

## 5

### *Yet Another Look at the Heuristics and Biases Approach*

**Gideon Keren and Karl H. Teigen**

#### **Introduction**

The research approach that has become to be known as the *heuristics and biases* research program, initially launched in the beginning of the 1970s by Amos Tversky and Daniel Kahneman (1974), has been highly influential in shaping the field of judgment and decision making. Its main aim was to study people's intuitions about uncertainty and the extent to which they were compatible with the normative probability calculus. It stimulated hundreds of articles designed to test the robustness as well as the limitations of this approach. Like any successful research program it did not escape critical evaluation. Indeed, several authors raised their doubts regarding the ecological validity and logical soundness of this approach (e.g., Cohen, 1981; Gigerenzer 1991, 1996). Even the originators of this highly successful program have, during the course of time as research results accumulated, changed their perspective and suggested new interpretations (e.g., Tversky and Kahneman, 1974 vs. Kahneman and Frederick, 2002). Indeed, given abundant new studies and an increasing list of heuristics and biases, the understanding of the term has gradually changed, and acquired some new interpretations.

The success of the heuristic and biases research program to attract so much attention and stimulate an ever increasing stream of studies can be explained on several grounds. First, having been launched shortly after the so called "cognitive revolution," it raised interest for two opposing reasons. On the one hand, this research program and its method of investigation matched well the principles underlying the cognitive paradigm and the belief that human behavior could (and should) be explained mainly in cognitive terms. It offered a new experimental methodology to the study of cognitive processes. At the same time it implicitly challenged some tacit assumptions about the abilities and

the limits of the cognitive system. It was probably this latter aspect that associated the heuristics and biases research program with the broad problem of rationality. The dispute concerning rationality, implied by the empirically exhibited biases, had implications not just for psychology. It challenged the fundamental assumptions underlying economic theory. Thus, the initial results reported by Kahneman and Tversky (1972, 1973; Tversky & Kahneman, 1972, 1974) carried an important message not just for psychology but for the social sciences in general.

Second, this research program evolved from previous investigations that laid the ground for the systematic study of how people cope with uncertainty and, in particular, the extent to which they obey the probability calculus. Precursors included the study of probability matching (Hake & Hyman, 1953), Meehl's (1954) essay on clinical versus statistical prediction, John Cohen's (1960, 1964) pioneering research on chance, skill, and luck, and the work of Ward Edwards and his colleagues who tried to assess the extent to which people behave as Bayesian statisticians (for a review, see Peterson & Beach, 1967). Kahneman and Tversky's heuristics and biases consolidated and in some respects challenged this previous work, and contained the outline of a novel, coherent, and meaningful framework.

Third, many of the demonstrations of biases were simple, easy to comprehend and thus very compelling. Indeed, for participants in these experiments the potential errors and inconsistencies were rather opaque (and some of the critics of these experiments argued that change of presentation may be sufficient to eliminate the observed biases). However, when presented in a transparent frame, to readers who were supposedly familiar with the basics of probability theory and who examined the experimental results analytically, the discrepancy between the intuitive and the analytical mode of reasoning became immediately evident.

Notwithstanding, simple introspection suggested to the honest reader that he or she might also be vulnerable to several of the observed biases. Consider for instance the letter frequency problem (Tversky & Kahneman, 1972) intended to demonstrate the availability heuristics. Participants were asked to estimate the likelihood that a given letter will appear in the first or third position of a word. For example, is the letter R more likely to appear in the first or the third position? Evidently, the majority of the participants judged the likelihood to be larger in the first position despite the fact that the letter R is more likely to be in the third position. Tversky and Kahneman suggested that people estimate the likelihoods of the two categories (first or third position) by roughly assessing the ease with which instances of the two categories come to mind. Taking a quick sample, it is mentally much easier to retrieve words with the letter R in the first rather than in the third place. Obviously, researchers reading the article were not more knowledgeable (than the average participant) about the frequency of different letters in different positions of a word. However, by placing oneself in the participants' role and attempting to simulate what participants in this task have done (in a way, using the simulation heuristic), it is easy to imagine that one would use exactly the same strategy supposedly used by the participants. Many of the problems used by Kahneman and Tversky were persuasive because they lent themselves easily to be imagined by the reader. So in a way, and perhaps paradoxically, the success of the heuristics and biases program could be partly attributed to a clever use of the simulation heuristic, whereby a conclusion appears

convincing by being easily constructed as a part of a good scenario (Kahneman & Tversky, 1982b). It is of course impossible to provide a complete and detailed treatment of this innovative and stimulating research program in a single chapter. An extensive coverage is provided in Kahneman, Slovic and Tversky (1982), and Gilovich, Griffin, and Kahneman (2002), both of which carry, not incidentally, the same title. In this chapter we cover a small selection of the existing literature and highlight what seems to us to be some of the more important facets of the area. We first examine more closely the meaning of the two key concepts of “bias” and “heuristics.” Subsequently, we offer a brief discussion in which the heuristic and bias program is related to perceptual processes on the one hand, and to the psychology of reasoning on the other hand. The following two sections contain a brief description of the three heuristics (representativeness, availability, and anchoring) and some more recent developments. Finally, a two-stage framework is proposed in which, borrowing from prospect theory, it is suggested that the processes underlying probability judgments consist of an *editing* and an *evaluation* phase.

## **What is a Bias?**

The heuristics and biases approach rests on the marriage between two key concepts, neither of which are unproblematic and unambiguous by themselves. We will discuss them in turn.

According to the *Oxford English Dictionary* (2002), the term “bias” was originally used to describe a slanting line (e.g., the diagonal in a square), and the oblique motion of a loaded bowling ball; it also referred to the asymmetric construction of the bowling ball achieved by loading it on one side with lead, as exemplified in a Shakespearean passage: “Well, forward, forward thus the bowle should run. And not unluckily against the bias” (Shakespeare, 1596, *The Taming of the Shrew* IV, v. 25).

These usages illustrate two distinctions still implied in various contexts of the modern term. First, biases are often used to describe deviations from a norm (as with Shakespeare’s bowl) but, in another more neutral sense, they can simply indicate a tendency to slant in one way rather than another (like the diagonal). For instance, the term “positivity bias” has been used to describe a preponderance of positive over negative evaluations in person perception and, more generally, in everyday language (Kanouse & Hanson, 1971; Peeters, 1971). This does not in itself indicate any errors of judgment, unless we believe that, in reality, positive and negative events should balance each other out. On the other hand, the concept of a “desirability bias” (Budescu & Bruderman, 1995) implies a tendency to assign exaggerated probability estimates to desired outcomes, not because of the amount of supporting evidence, but simply because we want them to come true. Such biases can be regarded as systematic, suboptimal judgments, sometimes labeled “errors,” or even “fallacies.”

