

Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition

Joachim I. Krueger

Department of Psychology, Brown University, Providence, RI 02912.

joachim_krueger@brown.edu

<http://www.brown.edu/departments/psychology/faculty/krueger.html>

David C. Funder

Department of Psychology, University of California, Riverside, CA 92506

funder@citrus.ucr.edu

<http://www.psych.ucr.edu/faculty/funder/rap/rap.htm>

Abstract: Mainstream social psychology focuses on how people characteristically violate norms of action through social misbehaviors such as conformity with false majority judgments, destructive obedience, and failures to help those in need. Likewise, they are seen to violate norms of reasoning through cognitive errors such as misuse of social information, self-enhancement, and an over-readiness to attribute dispositional characteristics. The causes of this negative research emphasis include the apparent informativeness of norm violation, the status of good behavior and judgment as unconfirmable null hypotheses, and the allure of counter-intuitive findings. The shortcomings of this orientation include frequently erroneous imputations of error, findings of mutually contradictory errors, incoherent interpretations of error, an inability to explain the sources of behavioral or cognitive achievement, and the inhibition of generalized theory. Possible remedies include increased attention to the complete range of behavior and judgmental accomplishment, analytic reforms emphasizing effect sizes and Bayesian inference, and a theoretical paradigm able to account for both the sources of accomplishment and of error. A more balanced social psychology would yield not only a more positive view of human nature, but also an improved understanding of the bases of good behavior and accurate judgment, coherent explanations of occasional lapses, and theoretically grounded suggestions for improvement.

Keywords: Bayesian inference; biases; normative models; personality; positive psychology; rationality; reasoning; social behavior; social judgment; social psychology

Odious behavior (“sin”) is at the heart of our most powerful research in social psychology.

(Elliot Aronson 1999a, p. 104)

How could people be so wrong?

(Lee Ross & Richard Nisbett 1991, p. 139)

Oops, I did it again.

(Britney Spears 2001)

1. Introduction

Although everyday social behavior and cognition includes both appalling lapses and impressive accomplishments, mainstream social psychology has for decades emphasized the negative side of this equation. A prevalent research strategy has been to propose a prescriptive norm for social behavior or cognition and then to demonstrate that human

performance falls short of it. Using this strategy, some of the most influential studies of social behavior documented conformity with false group judgments, obedience to malevolent authority, and failure to help those in need. Studies of social cognition showed how – among numerous other shortcomings – people misuse social information, perceive themselves erroneously, and are too quick to attribute attitudes and personality traits to others. The selective demonstration of negative phenomena is further compounded by the message that people’s intuitions regarding social behavior and cognition are also flawed. For example, people are said to believe that others, but not they themselves, are prone to bias (Friedrich 1996; Pronin et al. 2002). Some investigators have begun to revive interest in human strengths (Seligman & Csikszentmihalyi 2000; Sheldon & King 2001; Snyder & Lopez 2002) and cognitive accomplishments (Gigerenzer et al. 1999; Klein et al. 2002), but so far their influence on social psychology has been limited.

The purpose of the present article is to examine some of the causes and consequences of the prevalent negative research orientation and to sketch analytical and theoretical routes leading to a more balanced social psychology. The time for reform is ripe because the historically rooted paradigm of uncovering ever more behavioral and cognitive flaws may be approaching a dead end. It is becoming progressively less informative as it continues to proliferate, causing human strengths and cognitive skills to be underestimated and impairing the development of theory.

The persistent emphasis on the negative is problematic because research designed to uncover misbehavior or cognitive failures is sure to find some. Without efforts to also examine behavioral strengths and cognitive successes, a distorted view of human nature emerges that yields a cynical outlook on human nature rather than usable guidance for behavior and judgment. It is doubtful, for example, that people could function effectively if they refrained from all obedience, intervened in all apparent crisis situations, discarded judgmental heuristics, or suspended judgment altogether; yet, that is what research demonstrating human shortcomings in each of these domains would seem to recommend.

Studies of bad behavior and flawed reasoning often settle for rather simple demonstrations. The empirical section of the typical article shows that people can be induced to do something objectionable or to think in a way they should not. The discussion section may contain some speculation of how many social problems must be due to this tendency, and a call may be placed for research on how to reduce its prevalence. The analysis generally stops there, short of asking *why* such a behavioral or cognitive tendency exists, or what general purpose it might serve. As a result, the development of integrative theory and sensible advice is stymied (Katzko 2002).

The situation is reminiscent of the early days of vision research. When visual illusions were first discovered, they were considered mistakes produced by arbitrary design flaws (Gregory 1971). An early interpretation of the Müller-Lyer illusion, for example, was that it reflects a general ten-

dency to overestimate acute angles and to underestimate obtuse ones. Then, in 1896, psychologist A. Thiery proposed that this and other illusions reflect processes that permit accurate perception in real-life contexts. Today, optical illusions are no longer seen as failures of the visual system, and airline pilots are not taught that the Müller-Lyer and Ponzo illusions pose threats to their performance. In contrast, the pre-1896 view still dominates social-cognitive psychology. Behaviors and judgments that violate experimenter-imposed norms are interpreted as revealing flawed psychological processes that need to be fixed (Funder 1987).

The current state of social psychology has parallels in biomedical research, which is often based on problem-finding and indeed may be funded on the basis of the problem it seeks to alleviate. The search for a cure for a particular disease has popular and even political appeal. But ultimately, it is the systematic, theory-based research of basic physiology that explains how the human body usually functions well, and also how it malfunctions under certain conditions (Fields 1994; Skalka 1993). In a parallel manner, basic, theory-driven research on social psychological processes will most fully illuminate the peculiar shortcomings *and* the adaptive successes of the social animal.

2. Negativity in social psychology

Two traditions, a classic behavioral one and a more recent cognitive one, characterize the history of social psychology. The emphasis of both has been disproportionately negative.

2.1. Social behavior

The most remarkable fact about social behavior, according to the mainstream view, is how often it violates normative standards of conduct. In the words of one eminent researcher, “odious behavior (‘sin’) is at the heart of our most powerful research” (Aronson 1999a, p. 104). Historically, the concern with the odious began with analyses of human behavior in crowds (Le Bon 1895). With the possible exception of Floyd Allport (1924), the founders of social psychology worried that men (and women) could only be trusted to behave properly when left to their own devices, and that the social influence of the group would transform them into irrational, suggestible, and emotional brutes (see Asch 1952 for a review and critique).

In the 1950s and 1960s, a number of laboratory studies cemented the view that social influence has nefarious consequences on otherwise rational individuals. These studies demonstrated conformity with bizarre group behavior, obedience to destructive authority, and apathy among the bystanders of an emergency. Yielding to social influence was tantamount to violating behavioral norms of independence and empathy. Even research addressing seemingly positive aspects of human nature, such as interpersonal attraction, or neutral topics, such as attitude change, focused on the negative. One of the most widely cited studies of human attraction concluded that superficial cues of physical attractiveness overwhelmed cues to other personal qualities that people claim they value (Walster et al. 1966). The basic theme of attitude change research, whether from the cognitive dissonance tradition or the competing self-perception approach, has been that people are typically unaware of the degree to which their attitudes come from rational-

JOACHIM I. KRUEGER is Associate Professor of Psychology and Human Development at Brown University. He received his Ph.D. in psychology from the University of Oregon in 1988. A fellow of the American Psychological Society, he has published in the areas of self-perception, stereotyping, and statistical inference. The study of the conditions under which human judgment is more or less rational or accurate is a theme underlying much of his research.

DAVID C. FUNDER is Professor of Psychology and Chair of the Department of Psychology at the University of California, Riverside. He received his Ph.D. in psychology from Stanford University in 1979. A fellow of the American Psychological Association, the American Psychological Society, and the American Association for the Advancement of Science, he has published in the areas of person perception, personality assessment, and personality development. His primary line of research addresses the moderators of accuracy in judgments of personality, and the social processes that make accurate judgment possible.

ization rather than from rational thought (Aronson 1969; Bem 1972). But these conclusions are only implicitly negative. As we now illustrate, some of the most influential studies of social behavior and cognition have been explicitly interpreted as demonstrating surprising flaws in human nature.

2.1.1. Conformity. Solomon Asch (1956) pitted naïve participants against a unanimous group of confederates who, on occasion, rendered bizarre judgments concerning the relative lengths of lines. This situation included considerable social pressures to conform, but no incentives to resist. In the maximum-impact experiment, 90% of the participants gave two or more incorrect responses and about one-third of all responses were false. Conformity meant that participants violated the normative expectation that they should honor their own perceptions and be able to tolerate disagreement with others. Although Asch was also interested in – and empirically demonstrated – processes that allow resistance, the story of conformity carried the day (Friend et al. 1990).

