

during interest in, and support for, science is consistently grounded in utilitarian concerns; and although by no means the only way of arguing for the importance and relevance of studying social behavior and cognition, the field's focus on the negative and the surprising has been an important tool in promoting social psychology's professional and public support.

Social science is inescapably linked to persuasion (cf., Friedrich & Douglass 1998) – to convincing various audiences that certain questions or theories are worth pursuing, that particular interpretations are superior to others, and that developments in the field should matter to the public and to potential sources of funding. Other sciences face similar constraints but benefit from: (a) more visible track records of utilitarian benefits, and (b) an absence of lay theories the public assumes are as good as the theories researchers have to offer. Biases and misbehaviors have provided vivid ways to communicate to multiple audiences why the field matters at all, and why it provides a necessary counterpoint to lay explanation and folk psychology. True – important topics have been left unstudied as a result – but part of the solution seems to lie in our simply doing a better job of articulating for fellow scientists and the public the benefits of pursuing those topics.

The authors note that the significance or consequences of purported errors and biases may be exaggerated both by the experimental designs used and by the manner of presentation. But in many cases, the true importance remains largely unknown. The core concerns for the most prestigious journals in the field are (appropriately) in the area of theory development, and yet that means that papers providing real-world demonstrations or applications without breaking new theoretical ground are typically diverted to less prestigious “applied” journals. One should hardly be surprised at a relative dearth of studies demonstrating real-world magnitude and importance given the structure of professional rewards in the field.

The second, and more important, virtue of the focus the authors critique is the role it can play in stimulating theory development. The focus on bias and misbehavior in the research literature has clearly drawn us to anomalies and, in doing so, has called into question basic assumptions about human behavior and cognition. Consistent with the central role of falsification in theory testing, attention to “negative” anomalies often provides the counter-examples that fuel theory development.

Perhaps the most salient external validations of this notion are the recent Nobel Prizes for basic research relevant to “behavioral economics” (Dubner 2003), but there are numerous other examples of integrative theories of the sort that the authors advocate. The empirical literature is peppered with demonstrations of boundary conditions and moderator variables that qualify various claims, but it often takes the development of a rather large and diverse empirical literature before powerful syntheses emerge. Petty and Cacioppo's Elaboration Likelihood Model of persuasion (Petty & Wegener 1999) is an example of a theory that emerged at least in part as an effort to reconcile widely disparate findings, some of which showed that people were inconsistent in how thoughtfully or “centrally” they processed messages. Yet another example comes from work integrating a broad and often conflicting literature on the so-called confirmation bias (Friedrich 1993; Trope & Liberman, 1996) – work that has been invoked to counter “negative” views of human behavior in other disciplines (e.g., Mele 2001). Space prohibits an exhaustive review of the wide range of integrative efforts, but it is noteworthy that these two examples and others are broad in scope and seek to explain both strengths and weakness in human cognition within a parsimonious set of principles.

A continuing, serious problem to which the authors appropriately direct attention is the tendency to construct narrow, post hoc theories to account for specific effects. As Seidenberg (1993) has argued, a vital but often ignored component of a good theoretical account is that it is founded on “independently motivated principles.” That is, the principles used to account for a phenomenon should emerge from a broader understanding of human behavior

and its constraints and be justified outside of the effect to be explained. Evolutionary considerations (e.g., Friedrich 1993; Gigerenzer et al. 1999; Kenrick et al. 2002) and connectionist architectures (e.g., Smith 1996) are noteworthy sources of powerful, independently motivated principles that have been used in providing comprehensive accounts of strengths and weaknesses within a single framework.

The rewards of frequent empirical publication and narrowly focused methodological precision – not to mention the far more limited professional outlets for integrative theoretical work – might well play as large a role in the slow and skewed progress of theory development as researchers' fascination with the anomalous, the negative, and the statistically significant. Nevertheless, K&F's article is an important part of the self-correcting nature of our science. In keeping with much of the literature that the authors take to task, their work highlights a significant concern for the field with a potentially large but undetermined effect size. The “bias” bias itself reflects certain adaptive properties, and the challenge is to capitalize on its inherent strengths while limiting the harms that come from misapplication and overuse of certain strategies.