Another distinction concerns bias as a cause versus bias as an effect. The bias of the bowl can be its shape or loading, causing it to deviate from a straight run. It also designates its trajectory, resulting from the lopsided construction. In the psychology of judgment, biases were originally conceived as effects (to be explained, for instance, by

heuristics), rather than causes. But in many contexts, they have been used as explanations rather than phenomena to be explained. For example, in studies of logical tasks, Evans (1989) suggested that many errors of deductive reasoning can be explained on the basis of a more general “matching bias,” namely the tendency to endorse conclusions that are linguistically compatible with the premises (this may in turn be regarded as a manifestation of a more general principle of relevance). Similarly, “confirmation bias” in hypothesis testing can be conceived as a general strategy for testing hypotheses through verification rather than falsification procedures (Wason, 1960; Klayman & Ha, 1987), either by searching for positive instances rather than negative ones, or by finding observed confirmations more compelling than disconfirmations. It has alternatively been described as a general outcome of these and similar mechanisms (e.g., matching), reflecting the fact that hypotheses, for whatever reason, appear to be more easily retained than rejected.

The concept of a bias in the latter sense, namely as a systematic deviation from a norm (or as an inclination towards one judgment rather than another), does not in itself imply one specific kind of explanation. Biases can be the result of cognitive limitations, processing strategies, perceptual organizing principles, an egocentric perspective, specific motivations (e.g. “self-serving biases” in social psychology), affects, and cognitive styles. In the heuristics and biases tradition, the general approach has been to regard biases as a more or less regular by-product of some more general principles of judgment, labeled heuristics, to which we now turn.

### **What is a Heuristic?**

Paraphrasing William James, “everyone knows what heuristics are” or, at least, that is the impression given by the literature on heuristics and biases, where a definition of heuristics is rarely, if ever, attempted. The reason could also be that the term heuristics was, in this program, used in a deliberately imprecise way, more as a hint about the role of the psychological processes involved than as a description of their precise nature.

Following the Webster dictionary, the term heuristics implies inventing or discovering, and more specifically designates a method of education or a computer program that, searching for a solution or answer to a given question, proceeds along empirical lines using rules of thumb. It has been originally dubbed by Polya (1945) as a sort of reasoning “not regarded as final and strict but as provisional and plausible only, whose purpose is to discover the solution of the present problem” (p. 115). Being “provisional” rather than final, a heuristic approach will necessarily be incomplete and error prone. Einstein called his first Nobel Prize-winning paper on quantum physics (1905): “On a heuristic point of view concerning the generation and transformation of light,” using the term “heuristic” rather than “theory” to indicate that he regarded it at this stage only as a useful approximation to truth.

The term has been adopted and applied both in computer science and in the (psychological) domain of problem solving as a prescriptive method in which a problem solver (or a machine in the case of artificial intelligence) proceeds along empirical

guidelines to discover solutions or answers. Such procedures entail both advantages and risks, as they may lead us by a short cut to the goal we seek or they may lead us down a blind alley. Heuristics are, in this literature, often contrasted with algorithms, which are explicit and detailed rules that guarantee a correct result, but could be effortful and time-consuming, and hence impractical in situations characterized by limited cognitive resources.

The meaning of the term heuristics, as first used by Kahneman and Tversky, was highly similar to its use in the problem-solving literature, by being considered to be simplified methods intended to cope with humans' limited processing capacity. They were also error prone, leading generally to acceptable (although imprecise) estimates, but under certain circumstances, to systematic biases. Finally, they could be contrasted with normative, "algorithmic," procedures for estimating probabilities, which may require full statistical information of all outcomes involved, knowledge of the basic principles of probability theory (like combinatorial rules and Bayes' theorem), as well as cognitive capacity to carry out calculations based on these principles. However, one question remained: While heuristics in computer science and problem solving usually are explicit strategies, that can be applied (mostly with success) or not applied, it was not at all clear whether (or when) the judgmental heuristics described by Kahneman and Tversky were deliberate and under the control of the individual. Current views (Kahneman & Frederick, 2002) seem to suggest that the mechanisms underlying heuristics are essentially automatic, and supposedly do not operate under the individual's awareness. We elaborate on this point later.

## Two Metaphors

The psychology of judgment can be conceived as occupying a middle ground between the psychology of thinking and the psychology of perception. It may be slow and deliberate, like problem solving, and quick and immediate, like for instance distance perception, where we seemingly jump to the conclusion (e.g. "a car is approaching") without conscious knowledge of the premises, or "cues," on which this conclusion is based (for a discussion of these two metaphors within the framework of Brunswikian social judgment theory, see Chapter 3, this volume).

It has been known for a long time that the subjective conclusions drawn in both areas are sometimes nonveridical, or incorrect. In the literature on deductive reasoning, such errors have traditionally been called *fallacies*, whereas perceptual mistakes have typically been called *illusions*. Classic texts on logic have often included a chapter on fallacies (e.g., Mill, 1856), in many ways reminiscent of the "biases" apparently rediscovered in the heuristics and biases tradition. Similarly, treatises on sensation and perception have contained lists of visual (and other) illusions as an integral part. The traditional distinction between fallacies and illusions is nicely illustrated by two volumes appearing in the same "International Scientific Series" more than one hundred years ago, one by psychologist James Sully (1882), entitled *Illusions*, the second on *Fallacies*, by the logician Alfred Sidgwick (1883). However, both authors admitted that the distinction between illusions,

defined as errors of “immediate, self-evident, or intuitive knowledge,” and fallacies, denoting false inferences or errors of reasoning, is hard to draw. If one wants to draw attention to the *process* involved in drawing a conclusion (even a perceptual one), the reasoning or inferential metaphor seems particularly apt; if, on the other hand, emphasis is put on the immediate or inevitable *gut feeling* of what is the case, the perceptual metaphor will be more appropriate.

Indeed, Kahneman and Tversky often drew a parallel between heuristics and biases, and comparable perceptual processes. It is in this respect that the term “cognitive illusions” was introduced as an analog to visual illusions. Though rarely described in these terms, many of the Gestalt laws, such as grouping or closure, constitute non-deliberate automatic processing. In a similar vein, and congruent with current interpretations (e.g., Griffin, Gonzalez, & Varey, 2001; Kahneman & Frederick, 2002), we assume that many “heuristic” judgments are performed automatically and cannot be entirely controlled. In other contexts, the term “fallacies” has been employed (e.g., the conjunction fallacy, the gambler’s fallacy, and the “planning fallacy”), pointing more directly to the logical inconsistencies involved.

The perceptual metaphor, applied to subjective probability judgments, did not originate with Kahneman and Tversky, but can be traced back at least to Pierre Simon Laplace, one of the founders of probability theory. In his *Essai philosophique sur les probabilités* (1816) he included a chapter called “Illusions in probability estimation.” Here, the reader is told that “the mind has its illusions, like the sense of vision” (p. 182), which need to be corrected by “reflection and calculation.” Still, the subjective probabilities that are based on everyday experience, and exaggerated by hope and fear, are more striking than those that are merely a result of calculation. Subjective probabilities are, according to Laplace, governed by the principles of association, the main being contiguity (strengthened by repetition), and resemblance. These are, like heuristics, basically sound and helpful principles, but can occasionally be misleading. Indeed, the parallel between the laws of association and the heuristics suggested by Kahneman and Tversky is more than superficial, repetition frequency corresponding to availability, and resemblance corresponding to the representativeness heuristic. In a remarkable chapter on “Unphilosophical probabilities,” David Hume (1976[1739]) made the same point, by showing how people judge probability by how “fresh” an event is in memory; unfortunately memorability is not only affected by frequency, but also by recency and vividness. This is of course an early, but quite accurate, description of the currently popular “availability heuristic.”

