principles, and dismisses negative instances as due to obstacles like "cognitive illusions". If there exists an empirical argument that could lead him to reconsider his premisse, it is the argument of consistent logical inconsistencies. Today we can survey numerous experimental studies on adult statistical reasoning – from Hofstätter's (1939) pioneering work to the present avalanche. I take the main lesson of those studies to be that untutored reasoning is consistently inconsistent, *if* you look at reasoning from the point of view of logical structure. For instance, if we want to understand the experimental results reported above within a structural-logical conception of the mind, we would have to assume that our typical subject is a Baconian probabilist when she reasons about whether "Jack" is an engineer; uses a representativeness heuristic (but not Baconian probability) when she reasons about the undiagnostic description "Dick"; is a Pascalian probabilist when she encounters only base rates and no description; is an approximate Bayesian if she does the random sampling herself; and is

a proper Bayesian if soccer is the issue. This is why I do not believe that some logical structure alone is sufficient to explain how untutored reasoning functions.

The Helmholtzian alternative to both Cohen and the "general heuristics" view of Kahneman and Tversky would be to consider the task of reasoning as coming to terms with a world that presents itself to us only by uncertain cues, and to accept that this uncertainty forces us to rely on context and content (rather than only on formal structure) to improve our bet. But even this aspect of the analogy has its limits. Unlike reasoning, perception goes on to make the same bet again and again, despite better knowledge. As Egon Brunswik once put it, perception is like a stupid animal.

Cueing different interpretations of probability

What would a different research program look like if we took the final aspect of the analogy with visual illusions seriously? Despite Cohen's claim for the in-principle-rationality of untutored reasoning and its competence in applying statistical rules, his work contains the seeds of an interesting answer. In particular, Cohen suggested that the conception of probability, standard or Baconian, frequency or propensity, with which subjects construe their task is cued by the wording and content of the instructions. The question then becomes, what contextual cues and contents cue which notion of probability. This suggests the further question of whether experimenters and subjects have a similar understanding of the task, in particular of key concepts like "probability", "frequency", or "confidence". I consider this suggestion of Cohen's a practical and fruitful one. Curiously, contextually cued "schema" and "frames" are frequently used in cognitive psychology as an alternative view to the application of general systems of mental logic, but rarely to refer to cueing alternative (and equally legitimate) conceptions of "probability".

Let me give an example to show how Cohen's suggestion proves to be particularly fruitful. The "overconfidence" phenomenon is reported to be one of the "notorious biases of human reasoning" (Bar-Hillel & Margalit, 1983, p. 247) that resists numerous "debiasing methods" such as to raise stakes, clarify instruction, warn and inform about the phenomenon, or use different response modes (Fischhoff, 1982). It is classified as "a reliable, reproducible finding" (von Winterfeldt & Edwards, 1986, p. 539). All this makes it look like a good candidate for a stable reasoning error. What is the phenomenon? In a typical experiment, subjects are given two-alternative choice tasks, such as "Which city has more inhabitants? (a) Bonn, (b) Konstanz". After choosing what they believe to be the correct answer, subjects are asked for their confidence that the answer chosen is correct. These confidences were often obtained on a scale from 50% to 100% confident. After many such questions are given to many subjects, the relative frequency of correct answers in each confidence category is determined. A number of studies found that the confidence ratings were consistently larger than the relative frequencies of correct answers, a phenomenon which was called "overconfidence". For instance, when subjects said they were 100% confident, their average relative frequency of correct answers was only about 80%, for 90% confidence it was about 75%, for 80% confidence 65%, and so on. This discrepancy was considered a bias of reasoning, and explained by general cognitive heuristics such as, among others, search for confirming evidence ("confirmation bias"), or general motivational tendencies such as the tendency to overestimate one's knowledge ("illusion of validity").