The irrationality paradox

Gerd Gigerenzer

Max Planck Institute for Human Development, 14195 Berlin, Germany.
gigerenzer@mpib-berlin.mpg.de

Abstract: In the study of judgmental errors, surprisingly little thought is spent on what constitutes good and bad judgment. I call this simultaneous focus on errors and lack of analysis of what constitutes an error, the *irrationality paradox*. I illustrate the paradox by a dozen apparent fallacies; each can be logically deduced from the environmental structure and an unbiased mind.

The objective of demonstrating that people systematically underestimate or overestimate a quantity has become a paradigm in social psychology. Researchers tell us that we overestimate small risks and underestimate large risks, that our average confidence in our knowledge is larger than the actual proportion correct, and that we overestimate the long-term impact of emotional events, such as losing a child. This paradigm is part of a broader movement that emphasizes human irrationality and leads to a paternalistic attitude towards citizens, such as in behavioral economics and in behavioral law and economics (e.g., Sunstein 2000). I would not object to paternalism if the norms were well reasoned and argued. Yet, in the study of judgmental errors, surprisingly little thought is spent on the question of what actually constitutes good and bad judgment (Gigerenzer 1996b; 2000). Rather, researchers tend to take normative claims about irrationality at face value or accept these by authority, not by an analysis of the problem. I call this simultaneous focus on errors and lack of analysis of what constitutes an error, the *irrationality paradox*.

This commentary is about the missing study of ecological rationality in social psychology, an issue that I believe is sympathetic to Krueger & Funder's (K&F's) important article, but which they did not put into the foreground. The basic tenet of ecological rationality is that the rationality or irrationality of a judgment can only be decided by an analysis of the structure of the environment or the experimental task. Herbert Simon (1990) expressed this tenet once through the analogy of a pair of scissors: “Human rational behavior is shaped by a scissors whose blades are the structure of task environments and the computational capabilities of the actor” (Simon 1990, p. 7). By looking only at one blade, one cannot understand how minds work, just as one then cannot understand how scissors cut.

Environmental structures include statistical structures, such as the signal-to-noise ratio, the shape of distributions, and the size of samples, as well as social structures, such as the presence of com-

Table 1 (Gigerenzer). *Twelve examples of phenomena that were first interpreted as “cognitive illusions” (left), but later revalued as reasonable judgments given the environmental structure (right)*

Is a phenomenon due to a biased mind or to an environmental structure plus an unbiased mind?
Overconfidence bias (defined as miscalibration)	“Miscalibration” can be deduced from a unbiased mind in an environment with substantial unsystematic error, causing regression toward the mean (Dawes & Mulford 1996; Erev et al. 1994)
Overconfidence bias (defined as mean confidence minus proportion correct)	“Overconfidence bias” can be deduced from an unbiased mind in an environment with unrepresentative sampling of questions; this disappears largely with random sampling (Juslin et al. 2000)
Hard–easy effect	“Hard–easy effect” can be deduced from an unbiased mind in an environment with unsystematic error, causing regression toward the mean (Juslin et al. 2000)
Overestimation of low risks and underestimation of high risks	This classical phenomenon can be deduced from an unbiased mind in an environment with unsystematic error, causing regression toward the mean (Gigerenzer & Fiedler 2003)
Contingency illusion (based on prior beliefs or prejudices)	“Contingency illusion” can be deduced from an unbiased mind performing significance tests on samples with unequal sizes, such as minorities and majorities (Fiedler et al. 1999)
Most drivers say they drive more safely than average	The distribution of the actual numbers of accidents is highly skewed, which results in the fact that most drivers (in one study, 80%) have less accidents than the average number of accidents (Gigerenzer 2002; Lopes 1992)
Availability bias (letter “R” study)	“Availability bias” largely disappears when the stimuli (letters) are representatively sampled rather than selected (Sedlmeier et al. 1998)
Preference reversals	Consistent social values (e.g., don’t take the largest slice; don’t be the first to cross a picket line) can create what look like preference reversals (Sen 1993; 2002)
Probability matching	Social environments with $N > 1$ individuals competing over resources can make probability matching a more successful strategy than maximizing, whereas this would not be the case for an individual studied in isolation (Gallistel 1990)
Conjunction fallacy	“Conjunction fallacy” can be deduced from the human ability for pragmatic inference about the meaning of natural language sentences – an ability no computer program has so far (Hertwig & Gigerenzer 1999)
False consensus effect	This “egocentric bias” can be deduced from Bayes’s rule for problems of which a person is ignorant, that is, where a person has no knowledge about prior probabilities (Dawes & Mulford 1996)
Violations of logical reasoning	A number of apparent “logical fallacies” can be deduced from Bayesian statistics for environments in which the distribution of events (e.g., P, Q, and their negations) is highly skewed (McKenzie & Amin 2002; Oaksford & Chater 1994), and from the logic of social contracts (Cosmides & Tooby 1992; Gigerenzer & Hug 1992)