## **The Domain of Heuristics and Biases**

What kinds of phenomena lend themselves to “heuristic” approaches, and in which areas do we find “biased” outcomes of such an approach? The original focus of the heuristics and biases program was clearly within the field of prediction under uncertainty and estimation of probabilities and frequencies. In these areas many responses that are incompatible with normative considerations have been documented (as testified by Hume and Laplace), and the suspicion arose that people are not just inaccurate or lack the skills for

calculating probabilities, but that they use an entirely different approach from that of the mathematician.

Soon, the search for biases was generalized to the whole area of judgment and decision making (JDM), giving rise to decision biases like the status quo bias (e.g., Kahneman, Knetsch, & Thaler, 1991), omission bias (Spranca, Minsk, & Baron, 1991), and outcome bias (Baron & Hershey, 1988). We may also speak of choice heuristics (Frederick, 2002), and specific heuristics tailored to concrete judgment tasks (Gigerenzer, Todd, and the ABC Research Group, 1999).

In an even wider sense, the concepts of heuristics and biases have – separately or in combination – been applied to areas outside the JDM field, both within cognitive psychology (hypothesis testing, inductive and deductive reasoning) and by social psychologists studying issues of social cognition (Nisbett & Ross, 1980). In particular, biases are frequently discussed within the framework of attribution theory, as for instance “the correspondence bias” (Gilbert & Malone, 1995), referring to the tendency to draw inferences about a person’s dispositions from his or her behaviors (also called “overattribution,” and “the fundamental attribution error”), the “actor–observer bias”, and various “self-serving” biases, referring to patterns of attribution that tend to protect or boost the person’s self-esteem. Biases have also been found in the area of self–other comparisons, where people commonly judge themselves as better, more lucky, or more special than other people (above-average bias, illusory optimism, and false uniqueness effect). Pronin, Lin and Ross (2002) recently demonstrated that people are even biased to think that they are less biased than others!

The remaining part of this chapter will be devoted mainly to a discussion of predictions and probability judgments, being the original core area of the heuristics and biases approach, but also with an eye to related developments in judgment and decision making, more broadly conceived. Biases in other areas of cognitive and social psychology are beyond the scope of the present chapter.

### **Three Canonical Heuristics**

In their early work, Tversky and Kahneman (1974) described three judgmental heuristics for estimating probabilities, frequencies, and other uncertain quantities. These three, labeled representativeness, availability, and anchoring and adjustment, respectively, were not introduced as the *only* three, not even as the three most important heuristics, yet they have since the time of their introduction occupied a unique position as “prototypical” or canonical heuristics within the heuristics and biases approach.

#### *Representativeness*

Probability judgments are rarely completely unconditional. Some go from hypothesis to data, or from population to sample or, more generally, from a Model M to some instance or event X, associated with the model (Tversky & Kahneman, 1982). Such judgments

could be: what is the probability of getting five heads in a row from an unbiased coin; or what is more likely: that the best student in the class this year will perform equally well, less well, or even better next year? Another set of probability questions goes the opposite way, from data to hypothesis, sample to population, or more generally from X to M. We observe the five heads, and wonder whether the coin is unbiased or not; or, we observe that the student is performing less well the following year, and wonder about the most likely explanation. The first set of problems can be regarded as problems of prediction, the second as problems of diagnosis, or explanation.

In three early important papers, Kahneman and Tversky (1972, 1973; Tversky & Kahneman, 1971) demonstrated that both types of probability judgments are often performed as a simple comparison between X and M. If X looks like a typical instance of M, it will be regarded as a probable outcome. In such cases, predictions are said to be performed by a “representativeness heuristic.” Accordingly, we may think that five heads in a row is not a very likely outcome, because it does not fit our model of a random series; whereas we think it is likely that a good student will remain at the top of his class, because this looks like a typical thing for a good student to do.

M can also be diagnosed from X by the same mechanism. When five heads actually appear, we may suspect the coin of being loaded; if the student’s achievement is more mediocre next year, we look for causal rather than statistical explanations (perhaps he was overworked, or spoiled by his initial success). Such probability judgments by similarity, which are the essence of the “representativeness heuristic,” seemed well suited to explain several well-known biases of probability judgments, like the gamblers’ fallacy and the problem of non-regressive predictions. It could also make observers (including scientists) place undue weight on characteristics of small samples (facetiously termed “belief in the law of small numbers” by Tversky & Kahneman, 1971), and to neglect base rates in diagnostic judgments.

One of the more striking manifestations of representativeness reasoning is to be found in the so-called conjunction fallacy. Here the predicted outcome, X, is typically a combination of a high-probability and a low-probability event, where the first is a good and the second a poor match for the model (Linda as a feminist, and Linda as a bank teller). The conjunction (a feminist bank teller) is, by the logic of probability theory, less likely than both its components (the number of feminist bank tellers cannot exceed the number of bank tellers), but from a similarity point of view, the picture looks different. One typical and one atypical characteristic can give the conjunction an appearance of being neither likely, nor completely unlikely, but something in between (Tversky & Kahneman, 1983).

Representativeness captures an aspect of probability that, in many languages, is embedded in the probability vocabulary itself, namely its verisimilitude, or likeness to truth (cf. French: “vraisemblable,” German: “Wahrscheinlich,” Swedish: “sannolik,” Polish: “prawdopodobny”). It has been conceived as a very general mechanism, applicable both to singular and repeated events. It has also a high degree of ecological validity, since in most distributions, the central, or most typical value is at the same time the modal (most frequent) one. It is, at the same time, a quick and effortless type of judgment, requiring a minimum of cognitive resources. As a theoretical concept, critics have pointed out that it is underspecified and lends itself poorly to specific, falsifiable predictions (Olson, 1976; Gigerenzer, 1996).



### *Availability and simulation*

The second main heuristic, introduced by Kahneman and Tversky (1973), was termed *availability*. In this case, events are not compared to a model in terms of similarity, they are instead evaluated according to the ease by which they can be imagined or retrieved from memory. Again, this refers to a class of phenomena, rather than one specific process. In the most concrete case, instances of the target event are simply recalled; if a number of instances are readily recalled, the event is judged to be frequent, and predicted with a high probability to happen again in the future. Events that are harder to recall, are regarded to be less frequent and less probable. Unfortunately, recall can be influenced by factors other than frequency, such as public exposure, vividness, primacy and recency, leading people for instance to overestimate highly publicized and dramatic risks (like terrorism and airplane accidents) and underestimate less spectacular ones (like diabetes and tobacco smoking). Recall can also be affected by retrieval principles and memory organization, as illustrated by the case of words with R in the first, vs. third position, described earlier in this chapter. Recent research indicates, however, that people are more accurate in estimating letter frequencies than implied by this classic demonstration (Sedlmeier, Hertwig, & Gigerenzer, 1998). Research by Schwarz, Bless, Strack, Klumpp, Rittenauer-Schatka, & Simons (1991) suggests that “ease of recall” is a more important determinant than “number of instances” recalled. The availability principle is thus more than a simple generalization from the size of the sample of recalled instances to the whole population of events. It also, and perhaps primarily, refers to the feelings of effort and effortlessness of mental productions.