2.1.2. Obedience. Stanley Milgram (1974) led his participants to violate a norm of good behavior in particularly dramatic fashion. Believing they were assisting in a learning experiment, the participants faced an experimenter who relentlessly ordered them to deliver ever-increasing electric shocks to a faltering confederate. Overall, about 40% of the participants administered what they must have thought were lethal voltages. This was a surprising finding on the assumption that ordinary people would not hurt innocent others even when ordered to do so. A panel of psychiatrists predicted that only the rare sociopath would inflict mortal harm. Like Asch, Milgram went on to do under-appreciated work that demonstrated how various situational variables, such as the distance between participant, victim, and experimenter, caused the compliance rate to vary.¹

2.1.3. Bystander intervention. Inspired by the infamous murder of Kitty Genovese, a series of experiments by John Darley and his colleagues showed that people would fail to intervene in an emergency inasmuch as other witnesses were present (Darley & Latané 1968) and inasmuch as they were under time pressure (Darley & Batson 1973). Overall, about half the participants intervened,² although normatively all of them were supposed to. Quoting from Luke (10:29–37), Darley and Batson noted that the behavior of their research participants fell short of the example set by the Good Samaritan.

2.2. Sources of negativity

In each case, the aspect of the results that aroused the most interest was not the power of the situation per se, but the power of particular situations to elicit bad behavior. However, the same studies could also be construed as equivalently revealing the sources of nonconformity, independence, and helping. Asch, Milgram, and Darley and Latané showed that variations in the setting, such as the presence of allies or being held individually accountable, increased the prevalence of normative behavior, and many participants acted normatively even in the maximum-impact conditions. But this variation was seldom emphasized in the reviews and texts that made these studies famous. In 1997, a

segment of NBC's *Dateline* featured a re-enactment of the Asch experiment and excerpts from Milgram's obedience film, but made no mention of any of the moderator variables. Instead, the broadcast began with episodes from *Candid Camera* showing how strange situations can elicit strange behavior (e.g., the person who faces the back of the elevator because everyone else does). A spirited on-line discussion sponsored by the *Society of Personality and Social Psychology* did not reveal any concerns about this biased presentation. Instead, one commentator warned that the power of the situation to elicit compliance had not been emphasized enough.

The emphasis on negative outcomes leaves the powerful impression that the underlying psychological processes must be intrinsically maladaptive. Just as in visual perception, however, process and outcome are separate issues. Processes that can produce bad behavior in particular circumstances may yield desirable or adaptive results in other circumstances. The judgments of others can be informative (Deutsch & Gerard 1955), obedience to legitimate authority is important for social order (Hogan et al. 1994), and hesitation to get involved in someone else's struggle may save one's own neck. A more balanced recognition of the costs and benefits of conformity, obedience, intervention and other seemingly problematic behaviors would not only be more realistic, it would also ask that theories explain the range of human behavior, not just the negative end.

So why is the emphasis so unbalanced? At the most general level, it seems that negativity itself is a powerful motivator. Social psychological research has documented numerous ways in which negative information commands resources of perception, attention, and memory in ways that positive information cannot (Baumeister et al. 2001; Rozin & Royzman 2001). If the dominance of the negative evolved as an adaptation to life in uncertain and potentially hazardous environment, it may, in part, explain the negative bent of social research. This idea cannot explain, however, why other fields within psychology have a more positive outlook. Much developmental work, for example, is motivated by the search for capabilities among infants that no one expected they had. To move past the somewhat tautological idea that social-psychological research tends to be negative because of negativity bias, we consider four specific sources of this bias.

2.2.1. Zero-tolerance norms. In the classic studies, the frequency of misbehavior was considered to be surprisingly high even if it was limited to a minority of participants. But what is the smallest number sufficient to trigger surprise? Because the strict view that "nobody will be induced to behave badly" is too easily refuted, data analysis commonly proceeds along probabilistic lines. For example, an investigator might lay a confidence interval around the observed percentage of violations. As sample sizes increase, shrinking confidence intervals ultimately exclude zero. At that point, norm violations are considered established even if most participants acted properly. Consider the stock finding of ethnocentrism in laboratory groups. In the classic study, Henri Tajfel found that most participants distributed rewards equitably among ingroup and outgroup members. Only when the allocation matrix made fairness impossible did a significant number of participants reward outgroupers less than ingroupers (Tajfel et al. 1971). This finding led to the widely accepted conclusion that people discriminate

against outgroups without sufficient justification (but see Gaertner & Insko 2000 for a recent challenge of this view).

2.2.2. Informativeness of norm violation. Norm violations stand out as figures before a ground, just as they stand out statistically as signals against a background of noise. Because it is the expected behavior almost by definition, norm adherence does not demand explanation, and may even seem boring (Jones & Davis 1965). If people had behaved as they should, Asch's experimental design would have appeared ludicrous, Milgram's colleagues would have felt vindicated, Darley and Latané's research would have confirmed that people live according to scriptural precepts, and few readers would have heard of any of these studies. But that is not what happened. Instead, classic social psychology garnered great attention by exposing expectations of normative behavior as naïve.

Note that on purely numerical grounds, a persistent emphasis on norm violation ought to be self-eliminating. As demonstrations pile up, their surprise value should dissipate as the counter-norm becomes the new norm. This does not appear to have happened, probably because most of the social norms that are invoked are anchored not merely in statistical expectations, but in moral or ethical precepts.

2.2.3. Appeal of counterintuitive findings. Ordinary people know a great deal about human behavior, and this knowledge has helped to identify basic psychological principles (Kelley 1992). Nevertheless, as findings based on commonsense hypotheses risk being dismissed as "bubba psychology" (McGuire 1997), "psychology has often embraced counter-CS [i.e., common-sense] data as knowledge" (Kluger & Tikochinsky 2001, p. 411). Pressures to refute common sense arise from both inside and outside the field. From the inside, findings consistent with intuitive expectations seem uninteresting. From the outside, this view is reinforced by those who claim that they "always knew that." Senator William Proxmire once offered a "golden fleece" award to federally supported psychologists who obtained results he considered obvious. In contrast, demonstrations of norm violation are protected from ridicule, and may even gain a cachet of urgency and truth. To report that more people conform, obey, and fail to intervene than even one's grandmother (i.e., "bubba") would have expected, is an effective rebuttal to all those laypeople who feel they understand behavior as well as trained social psychologists do.

But some recent reviews question the robustness of counter-intuitive findings (Kluger & Tikochinsky 2001). A longstanding staple in the cupboard of counterintuitive knowledge has been that one's confidence in judgment has no relation to one's accuracy. However, recent research has shown that under realistic circumstances (e.g., when observational circumstances are varied), the confidence of eyewitnesses is quite closely related to their accuracy ($r = .59$; Lindsay et al. 1998; see also McCullough 2002). This is just one example. Kluger and Tikochinsky (2001) reported nine other cases in which an accepted counter-intuitive finding was reversed. Of these, eight lay within the areas of social and personality psychology.

2.2.4. Usefulness to society. Following Lewin (1948), many social scientists are concerned with the practical relevance of their findings (Aronson 1999a; Redding 2001). The goal of curing social ills requires first that a social prob-

lem be identified. Then, the critical undesirable behavior needs to be reproduced in the laboratory. Finally, the empirical demonstration can be interpreted as reflecting a typical – and thus potentially dangerous – human liability. At times, this sequence is followed by efforts to retrain individuals in order to alleviate their diagnosed flaws.

2.3. Social cognition

Over the last three decades, the cognitive reorientation of social psychology has shifted attention away from social behavior and towards social perception and judgment. Initially, normative models were believed to characterize how people ought to make inferences as well as how they actually do make inferences (Kelley's 1967; Peterson & Beach 1967). By 1980, this optimistic view had been displaced by a focus on inferential shortcomings and errors (Kahneman et al. 1982; Nisbett & Ross 1980). This emphasis continues today (Gilovich et al. 2002; Myers, 2002), and it has penetrated the literature on the application of psychological science to areas such as medicine, counseling, and management (Bazerman 2001; Heath et al. 1994).

The shift towards the negative followed a similar development in the field of judgment and decision making (Mellers et al. 1998), which, in turn, was stimulated by a series of articles by Tversky and Kahneman (e.g., Tversky & Kahneman 1974). Kahneman and Tversky (1973) challenged the axiomatic status of rationality in economic theories of choice, and social psychologists soon saw the potential of this new paradigm for the study of social cognition. Dawes (1976) reviewed the historical context for the psychological approach to irrationality. From Aristotle to the Catholic Church to Sigmund Freud, he argued, irrational thought and behavior was viewed as the result of capricious emotional forces disrupting the workings of an otherwise rational mind. In contrast, the modern view holds that the presumably rational apparatus of conscious thought is itself fraught with deficiencies. This is an even more depressing verdict than the traditional one. If conscious capacities cannot be counted on to detect and correct creeping errors, what can?

Ross (1977) took the next step by recasting the history of social psychology from the perspective of the heuristics-and-biases paradigm. He argued that the classic studies of social misbehavior gained their significance from the inability of everyday observers to predict their outcomes and to understand their implications. This inability, he suggested, betrayed a profound failure to think scientifically. The enthronement of the scientific method – as social psychological investigators understood it at the time – as the gold standard of good reasoning was a crucial step. It provided an uncompromising norm for human judgment, much as strict ethical demands had served as standards of behavior. But the metaphor of the intuitive scientist elevated the human mind only to denigrate it. And it had to be so because without certain key biases, "social psychology's most beloved phenomena would not have occurred and its most beloved experiments would have been mere platitudes" (Gilbert 1998, p. 130). Thus, a merger was achieved that tied the psychology of misbehavior to the psychology of flawed thinking.