I want to make two points. First, both the definition of the phenomenon and its interpretation as a bias of reasoning rests on the assumption that relative frequencies are the normative yardstick against which probabilities of individual events (such as that a given answer is correct) are to be evaluated as correct or incorrect. This position has yet to sweep the field of probability and statistics, winning over neither frequentists such as von Mises (1957) nor subjectivists such as de Finetti (1989). Nevertheless, little attention has been paid to this serious conceptual problem in the overconfidence literature, a problem, both for the definition of the phenomenon and its interpretation as a bias (an exception is May 1987). Cohen (1982, p. 256) noticed this widespread neglect of the various interpretations of formal probability in the experimental literature, and claimed that the possibility that subjective probabilities "may differ legitimately in value from the corresponding objective relative frequencies turns out to be vitally important for the interpretation of experimental data on probability judgment". Second, I shall argue that Cohen's conjecture is correct and fruitful for understanding the overconfidence phenomenon. Imagine an experimental task as described above, but where in addition the subjects are asked to estimate the *relative frequency* of correct answers in each of the confidence categories. This allows to compare true relative frequencies with estimated relative frequencies (in addition to comparing true relative frequencies with subjective probabilities for single events).

If the discrepancy called "overconfidence" were a reasoning bias due to limited processing abilities, confirmation biases or general motivational tendencies, then the discrepancy should occur independently of whether the subjects report their degrees of belief that a particular answer is correct, or estimate their relative frequency of correct answers.

If, on the other hand, the philosophers' distinction were relevant and meaningful to the untutored mind, then we might expect the discrepancy to change when we compare estimated with true relative frequencies. We were curious and actually performed such experiments (Gigerenzer, Kleinbölting & Hoffrage, 1990). Subjects were given several hundred two-alternative choice tasks, and were asked after each answer to estimate the probability that the answer was correct, as in previous studies and described above. In addition, after they finished, they were asked to estimate the relative frequency of correct answers in each of the confidence categories. When we compared estimated frequencies with true relative frequencies of correct answers the overestimation disappeared completely. For instance, in those cases where subjects said they were between 70% and 80% confident, only 65% of the answers were correct. But when their task was to estimate the relative frequency of correct answers (among those in which they said they were 70% to 80% confident), the average response was 65%. Different kinds of frequency judgments, such as asking the subject after 50 answers, "How many of the last 50 answers are correct?" gave a similar result. In all cases overestimation completely disappeared, and the discrepancy between estimated and true frequencies was either zero or even turned out to be an underestimation. This result cannot be explained by a general confirmation bias or general motivational tendencies to make ourselves appear better than we are. It suggests, however, that in evaluating one's knowledge, "probabilities for single events" and "relative frequencies" refer to different concepts in the minds of our untutored subjects, as they do in the minds of many philosophers. It also reveals that the previous interpretation of the discrepancy between degrees of belief and frequencies as a *bias in reasoning* meant comparing apples and oranges. Such results open up a quite different theoretical account of this allegedly reliable, reproducible "cognitive error". For one thing, more optimism is in order. Edwards and von Winterfeldt's (1986, p. 656) final word on overconfidence, after they asserted the stability of this "cognitive error", was: "Can anything be done? Not much". My answer is: Keep distinct meanings of probability straight, and much can be done — cognitive illusions disappear.[6]

However, many cognitive psychologists consider Cohen's conjectures of "little relevance to research on judgment and decision making" (Mynatt et al., 1983, p. 507), of no benefit for scientific work (Dawes, 1983) and at best, of merely philosophical interest. Fischhoff (1981, p. 337) even takes Cohen's work as evidence "that psychologists and philosophers would do well to ignore one another", since their pursuits "are somewhat irrelevant to the other". It is ironical that Fischhoff is among those who did the major studies on the overconfidence phenomenon.[7]

Psychologists can learn from philosophers. A sharp eye for conceptual distinctions, supplemented by experiments that can distinguish between them, is a major alternative to the traditional "heuristics and biases" program. It vindicates Cohen's point that even untutored reasoning is capable of making consistent conceptual distinctions of the sort philosophers make.

Institut für Psychologie
Universität Salzburg
Hellbrunnerstr. 34
5020 Salzburg
Austria

FOOTNOTES

[1] The recent debate on reasoning biases and rationality initiated by L. J. Cohen is documented in, among others, *The Behavioral and Brain Sciences* 1981, 1983, 1984, 1987, and *Cognition* 1979, 1980, 1982.