This is even more transparent in the *simulation heuristic*, sometimes described as a subspecies of availability, namely “availability for construction” in contrast to “availability for recall” (Kahneman & Tversky, 1982b). In prediction, we often compare causal scenarios of the future, and tend to be most convinced by the story that is most easily imaginable, most causally coherent, appears to be most “natural” or normal, and is most easy to follow. Mental simulation is also observed in instances of counterfactual reasoning, when we discuss the probability of events that did not actually occur, but “could” have happened (see Chapter 7, this volume). In some respects, the simulation heuristic can be regarded as an implication of an *a priori fallacy* described by John Stuart Mill, namely to believe that what is natural for us to think must also exist, and what we cannot conceive, must be non-existent. More specifically, “even of things not altogether inconceivable, that we can conceive with the greatest ease is likeliest to be true” (Mill, 1856, p. 312). As with representativeness, the concepts of “availability” and “simulation” do not in themselves specify the processes that bring instances of type X easy to mind, or make models of type M easy to run. Rather, they invite investigators to look for factors that make X and M more retrievable and plausible and hence, more likely.

It may be constructive to point out that both representativeness and availability could be viewed as instances of categorization. Smith, Patalano, and Jonides (1998) proposed that categorization of an instance can be carried out either by applying a category defining rule to an instance in question, or by determining the instance’s similarity to remembered exemplars of a category. Both representativeness and availability are supposedly based on

processes of the latter type. For example, the lawyers/engineers study (Kahneman & Tversky, 1973) that was intended to demonstrate base-rate-neglect can be viewed as a categorization task in which participants have to judge whether a person (briefly described in a personality sketch) should be classified as a lawyer or an engineer depending on the judged similarity between the person's description and the respective prototypes of the two categories. Similarly, regarding availability, when participants attempt to estimate the frequency of the letter R in the first and third place of a word, they supposedly retrieve a few exemplars from the relevant categories and base their estimates on these exemplars (Smith & Medin, 1981).

The interpretation of studies on representativeness may differ depending on whether they are viewed as experiments on probability judgments or whether the focus is on categorization. Probability theory, which serves as the benchmark for assessing representativeness experiments, is a formal theory based on computational principles and as such lends itself exclusively to what Sloman (1996) has termed the rule-based system of reasoning. Categorization, in contrast, in which similarity plays a major role, is more likely to be performed by what Sloman calls the associative system. Examining representativeness (and availability) from these two different perspectives, may provide some useful insights.

### *Anchoring and adjustment*

Judgments are also influenced by initial values, usually suggested by an external source. If asked whether I am willing to sell my old car for \$2,000, I will think of it as less valuable than if I am offered \$4,000, even if I find both offers "outrageously" low. In the first case, I may ask for \$5,000 rather than \$2,000, in the second I may ask for \$7,000, with little awareness about the extent to which my own "independent" estimates are, in fact, influenced by the original suggestions. The estimates can in such cases be regarded as upwards or downwards "adjustments" of the suggested values, whereas the initial suggested values serve as "anchors," towards which the estimates are pulled. This process of *anchoring and adjustment* (Tversky & Kahneman, 1974) thus creates estimates that tend to be biased, or assimilated, in the direction of the anchor.

Despite the inbuilt bias, anchoring and adjustment is clearly an adaptive heuristic whenever the anchor is informative and relevant. In the car sale example, the offer from a prospective buyer provides helpful information about the market value of my car, and should legitimately be taken into account. Without any external hint, my own price expectations might be less biased, but more variable and inaccurate. Sensible people anchor their predictions about the future based on the situation today, resulting in a conservative bias (by judging the future to be more similar to the past than warranted), but a conservative bias may be better than an estimate anchored on a sanguine wish, or simply coming out of the blue. There is, however, no such thing as a foolproof heuristic; when people are uncertain, they can be influenced by an irrelevant anchor value (Wilson, Houston, Etling, & Brekke, 1996) or a completely implausible one (Strack & Mussweiler, 1997).

The anchoring and adjustment heuristic is more general than representativeness and availability, describing a process that applies equally well to frequency judgments, value

judgments, magnitude judgments, and even causal attributions (Gilbert & Malone, 1995; Quattrone, 1982). In the area of probability judgment, anchoring phenomena have been used to explain the hindsight bias (where judgments about the past are biased by one's outcome knowledge), and various phenomena of overconfidence, for instance the tendency to produce too narrow confidence ranges in estimates of uncertain quantities (Alpert & Raiffa, 1982). In this case, the individual performs a guess about his or her most likely estimate, and makes (insufficient) adjustments upwards and downwards to incorporate the uncertainty involved. Alternatively, the lower estimate may function as an anchor for the higher estimate, or vice versa.

Despite the robustness of anchoring phenomena, there is no consensus about the mechanism behind them, not even whether actual adjustments are involved. Chapman and Johnson (2002) distinguish two main categories of explanations: insufficient adjustments (overweighing the anchor compared to other evidence), and selective activation and accessibility of evidence. In the first case, we could perhaps describe anchoring as a *primacy* effect; in the second case it functions as a special case of *priming* (Mussweiler & Strack, 2000). Epley (see Chapter 12, this volume) suggests that anchoring phenomena might be due to several, independent mechanisms.

## **Heuristics and Biases: A Current Evaluation**

The introduction of the heuristics and biases program was enthusiastically adopted by researchers and has been followed by 30 years of intensive research and corresponding disputes. This accumulating research was often guided by the question concerning the extent to which the heuristics and the associated biases should be considered as evidence for failures of rationality (e.g., Cohen, 1981, 1983; Evans & Over, 1996; Gigerenzer, 1996; Stanovich & West, 2002).

Much of the research consolidated previous findings and at the same time delineated the circumstances and conditions under which specific biases would appear, and sometimes disappear. For instance, a review paper by Koehler (1996) on the base-rate fallacy (one of the more prominent biases linked to the representativeness heuristic) provides overall evidence for the robustness of the phenomenon. Yet, at the same time, Koehler points out possible methodological shortcomings indicating that researchers have been too quick to conclude that people simply “neglect” the base rates.