Once this view was established, debates focused primarily on *which* of various negative metaphors explains the social perceiver's failings best. Following Dawes (1976), some

avored the metaphor of the “cognitive miser” by emphasizing limited mental resources, reliance on irrelevant cues, and the difficulties of effortful correction (Wilson & Brekke 1994). Others preferred the “totalitarian ego” metaphor (Greenwald 1980) to emphasize needs for self-esteem and control, as well as the self-deception necessary for the satisfaction of these needs (Ehrlinger & Dunning 2003). Despite their differences, both the cognitive and the motivational approach viewed distortions and errors as the fundamental and most informative aspects of social cognition. Whereas the early theorists regarded individual rationality as a haven from the irrationality of the group and a protector against collective excesses, the current view leaves little room for refuge. To illustrate, we now turn to three of the most widely studied biases.³

2.3.1. False consensus. In the study of perceptions of social groups, the preponderant bias became known as the false consensus effect (FCE). The FCE is understood as a projective tendency in which self-referent information serves as a judgmental anchor, from which other-referent or group-referent properties are inferred. In the paradigmatic study, undergraduates decided whether to assist in a study of communication by walking around campus wearing a sandwich board that read “Eat at Joe’s,” and then were asked to estimate how many other students would agree to wear the sign (Ross et al. 1977). The implicit model for the social perceiver was that of an ideal scientist who would discard idiosyncratic perceptions and appraise sample data with the cold eye of objectivity. To match this ideal, a social perceiver would need to ignore his or her own paltry contribution to the population average and base consensus estimates only on observations drawn from sufficiently large and unbiased samples (which happened to be unavailable to the participants in this study). To detect bias, investigators did what naïve social perceivers are incapable of doing. Comparing the average consensus estimate offered by compliant participants with the average estimate by noncompliant participants, they found a statistically significant difference.

2.3.2. Self-enhancement. When people evaluate themselves, they rely in part on comparisons with others. The bias of self-enhancement is said to occur when people think they are better than the average person. In the paradigmatic study, participants rated positive traits as more descriptive of themselves than of most others (Brown 1986). The verdict that this bias was ethically unwarranted stemmed from a comparison of the findings with the humanistic ideal that well-adjusted people feel as favorably about others as they do about themselves (Rogers 1961). The verdict that the bias violated a norm of rationality stemmed from the assumption that people should be able to make correct estimates of their relative standing in the population. It was further assumed that when self-descriptions differ from descriptions of the average other, the former must be wrong.⁴

2.3.3. The fundamental attribution error. According to attribution theory, people try to explain behavior by looking for its causes in the person, in the situation, or in both (Heider 1958). Research on the “fundamental attribution error” (FAE)⁵ maintains that they characteristically fail at this task by overestimating the importance of properties of the person. In the paradigmatic study, participants read essays that

either favored or opposed the regime of Fidel Castro (Jones & Harris 1967). Some participants learned that the authors of the essays had freely chosen which position to take, whereas others were told that the authors were assigned their positions. When participants estimated the true attitudes of the essay-writers, they concluded that those with free choice were more likely to believe what they said than were those who were coerced. This finding effectively demonstrated the “discounting principle,” which demands that a potential cause for a behavior is weakened to the extent that other plausible causes emerge (Kelley 1972). It was a different result, however, that captured the attention of a generation of researchers. Even in the coerced condition, participants attributed more pro-Castro attitudes to writers of pro-Castro essays than to writers of anti-Castro essays. The inference was much weaker than it was in the noncoerced condition,⁶ but it remained significantly larger than zero. Again, the normative ideal violated was that of an idealized social scientist, who, in this case, would completely refrain from dispositional inferences once an effective situational cause had been found.

2.4. Sources of negativity

A striking indication of social psychology’s preoccupation with the negative is the sheer number of published errors. Table 1 presents a partial list of errors reported over the past few years.⁷ Some of these refer to the same or a nearly identical phenomenon; others share the same label, although they refer to different phenomena (the “jingle-jangle” effect, Block 1995); and still others are contradictory. For now, suffice it to note their number and variety. Just as God has been said to have an inordinate fondness for beetles, having made so many (Evans 1996), social psychologists may have an inordinate fondness for errors, having found so many.

Table 1. *Some errors of judgment identified and labeled by social psychologists*

Overconfidence bias	Correspondence bias
Fundamental attribution error	Halo effect
False consensus effect	False uniqueness effect
Positivity bias	Negativity bias
Confirmation bias	Disconfirmation bias
Justice bias	Male bias
Hot hand fallacy	Gambler’s fallacy
Self-protective similarity bias	Hindsight bias
Self-serving bias	“Ultimate” self-serving bias
Optimistic bias	Pessimistic bias
Sinister attribution error	Conjunction fallacy
Ingroup/outgroup bias	Positive outcome bias
Hypothesis-testing bias	Diagnosticity bias
Durability bias	Vulnerability bias
Self-image bias	Labeling bias
Observer bias	External agency illusion
Systematic distortion effect	Intensity bias
Asymmetric insight illusion	Just world bias
Dispositional bias	Romantic bias
Clouded judgment effect	Bias blind spot
Empathy neglect	Empathy gaps

Note: Partial list of major topics of studies published since 1985.

Most of these errors are presented as lamentable shortcomings signaling dangerous flaws in the system (Funder 1992).⁸ Tversky and Kahneman (1971) characterized human judgment as “ludicrous” (p. 109), “indefensible” (p. 108), “self-defeating” (p. 107), and guilty of “a multitude of sins” (p. 110). Ross and Nisbett (1991) described the typical person as “oblivious” (p. 124) and “insensitive” (p. 82), as well as beset by “ignorance” (p. 69), “general misconceptions,” and a “whole range” of other “shortcomings and biases” (p. 86). The only question left, it seemed, was “How could people be so wrong?” (p. 139).

This condemnationist orientation heuristically equates bias with inaccuracy, and ignores the long-range outcomes of presumably non-normative judgments. Like the presumed social misbehaviors, however, many social-cognitive biases yield considerable benefits. Social projection increases the accuracy of social perception (Kenny & Acitelli 2001; Krueger 1998b) and satisfaction with a partner (Murray et al. 2002). Positive self-concepts not only are their own hedonic rewards, but they also tend to increase the accuracy of self-judgments whenever the attribute in question is distributed with a negative skew (Krueger 1998a). For example, relatively few people are deathly ill; most are in the range from medium to very good health, and these people will have health “scores” higher than the arithmetic mean.⁹ Drawing dispositional inferences even from situationally constrained behavior, often interpreted as a manifestation of the “fundamental attribution error,” may be a sign of the perceiver’s own social competence and adjustment (Block & Funder 1986). In sum, it appears that many social-perceptual biases signal the operation of an adaptive system of social perception much like certain visual illusions reveal the efficiency of the visual system under realistic circumstances (Funder 1987). Again we need to ask, why is the emphasis of social research so negative?

2.4.1. Norm violation, usefulness, and counter-intuitiveness. Some of the reasons for the negative tone of research on social cognition parallel the ones considered in the context of social behavior, including the apparent informativeness of norm violations, the desire to alleviate social problems, and the appeal of the counterintuitive. When judgments consistent with a norm of rationality are considered uninformative, only irrationality is newsworthy. Assuming that social problems can be traced to poor thinking (Jones & Roelofsma 2000), many researchers seek to identify “systematic irrationalities” (Stanovich & West 2000, p. 646) and ways to eliminate them (Baron 1998). The appeal of counter-intuitive findings is even stronger in the area of social cognition than in the area of social behavior. As one writer put it, “Mistakes are fun! Errors in judgment make humorous anecdotes, but good performance does not. It is fun to lean back in our chairs and chuckle about our goofs” (Crandall 1984, p. 1499).

This rhetoric of irrationality created the perception of a deep crisis in human cognition that could only be overcome if people learned to set aside heuristics and reason as normative models prescribe (Lopes 1991). The rhetoric continues even though some investigators maintain that they never meant to impugn the capabilities of human judgment in the first place (Kahneman 2000; Ross & Nisbett 1991).

2.4.2. Rationality as a null hypothesis. Much like earlier research on violations of behavioral norms, research on cog-

nitive biases has been beholden to the methodological ideal of experimentation as the *camino real* to causal inferences (Gigerenzer 1996a; Krueger 1998c; Rozin 2001). As part of this ideal, social-cognitive research inherited the analytical tool kit of null hypothesis significance testing (NHST). Whereas NHST can be used to detect causal relationships, its task in bias research often is merely to detect the presence of a significant difference between the average judgment and a normative standard. Thus, NHST is used in its weak, confirmatory form. Being identified with a complete lack of a difference, rationality at best remains a null hypothesis that has failed to be rejected.

As sample size increases, the precision of measurement is improved, more robust statistical tests are employed, and ever-smaller effect sizes pass the threshold of significance (Kirk 1996; Wilcox 1998). In some cases, this allows biases to reach significance level even when the modal response is identical with the demands of the normative model.¹⁰ The dichotomous decision rule of NHST – a bias either has been demonstrated, or it has not (yet) been demonstrated – leaves no room for bounded, or good-enough rationality, nor does it distinguish between biased and unbiased individuals. As the boundary between rationality and irrationality dissolves, any opportunity to learn how many respondents got it right is lost.

