The Rationality of Risky Decisions

A Changing Message

Anton Kühberger
UNIVERSITY OF SALZBURG

ABSTRACT. Over the years research in risky decision making has diagnosed variable degrees of irrationality in people's judgements and choices. In the 1960s an optimistic view dominated of a widely rational decision maker. The work of Tversky and Kahneman at the beginning of the 1970s led to a pessimistic view of basically flawed decision processes that frequently end up in 'cognitive illusions'. In the 1980s a movement gained strength that pointed to the adaptiveness of seemingly irrational decisions. Recent work demonstrates that seemingly irrational choices may be due to different task construal between experimenters and participants. The respective evaluative change in what the rationality issue is generally taken to show is overdue, however. The negative message of fundamentally flawed human decision making has to be replaced by a more positive picture that acknowledges that some reactions to task and context are advantages rather than disadvantages of human decision making. Recent work on systems of thinking shows that different task construals can be meaningfully related to different systems of thinking, thus enabling a more unbiased treatment of the rationality issue.

KEY WORDS: bias, choice, judgement, rationality, risk

Imagine I ask you: 'What is your favourite colour?' You answer: 'Blue.' A little later, I ask again: 'What is your favourite colour?' You answer: 'Red.' I am perplexed because you answer differently to the same question, and I may be tempted to conclude that you lack rationality in your preferences, since preferences should be stable to some degree. Imagine I ask you: 'Why did you buy the Sony TV set?' You answer: 'Oh, I do not know. It was just a feeling. The Toshiba might be better, but somehow I ended up with the Sony.' Again, I may be tempted to conclude that you lack rationality in your choice, maybe due to 'intuitive' information processing. Imagine I ask you: 'Why did you leave your job?' You answer: 'Something happened there that made me so upset that I could no longer think clearly. I acted on my emotions.' Again I may be tempted to conclude that you failed to make a rational choice, since undue emotional influences disturbed your otherwise orderly information processing.

THEORY & PSYCHOLOGY Copyright © 2002 SAGE Publications. VOL. 12(4): 427–452
[0959-3543(200208)12:4;427–452;026293]

The above examples show that the term ‘rationality’ has multiple meanings in ordinary language. In scientific decision theory, however, it is used in yet another sense. In its weakest formulation, it means that people act in accordance with the situation that they believe themselves to be in. A stronger meaning of rationality is implied in utility theory: (i) people act so as to maximize their utility, and (ii) there exists a preference order that allows the choice of the option that offers the maximal utility. This is formalized by multiplying the utility x of each option i by its probability of occurrence p: $EU = \sum p_i x_i$. In general, behaviour that is in accord with the axioms of utility theory (Savage, 1954; von Neuman & Morgenstern, 1947) is seen as rational. Rational decision makers will thus be indifferent between different combinations of probability and value (*axiom of continuity*); indifferent options can be substituted by one another without changing their preference (*substitution principle*); the addition or subtraction of options does not change their preference order (*axiom of independence*); and the preference order is invariant with respect to various formal operations, for instance to multiplication (*axiom of reduction*). Thus, utility theory typically defines rationality with respect to the adherence to certain formal axioms of choice.

There is a good reason for this narrow construal of rationality in risky decision making. People commonly use terms such as ‘probability’ and ‘chance’ to describe their beliefs and expectations about unique events, although some statisticians have denied any role for the rules of probability in describing events that cannot be replicated for an infinite series. For repeated events there are rules to calculate the correct probability (e.g. the chance of a given hand of cards), and thus to evaluate the appropriateness of a judgement. But for unique events, such as, for instance, for the prediction of the probability of a recession in the year 2015, the appropriateness of the judgement can only be tested by examining its *coherence* relative to other judgements, and by examining its *calibration* when aggregated together with several other judgements (i.e. events predicted with .70 probability will occur 70 per cent of the time). Only judgements that follow the von Neuman and Morgenstern axioms are coherent and well calibrated; coherence and calibration can thus be used to evaluate the appropriateness of the psychologically more interesting probabilities for unique events.

In what follows I will summarize the upshot of 40 years of decision research. First the attack on rationality emanating from the studies of Tversky and Kahneman is discussed, then different views are sketched which see the decision maker as the source of irrationality. Then a view follows that indicates that the source of irrationality resides in the choice of models, rather than in the decision maker. Finally, as a synopsis, an outline is given of the interdependencies between systems of thinking, types of problems and rationality.

Rationality: Setting the Stage

In a 'mechanistic' explanation, a phenomenon is explained by analysing its internal causal structure, that is, the 'design level' of analysis (Dennett, 1987). In the rational or 'purposive' way (called 'intentional' by Dennett), a phenomenon is explained in terms of what its purpose is and why. In psychology, a purposive explanation of people's behaviour is in terms of their beliefs and goals. In the cognitive sciences, and especially in decision theory, there has been a strong predominance of mechanistic explanation (Chater & Oaksford, 1999), since computational models, whether symbolic or connectionist, have focused on specifying architectures and algorithms for cognition. For example, experimental studies have carefully documented the structural elements of memory and reasoning but have displayed relatively little concern for why these processes work as they do.

The Message of Heuristics and Biases

From this mechanistic view on explanation a picture has emerged of the cognitive system as an assortment of apparently arbitrary mechanisms. A prominent message that emerged especially from judgement and decision-making research in the 1970s (one that was well heard outside this area, outside psychology, and even outside the scientific literature; see, for instance, the McCormich (1987) article in *Newsweek*, or the McKean (1985) article in *Discover*), was that this assortment of mechanisms is subject to capricious limitations that are a permanent source of irrationality in judgement and decision making. It is important to note that this message stands in stark contrast to what was generally accepted a few years before. Prior to 1970 or so, most researchers in judgement and decision making believed that people are pretty good decision makers, as is indicated by the title of the most frequently cited summary paper of that time: 'Man as an Intuitive Statistician' (Peterson & Beach, 1967). The use of the statistician metaphor was intended to convey a positive view of the proper use of normative rules in judgement and decision making. Since then, however, opinion has taken a decided turn for the worse. However, it was neither the case that participants suddenly became any less adept in experimental tasks, nor that experimenters began to grade performance against a tougher standard. Instead, as Lopes (1991) argues, researchers began to (i) emphasize some results at the expense of others and (ii) use different paradigmatic tasks.

Differential Emphasis of Research Results

Christensen-Szalanski and Beach (1984) attempted to demonstrate such a selection process and differential emphasis of research results. They surveyed all the experimental results appearing in *Psychological Abstracts* for the years 1972–81 using the key words *decision-making*, *judgement* and *problem-solving*. This yielded 84 papers of which 37 reported good performance (44 per cent) and 47 reported poor performance (56 per cent). A count of the citations of these papers as a measure of visibility in the sampled period yielded an average of 27.8 citations of reports claiming poor performance, whereas reports of good performance were cited only 4.7 times. This represents a ratio of about 1:6 citations for good and poor performance and goes to show a clear bias to cite reports that claim poor performance.