The continuous build up of the heuristics and bias research program extended in two ways. First, the number of newly identified biases has been constantly growing. For instance, in one of the more popular textbooks on judgment and decision making, Baron (2002) counts no less than 25 biases (see the term bias in his subject index). Second, new heuristics have appeared, but not at the same pace, and not as widely adopted as the three original ones. Among the newcomers are “the numerosity heuristic” (Pelham, Sumarta, & Myaskovsky, 1994), according to which the number of instances of a target is used to indicate its probability (regardless of the number of non-target instances); “the recognition heuristic” (Goldstein & Gigerenzer, 1999), which says that alternatives with known (recognized) labels are automatically believed to be a bigger, better, and safer

than alternatives with unknown labels; and “the affect heuristic” (Slovic, Finucane, Peters, & MacGregor, 2002), referring to people’s tendency to regard objects and activities with positive connotations as yielding positive outcomes with higher probability, and negative outcomes with lower probability, than objects with negative connotations. It has also been suggested that people often assess probabilities by heuristically comparing the target outcome only to its strongest competitor, rather than to the whole set of alternatives, creating the “alternative outcomes effect” (Windschitl & Wells, 1998), and that people, especially in hindsight, evaluate probabilities of a counterfactual outcome by their impression of how close it was to occurring, thus apparently adopting a “closeness” or “proximity” heuristic (Kahneman & Varey, 1990; Teigen, 1998).

In hindsight, it may have been unfortunate that heuristics and biases were introduced in unison, as a slogan or brand name, giving rise to the impression that the main task of heuristics was to produce biases, and that any bias was to be explained by a corresponding heuristic. Critics (e.g., Fiedler, 1983; Gigerenzer et al., 1999; Lopes, 1991) have pointed out that the proposed heuristics are vague and hence not readily testable, that they do not constitute a comprehensive model of probability judgments, and that they differ from problem-solving heuristics by being more often automatic than conscious and deliberate. Perhaps it is fair to say that they were introduced – like Einstein’s model alluded to previously – not as a theory, but as a heuristic [sic] device suggesting rather than dictating ways of thinking about subjective probabilities. From the amount of research inspired by this approach, the idea of heuristics appears to have been a fruitful heuristic.

## **Two Stages of Probability Judgments**

If probability judgments, and the possible biases associated with such judgments, are not to be explained by a finite set of concrete “heuristics,” how could the judgment process (alternatively) be conceived?

Recent developments in research on heuristics (e.g. Kahneman & Frederick, 2002) suggest that probability judgments may result from an interaction between two modes of thinking: one intuitive, automatic, and immediate (labeled System 1), and another more analytic, controlled, and rule-governed form of reasoning (System 2). In this scheme, spontaneous System 1 judgments may or may not be biased, and these biases may or may not be endorsed, corrected, or adjusted by System 2. Typical heuristic judgments (e.g. impressions of representativeness and priming effects caused by anchoring) can be explained by operations that are dominated by the first rather than the second of these systems. Responses induced by the first system are spontaneous and often irresistible, bearing some similarity to output from the perceptual system. Like the perceptual apparatus, System 1 may occasionally wind up with (cognitive) illusions. System 2 processes are, on the other hand, more slow and deliberate. This does not necessarily mean that they are always compatible with normative prescriptions. Extensive empirical evidence suggests that we are capable of being mistaken in different ways, leading to violations of the laws of logic or probability calculus. We may lack the proper rule (e.g., regression

towards the mean) leading to what has been termed errors of competence. We may strongly believe in rules that are irreconcilable with normative considerations (e.g., the gambler's fallacy). And, even if we are familiar with the proper rule, we are occasionally prone to make mistakes resulting in what has been termed errors of application.

Without necessarily endorsing the view that there are two distinct ways of thinking, as proposed by some models (Epstein, 1994; Slovic, 1996; Stanovich & West, 2002), we may profit from the two-phase analysis and posit that most instances of prediction and probability judgments include a phase in which candidate judgments are suggested or formulated, and a phase in which these proposals (or hypotheses) are evaluated. This is especially apparent in the case of anchoring and adjustment, where the anchor represents an externally suggested candidate value, to be modified and evaluated during the subsequent adjustment stage. In the case of representativeness, an initial prediction is made on the basis of how well a sample or a target outcome matches, or resembles, salient characteristics of the parent population, or outcome source. This prediction may subsequently be corrected and moderated by factors like base rates, beliefs about cue validity, or a record of previous prediction accuracy. Sometimes people use simple rules of thumb to ensure that some corrections are made. When asked about her confidence of testimony, an eyewitness (in a Norwegian murder case) recently claimed that she was "90 percent sure; when I do not say 100 percent, it is because I *never* say 100 percent." This witness evidently used a simple, deliberate principle to modify her immediate, perceptually based impression that the observed person was identical with the suspect. We could even call her use of a correction factor a "judgmental heuristic," with "heuristic" in this case indicating a consciously chosen strategy (to minimize errors of overconfidence) rather than an immediate, intuitive process. As proposed earlier, probability judgments are based on psychological principles of perception on one hand, and thinking and reasoning on the other. Supposedly, initial impressions and assessments (of a situation or an event) are mainly construed according to perceptual laws, whereas the subsequent evaluation phase is mainly based on deliberate conscious reasoning. Analogous to the two stages underlying choice behavior as postulated by prospect theory (Kahneman & Tversky, 1979), we suggest that probability judgments are governed by an initial *editing and encoding* phase followed by *evaluation*.

### *Phase 1: Editing (encoding)*

The initial editing phase is composed of structuring and arranging the available incoming information in a meaningful way, preparing it for the subsequent evaluative-computational phase. Given a limited processing and memory capacity, editing is designed to encode the information in the simplest and most meaningful way. The manner by which the perceptual system is tuned to encode the available information is based on what Bruner (1957) has referred to as perceptual readiness and is founded on some underlying (Gestalt) principles. In a broader context, editing is guided by what Pomerantz and Kubovy (1986) have termed the simplicity principle, according to which the perceptual system is geared up to find the simplest perceptual organization (what the Gestalt psychologists referred to as *prägnanz*).

Editing is responsible for selection of information and transforming it into an internal representation which, among other things, would depend on stimulus characteristics like concreteness and vividness. For instance, as originally proposed by Meehl (1954), and demonstrated in countless studies, people are evidently more tuned to the singular (clinical) than to statistical evidence. It has been proposed that the clinical singular case is more vivid, and therefore is given priority in the editing phase. The strength of this vividness effect would depend on how the available information (verbal or non-verbal) presents itself. For instance, in the well-known lawyer/engineer problem (Kahneman & Tversky, 1973), participants were presented with both a personality sketch (of either a lawyer or an engineer) and with base-rates regarding the number of lawyers and engineers respectively. Evidently, participants made their judgment mainly on the basis of the specific description ignoring the normatively important base-rate information. Kahneman and Tversky assert that participants evaluate the likelihood of a particular description to be that of an engineer or a lawyer by the degree to which the particular description resembles (or is representative) of the typical stereotype associated with these two occupations. In the framework proposed here, the editing phase is particularly sensitive to singular narrative information, which frequently grabs the major attention at this initial stage. In particular, the vivid character of the stereotypical sketch descriptions of the lawyer and the engineer draws immediate attention and is encoded as highly salient. This encoding, like the editing phase in general, is recognition based and to a large extent automatic. It is insensitive to the accuracy, validity, or diagnosticity of such descriptions which, if at all, are assessed only at the subsequent evaluation stage.

