
Peering Into the Bias Blind Spot: People's Assessments of Bias in Themselves and Others

Joyce Ehrlinger

Stanford University

Thomas Gilovich

Cornell University

Lee Ross

Stanford University

People tend to believe that their own judgments are less prone to bias than those of others, in part because they tend to rely on introspection for evidence of bias in themselves but on their lay theories in assessing bias in others. Two empirical consequences of this asymmetry are explored. Studies 1 and 2 document that people are more inclined to think they are guilty of bias in the abstract than in any specific instance. Studies 3 and 4 demonstrate that people tend to believe that their own personal connection to a given issue is a source of accuracy and enlightenment but that such personal connections in the case of others who hold different views are a source of bias. The implications of this asymmetry in assessing objectivity and bias in the self versus others are discussed.

Keywords: *bias; social judgment; self-perception; social perception*

The controversial outcome of the 2000 U.S. presidential election produced heated debate and accusations of bias as the courts handed down their rulings and political commentators offered their opinions on the issues at hand. Would hand counting of ballots in Florida produce a fairer and/or more accurate count than machine counting? Had a complicated ballot design robbed voters of the opportunity to make their vote count? Did the results suggest that it was time to abandon the Electoral College in favor of a direct vote? More pointedly, were the decisions of the Supreme Court and the Florida electoral officials fair, or an exercise in political partisanship?

Public opinion polls (Gallup, 2000) made it clear that the views Americans expressed on these issues were highly correlated with political affiliation. Ninety-four percent of George W. Bush supporters thought the

Supreme Court's ruling on the Florida recount was fair and justifiable, whereas only 17% of Al Gore supporters thought so. By contrast, 66% of Al Gore supporters, but only 31% of George Bush supporters, thought that members of the Supreme Court had been influenced by their "personal political views." When it came to the crucial actions of Florida secretary of state Katherine Harris, 79% of Bush supporters approved of the soundness of her performance—a sentiment shared by only 22% of Gore supporters. Ms. Harris herself expressed no doubts about her objectivity, insisting "I administered Florida's election law fairly, consistently, and evenhandedly throughout the controversial election" (Canedy, 2001).

In daily life, as in the aftermath of controversial elections, individuals must often assess the objectivity of opinions and judgments. They must determine whether a promotion decision was based on the relative merits of the rival candidates or the strength of their connections to the boss; whether a news report offered a balanced account of the issues at hand or reflected the news orga-

Authors' Note: Portions of this research were presented at the meetings of the Society for Personality and Social Psychology, Savannah, Georgia, 2002, and the American Psychological Society, Atlanta, Georgia, 2003. We thank Eric Hopson, Chris Minunni, and Rachel Smallman for collecting data and Jennifer Butler for comments on earlier drafts of this article. The research was supported by National Science Foundation grants SBR-9319558 and SES-0241638. Correspondence should be addressed to Joyce Ehrlinger, 271 Jordan Hall, Building 420, Department of Psychology, Stanford University, Stanford, CA 94305-2130; e-mail: joyce@psych.stanford.edu or tdl1@cornell.edu or ross@psych.stanford.edu

PSPB, Vol. 31 No. 5, May 2005 1-13

DOI: 10.1177/0146167204271570

© 2005 by the Society for Personality and Social Psychology, Inc.

nization's ideological leanings; whether a rosy economic forecast was based on steel-eyed analysis, wishful thinking, or even the self-interests of the forecaster. Individuals may be called upon to make assessments of objectivity not only about the views of others, but about their own as well. In deciding upon a future course of action, they must determine whether their assessments of a co-worker's efforts are tainted by envy, whether their reactions to a newscast are distorted by their own political leanings, or whether their guesses about how the future will unfold are biased by their hopes, fears, or formative experiences. In the four studies reported here, we expand upon previous research demonstrating that people are less likely to detect bias in themselves than in their peers and document two previously unexamined manifestations of this self-other asymmetry.

The Bias Blind Spot

A wealth of evidence suggests that judgments are often clouded by a number of cognitive and motivational biases (Gilovich, 1991; Kunda, 1990; Nisbett & Ross, 1980). Individuals consistently rate themselves above average across a variety of domains (Alicke, Klotz, Breitenbecher, Yurak, & Vredenburg, 1995; Dunning, Meyerowitz, & Holzberg, 1989), take credit for their successes but explain away their failures (Miller & Ross, 1975; Whitley & Frieze, 1985), assume they are more likely than their peers to experience the good things in life and avoid the bad (Weinstein, 1980), and tend to detect more support for their favored beliefs than is objectively warranted (Lord, Ross, & Lepper, 1979).