[2] For more on the stability of visual illusions see Robinson (1972).

[3] In fact, the likelihoods p("Blue"/Blue) and p("Blue"/Green) are not unambiguously specified in (ii). What is said is that the two errors, p("Blue"/Green) and p("Green"/Blue) add up to 20%, but this does not determine the size of either of them. But such ambiguities, which could be easily resolved, are not of interest here.

[4] In the present context, the meaning of the term "representativeness" can be reduced to saying that subjects use Bayes' theorem with uniform priors. This can be shown as follows, where UD, E, and L stand for Uninformative Description, Engineer, and Lawyer, respectively. What is to be explained is the subjects median judgment $p(E/UD) = .5$, that is the probability that Dick (the uninformative description, UD) is an eugineer (E). $p(E/UD)/(1-p(E/UD)) = p(UD/E)/p(UD/L)$ is Bayes' theorem with uniform priors. Since $p(UD/E) = p(UD/L)$ (the definition of uninformativeness), we get $p(E/UD) = .5$, which is the actual median answer. Baconian probability, in contrast, should assign $p_I(E/UD) = 0$. It is therefore not equivalent to the notion of "representativeness heuristic" in this context.

[5] Our results also contradict the interesting explanation of intuitive reasoning by reference to conversational maxims. Adler (1983, p. 246) argues that the experimenters' most salient contribution is the biographical sketch, and according to Grice's communicative maxim to "be relevant" it is socially healthy and more cooperative, if the subject judges solely on the basis of the description, which is of greater selective relevance, "rather than (less

cooperatively) integrating it with the base-rate data". Again, randomly drawing the descriptions out of an urn should not matter if the real issue was that conversational maxims such as being cooperative were in conflict with abstraction, as Adler believes. The elimination of base rate neglect by making random sampling explicit, however, indicates that, at least in the present context, cognitive representations are more powerful for probabilistic reasoning than communication maxims.

[6] The same conceptual distinction is crucial for whether people commit the so-called conjunction fallacy (Fiedler, 1988; Tversky & Kahneman, 1983).

[7] Why many psychologists and social scientists neglect conceptual distinctions within probability theory is a different story (see Gigerenzer et al. 1989, chapters 3 and 6, and Gigerenzer, 1987).

REFERENCES

Adler, J. E. (1983). Human rationality: Essential conflicts, multiple ideals. *The Behavioral and Brain Sciences*, **6**, 245-246.

Adler, J. E. (1984). Abstraction is uncooperative. *Journal for the Theory of Social Behaviour*, **14**, 165-181.

Bar-Hillel, M. & Margalit, A. (1983). To err is human. *The Behavioral and Brain Sciences*, **6**, 246-248.

Berkeley, G. (1709). *An essay towards a new theory of vision.* Dublin.

Birnbaum, M. H. (1983). Base rates in Bayesian inference: Signal detection analysis of the cab problem. *American Journal of Psychology*, **96**, 85-94.

Cohen, L. J. (1979). On the psychology of prediction: Whose is the fallacy? *Cognition*, **7**, 385-407.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *The Behavioral and Brain Sciences*, **4**, 317-331.

Cohen, L. J. (1982). Are people programmed to commit fallacies? Further thoughts about the interpretation of experimental data on probability judgment. *Journal for the Theory of Social Behaviour*, **12**, 251-274.

Cohen, L. J. (1983). The controversy about irrationality. *The Behavioral and Brain Sciences*, **6**, 510-517.

Cohen, L. J. (1986). *The dialogue of reason*. Oxford: Clarendon Press.

Daston, L. (1988). *Classical probability in the enlightenment.* Princeton NJ: Princeton University Press.

Dawes, R. M. (1983). Is irrationality systematic? *The Behavioral and Brain Sciences*, **6**, 491-492.

De Finetti, B. (1989). Probabilism. *Erkenntnis*, **31**, 169-223 (original work published in 1931. New York: Springer).