The operations of the editing and the corresponding initial impressions are highly dependent on the order of the incoming information and the manner by which it is structured and arranged. Studies of anchoring show the importance of order (primacy effects). Studies of framing effects reveal how the same, objective, facts can have different impact dependent upon how they are presented (Tversky & Kahneman, 1981; Levin, Schneider, & Gaeth, 1998). For instance an 80 percent chance of success (positive frame) appears more encouraging than a 20 percent chance of failure (negative frame), by directing our attention towards a positive versus a negative target outcome. Framing can also be achieved by the choice of probability terms, the 80 percent probability of success can be described as a “highly probable” success or a “not completely certain” success, the first description being more optimistic than the second (Teigen & Brun, 2003).

Similar to framing, editing is also vulnerable to all sorts of format effects. For instance, much of the controversy concerning base-rate neglect (the tendency to overweight singular narrative information and undermine corresponding statistical information) is directly linked to how the information is presented. The difference between studies that demonstrate base-rate neglect compared with those that fail to find the effect (Koehler, 1996), is largely dependent on how the two types of information are presented. Different presentation formats enhance some aspects more than other, resulting in a different structure of the internal representation. Note that framing is not necessarily restricted to verbal descriptions. Perceptual stimuli (and situations) can be equally presented and perceived in more than one way.

Descriptions of target outcomes can also differ by specificity, or amount of detail. This is a central point in support theory (Tversky & Koehler, 1994; Rottenstreich &

Tversky, 1997), where it is claimed that people do not allocate probabilities to events, but to descriptions of events. Events that are described in such a way that they will generate a large amount of support (positive evidence and favorable arguments) will be estimated as more probable than those that are described in such a way that they will be more sparsely supported. The most important corollary of this view is that an “unpacked” outcome (for instance deaths by traffic accidents, natural disasters, terrorism, homicide, or suicide) is believed to be more probable than the corresponding “packed” outcome (“death from unnatural causes”), even if the latter include the former. Such “subadditivity” has been documented in many domains.

### *Phase 2: Evaluation*

The editing phase determines which aspects of the incoming information will receive more or less attention, and arranges (structures) the information preparing it for the subsequent evaluation phase. This latter phase consists of assessing the different aspects of the available information obtained from the editing phase, eventually combining them into a probabilistic estimate (in a numerical or verbal form). The evaluation phase supposedly consists of deliberate cognitive processes that are, at least to some extent, based on what Bruner (1984) has termed the paradigmatic or logico-scientific mode of reasoning. This mode is regulated by requirements of consistency and non-contradiction, and in its most developed form fulfills the ideal of a formal mathematical system of description and explanation. However, there is overwhelming empirical evidence (much of which has been stimulated by the heuristics and biases approach) suggesting that the evaluation phase can also be prone to systematic errors and reasoning faults. Failures at the evaluation phase may be due to different reasons.

First, in many cases people are familiar with the appropriate (paradigmatic) way of thinking yet fail to apply it to the particular case thus resulting in what has been termed errors of application (Kahneman & Tversky, 1982a). For example, they presented (p. 127) participants with the following question: “As you know, a game of squash can be played either to 9 or to 15 points. Holding all other rules of the game constant, if A is a better player than B, which scoring system will give A a better chance of winning.”

Most participants believed that the scoring rule should not make a difference, yet (with few exceptions) they were convinced after being told that A (the better player) would be better off with a scoring rule of 15 because an atypical outcome is less likely to occur in a large sample. The likelihood of “correct” applications at the evaluation phase depends on the extent to which the problem structure is transparent, and in turn on the manner by which it is encoded at the initial editing phase.

Second, the principles underlying statistical theory are neither easy to grasp nor always compatible with natural intuitions (Lewis & Keren, 1999). Indeed, themes like regression toward the mean or inverse probabilities are not just difficult to comprehend, but (or because) they are not part of our natural reasoning tools. Hence, the evaluation phase fails in those instances in which the proper rule, procedure, or more generally way of thinking, is unknown or not recognized resulting in what is referred to as errors of comprehension.



Third, there are several statistical and probabilistic phenomena on which we possess deeply rooted misconceptions, that may dominate the evaluative phase. By misconceptions is meant beliefs that are neither compatible with the physical world nor with normative considerations based on the paradigmatic mode of reasoning. Two of the most pervasive ones are a deficient understanding of randomness (e.g., Bar-Hillel & Wagenaar, 1993) as exemplified, for instance, by the belief in the “hot hand” (Gilovich, Vallone, & Tversky, 1985), and the failure to understand statistical independence as exhibited in the gambler’s fallacy (Keren & Lewis, 1994).

When probability evaluations, even analytical and deliberate ones, sometimes differ from the normative rules, it could be due to the kind of probability concept people endorse. Even among probability theorists, there is no consensus about what is the true reference for a probability statement. Should probability statements be reserved for repeatable events, as claimed by proponents of the frequentistic approach, or are probability statements fundamentally statements about a person’s ideal degree of confidence, as claimed by the personalistic school (de Finetti)? Can probability statements legitimately refer to unique situations by being descriptive of the causal propensities involved (Popper). Lay people may, in different contexts, endorse versions of all these views, albeit in a less stringent and explicit form. For instance we may distinguish between “external” (sometimes called aleatory) and “internal” (epistemic) probabilities (Kahneman & Tversky, 1982c). In daily life, probability is for most of us a “polysemous” concept (Hertwig & Gigerenzer, 1999), referring on some occasions to relative frequencies, and in other situations simply to “plausibility.” In many cases, people seem to think of probabilities as a kind of causal forces, or dispositions, manifesting themselves not only in outcome frequencies but also in the strength and latency of target outcomes. For instance, when people are told about the risk of an earthquake in a particular region during the next three years, they will believe that it will come sooner and be stronger if  $p = .8$  than if  $p = .6$  (Keren & Teigen, 2001). With such interpretations, probabilities tend to become viewed as characteristics of causal systems, with no urgent need to obey formal axioms of distributive probabilities.

## **Closing Comments**

It is naturally impossible to cover, in a single chapter, all the aspects of the heuristic and bias research program and its implications for decision-making research. In this final section we briefly assess the achievements and the limitations of this research and the possible directions in which it may evolve in the future.