However, a re-analysis of these data by Robins and Craik (1993) weakens the claim of a differential emphasis considerably. Using the number of citations per year rather than the mean total citations, Robins and Craik found a citation ratio of about 1:3 citations for good and poor performance. More importantly, they identified a variety of qualitatively different standards for appraising good versus poor performance. For instance, the majority of 'good performance' articles featured evidence of psychological consistency in judgement while the majority of 'poor performance' articles documented deviations from a statistical model. They concluded that 'empirical support for a citation-rate disparity is modest at best and possibly unreliable' (p. 242).

A second factor for moving the focus of attention toward poor performance was that articles reporting poor performance tended to be published in more important journals. This is surely true for four empirical articles by Amos Tversky and Daniel Kahneman that constitute the cornerstone of what today is called the 'heuristics-and-biases'literature (Kahneman & Tversky, 1972, 1973; Tversky & Kahneman, 1971, 1973) and a summary article in *Science* (Tversky & Kahneman, 1974).[1] The central tenet of the heuristics-and-biases programme is that human reasoning is fundamentally irrational (for an overview see Kahneman, Slovic, & Tversky, 1982)

Although the notion of heuristics does not necessarily imply bias (heuristic processes can be less accurate due to random error rather than systematic error), heuristics were actually understood as indicating bias. It was especially the *Science* article (Tversky & Kahneman, 1974) that exported this idea to areas outside psychology. This article was cited a total of 227 times in 127 different journals between the years 1975 and 1980, and about 20 per cent of the citations were from sources outside psychology (Berkeley & Humphreys, 1982). Although some scholars (e.g. Hertwig & Ortmann, 2001), most notably economists (e.g. Harrison, 1994; Smith & Walker, 1993), still resist, the ease with which the message of bias and irrationality made its way into other fields of the social sciences like sociology, political science, law, economics, business and anthropology is surprising, since these fields tend to look with suspicion on the tasks that psychologists study in their laboratories, particularly when the studies are carried out with relatively unmotivated volunteers. However, this was not so in the heuristics-and-biases case, where the issue of generalizability was seldom raised and rarely mentioned: ‘Human incompetence is presented as a fact, like gravity’ (Lopes, 1991, p. 67).

It is noteworthy that the scientific community did welcome the negative bias message rather than resisting it. We can understand why some disciplines welcomed it. For instance, the then growing field of Artificial Intelligence (AI) could possibly profit from demonstrating an advantage of machine reasoning over human reasoning. However, as was subsequently seen, it was just the other way round. AI implemented the use of heuristics because it became clear that heuristics can have marked advantages over purely algorithmic processing. In AI the term ‘heuristics’ now has a positive rather than negative meaning. However, for other disciplines we see no such reason and we are left with the suspicion of a hidden motive to readily accept the negative message.

Different Paradigmatic Task

Kahneman and Tversky introduced a new task type in decision research. Before 1970, probabilistic thinking was studied by asking participants to give intuitive estimates of statistical indices, for example the variance of a set of scores or the correlations between scores, and researchers used urn-and-balls problems, a paradigm now termed the ‘bookbag-and-poker-chip paradigm’. Although it is somewhat odd to expect naïve people to have, for instance, intuitions about the average squared deviation of scores about a mean (i.e. the variance of scores), participants did quite well. To be sure, it was not assumed that people have statistical equations inside their heads, but rather the analysis bypassed the head altogether. Kahneman and Tversky introduced a new focus of attention: the *process* of judgement. In accordance with what was dominant practice in psychology then, they began to formulate hypotheses about how people come to know their answers. These hypotheses pointed to the use of cognitive short-cuts, termed ‘heuristics’, which people were assumed to use in the process of judgement. Put broadly, heuristic methods are quick-and-not-too-dirty procedural tricks that usually lead to acceptable solutions to problems at noticeably less cost than is required by alternative methods (called algorithms) that guarantee optimal solutions. Put differently, heuristics are methods that achieve efficiency by risking failure. Kahneman and Tversky originally proposed three such heuristics: *representativeness* (the tendency to judge the probability of a sample by the degree to which it resembles the parent distribution or displays the characteristics of the generating process), *availability* (the tendency to estimate the probability of an event by the ease with which instances of the event can be remembered or constructed in imagination) and *anchoring and adjustment* (the process of generating estimates by taking a value suggested by the statement of the problem or some partial computation and by adjusting it upward or downward to account for other relevant information). Additional biases were added later, like the framing bias (the tendency of risk attitude to depend on the formulation of problems), the hindsight bias (the tendency to correct past judgements in view of present knowledge) or the overconfidence bias (the tendency of people to be overconfident with respect to the correctness of own knowledge) (for an overview see Kahneman et al., 1982; McFadden, 1999).

In their experiments, Kahneman and Tversky replaced the parametric study with the 'problem study'. As an example, take the demonstration of the use of the availability heuristic in the judgement of the probability of an uncertain event: 'Consider the letter R. Is R more likely to appear in the first position of a word or in the third position of a word?' (Tversky & Kahneman, 1973, p. 211). This problem has only two possible answers. One answer is correct if people reason in accordance with probability theory (R is more frequent in the third position), and the other if people reason heuristically by availability (R more is frequent in the first position). This 'strong inference method' (Platt, 1964) guarantees that one, and only one, of the hypotheses will be supported by the data. This method showed heuristic reasoning in many different problems. Note, however, that the point in the problems used in the heuristics-and-biases tradition is that the use of heuristics usually leads to errors, while the probability mode predicts correct answers. That is, the method is biased against the heuristic mode (Krueger, 1998). To see this, recall the letter-estimation task: of the 20 possible consonants (excluding X), 12 are more common in the first position and 8 are more common in the third position. All of the consonants that Tversky and Kahneman studied were taken from the third-position group even though there are more consonants in the first-position group (Lopes, 1991). That is, problems are not chosen arbitrarily but deliberately so that the use of the heuristic mode leads to errors, while the use of the probability mode leads to correct prediction. The upshot of all this was a strong conviction in the view that 'people-are-irrational-in-dealing-with-probabilities-and-science-has-proven-it' (see Lopes, 1991).

Rationality as a Standard for Comparison

The issue of irrational choice or judgement began to dominate in decision theory following the influential papers of Kahneman and Tversky. Unlike their ancient predecessors, who revised the rules of logic and probability when human intuition was observed to deviate from those rules (Daston, 1988), a view began to dominate of the rules of probability and logic as norms against which human reasoning has to be evaluated rather than as codifications of it: when the two diverge, there is something wrong with the reasoning, not with the norms (Chase, Hertwig, & Gigerenzer, 1998). Note a striking mismatch here: contemporary biologists who observed similar nonnormative behaviour in animals have faulted the assumptions of their normative models rather than challenge that adaptiveness of the animals' choices. For example, biologists who model the choices of foraging animals initially started with the strong assumptions of expected utility theory that animals maximize expected returns (Charnov, 1976). However, when it became clear that this model failed adequately to describe some choices, the model was changed fundamentally: the assumption of maximization of expected utility was replaced by the assumption that animals maximize the probability of reaching some goal. The foraging behaviour of some birds, for instance, maximizes the probability of surviving for a given time period, rather than maximizing the amount of food intake (Stephens, 1981). Thus, animals were found to be risk-sensitive, following a rule like: 'Be risk-averse if expected energy budget positive, be risk-prone if expected energy budget negative.'