When efforts to detect bias fail, nothing positive can be said about the presence of rationality because the null hypothesis not only represents rational judgment but also chance variation (Gigerenzer & Goldstein 1996). As the negation of causation, chance can neither be produced nor explained (Hayes 2001). Therefore, psychological mechanisms are more readily invoked to explain bias than to explain the absence of bias. With no direct way of explaining accuracy, the absence of bias, when it occurs, might even be explained by the mutual cancellation of opposite errors (see Epley et al. 2002 for a case of a co-occurrence of social projection, i.e., the “spotlight effect” and the FAE).

3. Problems with the negative emphasis

The view that people characteristically violate norms of good behavior and rational thought raises two further problems. First, some of the imputations of misbehavior and error are themselves difficult to justify, and second, the problem-seeking approach tends to be atheoretical. The lack of attention to behavioral and judgmental accomplishments not only prevents understanding of adaptive behavior or accurate judgment, but it also retards a full understanding of the sources of the misbehaviors and errors when they do occur.

3.1. Rash imputations of error

The imputation of irrationality should demand a high standard of proof. The average human, no less than the common criminal, deserves to be considered innocent until proven guilty. More importantly, a false imputation of incorrect thinking threatens the validity of subsequent empirical or theoretical analysis. The typical lack of such a high standard opens the door to incoherent findings and contradictory conclusions.

3.1.1. Contradictory errors. In many studies of social judgment, the null hypothesis of rationality is sandwiched be-




Figure 1. Pointed rationality between ranging biases.

tween opposite biases. Consider the three paradigmatic areas of judgment. As shown in Figure 1, consensus estimates are unbiased only if they do not covary with the perceivers' own responses. Most studies show projective bias, but scattered reports of false uniqueness raise the unsettling possibility that opposite biases might co-exist (Krueger 2000a). In self-perception, people are typically found to self-enhance, but there are also reports of self-diminishment (John & Robins 1994; Yik et al. 1998). In attribution studies, the correspondence bias (or FAE) is the standard finding, but its inverse, the insufficient discounting of situational causes, has also been found. When the experimental design demands that the situational cause be discounted, participants overestimate its effect (Krull & Dill 1996; Quattrone 1982).

The co-existence of contradictory biases is not limited to the three paradigmatic areas of social judgment. Intuitive predictions have been found both to overstate and understate the probability that past events will recur. Belief in "the hot hand in basketball" exemplifies the former finding (Gilovich et al. 1985), whereas "the gambler's fallacy" exemplifies the latter (Keren & Lewis 1994). Similarly, many studies show that people neglect base rates when making predictions (Bar-Hillel 1980; Kahneman & Tversky 1973), whereas others suggest that they use them too much (Edwards 1982). Such contradictions can escape notice when opposite biases are presented as part of different topics using different terminologies. Textbook chapters on social cognition maintain that people make faulty predictions by relying too much on specific case information while under-using category (i.e., base rate) information. Chapters on intergroup relations maintain the opposite, namely that people overuse their categorical (i.e., stereotypic) beliefs while neglecting individuating information (Funder 1995b).

Opposite biases can even emerge in the same study. When this happens, *ad hoc* assumptions may take the place of theoretical integration. Kruger and Dunning (1999) reported that participants who scored low on a test of ability grossly overestimated how well they did relative to other test takers. In contrast, participants who scored high on the test slightly underestimated their relative standing. Kruger and Dunning dealt with this apparent inconsistency by attributing each error to a distinct cognitive failure. Poor performers, they argued, overestimated their ability because they lacked the meta-cognitive insight into their own weaknesses. They were "unskilled and unaware of it." The opposite bias displayed by the high performers was attributed to their falling "prey to the false consensus effect" (p. 1126) – one bias to which the unskilled were apparently immune.

When the rituals of NHST are suspended, it is no longer necessary to interpret all observed differences between estimates and normative values as distinct biases with correspondingly distinct mechanisms. Instead, the asymmetry in estimation errors can be explained by regression to the mean in conjunction with an overall, group-level, better-than-average effect. Estimated and actual performance can be expected to be positively but imperfectly correlated – hence regression – and overall, people can be expected to be optimistic rather than pessimistic – hence asymmetric "errors" (Krueger & Mueller 2002; see also Ackerman et al. 2002).

The debate over the putative "unskilled-and-unaware" effect was but a replay of an earlier controversy over asymmetric errors in consensus estimation. Meta-analyzing studies on consensus bias, Mullen and Hu (1988) noticed that people holding minority attributes grossly overestimated the prevalence of these attributes, whereas people holding majority attributes slightly underestimated the prevalence of majority attributes. Again, it was not necessary to associate different errors with different flaws of thinking (as Mullen and Hu did). Simply assuming that most people expect to be in the majority and noting that estimates are inevitably imperfect, a regression model replicated the pattern of asymmetric errors (Krueger & Clement 1997). Over the past decade, several other biases have been reinterpreted as the result of such random imperfection and regression. Among them are such core phenomena as over- (and under-) confidence (Klayman et al. 1999) and illusory correlations (Fiedler & Armbruster 1994).

As soon as one asks whether changes in one bias may result in changes in others, one moves towards a more comprehensive model. Since Floyd Allport's original exposition (Allport 1924), people have been charged both with "pluralistic ignorance" and social projection. Pluralistic ignorance reflects the perception of a difference between the self and the group. It is said to occur when people in general underestimate the prevalence of a certain (usually socially desirable) attitude. In contrast, social projection (or the FCE) reflects a perception of similarity between the self and the group. It is said to occur when those who hold a certain attitude believe it to be more common than those who do not hold it.

Studying attitudes toward alcohol consumption, Prentice and Miller (1993) found that on average, college students felt that others were more permissive toward drinking than they themselves were (pluralistic ignorance). At the same time, those who expressed a less permissive attitude thought

there was less permissiveness on campus than did those students whose own attitudes were more permissive (social projection). Prentice and Miller deflected the idea that these two errors might be opposites by noting that both can co-occur empirically. Indeed they can, because pluralistic ignorance is the difference between the true prevalence of an attitude and the mean estimate of that prevalence, whereas projection is the correlation between estimates and people's own individual attitudes. It is easily shown that as projection increases, pluralistic ignorance decreases (Krueger 2002). Once again, projection is more beneficial than irrational.

Another example involves the relationship between social projection and self-enhancement. When self-enhancement is construed as the "better-than-average effect," it is easily misinterpreted as the opposite of projection, that is, as a "false-uniqueness effect" (Fiske et al. 1998; Markus & Kitayama 1991). Whereas seeing oneself as different from the average suggests a psychological contrast (which need not be a false one; Krueger 1998a), social projection suggests assimilation. Again, however, a closer look at the units of analysis dissolves the paradox. Whereas self-enhancement is a mean-level effect, social projection is a correlational effect. For an individual judgment item, both effects tend to emerge, but they are negatively related across items. The more people assume others to be similar, the harder it is to feel superior. In this case, social projection serves as a brake against narcissistic over-evaluation of the self (Krueger 2000b).

3.1.2. Wrong or misapplied norms. Some putative demonstrations of error are themselves erroneous because the norm against which behavior or judgment is compared is incomplete, wrong, or misapplied. In each of the three paradigmatic social-judgment tasks, the norm of zero difference can no longer be taken for granted. As we have seen, social predictions and self-perceptions would be respectively less accurate if people ignored their own responses and if they rated themselves as being average on negatively skewed attributes. If they attributed coerced behavior entirely to the situation, they would concede that experimenters always manage to secure compliance from their participants, arguably an overgeneralization (Morris & Larrick 1995).

Research on the classic base-rate integration problem yields a similar conclusion. Participants who seemingly fail to produce a Bayesian probability estimate may be doing rather well if one assumes that they approach the task as a matter of signal detection (Birnbaum 1983). In this as in many other cases, the evaluation of performance depends on which of several plausible normative standards is brought to bear. And why should human performance be asked to measure up against the particular normative standard to which an experimenter happens to subscribe? As Nozick (1996) noted, "theorists of rationality have been intent upon formulating the one correct and complete set of principles to be applied unreservedly in all decision situations. But they have not yet reached this – at any rate, we do not have complete confidence that they have" (p. 46). Perhaps it is more sensible to ask whether research participants fail (or succeed) on their own terms (Ayton 2000; Moldoveanu & Langer 2002). Such an empowerment of the participants implies, of course, that "normative theories will be drained of all their evaluative force" (Stanovich & West 2000, p. 655). At any rate, this proposal does not mean

that judgments cannot be evaluated. Instead, multiple norms may need to be considered, and an effort should be made to understand which one best represents the meaning and goals of the participants.

3.1.3. Incoherent explanations of misbehavior and error.

A minimum requirement for rational judgment is to avoid outright contradictions (Dawes 1998; Krueger 2000a). Because the coherence criterion is such a powerful device, it ought to be applied to explanations of misbehavior and bias, too. When this is done, we find that many accounts of human judgment lack the very coherence they demand of naïve research participants. With regard to social projection, Dawes (2000) observed that "it is not the research subjects or intuitive experts that reach an irrational conclusion, but the psychological theorists analyzing them" (p. 134). In the following, we pursue Dawes' argument to examine three prominent accounts of the FAE, which deserves close inspection because of its flagship status as the self-proclaimed "fundamental" flaw of social intuition.