The heuristic and bias research program made several important contributions. First, it successfully combined perceptual principles with the psychology of thinking and reasoning, offering a new perspective on judgment under uncertainty. Second, it provided irrefutable evidence that humans’ reasoning and decision-making capabilities, though certainly remarkable, are prone to systematic errors. Third, and as a consequence, it challenged the rigid assumptions of economic theory regarding “Homo economicus” and human rationality associated with it. Evidently, people are not always able to follow



the prescriptions of normative theories (despite the fact, that these were originally constructed by the human mind) as is assumed by standard economic theory. Finally, it offered simple and clever methods for the study of probability judgments. Not undermining its inspiring achievements, a comprehensive theory that can encompass the different heuristics under one framework is still lacking. Different heuristics are explicated by different processes which are only partially linked. Given that the different heuristics are based on a wide range of perceptual and cognitive mechanisms, it is questionable whether an all-inclusive theory of heuristics and biases is feasible. One promising step has been the development of support theory (Rottenstreich & Tversky, 1997; Tversky & Koehler, 1994) according to which probability judgments correspond to an assessment of the relative balance of evidence for and against competing hypotheses. Though the theory can serve as a global framework for the heuristic approach, it does not explain how, and under what conditions, the different heuristics would be operating.

Most of the empirical demonstrations regarding the different heuristics are based on explicitly eliciting people's probability judgments. An open question is how different elicitation procedures induce different heuristics, leading to different biases. Are different heuristics deeply rooted facets of the cognitive system, or are they mainly brought to mind (online) by the specific elicitation method employed? This question has both theoretical and practical implications. Attempting to answer this question may provide a useful guideline for future theoretical research. From a more practical viewpoint, it may have an important contribution to the development of enhanced corrective (often referred to as debiasing) methods.

## References

- Alpert, M. & Raiffa, H. (1982) A progress report on the training of probability assessors. In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 294–305). Cambridge: Cambridge University Press.
- Bar-Hillel, M. & Wagenaar, W. A. (1993) The perception of randomness. In G. Keren and C. Lewis (eds.), *A Handbook for Data Analysis in the Behavioral Sciences: Methodological Issues*. Hillsdale, NJ: Lawrence Erlbaum.
- Baron, J. (2002) *Thinking and Deciding* (3rd edn.). Cambridge: Cambridge University Press.
- Baron, J. & Hershey, J. C. (1988) Outcome bias in decision evaluation, *Journal of Personality and Social Psychology*, 54, 569–79.
- Bruner, J. S. (1957) On perceptual readiness, *Psychological Review*, 54, 123–49.
- Bruner, J. S. (1984) Narrative and paradigmatic modes of thought. Paper presented at the annual APA meeting, Toronto (August 25, 1984).
- Budescu, D. V. & Bruderman, M. (1995) The relationship between the illusion of control and the desirability bias, *Journal of Behavioral Decision Making*, 8, 109–25.
- Chapman, G. B. & Johnson, E. (2002) Incorporating the irrelevant: Anchors in judgments of belief and value. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 120–38). Cambridge: Cambridge University Press.
- Cohen, J. (1960) *Chance, Skill, and Luck*. Baltimore: Penguin.
- Cohen, J. (1964) *Behavior in Uncertainty*. New York: Basic Books.
- Cohen, L. J. (1981) Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences*, 4, 317–70.

- Cohen, L. J. (1983) The controversy about irrationality, *Behavioral and Brain Sciences*, 6, 510–17.
- Epstein, S. (1994) Integration of the cognitive and the psychodynamic unconscious, *American Psychologist*, 49, 709–24.
- Evans, J. St. B. T. (1989) *Bias in Human Reasoning: Causes and Consequences*. Hove, UK: Erlbaum.
- Evans, J. St. B. T. & Over, D. E. (1996) *Rationality and Reasoning*. Hove, UK: Psychology Press.
- Fiedler, K. (1983) On the testability of the availability heuristic. In R. W. Scholz (ed.), *Decision Making Under Uncertainty* (pp. 109–19). Amsterdam: North-Holland.
- Frederick, S. (2002) Automated choice heuristics. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 548–58). Cambridge: Cambridge University Press.
- Gigerenzer, G. (1991) How to make cognitive illusions disappear: Beyond “heuristics and biases,” *European Review of Social Psychology*, 2, 83–115.
- Gigerenzer, G. (1996) On narrow norms and vague heuristics: A reply to Kahneman and Tversky, *Psychological Review*, 103, 592–6.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999) *Simple Heuristics that Make Us Smart*. Oxford: Oxford University Press.
- Gilbert, D. T. & Malone, P. S. (1995) The correspondence bias, *Psychological Bulletin*, 117, 21–38.
- Gilovich, T., Griffin, D., & Kahneman, D. (2002) *Heuristics and Biases*. Cambridge: Cambridge University Press.
- Gilovich, T., Vallone, R., & Tversky, A. (1985) The hot hand in basketball: On the misperception of random sequences, *Cognitive Psychology*, 17, 295–314.
- Goldstein, D. G. & Gigerenzer, G. (1999) The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. Todd, & the ABC Research Group, *Simple Heuristics that Make Us Smart* (pp. 37–58). Oxford: Oxford University Press.
- Griffin, D., Gonzalez, R., & Varey, C. (2001) The heuristics and biases approach to judgment under uncertainty. In A. Tesser and N. Schwarz (eds.), *Blackwell Handbook of Social Psychology: Intra-individual Processes* (pp. 207–35). Oxford: Blackwell.
- Hake, H. W. & Hyman, R. (1953) Perception of the statistical structure of a random series of binary symbols, *Journal of Experimental Psychology*, 45, 64–74.
- Hertwig, R. & Gigerenzer, G. (1999) The “conjunction fallacy” revisited: How intelligent inferences look like reasoning errors, *Journal of Behavioral Decision Making*, 12, 275–305.
- Hume, D. (1976) *A Treatise on Human Nature*. Oxford: Clarendon Press (Original published in 1739).
- Kahneman, D. & Frederick, S. (2002) Representativeness revisited: Attribute substitution in intuitive judgments. In T. Gilovich, T. D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 49–81). Cambridge: Cambridge University Press.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991) The endowment effect, loss aversion, and status quo bias, *Journal of Economic Perspectives*, 5, 193–206.
- Kahneman, D., Slovic, P., & Tversky, A. (1982) *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. & Tversky, A. (1972) Subjective probability: A judgment of representativeness, *Cognitive Psychology*, 3, 430–54.
- Kahneman, D. & Tversky, A. (1973) On the psychology of prediction, *Psychological Review*, 80, 237–51.
- Kahneman, D. & Tversky, A. (1979) Prospect theory: An analysis of decision under risk, *Econometrica*, 47, 263–91.
- Kahneman, D. & Tversky, A. (1982a) On the study of statistical intuitions, *Cognition*, 11, 123–41.