Accepted Wisdom by the 1980s: Irrational Decision Makers

Out of five different sources of irrationality that are discussed in a summary paper by Abelson and Levi (1985), four were related to the individual decision maker: the corrigible rationalist, the bounded rationalist, the error-prone intuitive scientist and the slave to motivational factors.

The Corrigible Rationalist

The corrigible rationalist view holds that violations of rationality follow from people being insufficiently trained in the appropriate probability calculations or other logical tools. Thus, by providing the appropriate task information and training, these violations of optimality can be corrected. However, research has shown that formal training helps only in limited situations. For instance, Cheng, Holyoak, Nisbett and Oliver (1986) showed that training on the rules of propositional logic, particularly *modus tollens*, does not lead to any improvement in performance. Similarly, training and guidance in drawing a decision tree did not result in more effective decision making (Huber & Kühberger, 1996). More generally, an overview over efforts at debiasing by Fischhoff (1982) concludes that 'biases have proven moderately robust, resisting attempts to interpret them as artifacts and eliminate them by "mechanical" manipulations, such as making subjects work harder. Effective debiasing usually has involved changing the psychological nature of the task' (p. 440). In sum, the success of debiasing efforts has proven elusive and unpredictable, thus rendering the overly general corrigible rationalist view implausible.

The Bounded Rationalist

The bounded rationalist position was formulated by Simon (1955) and holds that decision makers do not necessarily optimize, but *satisfice*—that is, they choose options that are good enough given the limitations of the task and the cognitive system. The picture is of:

> . . . a creature of bounded rationality who copes with the complexity that confronts him by highly selective serial search of the environment, guided and interrupted by the demands of his motivational system, and regulated, in particular, by dynamically adjusting multidimensional levels of aspiration. (Simon, 1979, p. 4)

This conception of rationality contrasts with the classical view of rationality as maximization of expected utility that demands unlimited memory capacity and computational complexity.[2] The main shortcoming of the bounded rationalist position is that it suffers from being too glib. All departures from rationality can be dismissed as being reasonable enough.

The Error-Prone Intuitive Scientist

At the heart of the heuristics-and-biases programme is the view of irrationality as a consequence of an error-prone intuitive scientist. The decision maker is seen as trying vaguely to be rational but as often falling short if the situation gets subtle. This view is similar to the view of bounded rationality, but here the boundedness comes not from the impracticability of spending time and effort to make truly optimal decisions but from a genuine failure to appreciate normatively appropriate strategies.

This view assumes that there exist accepted criteria of valid inference and that psychological factors may bias judgements away from these criteria. Such ‘normative’ criteria of valid inference imply methods of judgement and decision that lessen the incidence of inaccuracies. However, the existence of normative models is restricted to certain domains of judgement (e.g. the revision of probabilities by Bayes’ rule) while other domains (e.g. causal attributions) are assumed to lack equally normative formulations (Krueger, 1998; Nisbett & Ross, 1980). In the absence of accepted normative criteria, inferential errors are sometimes identified by comparing intuitive judgements to generally trusted procedures such as, for instance, counting (which is what Tversky and Kahneman [1973] did in their letter R task). Doing this may be seen as some sort of ‘direct verification’ (Kruglanski & Ajzen, 1983). Another standard of the veridicality of inferences, when neither normative models nor direct verification seems available, is that investigators use their own tutored judgement. For instance, it has been argued that an actor’s behaviour should not be used to infer dispositional attributions when it is performed under strong situational constraints (Jones, 1979). A finding that the actor’s behaviour does affect attributions under these conditions is then taken as evidence of error.

These examples show that it makes doubtful sense to ask 'Are humans rational?' Rather, similar to the question 'Are humans tall?' it is incomplete without some specification of either a standard or a purpose. Thus the view of the decision maker as an error-prone intuitive scientist is fraught with problems as to what constitutes the appropriate standard for comparison between rational and biased judgement and choice.

The views of the decision maker as a corrigible rationalist, as a bounded rationalist and as an error-prone intuitive scientist can be classified as cognitive in origin, since irrationality originates in the limitations of otherwise reasonable information-processing strategies. Individuals are said to possess a considerable repertory of such sub-optimal strategies (Kahneman et al., 1982; Nisbett & Ross, 1980). Broadly speaking, these strategies are assumed to direct people's attention to some types of information and hypotheses at the expense of others. In addition to biases that have their origin in cognition, there are also biases that have a motivational origin. From a motivational point of view, biases are characterized by a tendency to hold and form beliefs that serve the individual's needs and desires.

The Slave to Motivational Forces

Motivationally based irrational tendencies are said to ensue from two broad needs: ego enhancement and control (Kruglanski & Ajzen, 1983). As a consequence of the presumed need of ego enhancement, people are eager to accept credit for success but are reluctant to accept blame for failure. This self-serving bias leads, for instance, to a 'false consensus effect', whereby people tend to perceive their own behaviour and judgements as common and appropriate, and at the same time tend to perceive alternative responses as uncommon, deviant or inappropriate (Ross, 1977). The desire to exercise control may bias individuals' attributions to controllable factors rather than to factors over which they have no control (Kelley, 1971). Alternatively, people may entertain an 'illusion of control' because they expect personal success with a higher probability than the objective probability warrants (Langer, 1975; but see Kühberger, Perner, Schulte, & Leingruber, 1995).

The view of the irrational decision maker as the slave to motivational forces comes closest to the understanding of irrationality in laypeople. It is the common-sense view that irrationality ensues when people fall into the grip of emotional forces that they cannot or are unwilling to control. Since emotion is commonly viewed as a disorganizer of ongoing deliberative behaviour, irrationality produced by emotional arousal will have a diffuse and impulsive character. However, as recent research has demonstrated, emotion is better treated as a precondition for deliberative behaviour, rather than as a disorganizer (e.g. Damasio, 1994).

In conclusion, the above summary of the accepted wisdom by the 1980s as portrayed by Abelson and Levi (1985) shows a prevailing tendency to (i) accept the ubiquity of irrational decision making, and (ii) accuse the decision maker of being the primary source of this irrationality. A growing collection of biases (insensitivity to sample size, ignoring of base rates, insufficiently regressive predictions, insufficient adjustments, misperceptions of chance, illusions of validity, illusory correlations, overconfidence, conjunction fallacy, framing bias, hindsight bias, illusion of control, etc.) bolstered a general interest in biases of decision making, and in probabilistic judgement. Biases aplenty. This message subsequently caught the attention of social psychologists, and with their contribution (see Nisbett & Ross, 1980), the field reached a very high level of visibility within psychology and beyond.

The Decision Maker Exonerated

The widely shared assumption that the irrationality resides within the decision maker soon came under heavy attack, however, due to several arguments.