3.1.3.1. Logical incoherence. Recall that the classic demonstrations of non-normative behavior aimed to show that situational forces overwhelm people's intentions or dispositions to behave responsibly (e.g., Conner 2000). At the same time, however, their surrender to situational pressures has been taken to indicate negative dispositions, such as lack of autonomy, empathy, or ethical fiber. If people conform with patently false judgments, obey malevolent masters, and fail to help those in need, they might also readily accept anti-democratic values, ignore the suffering of others, and participate in genocide. When evil becomes banal, there must be something wrong with people. Inferences of this kind amount to precisely the kind of dispositional attributions in which untutored people are said to overindulge.

The most influential account of the FAE rests on Heider's (1958) classic distinction between internal and external sources of behavior. It assumes that

the human skin [is] a special boundary that separates one set of "causal forces" from another. On the sunny side of the epidermis are the external or situational forces that press inward upon the person, and on the meaty side are the internal or personal forces that exert pressure outward. (Gilbert & Malone 1995, p. 21)

This version of attribution theory assumes that behavioral causation is a zero-sum game. Dispositional causes must be discounted when situational causes are shown to be effective. It follows that perceivers are mistaken to appeal to dispositional causes when a change in the situation explains the observed behaviors. What are the implications of this hydraulic person-situation model for the researchers' own chain of inference? To preserve coherence, they need to argue that perceivers' dispositional judgments were elicited (and thus caused) by specific experimental conditions. If so, the conclusion that perceivers' dispositional inferences were reflections of their own disposition to commit the FAE would be an expression of the very error it is meant to explain. This paradox is so delicious that it deserves to be savored like a fine Merlot. To claim that only situational effects are real while bemoaning participants' dispositional lack of insight into this important truth is incoherent. *The FAE cannot be discovered unless investigators, by their own criteria, commit it themselves!*

The hydraulic model does not do justice to the dynamic interplay of personal and situational variables (Sabini et al. 2001). Lewin (1951) famously observed that behavior is a function of the person *and* the situation, which specifically means it does not have to be one or the other. The Milgram setting, for example, can be understood as a situation in which two opposing external forces interacted with two opposing internal forces. The implacable experimenter provided a situational force towards obedience, whereas the complaining victim provided a situational force towards compassion. At the same time, the individual's disposition toward obedience was opposed by whatever disposition he or she had towards compassion (Sabini & Silver 1983). When observers predicted less obedience than Milgram obtained, they not only underestimated the experimenter's situational influence, they also *overestimated* the victim's situational influence. It is also correct to observe that observers overestimated the dispositional tendency towards compassion at the same time that they *underestimated* the dispositional tendency towards obedience. The equivalence of these two superficially different accounts underlines how behavior makes little sense without the interplay of both situation and disposition (Wright & Mischel 1987).¹¹

3.1.3.2. Statistical incoherence. The second interpretation of the FAE was developed in an effort to sidestep the difficulty of justifying the a priori separation of situational and dispositional causes. According to this statistical account, behavior can be assumed to be under situational control when behavioral variance is low. When the variance is high, it is assumed to be under dispositional control (Jones & Davis 1965; Kelley 1967). Thus, a religious service is a more powerful situation than the opening of a department store. In the former, nearly everyone does as required (e.g., pray), whereas in the latter, many people do not do as hoped (e.g., shop; Snyder & Ickes 1985). Note that this interpretation still presupposes a hydraulic model because it conceives of situations and dispositions as competing accounts of behavioral variance (Ross 1977).

Even as the statistical argument has become accepted as a “logical standard [that] does not seem . . . to have any serious competition” (Gilbert 1998, p. 135), its implications for the FAE have been overlooked. Indeed, the statistical criterion *reverses* the standard interpretation of situational power (Sabini et al. 2001). Because many participants disobeyed Milgram's experimenter, the situation bore a closer resemblance to the opening of a department store than to a prayer service. When estimating that hardly anyone could be induced to administer severe shocks, psychiatrists predicted the opposite, namely that the situation would have an extremely strong effect by evoking uniform independence. Thus, they *overestimated* the power of the situation by underestimating interpersonal variation. Predictions regarding behavior in the Asch conformity experiment or in Darley and Latané's studies on bystander apathy can be interpreted in a parallel manner. Because it so often implies that people overestimate instead of underestimate situational forces, the statistical interpretation of the FAE is incoherent.

3.1.3.3. Empirical incoherence. The third argument refers to the presumably large effect sizes obtained from situational variations. These effects are often portrayed as much larger

than the effects of individual differences in personality, a difference that naïve perceivers presumably fail to appreciate. This argument also presupposes a hydraulic causal model. If, for example, a personality variable correlates .40 with a behavioral outcome, it is frequently assumed that the situation must explain the remaining 84% of the variance (Kunda & Nisbett 1986; Mischel 1984). Because there is no well-developed taxonomy of situations or accepted set of situational variables, variance is assigned to situations by default – they get whatever is left after the effects of personality variables are accounted for (Funder 2001a). But it is as plausible to assign variance not explained by any particular personality variable to other personality variables that were not measured, as it is to assign it to situational variables that were also not measured (Ahadi & Diener 1989). Despite the rhetoric about the “power of the situation,” very little is known about the basis of that power or its real amount.

It is not even clear that the effects of situational variation are greater than the effects of dispositional variation. When Funder and Ozer (1983) recalculated the effect sizes for situational variables such as the distance of the experimenter and the victim in the Milgram experiment, the number of bystanders and degree of hurry in the studies on bystander intervention, and the level of incentive offered to induce attitude change through cognitive dissonance, they found correlations ranging between .30 and .40. These values were similar to the correlations typically found between individuals' behaviors across different situations (Funder 1999; Funder & Colvin 1991) and the notorious “personality coefficient” that situationists consider to be the upper limit for the effect of personality on behavior (Mischel 1968; Nisbett 1980). In a sophisticated analysis, Kenny et al. (2001) compared person, situation, and interaction effects directly, and found the person effect to be the largest. Thus, the third account of the FAE violates the coherence criterion in that it relies on empirical data that either do not support or that reverse the commonly assumed direction of the error.

3.2. Theoretical shortcomings

An exclusive focus on norm violations discourages cumulative research and theoretical development. Misbehaviors and errors tend to be viewed in isolation; they have narrower implications than is often assumed; and they do not contribute much to theories of the whole range of behavior.

3.2.1. Isolation. To the extent that behavioral social psychology becomes the study of misbehavior, and cognitive social psychology becomes the study of judgmental shortcomings, the field is reduced to a catalog of things people do badly. Each misbehavior generates its own explanation, but these explanations are seldom integrated, much less drawn from broader theories of behavior or cognition. Many of the errors listed in Table 1 are associated with particular groups of investigators or even single psychologists. This isolation facilitates a profusion of overlapping labels, it allows the discovery and survival of mutually contradictory errors, and it discourages the development of overarching theory (Kruglanski 2001). In Asch's (1987) words, “the current expansion [comes with] a shrinking of vision, an expansion of surface rather than depth” (p. x).

3.2.2. Limited implications. Errors in judgment are studied because of the belief, often explicitly expressed, that they have important implications for evaluating human reasoning. But some errors reveal little more than the difficulty of the presented task (Funder 2000). People who are good at solving problems on the Scholastic Aptitude Test (the SAT-1, which is highly saturated with conventional IQ) are also good at solving many of the standard problems of judgment and decision making (Stanovich & West 2000). Indeed, many problems used in heuristics-and-biases studies would be suitable for use as SAT items because they correlate as well with total scores as do individual SAT items themselves.

Consider the implications of this psychometric finding. To detect differences among highly able test takers, the Educational Testing Service (ETS) has written many difficult SAT items without claiming to uncover systematic and discrete cognitive deficits (cf. Stanovich & West 2000, p. 646). By the standards of heuristics-and-biases research, however, each difficult SAT item should merit a separate article presenting the discovery of a cognitive flaw.

3.2.3. Incompleteness. Because findings of error are seen as demanding of explanation, whereas rationality may merely be assumed, theoretical explanations of error do not even seek to explain the entire range of performance. They concern the participants who get the answer wrong, not the ones who get it right. One way to overcome this limitation is to examine the relationship between bias and accuracy. Often this relationship is positive. Projected consensus estimates (Hoch 1987), self-enhancement (Krueger & Mueller 2002), and overattribution (Block & Funder 1986) benefit the perceiver most of the time. These examples are no exceptions. The hindsight effect (Hoffrage et al. 2000), the positive testing strategy (Klayman & Ha 1987; Oaksford & Chater 1994), the halo effect (Borman 1975), overconfidence (Dawes & Mulford 1996; Erev et al. 1994), and various heuristics in probability estimation (McKenzie 1994; Moldoveanu & Langer 2002) have similar advantages.

The idea that judgmental biases serve adaptive functions vindicates Egon Brunswik's (1956) view that social perception operates through a lens of probabilistically valid cues and probabilistically correct use of these cues. By and large, cues are valid enough and perceivers use them well enough to achieve a fair degree of judgmental accuracy. Brunswik's approach distinguishes between adaptive errors and harmful mistakes (see Funder 1987 for details on this distinction). As noted earlier, visual illusions are also erroneous interpretations of experimental reality, but they reveal underlying mechanisms of the visual system that yield accurate and adaptive results under most ecologically representative conditions (e.g., Vecera et al. 2002). If these illusions were eliminated from human perception, perceptual accuracy would surely get worse, not better.