Edwards, W. & von Winterfeldt, D. (1986). On cognitive illusions and their implications. In H. R. Arkes & K. R. Hammond (Eds.), *Judgment and decision making.* Cambridge: Cambridge University Press.

Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research*, **50**, 123-129.

Fischhoff, B. (1981). Can any statements about human behavior be empirically validated? *The Behavioral and Brain Sciences*, **4**, 336-337.

Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slovic & A. Tversky (Eds.), *Judgment under uncertainty. Heuristics and biases*. Cambridge: Cambridge University Press.

Gigerenzer, G. (1987). Probabilistic thinking and the fight against subjectivity. In L. Krüger, G. Gigerenzer & M. S. Morgan (Eds.), *The Probabilistic Revolution*. Vol. 2: *Ideas in the sciences*. Cambridge, MA: MIT Press.

Gigerenzer, G., Hell, W. & Blank, H. (1988). Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance*, **14**, 513-525.

Gigerenzer, G., Kleinbölting, H. & Hoffrage, U. (1990). Confidence in one's knowledge. A Brunswikean view. Manuscript submitted for publication.

Gigerenzer, G. & Murray, D. J. (1987). *Cognition as intuitive statistics*. Hillsdale, NJ: Erlbaum.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L. J., Beatty, J. & Krüger, L. (1989). *The empire of chance: How probability changed science and everyday life*. Cambridge: Cambridge University Press.

Ginossar, Z. & Trope, Y. (1980). The effects of base rates and individuating information on judgments about another person. *Journal of Experimental Social Psychology*, **16**, 228-242.

Gregory, R. L. (1974). *Concepts and mechanisms of perception*. New York: Charles Scribner's Sons.

Hofstätter, P. R. (1939). Ueber die Schätzung von Gruppeneigenschaften. *Zeitschrift für Psychologie*, **145**, 1-44.

Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986). *Induction*. Cambridge, MA: MIT Press.

Kahneman, D. & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, **80**, 237-251.

Kahneman, D. & Tversky, A. (1982a). On the study of statistical intuitions. In D. Kahneman, P. Slovic & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Kahneman, D. & Tversky, A. (1982b). Variants of uncertainty. In D. Kahneman, P. Slovic & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

May, R. S. (1987). *Realismus von subjektiven Wahrscheinlichkeiten*. Frankfurt/Main: Lang.

Mynatt, C. R., Tweney, R. D. & Doherty, M. E. (1983). Can philosophy resolve empirical issues? *The Behavioral and Brain Sciences*, **6**, 506-507.

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

Robinson, J. O. (1972). *The psychology of visual illusion*. London: Hutchinson University Library.

Slovic, P., Fischhoff, B. & Lichtenstein, S. (1976). Cognitive processes and societal risk taking. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior*. Hillsdale, NJ: Erlbaum.

Tversky, A. & Kahneman, D. (1971). Belief in the law of small numbers. *Psychologic Bulletin*, **76**, 105-110.

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biase *Science*, **185**, 1124-1131.

Tversky, A. & Kahneman, D. (1977). Causal thinking in judgment under uncertainty. I R. Butts & J. Hintikka (Eds.), *Basic problems in methodology and linguistics* Dordrecht-Holland: Reidel.

Tversky, A. & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. I M. Fishbein (Ed.), *Progress in social psychology*. Vol. 1. Hillsdale, NJ: Erlbaum.

Tversky, A. & Kahneman, D. (1983). Extensional versus intuitive reasoning: The con junction fallacy in probability judgment. *Psychological Review*, **90**, 293-315.

von Mises, R. (1957). *Probability, statistics, and truth*. New York: Dover.

von Winterfeldt, D. & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge: Cambridge University Press.

Wells, G. L. & Harvey, J. H. (1978). Naive attributors' attributions and predictions: What is informative and when is an effect an effect? *Journal of Personality and Social Psychology*, **36**, 483-490.

Zukier, H. & Pepitone, A. (1984). Social roles and strategies in prediction: Some determinants of the use of base-rate information. *Journal of Personality and Social Psychology*, **47**, 349-360.