A Misrepresented Criterion of Optimality

One line of critique argued that the portrayal of irrational heuristics is misunderstood as negative, but is in fact much more positive. For instance, Oaksford and Chater (1992) argued that the idea of bounded rationality implies that people *could not* use optimal strategies. People cannot apply Bayesian inference, for example, since this makes exponentially increasing demands on computational resources even for relatively simple problems. Since the brain is a limited information processor, the processes of risky decision making cannot be based on optimal, algorithmic procedures, because these are too complex.

> This means that the only rationality to which we can aspire, as individual decision-makers, is one bounded by our limited computational resources. In consequence, the observation that we do not behave in accordance with Bayes's theorem could not impugn our rationality. Our rationality could be questioned only if we were capable of using the optimal strategy but failed to do so. Thinking otherwise is akin to condemning us because we do not fly even though we do not possess wings. (Oaksford & Chater, 1992, pp. 226–227)

One could argue against this view that the heuristics-and-biases tradition frequently uses very simple tasks such that computational complexity is kept well within manageable margins—you can do Bayes' rule on every pocket calculator. Oaksford and Chater (1992) dismiss this argument on the grounds that experimental participants typically lack the appropriate schooling in statistics and thus rely on strategies that they use in the normally more complex tasks. However, in staying with the wings metaphor, this is to continue running even if one could—and better should—fly.

The Corrigible Rationalist Revisited

Another line of evidence against the view that humans are fundamentally irrational decision makers emanated from the effort model of decision making. The essence of the effort model of decision making is that more cognitive effort leads to better decisions.[3] That is, people are seen to be corrigible if provided with enough incentive. This new corrigible rationalist view puts the focus on incentives rather than on training. The intuition seems to be that participants will calculate more, think harder or somehow see the appeal of axioms when faced with important stakes. There is evidence that incentives do prolong deliberation (e.g. Paese & Sniezek, 1991; Tversky & Kahneman, 1986; Wilcox, 1993) but in many cases providing incentives does not lead to better decisions (Camerer & Hogarth, 1999; Kühberger, Schulte-Mecklenbeck, & Perner, 2000).

In sum, the evidence suggests that incentives may not be the appropriate means to infuse rationality into decision making, though they routinely infuse additional cognitive effort. That is, a relevant change in the decision maker's motivation fails to result in a relevant change in behaviour with respect to rationality. This makes it unlikely that a source within the person is mainly responsible for irrational behaviour. Staying with the cognitive illusions metaphor, we might say that you cannot make an illusion disappear by providing an incentive. Working harder is not working smarter.

Faulty Normative Models

Some researchers began to follow the biologists in maintaining that when people do not conform to a normative model, the normative model rather than the people may be at fault (e.g. Oaksford & Chater, 1996). In a review of new trends in decision making in the 1990s, Mellers, Schwartz and Cooke (1998) side with this view. They hold that, within decision theory, definitions of errors are often based on three faulty assumptions. First, it is assumed that there is a single correct response. But tasks are often not fully specified and therefore can have many correct answers. For instance, Birnbaum (1983) has shown that the famous cab problem has many correct answers, Kühberger (1995) has shown that the Asian disease problem (see below) is incompletely specified, and Gigerenzer, Hoffrage and Kleinbölting (1991) have argued that overconfidence need not necessarily be a reasoning error. Even in relatively well-defined areas like the updating of ambiguous conditional probabilities, mathematicians come up with different 'normative' models (Walley, 1991, see esp. chap. 5).

Second, rationality is expressed as internal coherence and logical consistency within a system of beliefs and preferences. However, this means that the empirical accuracy of decisions is largely missing in the rationality discussion. Similarly, the classical definition of rationality is blind to content and context. Good judgement, however, is domain-specific and should reflect the basic principles of survival and adaptation (Cosmides & Tooby, 1996; Gigerenzer, 1996).

Third, it is assumed that rationality is the same for subjects and experimenters. But the metaphor of people as intuitive statisticians may be misleading. People may be better characterized as intuitive politicians who balance pressures from competing constituencies, intuitive prosecutors who demand accountability, or intuitive theologians who protect sacred values from contamination (Tetlock, 1992).

Domain-Specific Rationality

Many researchers, most notably the group around Gigerenzer (e.g. Chase et al., 1998; Gigerenzer, Todd, & the ABC Research Group, 1999), subscribe to a view of domain-specific rationality. Gigerenzer's group exploits the analogy between inference and perception behind the cognitive illusions metaphor. If they are illusions, violations of rationality should be the exception rather than the rule. Just as vision researchers construct situations in which our perceptions lead to incorrect inferences about the world (e.g. about the length of the lines in the Müller–Lyer illusion), and just as motor researchers generate visual flow patterns that lead to motor illusions (e.g. physically impossible step frequencies in the walking-in-the-drum paradigm; Lackner & DiZio, 2000), researchers in the heuristics-and-biases programme select problems in which reasoning by cognitive heuristics leads to violations of rationality. However, what is different between vision and motor researchers, on the one side, and decision researchers, on the other side, is the conclusions drawn from such demonstrations. Vision scientists do not conclude from the robustness of the Müller–Lyer illusion that people are generally poor at inferring object lengths, and motor researchers do not conclude from the demonstration of stepping illusions that people do walk improperly. However, many advocates of the heuristics-and-biases programme conclude from the cognitive illusions found in laboratory tasks that human judgement is severely handicapped.

In contrast, 'fast and frugal heuristics' (Gigerenzer et al., 1999) can perform about as well as algorithms that require much more information. Their secret is their 'ecological rationality': that is, they exploit environmental regularities to make smart inferences. For instance, the recognition heuristic exploits the fact that our ignorance is often systematically related to variables that we want to infer (e.g. we may be more likely to recognize big cities, companies and universities than small ones). Another aspect of fast and frugal heuristics is that they are not indifferent to how the information is presented. For probability theory it does not matter whether it deals with probabilities, percentages or frequencies, but ecologically rational heuristics have evolved to represent information in the format of natural frequencies: that is, absolute frequencies that have not been normalized with respect to base rates (Kleiter, 1994).

Researchers endorsing the view of rationality as domain-specific argue for bounded rationality in an effective, but predictable, way. They reject any internalist criteria for rationality, primarily consistency and coherence, but endorse an external alternative: it is 'rationality that is defined by its fit with reality' (Gigerenzer et al., 1999, p. 5). Cognitive processes are rational in this sense if they yield accurate results. The proposal is that human cognitive processes accurately represent the world by means of fast and frugal heuristics that work well in their natural environments. These fast and frugal heuristics are domain-specific; the mind contains no domain-general learning or logical process, but it only consists of domain-specific modules that are designed by natural selection to solve particular adaptive problems. The metaphor is of the mind being like a Swiss army knife, with many blades dedicated to solving special problems, but with no general-purpose blade. This view of domain-specific rationality as reflected by fast and frugal heuristics holds that the standards of rationality should be relativized to our needs as situated, finite enquirers after truth. Thus, the rationality of processes has to be assessed in relation to reliability *and* fecundity (truth and speed).

Irrationality: A Misunderstanding?