A frequent charge is that people "over- or under-apply particular rules or use shortcut 'heuristics' instead of relying on normative rules" (Jacobs & Klaczynski 2002, p. 146). The explanatory power of this charge depends on whether people can be expected to know when and how to switch from a heuristic mode to a more formal mode of thinking. Often, no such meta-decision can be made without running into the paradox that the switch cannot be made without foreknowledge of the answer (Krueger et al. 2003). Suppose people know that most distributions of social (e.g.,

self-esteem) or academic (e.g., grades) characteristics are negatively skewed. The heuristic expectation of being better than average would minimize the aggregated error, although it would produce some false positives (Einhorn 1986). To avoid overgeneralization, the person would have to know by non-heuristic means on which dimensions he or she is merely average or worse. If the person effortfully recruits such knowledge for each dimension, the need to think heuristically never appears in the first place, but neither do its effort-saving advantages.

The heuristics-and-biases paradigm makes any benefits of heuristic strategies impossible to detect. When the stimulus is held constant, the data cannot show how accurate a person would be across stimuli or under more realistic circumstances. When the stimuli selected for research are limited to those for which use of the heuristic yields inaccurate results, it is tempting – and rather typical – to conclude that the heuristic represents a systematic flaw (Kühberger 2002).

When multiple stimuli are employed, statistical analysis typically focuses on bias to the exclusion of accuracy. When bias is expressed by a partial correlation between heuristic cues and judgment after the reality criterion is controlled, it is impossible to estimate how much a heuristic contributes or detracts from accuracy. All that can be said is that "all heuristics – by mathematical necessity – induce weighting biases" (Kahneman 2000, p. 683). If only the partial correlation between the bias cue and the prediction is tested for significance (with the reality criterion being controlled), the utility of the bias cue for the improvement of accuracy necessarily remains unknown.

4. Back to balance

We do not question all research on problematic behaviors or flawed reasoning. We do suggest, however, that social psychology is badly out of balance, that research on misbehavior has crowded out research on positive behaviors, that research on cognitive errors has crowded out research on the sources of cognitive accomplishment, and that the theoretical development of social psychology has become self-limiting. We now offer empirical, analytic, and theoretical recommendations to redress the current imbalance.

4.1. Empirical suggestions

4.1.1. De-emphasize negative studies. If current trends continue, new entries for Table 1 will continue to appear, and indeed several have been added since this article was first submitted. Many of these will be old biases resurfacing under new names or new biases contradicting old ones. Even to the extent that new biases are discovered, one could question what new exhibits in the Museum of Incompetence will contribute to our understanding of social inference. A slowing rate of output of error-discovery would not only stem the fragmentation of the literature, but also free journal space for studies that examine errors in the context of accomplishments and vice versa.

Not all current research is negative. A "positive psychology" movement has begun to focus research on human strengths and abilities to cope and develop (e.g., Diener & Biswas-Diener 2002; Diener & Seligman 2002; Lyubomirsky 2001). Though important, increased research on positive topics will be an insufficient remedy. A one-sided re-

search emphasis on positive behavior, perhaps complete with null hypotheses where *bad* behavior represents the null to be disconfirmed, might eventually generate problems parallel to those besetting the one-sided emphasis on negative behavior. We recommend that the *range* of behavior be studied, rather than showing that behavior is bad – or good – “more often than most people would expect.”

In the area of judgment and decision making, Gigerenzer and colleagues (Gigerenzer et al. 1999) find that heuristics can “make us smart” as well as produce error. A movement in cognitive psychology, parallel in some ways to positive psychology, has also begun to question the logical and empirical bases for studying errors (Cohen 1981; Moldoveanu & Langer 2002). Some of the arguments have been “Panglossian” by suggesting that psychologists have no grounds for evaluating the judgments of their fellow humans (Cohen 1981), whereas others have suggested that certain imputations of error are themselves erroneous (Dawes 2000; Lopes 1991).

4.1.2. Study the range of behavior and cognitive performance. As an example of a more comprehensive approach, consider Stanovich and West’s (2000) study of individual differences in judgmental performance and general cognitive ability. Even smart research participants get certain problems wrong, which suggests that these were simply too difficult or perhaps even incorrectly keyed. More importantly, Stanovich and West placed both normative and counter-normative decision making in a common framework to explain when normative decisions might be expected, what psychological processes produce them, and the prescriptive status of the normative model employed.

For another example, Paul Ekman (1991/1992) examined people’s ability to discriminate between spontaneous and staged nonverbal behavior. In one study, observers were better able to detect concealed emotions by attending to postural cues instead of facial expressions (Ekman & Friesen 1974). Our point is not that this work was “positive” in any particular sense, but rather that it examined the conditions under which both failures and successes occur. Other studies have shown that people can form accurate impressions on the basis of minimal information (see Hall & Bernieri 2001 for a survey). Short soundless videos suffice (Ambady et al. 1995), as do handshakes (Chaplin et al. 2000), or even a mere peek at a person’s office or bedroom (Gosling et al. 2002). Studying “empathic accuracy,” Ickes (1997) explored the conditions under which people can intuit the thoughts and feelings of their interaction partners. Funder (1995a) and Kenny (1994) have evaluated the accuracy of judgments of personality with criteria such as interjudge agreement and correct behavioral prediction (see also Diekmann et al. 2002 for an innovative study on the accuracy of gender stereotypes).

An important property of these research programs is that they allow the possibility of accurate judgment. There is a criterion – a target is lying or not, thinking a particular thought or not, characterized by a particular trait or not – that the participants might successfully predict. This contrasts with the artificial design of many error studies where nothing true can possibly be said about the target. Consider the study of expectancy effects. In the paradigmatic (non-ecological) error study, a participant such as a teacher receives false information about the potential ability of some students. The classic result is that such false expectancies

predict modest increases in test scores (Rosenthal 1994). Arguing from an ecological perspective, however, Jussim (1991) asked how expectancies typically arise, and whether their predictive utility is necessarily false in the sense of being self-fulfilling prophecies. Indeed, most teachers’ expectancies are based on valid information, and the effect of erroneous expectations is comparatively small (see also Brodt & Ross 1998 for the predictive utility of stereotypic expectancies). Again, we do not merely wish to emphasize the positive conclusion, but the availability of research designs that allow participants to be correct.

4.2. Analytic suggestions

4.2.1. Handling NHST with caution. The proliferation of documented errors owes much to the ritual use of NHST. Skepticism about the value of NHST has a long history (Harlow et al. 1997), and these concerns apply a fortiori to a value-laden field such as social psychology. The method’s most serious shortfall is that by misapplying *modus tollens* to inductive inferences, NHST misses its own ideal of rationality. According to this ideal, a null hypothesis (e.g., of rational thinking) may be rejected if the data are improbable under that hypothesis. Logically, knowing that P implies Q means that $\neg Q$ implies $\neg P$. When the consequence is denied, the antecedent cannot be true. If the null hypothesis suggests that certain data are improbable, however, finding such data does not guarantee that the null hypothesis is improbable (Cohen 1994). Because knowledge of just that improbability is needed for the rejection of the hypothesis, NHST does not deliver what researchers want. It does not provide the inverse leap from data to hypothesis. As a consequence, reliance on NHST can generate contradictory claims of bias, each apparently supported by improbable data.

Because we do not expect NHST to fall out of favor, we emphasize the need to understand its limitations and to use additional data-analytic strategies. Several commentators have proposed an integration of NHST with Bayesian concepts of hypothesis evaluation (Krueger 2001; Nickerson 2000; Task Force on Statistical Inference 2000). The Bayesian approach acknowledges that data do not speak directly to the truth or falsity of a hypothesis unless there is a prior theory or expectation about the chances of the hypothesis to be true. If such expectations are specified, Bayes’ Theorem gives a posterior probability for each hypothesis. The differences between prior and posterior probabilities then reflect how much has been learned from the evidence, and research becomes an incremental learning process. The following examples illustrate how the Bayesian approach combines prior expectations with significance levels to allow the estimation of the probabilities of the hypotheses in the light of the data.

4.2.2. Bayesian inferences. Consider the simplest case, in which a researcher is not able (or willing) to advance any hypothesis but the null hypothesis, H_0 . The alternative hypothesis, H_1 , may then be the empirically observed effect size. The probability of the data under the null hypothesis, or data more extreme, is the familiar significance level derived from the statistical test, $p(D|H_0)$. The probability of the data, or data more extreme, under the alternative hypothesis is $p(D|H_1)$. If the observed effect size stands in as the alternative hypothesis, this probability is .5 because the distribution is centered on the observed effect after the fact (other cases

Table 2. Bayesian evaluation of hypotheses: The effects of varying priors

Hypothesis	Case 1: Uniform Priors		
	p(H _i)	p(D H _i)	p(H _i D)
-.4	.2	.00000005	8.33E-08
-.2	.2	.00003	.000005
0	.2	.05	.083
.2	.2	.5	.833
.4	.2	.05	.083

Hypothesis	Case 2: Bias as a “risky” hypothesis		
	p(H _i)	p(D H _i)	p(H _i D)
-.4	.125	.00000005	6.67E-08
-.2	.125	.00003	.000004
0	.5	.05	.267
.2	.125	.5	.667
.4	.125	.05	.067

Hypothesis	Case 3: Bias as a “safe” hypothesis		
	p(H _i)	p(D H _i)	p(H _i D)
-.4	.125	.00000005	2.38E-08
-.2	.125	.00003	.000001
0	.125	.05	.024
.2	.5	.5	.952
.4	.125	.05	.024

are discussed below). Bayes’ Theorem then combines expectations with evidence to yield what the researcher wants, namely, the probability of each hypothesis, given the data. This probability is the product of the probability of the data, given the hypothesis and the prior probability of the hypothesis divided by the overall probability of the data, or

$$p(H_0|D) = \frac{p(D|H_0)p(H_0)}{\sum p(D|H_i)p(H_i)}$$

When studies are selected for their statistical significance – as they often are when articles on bias are published – Bayesian posterior probabilities tend to be higher than significance levels. The reason for this is twofold. First, good studies are supposed to be risky, which means that the prior probability of H₀ is assumed to be high. Second, significance levels are positively but imperfectly correlated with the posterior probabilities of the null hypothesis across studies. By regression to the mean, the posterior probability of the null hypothesis is less extreme than the significance level. Thus, Bayesian inferences are conservative because they take prior expectations into account.