Given the pervasiveness of such optimistic and self-serving biases, it is hardly surprising that assessments of bias in the self versus others are often biased as well. People seem to suffer from a "bias blind spot," or the conviction that one's own judgments are less susceptible to bias than the judgments of others (Pronin, Gilovich, & Ross, 2004; Pronin, Lin, & Ross, 2002; see also Armor, 1999; Friedrich, 1996). People also believe they are more moral (Epley & Dunning, 2000) and less self-interested (Miller & Ratner, 1998) than others. Indeed, there is evidence that people may be *overly* cynical in this regard, predicting greater bias in other's judgments, on average, than proves to be the case (Kruger & Gilovich, 1999).

Some of the reasons why people might be more prone to see bias in others than in themselves are obvious. First, the desire to see oneself as above average on desirable attributes is likely to lead people to believe they are less subject than others to the influence of motives that might taint their judgments. Second, individuals tend to be naïve realists, believing their own understandings of the world are direct, unmediated perceptions of the way things are—and hence are inclined to view beliefs and

assessments that differ from their own as uninformed or biased (Pronin et al., 2004; Ross & Ward, 1996).

Two Criteria for Detecting Bias in Oneself and Others

The present research derives from a third source of the asymmetry in assessments of bias in oneself versus others—namely, differences in the strategies people use to determine whether a given judgment may have been unduly influenced by various motivational factors. Previous research has established the existence of two broad strategies for assessing bias (Pronin et al., 2004). First, individuals frequently consult their abstract theories of bias. For example, understanding that individuals with a vested interest in an issue are likely to view it through the prism of self-interest, people are apt to insist on full disclosure of such interests and to treat self-interested judgments with a healthy dose of skepticism. More generally, people are aware that human beings are motivated to seek pleasure and avoid pain, give heavy weight to their own needs and preferences, and have an arsenal of psychological defense mechanisms at their disposal. Such insights guide attributional analyses whenever contextual or situational cues suggest the possibility of bias. A claim that a contested electoral outcome was fair, a procedure for counting votes appropriate, or a judicial decision wise thus becomes subject to suspicions of bias when the claim is made by someone who was well satisfied with the outcome.

A second strategy for assessing whether a given judgment has been tainted by bias is to rely on the phenomenology of the person making that judgment. This strategy rests on the assumption that the various sources of bias leave some detectable trace—that one would have access to (and faithfully recognize and report) the tug of wishful thinking or the taint of self-interest. Such an assumption, even in the case of self-assessments, is shaky at best. The processes that give rise to most biases may leave no trace of their operation (Wilson, Centerbar, & Brekke, 2002).

In fact, introspection is as likely as not to lead to the conclusion that one acted *in spite* of one's preferences, not because of them. By contrast, protestations of objectivity by others are taken with a grain of salt. The lay psychologist knows that people are capable not only of trying to deceive others about their objectivity, but also of deceiving themselves.

Two consequences follow from these two strategies for assessing bias in the self versus others. First, introspecting in search of evidence of bias in the self is unlikely to yield such evidence. Thus, people will tend to believe that their assessments are relatively free of bias. Second, because people cannot, by definition, introspect into the minds of others, and because they regard others' introspection with skepticism, they will rely on

abstract theories to determine whether others are biased. As a result, people will readily infer bias when they see individuals make judgments that serve or coincide those individuals' self-interests.

The Present Research

This research was designed to investigate two important implications of the different strategies people employ to detect bias in themselves and others. First, because people are bound to rely more on introspective evidence when assessing whether they have been biased in a given instance than when assessing their more general susceptibility to bias, we predict that people will be more apt to concede that their judgment *in general* is subject to bias than they are to concede that any *specific* judgment has been tainted by bias. Studies 1 and 2 investigate this hypothesis.