The rationality issue in risky decision making up to here was pinpointed as a discussion with various 'losers'. The 1960s had no loser, since the view prevailed of the decision maker as being fairly rational in his or her choices. The 1970s and 1980s portrayed the individual as the loser: the decision maker, due to his or her mental make-up, was seen as the easy prey of a wealth of cognitive illusions. The end of the 1980s and the 1990s witnessed a movement to blame the models that were used as yardsticks to judge rational behaviour.

From a somewhat different perspective, however, one might also conceive of the rationality debate, especially of the heuristics-and-biases stance and of the adaptive rationality stance, as similar positions differing only in focus. The message of the heuristics-and-biases tradition can be paraphrased as: 'The world is too complicated, thus the mind simplifies it—this leads frequently to error that could be avoided by using complex normative rules.' Alternatively, the adaptive rationality position can be paraphrased as: 'The world is too complicated, thus the mind simplifies it—this frequently leads to correct solutions at a much lower cost than using complex normative rules.' Thus, both views agree that decision making and judgement are heuristic, but they differ in their evaluation of the consequences of the use of heuristics: the heuristics-and-biases position emphasizes the contras, while the adaptive rationality position emphasizes the pros.

Research beginning in the 1990s pinpoints a new movement: the acknowledgement that mental processes are governed by higher-level assumptions about the social context of the information to be processed. There is more to the simple laboratory tasks than experimenters might think. This time the loser is the experimenter him- or herself who does not fully know what s/he communicates.

A good example is given by Hilton (1995, p. 248) of how the word 'family' can be differently interpreted and thus can lead to seemingly inconsistent judgements expressed in a conversational exchange:

> *Question:* How is your family?
> *Answer:* Fairly well, thank you.

A married man might reply this way if he considers that his wife has recently suffered the loss of a close friend but that his children are in good form. The interpretation of 'family' is 'wife and kids'. Now consider this exchange:

> *Question:* How is your wife?
> *Answer:* Not too good, I'm afraid.
> *Question:* And how is your family?
> *Answer:* Extremely well, thank you.

Here different responses are given to the same question, something that usually goes to show irrationality. However, this is premature. Specifically, in the case where the hearer has just been told about the health of the speaker's wife, interlocutors are entitled to assume that the reference to 'family' meant 'just the kids', rather than the 'wife and kids'. This interpretation of the word 'family' is conversationally rational, since norms of rational communication (e.g. Grice, 1975) require a speaker to be cooperative with a hearer.

> Thus, survey researchers who attribute inconsistent patterns of response to questions to cognitive shortcomings, or experimental psychologists who explain patterns of judgment in terms of purely intrapersonal variables—such as memory capacity, attention factors, memory activation levels, search strategies, and judgmental heuristics—may be in danger of . . . misattributing patterns of inferential behavior to features of the person and overlooking how it is constrained by its interpersonal context. (Hilton, 1995, p. 249)

Experiments as Social Interaction Situations

Hilton (1995) shows that in experimental tasks the social context can influence the response in two ways. The first way comprises the interpretation of the task, and means that participants choose the most rational interpretation using the criterion of consistency with higher-order assumptions about conversation and knowledge about the discourse context. The second way involves applying a normative model of reasoning to the representation formed.

Most research in judgement and decision making has focused on the second way. However, there are examples of how the general assumption of cooperativeness (one of Grice's axioms) can make information relevant that is normally considered incidental to the experimental task. For instance, one of the most fundamental assumptions that people make in a conversation is that utterances are intentionally produced by the speaker. Take the influential engineers-and-lawyers problem introduced by Kahneman and Tversky (1973), where the authors found that participants were more likely to rely on individuating information about the target than on base-rate information. Krosnick, Li and Lehman (1990) showed that the use of base-rate information can be influenced by various pragmatic factors, for example the order of presentation. Normally, individuating information is presented first in the engineer-and-lawyer problem and base-rate information is presented later. If participants use the order of presentation as a cue to significance, base rates will be used more if they are presented first. Krosnick et al. found exactly this. In addition, they showed that the order cue is invalidated when participants know that the order has not been produced intentionally by the experimenter.

Another example of the importance of implicit assumptions that are conveyed by the cover story of problems is the demonstration of the conjunction fallacy (Tversky & Kahneman, 1983). The problem involves the description of a woman who appears to be a social activist. Participants are then asked to rate or rank the probability that the woman is a bank teller as opposed to a bank teller who is active in the feminist movement:

> *Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.*
>
> *Which of the following is more probable?*
>
> *Linda is a bank teller (t)*
>
> *Linda is a bank teller and is active in the feminist movement (t&f)*

Usually, about 80 per cent of the participants rate the conjunction of constituents (*t&f*) as more probable than the single constituent (*t*). Tversky and Kahneman argued that this error is due to the use of the representativeness heuristic, because *t&f* resembles the personality description of the woman better than *t* alone. Thus, *t&f* is judged more likely. Dulany and Hilton (1991) challenged the characterization of this conjunction effect as a fallacy. They argued for the influence of conversational norms since they found that participants who were most likely to commit the conjunction fallacy were also most likely to draw negation implications ('Linda is a bank teller' is taken to mean 'Linda is a bank teller *and* is *not active* in the feminist movement' [*t&~f*]). As Tetlock (1992) states it, experimental participants:

> . . . assume that the experimenter is trying to make the most informative statement (thus observing the Grician maxim of quantity), would only make statements relevant to the topic under discussion (thus observing the Grician maxim of relevance), would not say anything that he or she knows to be false (thus observing the maxim of truthfulness), and would refrain from making statements for which he or she lacks adequate evidence (thus observing the maxim of quality). (p. 366)

Another illuminating example comes from Kühberger (1995) with the framing phenomenon (Tversky & Kahneman, 1981):

> Problem 1:
>
> *Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programmes to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programmes are as follows:*
>
> *If Program A is adopted, 200 people will be saved.*
>
> *If Program B is adopted, there is 1/3 probability that 600 people will be saved and 2/3 probability that no people will be saved.*

Now consider this problem with a slightly different verbal description of the outcomes:

> Problem 2:
>
> *If Programme C is adopted, 400 people will die.*
>
> *If Programme D is adopted, there is 1/3 probability that nobody will die and 2/3 probability that 600 people will die.*

In the positively framed version of the Asian disease problem, a clear majority of respondents preferred saving 200 lives for sure (72 per cent), over the option that offered a 1/3 chance of saving 600 lives (28 per cent). In the negatively framed version, however, most people preferred the 1/3 chance of losing no lives (78 per cent) to the sure loss of 200 lives (22 per cent). From a formal point of view, options A and B in problem 1 are indistinguishable from options C and D in problem 2; all four options yield either 200 lives for sure or an expected value of 200 lives for the risky options. Thus, there should not be any systematic preference. The finding is different: there seems to be a general tendency of risk aversion for positively framed problems and a general tendency of risk seeking for negatively framed problems. This framing effect is fairly robust (Kühberger, 1998).