A crucial advantage of the Bayesian approach is that data analysis in an individual study can reflect the maturity of the field. Theoretical advances and past research evidence can inform the selection of contending hypotheses and the prior probabilities assigned to them. Research can then move away from rejecting individual hypotheses with data that depart from it in whichever direction, and thus away from incoherent claims. Consider three examples with multiple hypotheses and varying priors (Table 2). The examples

share the assumption that the observed effect size is .2 in standard units, and that the significance level, p(D|H₀) is .05. The other values of p(D|H_i) can be found in tables for cumulative probabilities in normal distributions. Bayes’ Theorem then gives the posterior probability for each hypothesis (see rightmost column in Table 2).

In the top panel of Table 2, the five possible hypotheses start out equiprobable, suggesting a novel area of investigation where theory is tentative and the empirical base is thin. The posterior probability of the null hypothesis (.08) is not as low as the significance level, and the posterior probability of the observed effect is .83 instead of .92 because there are several alternative hypotheses. If applied to the study of bi-directional bias, this example shows that data indicating, say, a positive bias, also reduce the probability of negative bias. Here, the probability that the true effect is -.2 or -.4 has decreased drastically.

The probabilities in the center panel reflect the assumption that the hypothesis of no bias is as probable a priori as the combined hypotheses of bias. This assumption is implicit in much error research. The demonstration of an error is considered important because the implicit priors suggest that such a demonstration would be difficult to obtain. If such an expectation were made explicit, however, one would have to acknowledge that the posterior probability of rationality did not shrink much (here it dropped from .5 to .267). A Bayesian approach prevents the researcher from having it both ways. *A Bayesian cannot claim that rationality is a strong contending hypothesis and then reject it on the grounds of significance alone.*

The probabilities in the bottom panel reverse the situation presented in the center panel. The observed effect size has a prior of .5, and the remaining priors are equally shared by the other four hypotheses. This example reflects a more mature area of study because researchers already expect to find what they end up finding. The incremental benefit of each new study diminishes as a field matures. Looking back at the three sample cases, the average difference between prior and posterior probabilities was highest for uniform priors (M = .25), intermediate for the high prior of rationality (M = .22), and lowest for the case in which a certain bias was already expected (M = .18).

4.2.3. Advantages of the Bayesian approach. The Bayesian approach encourages investigators to be clear about their expectations. They can no longer use NHST as a surrogate for theory (Gigerenzer 1998), knock down the null hypothesis as a straw man, or treat bias as a foregone conclusion (Krueger 1998c). Bayesianism permits the integration of new evidence with theory and past research even at the level of the individual study. This may prove to be a crucial advantage because some critics of NHST have proposed that all evaluation of hypotheses be ceded to meta-analyses (Schmidt 1996). However, this suggestion creates a social dilemma for individual researchers. If a final judgment regarding the existence of a phenomenon can only be reached by aggregating the results of multiple studies, there is no incentive for a researcher to gather data. Rather, the most effective strategy would be to hope that others will do the studies, wait until enough studies have accumulated, and then do the meta-analysis before anyone else does. With the Bayesian approach, the lessons from the past are not set aside to be rediscovered by meta-analysts. Individual researchers who replicate their work can quantify its di-

minishing returns and reach a rational decision for when to stop.

4.3. Theoretical suggestions

4.3.1. Explanations for the complete range of performance. Our central recommendation is that empirical work and theoretical modeling address the whole range of performance, seeking to understand and explain how both positive and negative phenomena may arise and how they are interrelated. Of course, this recommendation is not new. In the preface to the first edition to his social psychology text, Asch (1952) noted that “social relations at the human level, even the simplest, require emotional and intellectual capacities of a high order. I have tried to call attention to their productive character, which is to my mind also necessary for the understanding of destructive social processes” (p. x).

4.3.2. Social behavior. With Asch’s evenhanded stance in mind, the classic behavioral studies can be recast in their original light. We note the high rate of independence in Asch’s own work, and the sometime successful struggle towards resistance in Milgram’s studies. Milgram emphasized these successes by opening his classic movie with the portrayal of several disobedient subjects. His theoretical account referred to Lewin’s idea of competing “force fields” emanating from the victim and the experimenter. Each participant’s ultimate behavior then revealed the relative strength of these external forces as well as the relative strength of the competing dispositions internal to *that* person. The notable virtue of this approach was that it aimed not at the phenomenon of obedience per se, but at the differences between circumstances under which behavior is more likely to be influenced by the wishes of the experimenter or by the needs of the victim.

In Darley and Latané’s work on bystander intervention, the tone was not as balanced. It emphasized how people’s behavior often violated Biblical norms, but the research also included conditions that increased the rate of intervention (Darley & Latané 1968). Following the lead of the pioneers, comprehensive theoretical accounts of behavior should address its moderators, and not just the surprising, attention-getting, and simplistic message that people can be made to behave badly.

No theoretical account of a range of behavior is complete without a cost-benefit analysis. For example, most readers of this article probably would stop at an intersection if told to do so by a traffic officer, perhaps even if the reasons for the stop are not obvious or appear to be wrong (e.g., when no traffic is on the cross street, while impatient drivers accumulate behind you). Why is this? A superficial reading of Milgram’s work might suggest that people blindly obey anyone who looks like an authority figure. But at least two other reasons suggest themselves. First, some instructions from authority are based on expert knowledge. The traffic officer may know that a fire truck is on the way. Second, and more generally, obedience to legitimate authority is an important part of social order and should be withdrawn only under compelling circumstances. Surely, a command to give lethal shocks to an innocent victim is one of those circumstances but notice that the issue is not that obedience is a bad thing, but rather of where to draw the line (and how to know where that line is). Similar accounts could be of-

fered for conformity (as the behavior of others can be an important source of information about what is safe and appropriate to do), and bystander intervention (as it may be rational to hold back from immediate intervention while assessing the legitimacy of the need and one’s own capabilities to help).

After social psychology began to focus its interest on cognitive processes, few modern classics have been added to the canon of misbehavior. But to the extent that such studies are done, it is important to include opportunities for participants to do the right thing, to interpret the findings in terms of the circumstances that produce the whole range of behavior, and to evaluate the costs and benefits implicit in the behavioral choices. Admittedly, this strategy will produce fewer counter-intuitive or “cute” findings, but it would yield more solid and informative research. In the meantime, it will be helpful for social psychologists to broaden how they think about, and teach, the landmark studies in their field. It might even be salutary if occasionally a social psychologist were to object when NBC tries to turn classic research into *Candid Camera*-like demonstrations of how people are funny.

4.3.3. Social cognition. The road to reform may be especially difficult in the field of social cognition, which suffers from a particular addiction to counter-intuitive findings. All inferential errors are counter-intuitive in the sense that they show ordinary inferences to be wrong. This is the most important reason, we suggest, why lists like Table 1 continue to grow, even as entries duplicate and contradict each other. Overcoming this addiction will be difficult and will require two kinds of reform.

4.3.3.1. Consider adaptive mechanisms underlying error. First, as in the case of behavioral social psychology, some of the classic studies could be appreciated in a new light. Researchers might follow the example of research on visual perception, as conducted after 1896, and entertain the possibility that the psychological mechanism underlying an apparent inferential error might lead to adaptive results outside of the laboratory (Evans & Over 1996; Funder 1987).

Although some researchers in the heuristics-and-biases tradition have acknowledged this idea, experimental demonstrations rarely show how heuristics can produce accurate judgments. Even more notable is the way that their research is so widely interpreted as implying that human judgment is fundamentally erroneous (e.g., Shallice 1991). We submit that few readers of this literature have carried away the dominant message that representativeness, availability, or the fundamental attribution error are essential components of adaptive social cognition. But of course, to the extent that these and other heuristics have been correctly characterized, they probably are. Like processes underlying the Müller-Lyer illusion, the heuristics that drive human inference are more likely to be part-and-parcel of adaptive cognition than arbitrary design flaws.

In an extensive research program, Gigerenzer and his research group (Gigerenzer et al. 1999) have explored the conditions under which “fast and frugal” heuristics can, like mechanisms of visual perception, lead to interesting errors while yielding many adaptive and accurate results in the complex, chaotic, and consequential settings of the real world. This position has met with some resistance (e.g., Margolis 2000), despite the assertions elsewhere in the er-

ror literature itself that heuristics are not necessarily maladaptive.