petitors, social values, and contracts (Anderson 1990; Gigerenzer 1996c; Gigerenzer et al. 1999). Table 1 (left side) shows a dozen phenomena that have been claimed to be cognitive illusions – by mere assertion, not ecological analysis. Table 1 (right side) shows that as soon as researchers began to study the structure of information in the environment, what looked like a dull cognitive illusion often turned out to be a sharp scissors. Note that these researchers listed on the right side of the table provided formal arguments, not just an optimist’s story set against a pessimist’s saga. Their argument is a deductive one: The environmental structure plus an *unbiased* mind is *logically sufficient* to produce the phenomena. Note that a sufficient reason does not exclude the possibility of alternative explanations (this can be decided by empirical test; see below).

The general argument is that environmental structure (such as unsystematic error, unequal sample sizes, skewed distributions) plus an unbiased mind is *sufficient* to produce the phenomenon. Note that other factors can also contribute to some of the phenom-

ena. The moral is not that people never err, but that in order to understand good and bad judgments, one needs to analyze the structure of the problem or the structure of the natural environment.

For example, consider task environments with substantial but unsystematic error, such as when people are confronted with general-knowledge questions, the answers to which they do not know. A typical finding is that when participants were 100% confident of giving a correct answer, the average number correct was lower, such as 80%. This phenomenon was labeled “overconfidence bias” or “miscalibration” and was attributed to confirmation biases or wishful thinking. An analysis of the environmental structure, however, reveals substantial unsystematic error, which *in the absence of any cognitive bias* leads to regression towards the mean: The average number correct is always lower than a high confidence level. Therefore, the environmental structure is a logically sufficient condition for the phenomenon. Now we can ask if there is, in addition, a trace of a real cognitive bias? When Erev et al. (1994) and Dawes and Mulford (1996) plotted that data the other way

round, regression towards the mean produced a mirror pattern that looked like *underconfidence bias*: When participants answered 100% correctly, their mean confidence was lower, such as 80%. They found no real bias. The same unsystematic error is a sufficient condition for two other phenomena listed in Table 1, people's apparent error of overestimating low risks and underestimating high risks (Lichtenstein et al. 1978), as well as the hard-easy effect (see Gigerenzer & Fiedler, 2003; Juslin et al. 2000).

Consider next how stimulus objects are sampled from an environment and a class of phenomena known as "contingency illusions," which were attributed to irrelevant prior beliefs or prejudices against minorities. Versions of the contingency illusion have been claimed in research on self-fulfilling prophecies (Jussim 1991; Kukla 1993), on confirmation biases in hypothesis testing (Snyder 1984), and on alleged memory advantages for negative behaviors in minorities (Hamilton & Gifford 1976; Hamilton & Sherman 1989).

Let me use evaluative judgments of minorities as an illustration. It is an ecological truism that minorities are smaller than majorities, and a recurrent property of social environments is that the rate of positive, norm-conforming behaviors is higher than the rate of negative, norm-violating behaviors (Fiedler 1991; Parducci 1968). When these two ecological assumptions are built into the stimulus distribution presented in a social psychological experiment, participants may be exposed to the following description:

Group A (Majority):	18 positive	and	8 negative behaviors
Group B (Minority):	9 positive	and	4 negative behaviors

Note that the same ratio of positive to negative behaviors (18:8 = 9:4) holds for both groups, but people nevertheless tend to conclude that there is significantly more positive behavior in the majority than in the minority: a "contingency illusion." Given the unequal sample sizes, however, an unbiased mind using an (unconscious) binomial test would infer that there are significantly more positive than negative behaviors in the majority group ($p = .038$), but not in the minority group ($p = .13$). Thus, unequal sample size is a sufficient condition for a class of phenomena labeled "contingency illusions." Again, one can empirically test whether an additional bias exists because of prior knowledge, such as by replacing real groups by neutral labels, to rule out any influence of prior knowledge (Fiedler et al. 1993; 1999).