One of the most important criticisms against the standard rational choice model has emerged from the framing effect. After all, choosing differently in identical situations is clearly irrational. However, from a conversational norm perspective, doing so in framing experiments is not necessarily irrational. Closer inspection of the Asian disease problem shows that programme A in problem 1 and programme C in problem 2 are not completely described. Programme A states that 200 people will be saved; a more complete description should then specify that 400 people will not be saved, since a total of 600 lives are at stake. Similarly, programme C in the negative framing condition specifies that 400 people will die and ought, in order to be complete, specify that 200 people will not die. If people do not assume this to be the case, the two problems are not equivalent and therefore different choices are to be expected.

In sum, experimental participants may perceive problems differently from what experimenters expect.

> As we amply illustrate, the context is not given but chosen. Moreover, humans are not in the business of simply assessing the relevance of new information. They try to process information as relevantly as possible; that is, they try to obtain from each new item of information as great a contextual effect as possible for as small as possible a processing effort. For this, they choose a context which will maximize relevance. (Sperber & Wilson, 1987, p. 703)

Experiments as Machineries that Generate False Collective Judgements

One important observation about cognitive biases is that participants do not only arrive at some wrong answer, but they frequently arrive at the *same* wrong answer. Cognitive biases can thus be regarded as machineries generating false collective judgements. It is plausible to interpret them as biases that have a cultural or biological origin leading ordinary inference along false paths (this is the usual interpretation). Alternatively, they can be interpreted as indicating that respondents give wrong answers because they have strong reasons to do so (Boudon, 1996). Hilton’s (1995) analysis deals mainly with conversational implicatures that follow from the use of the Grician maxims of cooperativeness, quality, quantity, relation and manner. However, there are examples that are suggestive of another instance of going beyond the information given. For instance, Bless, Betsch and Franzen (1998) have shown that framing effects tend to disappear if the task is embedded in a statistical context, rather than in a disease context. This type of framing via the context can exert a powerful influence on participants’ construal of what the problem is about. Even if participants know the

TABLE 1. A negative correlation perceived as positive

	Disease		
Symptom	Ill	Not Ill	Total
Yes	37	33	70
No	17	13	30
Total	54	46	100

statistically correct way to deal with a problem, they may prefer to do otherwise, because they feel that a statistical treatment is inadequate.

Another example for going beyond the information given from cognitive psychology is a study by Shweder (1977), who gave a population of nurses 100 cards, each representing a hypothetical patient. Two fictitious pieces of information were presented on each of these cards: whether or not the patient has a given symptom, and whether or not the patient has a given disease. The nurses were then asked whether the symptom indicates the disease. Table 1 gives the summary information of the study. (This table was not presented to the nurses.)

The result was that the nurses perceived the symptom to be correlated with the disease. To arrive at this conclusion, the nurses used, presumably, only the proportion of cases where the patient has contracted the disease and displays the symptom. These cases are relatively frequent (37), thus leading to a perceived positive correlation between symptom and disease. To verify such a relationship, normatively, one has to compare the proportion of patients who have the disease and do show the symptom (37/70 = .53) with the proportion of patients who have the disease and do not show the symptom (17/30 = .57). Since the former value is a little lower than the latter, the modal answer of the nurses is wrong.

However, the nurses may have good reasons for their judgement. The argument (Boudon, 1996) is as follows: It is true, in principle, that to derive a causal statement from a contingency table one has to use all four pieces of information. Practically, however, one piece may be sufficient, notably when there is some implicit knowledge about the order of magnitude of the other pieces of information. For instance, one may know that the frequencies are asymmetrically distributed, as is usually the case for diseases and symptoms: to be ill and to show symptoms is less frequent than not to be ill and not to show symptoms, since pathologies are less frequent than normal phenomena. Thus the nurses may have implicitly assumed that the frequency of the disease should be low, as is the frequency of most diseases they are confronted with. Similarly for the frequency of the symptom. This makes the whole situation unrealistic. In other words, the data are implausible, given

TABLE 2. A more realistic distribution of frequencies (after Boudon, 1996)

	Disease		
Symptom	Ill	Not Ill	Total
Yes	4	16	20
No	16	64	80
Total	20	80	100

their supposed meaning. Table 2 presents data as they appear plausible given the marginals by Boudon (1996).

With realistically low frequencies of symptom and disease, a finding of 37 out of 100 people having both the symptom and the disease is a serious indicator of the existence of a causal relationship between the two variables.

This example brings us back to fast and frugal heuristics that work well in their natural environments. It is just another instance showing that otherwise reliable statistical intuition, if applied to an artificial situation, can be incorrectly evaluated as deficient. In sum, acknowledging the inductive nature of conversational inference, and acknowledging that implicit meta-cognitive assumptions may be brought to bear, suggests that many of the experimental results that have been attributed to faulty reasoning may be reinterpreted as being due to rational interpretations of experimenter-given information.

Synopsis

A clarification of the meaning of rationality that is in concord with the development of the issue in decision theory has been brought forward by Fillieule (1996). It is the distinction between *rationality of frame* and *rationality of inference*. The study of the rationality of frame deals with the choice between contextual frames, while the study of the rationality of inference deals with the choice between rules of inference. Take again the framing effect. We depart from the general assumption that participants may have good reasons for their preferences, given the information that is communicated to them. Since this information conveys different aspects of reality, it breeds different preferences. The rationality of inference is taken into account here, since each participant presumably chooses the best option given his or her preferences. If we admit that the economic model of rational choice is strictly limited to the rationality of inference, then this model is in no way refuted by the experiment. The anomaly originates in the choice of the conceptual frame. That is, if we want to explain choice anomalies, we must transcend the rationality of inference and state the problem in terms of the rationality of frame. Tversky and Kahneman do not state the problem in these terms, maybe because they are under the influence of the economic model, which does not aim at explaining conceptual frames.

The issue thus is to clarify the rationality of frames before talking about the rationality of inferences. Stanovich and West (2000) distinguish between four interpretations that preserve the assumption of rationality in human cognition. These explanations posit that the so-called 'normative–descriptive gap' is due to (i) performance errors, (ii) computational limitations, (iii) the wrong norm being applied by the experimenter, and (iv) a different task construal by the subject. In an impressive research enterprise (Stanovich & West, 1998, 1999) they measure the correlation between various tasks that have exhibited bias and SAT as a measure of cognitive capacity. They report significant cross-task correlations between tasks and conclude that this renders the performance error view unlikely while being (weak) evidence for the computational limitations view. That is, to a moderate extent, discrepancies between actual performance and normative models can be accounted for by individual differences in cognitive capacity.

With respect to norm application and task construal, Stanovich and West find that cognitive ability differences are strong in cases where there is a dispute about the proper construal of the task, while differences are markedly attenuated in cases with little controversy about alternative task construal. That is, participants high on SAT tend (i) to construe tasks similar to experimenters, and (ii) to stay within a frame once construed. Stanovich and West champion an interpretation of this pattern in terms of two-system theories of reasoning (e.g. Epstein, 1994; Sloman, 1996). System 1 reasoning is characterized as automatic, largely unconscious and relatively undemanding of computational capacity. It entails properties of heuristic processing. System 2 reasoning conjoins characteristics that are seen as typifying controlled processing (e.g. rule-based, demanding of cognitive capacity, relatively slow). With respect to rationality, the most important difference between the two systems is that they tend to lead to different task construals. Construals triggered by System 1 are highly contextual, personalized and socialized. They are mainly driven by considerations of relevance and aim at inferring intentionality by the use of conversational implicatures even in situations that are devoid of conversational features. Consequently, we see a tendency of automatic contextualization of problems. On the other hand, System 2 triggers decontextual and depersonalized task construals, more akin to the 'deep structure' of problems.