- Kahneman, D. & Tversky, A. (1982b) The simulation heuristic. In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 201–8). Cambridge: Cambridge University Press.
- Kahneman, D. & Tversky, A. (1982c) Variants of uncertainty, *Cognition*, 11, 143–57.
- Kahneman, D. & Varey, C. A. (1990) Propensities and counterfactuals: The loser that almost won, *Journal of Personality and Social Psychology*, 59, 1101–10.
- Kanouse, D. E. & Hanson, L. R. (1971) Negativity in evaluations. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, and B. Weiner (eds.), *Attribution: Perceiving the Causes of Behavior* (pp. 47–62). Morristown NJ: General Learning Press.
- Keren, G. & Lewis, C. (1994) The two fallacies of gamblers: Type I and Type II, *Organizational Behavior and Human Decision Processes*, 60, 75–89.
- Keren, G. & Teigen, K. H. (2001) The probability-outcome correspondence principle: A dispositional view of the interpretation of probability statements, *Memory & Cognition*, 29, 1010–21.
- Klayman, J. & Ha, Y.-W. (1987) Confirmation, disconfirmation, and information in hypothesis testing, *Psychological Review*, 94, 211–28.
- Koehler, J. (1996) The base-rate fallacy reconsidered: Descriptive, normative, and methodological challenges, *Behavioral and Brain Sciences*, 19, 1–53.
- Laplace, P. S. (1816) *Essai philosophique sur les probabilités*. Paris: Courcier.
- Levin, P., Schneider, S. L., & Gaeth, G. J. (1998) All frames are not created equal: A typology and critical analysis of framing effects, *Organizational Behavior and Human Decision Processes*, 76, 149–88.
- Lewis, C. & Keren, G. (1999) On the difficulties underlying Bayesian reasoning: A comment on Gigerenzer and Hoffrage, *Psychological Review*, 106, 411–16.
- Lopes, L. L. (1991) The rhetoric of irrationality, *Theory and Psychology*, 1, 65–82.
- Meehl, P. E. (1954) *Clinical Versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*. Minneapolis: University of Minnesota Press.
- Mill, J. S. (1856) *A System of Logic*. London: Parker.
- Mussweiler, T. & Strack, F. (2000) Comparing is believing: A selective accessibility model of judgmental anchoring. In W. Stroebe and M. Hewstone (eds.), *European Review of Social Psychology*, 10 (pp. 135–67). Chichester, UK: Wiley.
- Nisbett, R. E. & Ross, L. (1980) *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Olson, C. L. (1976) Some apparent violations of the representativeness heuristic in human judgment, *Journal of Experimental Psychology: Human Perception and Performance*, 2, 599–608.
- Oxford English Dictionary* (2002). Oxford: Oxford University Press, online at: <http://dictionary.oed.com>
- Peeters, G. (1971) The positive-negative asymmetry: On cognitive consistency and positivity bias, *European Journal of Social Psychology*, 1, 455–74.
- Pelham, W. B., Sumarta, T. T., & Myaskovsky, L. (1994) The easy path from many to much: The numerosity heuristic, *Cognitive Psychology*, 26, 103–33.
- Peterson, C. & Beach, L. R. (1967) Man as an intuitive statistician, *Psychological Bulletin*, 68, 29–46.
- Polya, G. (1945) *How to Solve it: A New Aspect of Mathematical Method*. Princeton: Princeton University Press.
- Pomerantz, J. R. & Kubovy, M. (1986) Theoretical approaches to perceptual organization: Simplicity and likelihood principles. In K. R. Boff, L. Kaufman, and J. P. Thomas (eds.), *Handbook of Perception and Human Performance: Volume II. Cognitive Processes and Human Performance* (pp. 1–45). New York: Wiley.

- Pronin, E., Lin, D. Y. & Ross, L. (2002) The bias blind spot: Perception of bias in self versus others, *Personality and Social Psychology Bulletin*, 28, 369–81.
- Quattrone, G. A. (1982) Overattribution and unit formation: When behavior engulfs the person, *Journal of Personality and Social Psychology*, 42, 593–607.
- Rottenstreich, Y. & Tversky, A. (1997) Unpacking, repacking, and anchoring: Advances in support theory, *Psychological Review*, 104, 406–15.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991) Ease of retrieval as information: Another look at the availability heuristic, *Journal of Personality and Social Psychology*, 61, 195–202.
- Sedlmeier, P., Hertwig, R., & Gigerenzer, G. (1998) Are judgments of the positional frequencies of letters systematically biased due to availability? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 754–70.
- Sidgwick, A. (1883) *Fallacies: A View of Logic from the Practical Side*. London: Kegan Paul, Trench & Co.
- Slooman, S. A. (1996) The empirical case for two systems of reasoning, *Psychological Bulletin*, 119, 3–22.
- Slovic, P., Finucane, M., Peters, E., & MacGregor, D. (2002) The affect heuristic. In T. Gilovich, T. D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 397–420). Cambridge: Cambridge University Press.
- Smith, E. E. & Medin, D. L. (1981) *Categories and Concepts*. Cambridge, MA : Harvard University Press.
- Smith, E. E., Patalano, A. L., & Jonides, J. (1998) Alternative strategies of categorization, *Cognition*, 65, 167–96.
- Spranca, M., Minsk, E., & Baron, J. (1991) Omission and commission in judgment and choice, *Journal of Experimental Social Psychology*, 27, 76–105.
- Stanovich, K. E. & West, R. F. (2002) Individual differences in reasoning: Implications for the Rationality debate? In T. Gilovich, T. D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 421–40). Cambridge: Cambridge University Press.
- Strack, F. and Mussweiler, T. (1997) Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility, *Journal of Personality and Social Psychology*, 73, 437–46.
- Sully, J. (1882) *Illusions: A Psychological Study*. London: Kegan Paul, Trench & Co.
- Teigen, K. H. (1998) When the unreal is more likely than the real: Post hoc probability judgments and counterfactual closeness, *Thinking and Reasoning*, 4, 147–77.
- Teigen, K. H. & Brun, W. (2003) Verbal probabilities: A question of frame? *Journal of Behavioral Decision Making*, 16, 53–72.
- Tversky, A. & Kahneman, D. (1971) Belief in the law of small numbers, *Psychological Bulletin*, 76, 105–10.
- Tversky, A. & Kahneman, D. (1972) Availability: A heuristic for judging frequency and probability, *Cognitive Psychology*, 5, 207–32.
- Tversky, A. & Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases, *Science*, 185, 1124–31.
- Tversky, A. & Kahneman, D. (1981) The framing of decisions and the psychology of choice. *Science*, 211, 453–8.
- Tversky, A. & Kahneman, D. (1982) Judgments of and by representativeness. In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 84–98). Cambridge: Cambridge University Press.
- Tversky, A. & Kahneman, D. (1983) Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment, *Psychological Review*, 90: 293–315.
- Tversky, A. & Koehler, D. J. (1994) Support theory: A nonextensional representation of subjective probability, *Psychological Review*, 101, 547–67.

- Wason, P. (1960) On the failure to eliminate hypotheses in a conceptual task, *Quarterly Journal of Experimental Psychology*, 12, 129–40.
- Wilson, T. D., Houston, C. E., Etling, K. M., & Brekke, N. (1996) A new look at anchoring effects: Basic anchoring and its antecedents, *Journal of Experimental Psychology: General*, 125, 387–402.
- Windschitl, P. D. & Wells, G. L. (1998) The alternative-outcomes effect, *Journal of Personality and Social Psychology*, 75, 1411–23.