Table 1 lists two other phenomena that can be deduced from sampling. One has been called "overconfidence bias," and is defined as mean confidence minus proportion correct (many different phenomena have been labeled overconfidence). Note that "miscalibration" does not imply this phenomenon. It can be logically deduced from unrepresentative sampling of stimulus items and an unbiased mind (Gigerenzer et al. 1991). An analysis of 135 studies showed that "overconfidence bias" practically disappears when stimuli are randomly selected from an environment (Juslin et al. 2000). The second phenomenon is that people erroneously judge that there are more English words with a letter (such as "R") in first position than in third position, which has been attributed to "availability" (Tversky & Kahneman 1973). When one uses a representative sample of letters, rather than the five letters selected by Tversky and Kahneman (which are among the few that are more frequent in the third position), people's apparently systematic bias disappears (Sedlmeier et al. 1998).

The other "cognitive illusions" listed in Table 1 can be deduced in the same way for the task structure, including that of social environments (see also Gigerenzer 2000; 2001; Gigerenzer & Fiedler 2003; Krueger & Mueller 2002). An objection to my general argument is, "But people do commit errors!" No doubt, people commit errors; but I am talking about a blunder committed by a research program. The fact that little attention is paid to establishing what is good and bad reasoning cannot be excused by blaming John Q. Public.

Errors might be a window to cognitive processes, but falsely identified errors do not seem to be so, which is consistent with the fact that after 30 years of collecting errors, no model of cognitive

processes has emerged from overconfidence bias, the conjunction fallacy, or any of the other celebrated errors – only vague labels. In contrast, the study of amazing performance seems to be a better window to cognitive processes, such as the less-is-more effect, which led to the discovery of the recognition heuristic (Goldstein & Gigerenzer 2002).

The story is told that there are two personalities among psychologists, optimists and pessimists, who see the glass as half full or half empty, respectively. According to this legend, people like Funder, Krueger, and myself are just kinder and more generous, whereas the pessimists enjoy a darker view of human nature. This story misses what the debate about human irrationality is about. It is not about how much rationality is in the glass, but what good judgment is in the first place. It is about the kinds of questions asked, not just the answers found.

Null hypothesis statistical testing and the balance between positive and negative approaches

Adam S. Goodie

Department of Psychology, University of Georgia, Athens, GA 30602-3013.
 goodie@egon.psy.uga.edu
<http://www.uga.edu/psychology/faculty/agoodie.html>

Abstract: Several of Krueger & Funder's (K&F's) suggestions may promote more balanced social cognition research, but reconsidered null hypothesis statistical testing (NHST) is not one of them. Although NHST has primarily supported negative conclusions, this is simply because most conclusions have been negative. NHST can support positive, negative, and even balanced conclusions. Better NHST practices would benefit psychology, but would not alter the balance between positive and negative approaches.

Balance is hard to achieve, but Krueger & Funder (K&F) are right to try. Several of their concrete suggestions for the better balancing of positive and negative approaches in social cognition are constructive and likely to be beneficial. First, the canonical literature is not likely to change overnight, but a more balanced reading of it would help, and the canon is indeed more balanced than the treatment it has received in recent decades. Second, K&F avoid the temptation to implore research that is more likely to support positive conclusions. Their formulation, which is entirely reasonable, is to encourage research in which normative responding is at least possible. Third, it would be of salutary value if more researchers examined modal as well as mean responding. When most participants respond normatively, it is appropriate to model this behavior. This is not intended as a complete list of their constructive suggestions: there are more. However, one of their recommendations, although helpful in other ways, is likely to have no impact on the balance between negative and positive approaches. This is the advocacy of increased skepticism with regard to null hypothesis statistical testing (NHST).

To be sure, the limitations of NHST, as it is currently practiced, are real, and the widespread lack of understanding of these limitations has led to a great deal of ritualistic, often inappropriate, application of NHST (Gigerenzer 1993). Anything that may diminish the misunderstanding and ritualistic, inappropriate application of NHST in psychology is to be applauded.

However, although a more thoughtful approach to NHST by more researchers would help psychology in many ways, it would probably have little impact on the balance between positive and negative approaches in psychology. This is because NHST, whether used thoughtfully or ritualistically, may be used equally in support of positive conclusions as in support of negative conclusions.

It is true that the tools of NHST have mostly been used in social cognition research to promote negative conclusions. But this