Stanovich and West's research represents progress toward understanding irrationality in decision making by providing cues to predict different task construals. First, it predicts that problems will only appear if System 1 and System 2 lead to different task construals. Second, owing to the automaticity of System 1 processes, task construal will always be influenced by context and conversational principles. What is more likely to be possible for individuals of higher cognitive ability, however, is that the interactional intelligence of System 1 will be overridden by the analytic intelligence of System 2. In most cases, this leads to a task construal that is shared with the experimenter.

In a broader scheme of things the distinction between systems of thinking can be related to recent research on consciousness in cognitive neuroscience. Dehaene and Naccache (2001) report three fundamental empirical findings on consciousness: (i) that cognitive processing is possible without consciousness, (ii) that attention is a prerequisite of consciousness, and (iii) that consciousness is required for specific mental operations. System 1 thinking has very much in common with finding (i): current evidence suggests that there are many such processes and that they can occur not only on a perceptual, but also on a semantic level. System 2 processing is related to findings (ii) and (iii) in that attention is central to these processes. In addition, Dehaene and Naccache identify a structural criterion, active representation, which seems to be a precondition for consciousness and maybe also for System 2 thinking. Active representation means that the information must be represented in an active manner in the firing of one or several neuronal assemblies. This excludes the wealth of information that is present in the nervous system only in a latent form as patterns of anatomical connections, or in the form of strengthened memory traces. In sum, the distinction between systems of thinking gains relevance through recent research in cognitive neuroscience.

All this should not be read as indicating that System 1 is inferior to System 2, but rather that our systems of recruiting prior knowledge and contextual information to solve problems—even problems with formal solutions—are receptive to different types of information. However, in the case of decision making and reasoning, some people show the cognitive flexibility to decouple unneeded systems of knowledge and some people lack it. The distinction between systems of thinking also addresses the adaptive view of rationality. Stanovich and West (2000) argue explicitly that the features of System 1 are designed to track increases in the reproduction probability of the genes very closely, while System 2 works primarily as a control system focused on the interests of the whole person. In cases where maximizing personal utility will sacrifice genetic fitness, System 2 will seek to fulfil the individual's goals. That is, ecological rationality in the sense of Gigerenzer has to be discussed in relation to System 2 processing.

In conclusion, we have to diagnose a shortcoming of decision theory to adequately acknowledge the problem of rationality of frame and an overemphasis on rationality of inference. The rationality of inference can only be evaluated with respect to a given task construal. The topic of different task construal may have been overlooked for so long because it is strongly influenced by unconscious System 1 thinking. Unconscious thinking by its very nature tends to go unnoticed. Research in decision theory should now put more effort into research on the construal of tasks. We know very little about the features of tasks that lead to a discrepancy in task construals due to System 1 and System 2 processing. For instance, as Hilton (1995) notes, it is not clear why incentives should cause respondents to abandon an interpretation that is pragmatically correct. However, from the perspective of System 1 and System 2 processing, it may well be that the question of actual motivation is important. One prediction is that highly motivated individuals tend to do more System 2 processing. In addition, whether people activate System 2 processing (System 1 is activated automatically) may depend on what people consider to be appropriate processing for the task at hand. For instance, in face-to-face interaction, contextual thinking will possibly be considered more appropriate than decontextualized thinking.

As Overton and Newman (1982) argue, two distinct components are required for a complete psychological theory. One is a competence component that is an idealized model of the abstract knowledge possessed by an individual in a given domain—decision theory has much to say about this. The other, the activation-utilization component, encompasses the psychological procedures and situational factors that determine the manifestation of the competence—here decision research has much left to learn. My take-home message is this: the time is overdue for a change of what the rationality issue is generally taken to show. The negative message of fundamentally flawed decision makers must be replaced by a more positive picture of a decision maker who reacts to task and context and to the experiment as a social situation. The bet on bias is risky and is likely to be lost.

Notes

1. Two of these papers were ranked among the top 100 most influential works in cognitive science from the 20th century, as selected by a panel of esteemed judges from the Center for Cognitive Sciences, University of Minnesota. The ranking lists Kahneman and Tversky (1973) on 33, and Tversky and Kahneman (1974) on 68 (see http://cogsci.umn.edu/millennium/top100.html).
2. Bounded rationality also contrasts with the concept of quasi-rationality (Hammond, 1996). Quasi-rationality distinguishes between intuitive and analytical mechanisms of information processing, and does not always predict that analysis will outperform other modes of thought.
3. Better decisions are to be understood in the sense that decisions conform better to a normative model mostly by reducing subject-specific error variance (Smith & Walker, 1993).

References

Abelson, R.P., & Levi, A. (1985). Decision making and decision theory. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology* (3rd ed.; pp. 213–309). New York: Random House.

Berkeley, D., & Humphreys, P. (1982). Structuring decision problems and the ‘bias heuristic’. *Acta Psychologica*, *50*, 201–252.

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

Bless, H., Betsch, T., & Franzen, A. (1998). Framing the framing effect: The impact of context cues on solutions to the Asian disease problem. *European Journal of Social Psychology*, *28*, 287–291.

Boudon, R. (1996). The ‘cognitivist model’: A generalized ‘rational-choice model’. *Rationality and Society*, *8*, 123–150.

Camerer, C.F., & Hogarth, R.M. (1999). The effects of financial incentives in experiments: A review and capital–labor-production framework. *Journal of Risk and Uncertainty*, *19*, 7–42.

Charnov, E. (1976). Optimal foraging: The marginal value theorem. *Theoretical Population Biology*, *9*, 129–136.

Chase, V.M., Hertwig, R., & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences*, *2*, 206–214.

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*, 57–65.

Cheng, P.W., Holyoak, K.J., Nisbett, R.E., & Oliver, L.M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology*, *17*, 391–416.

Christensen-Szalanki, J.J., & Beach, L.R. (1984). The citation bias: Fad and fashion in the judgement and decision literature. *American Psychologist*, *39*, 75–78.

Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition*, *58*, 1–73.

Damasio, A.R. (1994). *Descartes’ error: Emotion, reason, and the human brain.* New York: Putnam.

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

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

Dennett, D. (1987). *The intentional stance.* Cambridge, MA: MIT Press.

Dulany, D.E., & Hilton, D.J. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition*, *9*, 85–110.

Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist*, *49*, 709–724.

Fillieule, R. (1996). Frames, inferences, and rationality. *Rationality and Society*, *8*, 151–165.

Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement under uncertainty: Heuristics and biases* (pp. 422–444). Cambridge: Cambridge University Press.