A related theoretical development is evolutionary psychology (e.g., Buss & Kenrick 1998; Klein et al. 2002), which assumes that the basic mechanisms of human cognition, like the basic mechanisms of human anatomy and physiology, evolved as adaptations to life in the ancestral environment. The evolutionary approach suggests, for example, that it is adaptive to predict future events on the basis of apparent similarities with current circumstances (one version of the representativeness heuristic) or to pay attention to salient and vivid information (the availability heuristic). A theoretical challenge for both social cognition and evolutionary psychology is to work towards greater convergence, and we expect that this challenge will be met. As we have shown, important “errors” such as consensus bias, self-enhancement, and even the so-called fundamental attribution error can lead to accurate judgments and positive outcomes.

4.3.3.2. Explain error and accuracy in the same framework.

The other theoretical reform for cognitive social psychology is that models be constructed which not only explain how errors occur but which also account for accurate judgment. For the area of personality judgment, Funder (1995a) called for an “accuracy paradigm” to complement the dominant “error paradigm.” The accuracy paradigm identifies accomplishments of social judgment by using correspondence criteria (Hammond 1996) rather than departures from normative models by coherence criteria. For example, rather than focusing on how people distort artificially input stimulus information, a study might evaluate the circumstances under which participants manifest inter-rater agreement in their personality judgments of real acquaintances (e.g., Funder et al. 1995), or are able to predict the behavior of themselves or others (e.g., Kolar et al. 1996; Spain et al. 2000).

Kenny (1994) presented a Weighted Average Model of social judgment (WAM) to explain the basis of inter-judge agreement in personality rating. The components of agreement include culture, stereotypes, communication, common observation, and personality, and their sum determines the degree of inter-judge agreement that may be found. A related approach, the Realistic Accuracy Model (RAM; Funder 1995a; 1999), assumes that personality characteristics are real, and it seeks to explain how humans

manage to evaluate the attributes of others correctly at least some of the time. Figure 2 shows the process and reveals the ancestry of this theory in Brunswik’s (1956) lens model of perceptual judgment.

First, the person who is the target of judgment must emit a cue, usually a behavior, that is *relevant* to the existence of the trait in question. A courageous person must do something brave, a smart person must do something intelligent, and so on. Unless an attribute of personality is manifested in behavior, an observer cannot judge it accurately. Second, this relevant behavior must be *available* to the judge. A highly relevant behavior performed in the next room, with the door closed, obviously is of no help to the judge’s accuracy. Less obviously, different behaviors are available to co-workers than to spouses, to parents than to children, and therefore different others will be differentially accurate across the traits that vary in their availability across these contexts. Third, the relevant, available behavior must be *detected*. A social perceiver may be perceptually acute and paying close attention, or distracted, preoccupied, or simply imperceptive. Finally, the relevant, available, and detected information must be correctly *utilized*. The judge must interpret the information in light of past experience and general knowledge. An intelligent and experienced judge can be expected to do this well, but if past experience or knowledge is misleading or if the judge applies it poorly, accuracy will be low.

This simple model has several implications. First, it describes not just a cognitive process of judgment, but rather, the interpersonal process necessary for *accurate* judgment. Second, it implies that accuracy is a difficult and remarkable achievement. A failure at any of the four stages of accurate judgment will dramatically reduce accuracy as failures at each stage combine multiplicatively. Third, the model implies that the traditional paradigm of social psychology addresses, at most, half of what is required to understand accuracy. The paradigmatic study presents social stimuli directly to participants, thus bypassing relevance and availability completely, and largely bypassing the task of cue detection. Traditional studies of social cognition concern the utilization stage.

A fourth implication is that although the RAM is optimistic in the sense that it describes the route to successful social judgment, it is not Panglossian because it recognizes the barriers between judgment and accuracy. In particular, it can incorporate the four moderators of accurate judg-




Figure 2. Realistic Accuracy Model (RAM).

ment identified in prior research – properties of the judge, target, trait, and information (Funder 1995a) – and suggest new ones. Some judges are inattentive or cognitively inept, for example. Some targets emit few relevant behaviors – because they are inactive or, in some cases, even deceptive. Some traits are difficult to judge because they are available in only a few contexts, or because their cues are difficult to detect. The information itself may be inadequate in the sense that the judge has not had enough or varied enough experience with the target for sufficient cues to be available on which to base a reasonable judgment. In sum, the RAM describes how people manage to make some of their most difficult judgments, addresses the many ways the process can go awry, and points to four specific stages where efforts to improve accuracy might productively be directed.

5. Conclusion

Discerning the pathological element in the typical is the social psychologist's privilege.
—Alexander Mitscherlich

I have made a ceaseless effort not to ridicule, not to bewail, not to scorn human actions, but to understand them.
—Baruch Spinoza

For decades, social psychology has emphasized how human behavior falls short of ethical standards and moral imperatives. When research attention shifted to judgment and inference, violations of norms for rational thought took center stage. As a final, and perhaps inevitable merger of these two intellectual strands, we now find that ordinary people's moral judgments are being exposed as both hypocritical (Batson et al. 2002) and irrational (Carlsmith et al. 2002). As psychoanalyst Alexander Mitscherlich observed, the fascination with the negative has turned social psychology into a psycho-pathology of everyday life. A more balanced, full-range social psychology, as we tried to sketch it, would be more sensitive to Spinoza's perspective. While seeking not to pass rash or harsh judgments on research participants (and the populations they represent), research in the spirit of Spinoza would seek to understand how people master difficult behavioral and cognitive challenges, and why they sometimes lapse. Ultimately, a more realistic and thus a more compassionate view of human nature may result.

This shift in perspective need not entail a return to the chimerical ideal of a value-free science. Social psychologists have always, at least implicitly, acknowledged which phenomena within their domain they consider desirable and which they consider undesirable. We think that this is as it should be. But we propose that alternative models of rationality be compared carefully, that values be discussed openly, and that the social animal not be judged in trials designed to establish guilt.

In our effort to bring about a shift in perspective, we presented a critique of the current negative paradigm in rather pointed fashion. We realize that social psychologists are often accused (and accuse themselves) of being overly harsh when evaluating work in their own field. Kelley (2000), for example, attributed the marginal status of social psychology among the sciences in part to its tendency for self-loathing. In contrast, our critique was meant to be constructive. We doubt that the traditional pan-critical approach to human

behavior and thinking can sustain the field. Instead of demanding respect, it is likely to backfire.

ACKNOWLEDGMENT

The research of the second author is supported by National Institute of Mental Health grant R01-MH42427.

NOTES

1. In 18 experimental conditions, compliance ranged from 93% (when the participant did not have to administer shocks personally) to 0% (when two authorities gave contradictory orders, when the experimenter was the victim, and when the victim demanded to be shocked). In the two best-known and most frequently portrayed conditions, when the experimenter was present and the victim could be heard but not seen, the obedience rates were 63% (at Yale) and 48% (at "Research Associates of Bridgeport"). Across all conditions the average rate of compliance was 37.5% (Milgram 1974, Tables 2, 3, 4 and 5).

2. The proportion of interveners in the lowest-helping conditions of the Darley and Latané (1968) and Darley and Batson (1973) studies were, respectively, 31% and 29%; across conditions the average proportions were 59% and 40%.

3. Over the years, cognitive-social psychology and the psychology of judgment and decision-making (JDM) progressively interpenetrated each other. In the classic collection of papers on heuristics and biases (Kahneman et al. 1982), eight of the 35 contributions were authored by social psychologists. Twenty years later, they account for half of the 42 contributions (Gilovich et al. 2002).

4. Without this further assumption the phenomenon would have to be renamed.

5. The same effect was earlier labeled the "correspondence bias" (see Gilbert & Malone 1995) but the more evocative "fundamental" label has come to predominate.

6. In the noncoerced condition the attributed difference in attitudes was 42.2 points; in the coerced condition the difference was 21.2 points.

7. Sheldon and King (2001) reported that an OVID search on the terms *error* and *bias* yielded more than twice as many hits as the terms *strength* and *virtue*.

8. A recent exception and a potential harbinger of change was an article that included a thorough examination of boundary conditions under which biases are found. As the authors commented, they did "not wish to argue that the [bias under discussion] is a poor strategy" (Goodwin et al. 2002, p. 241). It remains to be seen whether secondary accounts of this research will emphasize these boundary conditions and adaptive possibilities, or simply the finding of bias itself.

9. Intelligence and general well-being are two more examples of variables that probably have modal values higher than the arithmetic mean.

10. See, for example, Klar and Giladi's (1997) report on the "Everyone-is-better-than-average effect." Although most participants recognized the definitional truth that on average, people are average, the significant minority that erred, erred in the same direction, thereby yielding a difference between the average judgment and the modal judgment.

11. Machiavelli (1513/1966) noted that "Without an opportunity [the abilities of Moses, Cyrus, Romulus, and Theseus] would have been wasted, and without their abilities, the opportunity would have arisen in vain" (p. 26). According to an apocryphal story, a jealous Serifotis once told Themistokles that he, Themistokles, was famous only because he was an Athenian. The great strategist concurred and observed that he would be as obscure if he were a Serifotis as the Serifotis would be if he were an Athenian.