Gigerenzer, G. (1996). Rationality: Why social context matters. In P. Baltes & U.M. Staudinger (Eds.), *Interactive minds: Life-span perspectives on the social foundation of cognition* (pp. 319–346). Cambridge: Cambridge University Press.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, *98*, 506–528.

Gigerenzer, G., Todd, P.M., & the ABC Research Group (1999). *Simple heuristics that make us smart.* Oxford: Oxford University Press.

Grice, H.P. (1975). Logic and conversation. In P. Cole & J.L. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41–58). San Diego, CA: Academic Press.

Hammond, K.R. (1996). *Human judgement and social policy*. New York: Oxford University Press.

Harrison, G.W. (1994). Expected utility theory and the experimentalists. *Empirical Economics*, *19*, 223–253.

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists. *Behavioral and Brain Sciences, 24*, 383–451.

Hilton, D.J. (1995). The social context of reasoning: Conversational inference and rational judgement. *Psychological Bulletin*, *118*, 248–271.

Huber, O., & Kühberger, A. (1996). Decision processes and decision trees in gambles and more natural decision tasks. *Journal of Psychology*, *130*, 329–339.

Jones, E.E. (1979). The rocky roads from acts to dispositions. *American Psychologist*, *34*, 107–117.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgement under uncertainty: Heuristics and biases.* Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgement of representativeness. *Cognitive Psychology*, *3*, 430–454.

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

Kelley, H.H. (1971). *Attribution in social interaction.* New York: General Learning Press.

Kleiter, G. D. (1994). Natural sampling: Rationality without base rates. In G. Fischer & D. Laming (Eds.), *Contributions to mathematical psychology, psychometrics, and methodology* (pp. 375–388). New York: Springer.

Krosnick, J.A., Li, F., & Lehman, D.R. (1990). Conversational conventions, order of information acquisition, and the effects of base-rates and individuating information on social judgements. *Journal of Personality and Social Psychology*, *59*, 1140–1152.

Krueger, J. (1998). The bet on bias: A foregone conclusion? *Psycoloquy, 9.46* (ftp://ftp.princeton.edu/pub/harnad/Psycoloquy/1998.volume.9/psyc.98.9.46.social-bias.1.krueger).

Kruglanski, A.W., & Ajzen, I. (1983). Bias and error in human judgement. *European Journal of Social Psychology*, *13*, 1–44.

Kühberger, A. (1995). The framing of decisions: A new look at old problems. *Organizational Behavior and Human Decision Processes*, *62*, 230–240.

Kühberger, A. (1998). The influence of framing on risky decisions: A meta-analysis. *Organizational Behavior and Human Decision Processes, 75*, 23–55.

Kühberger, A., Perner, J., Schulte, M., & Leingruber, R. (1995). Choice or no choice: Is the Langer effect evidence against simulation? *Mind & Language*, *10*, 423–436.

Kühberger, A., Schulte-Mecklenbeck, M., & Perner, J. (in press). Framing decisions: Real and hypothetical. *Organizational Behavior and Human Decision Making*.

Lackner, J.R., & DiZio, P.A. (2000). Aspects of body self-calibration. *Trends in Cognitive Sciences*, *4*, 279–288.

Langer, E.J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, *32*, 311–328.

Lopes, L.L. (1991). The rhetoric of irrationality. *Theory & Psychology*, *1*, 65–82.

McCormich, J. (1987, August 17). The wisdom of Solomon. *Newsweek*, pp. 62–63.

McFadden, D. (1999). Rationality for economists? *Journal of Risk and Uncertainty*, *19*, 73–105.

McKean, K. (1985, June). Decisions, decisions. *Discover*, pp. 22–31.

Mellers, B.A., Schwartz, A., & Cooke, A.D.J. (1998). Judgement and decision making. *Annual Review of Psychology*, *49*, 447–477.

Nisbett, R.E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgement.* Englewood Cliffs, NJ: Prentice Hall.

Oaksford, M., & Chater, N. (1992). Bounded rationality in taking risks and drawing inferences. *Theory & Psychology*, *2*, 225–230.

Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, *103*, 381–391.

Overton, W.F., & Newman, J.L. (1982). Cognitive development: A competence-activation/utilization approach. In T. Field, A. Huston, H. Quay, L. Troll, & G. Finley (Eds.), *Review of human development* (pp. 217–241). New York: Wiley.

Paese, P.W., & Sniezek, J.A. (1991). Influences on the appropriateness of confidence in judgement: Practice, effort, information, and decision-making. *Organizational Behavior and Human Decision Processes*, *48*, 100–130.

Peterson, C.R., & Beach, L.R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, *68*, 29–46.

Platt, J.R. (1964). Strong inference. *Science*, *164*, 347–353.

Robins, R.W., & Craik, K.H. (1993). Is there a citation bias in the judgement and decision literature? *Organizational Behavior and Human Decision Processes*, *54*, 225–244.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in Experimental Social Psychology*, *10*, 173–221.

Savage, L.J. (1954). *The foundations of statistics.* New York: Wiley.

Shweder, R.A. (1977). Likeliness and likelihood in everyday thought: Magical thinking in judgements about personality. *Current Anthropology*, *18*, 637–659.

Simon, H.A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Simon, H.A. (Ed.). (1979). *Models of thought.* London: Yale University Press.

Sloman, S. (1996). The empirical case for two systems of reasoning. *Psychological Review*, *119*, 3–22.

Smith, V.L., & Walker, J.M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, *31*, 245–261.

Sperber, D., & Wilson, D. (1987). Précis of relevance: Communication and cognition. *Behavioral and Brain Sciences*, *10*, 697–754.

Stanovich, K.E., & West, R.F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, *127*, 161–188.

Stanovich, K.E., & West, R.F. (1999). Discrepancies between normative and descriptive models of decision making and the understanding/acceptance principle. *Cognitive Psychology*, *38*, 349–385.

Stanovich, K.E., & West, R.F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–726.

Stephens, D.W. (1981). The logic of risk-sensitive foraging preferences. *Animal Behavior*, *29*, 628–629.

Tetlock, P.E. (1992). The impact of accountability on judgement and choice: Toward a social contingency model. *Advances in Experimental Social Psychology*, *25*, 331–376.

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

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*, 207–232.

Tversky, A., & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgement. *Psychological Review*, *90*, 293–315.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, *59*, 251–278.

von Neuman, J., & Morgenstern, O. (1947). *Theory of games and economic behavior.* Princeton, NJ: Princeton University Press.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. New York: Chapman & Hall.

Wilcox, N. (1993). Lottery choice: Incentives, complexity and decision time. *Economic Journal*, *103*, 1397–1417.

ANTON KÜHBERGER is Associate Professor of Psychology at the University of Salzburg. His current research includes judgement and decision making, particularly the role of risk for judgement and choice; and the role of theory and simulation in the reading of minds. ADDRESS: Department of Psychology, University of Salzburg, Hellbrunner str. 34, A-5020 Salzburg, Austria. [email: anton.kuehberger@sbg.ac.at; http://www.sbg.ac.at/psy/people/kuehberger.htm]