Second, because introspection typically yields little evidence of bias in the self, we propose that such an "inward search" will lead people to believe they have acted independently of, or even in spite of, their self-interest or other motives, and not because of them. Thus, whereas others' identities and identity-based experiences will be seen as a source of *bias*, their own identities and experiences will be seen as a source of *enlightenment*. Studies 3 and 4 investigate this hypothesis.

STUDY 1

Past research indicates that people consider positive events more likely and negative events less likely to happen to them than to the average person (Weinstein, 1980). This unrealistic optimism is a robust phenomenon present across a broad range of content domains such as health (Price, 2001), academics (Eiser, Pahl, & Prins, 2001), relationship prospects (Murray & Holmes, 1997), gambling outcomes (Brandstaetter & Schwarzenberger, 2001), and the likelihood of completing tasks on time (Buehler, Griffin, & Ross, 1994). In this study, we investigate the extent to which people believe such assessments are swayed by wishful thinking.

Using a 2 × 2 design, we asked participants to consider either their own predictions or those of the average Cornell student and to consider either a specific set of predictions that had been made or predictions of this sort that are generally made. We predicted that the more participants were likely to rely on introspection, the less likely they would be to conclude that the judgments in question were biased. Because participants are most likely to introspect when considering their own judgments they just made, we expected these judgments to be seen as less tainted by bias than either their own general judgments or the judgments, of either sort, made by others.

A second, ancillary prediction follows directly from the previously documented difference in the reliance on introspection versus abstract theories of bias. One abstract procedure to determine whether a judgment is biased involves considering the extent to which it strays from the judgments others would make. Imagine, for example, that Susan thinks her daughter is unusually creative. One cue as to whether Susan is being objective is to consider what others think of her daughter's creativity. If Susan's judgment is similar to that made by others, one can be more confident that she is being objective than if her judgment is dramatically different from everyone else's. Because individuals tend to rely on abstract theories when introspection is a less viable strategy, they should pay greatest attention to whether a judgment strays from expectations when assessing bias in the judgments of others.

METHOD

Participants. Participants were 99 Cornell University undergraduates who were recruited from lecture courses in psychology and received extra credit for their efforts.

Procedure. Participants were randomly assigned to one of four conditions in a 2 × 2 design, being asked either to assess bias in the self or in another student and with respect to the sort of predictions that might be made or a specific set of predictions that had been made. Participants in the *self-specific group* were first asked to rate the likelihood that each of six events would happen to them relative to other Cornell students—receiving an attractive job offer before graduation, not finding a job for 6 months, suffering from lung cancer, owning a home, having an appealing job, and living past 80. All ratings were made on a scale ranging from "100% less likely than the average Cornell student" to "500% more likely than the average Cornell student." Participants in the *other-specific* condition were each yoked to a student in the self-specific condition and shown the predictions made by that student. Those in the *self-general* and *other-general* conditions considered the predictions that might be made if they or if the average Cornell student were asked to predict their future on these six dimensions.

Participants then read a description of how a person might be influenced by competing factors when making such predictions. It was explained that predictions of one's future might be influenced by a desire to make accurate, honest assessments, but also by a desire to think positive events likely and negative events unlikely to happen. Participants were asked to rate the degree to which the predictions they had just considered might be influenced by these competing pressures by placing an X on a 14.5-centimeter (cm) line anchored by "affected mainly by a desire to give an honest appraisal" and




Figure 1 Perceived bias in predictions of one's future, Study 1.

“affected mainly by a desire to think positive events likely and negative events unlikely to happen.”

To examine our ancillary prediction that assessments of bias based largely on abstract theories should be influenced by the extent to which a response is seen as normative, we collected data from an additional 18 participants who indicated how they expected the average Cornell student to rate the likelihood of the six future life events. They did so using the same scale as the participants themselves (i.e., from 100% less likely than the average Cornell student to 500% more likely).

RESULTS

Perceptions of bias. As Figure 1 illustrates, participants' ratings of bias conformed largely to our predictions. Participants thought their own specific assessments were the least biased and the hypothetical assessments of the average Cornell student were the most biased. A 2 (self vs. other) \times 2 (abstract vs. specific) ANOVA of participants' ratings of the degree to which the pertinent forecasts were biased yielded two significant main effects. As expected, participants thought their own forecasts were (or would be) less tainted by wishful thinking than the predictions made by the average Cornell student, $F(1, 95) = 15.70$, $p < .0001$.¹ Participants also thought that forecasts in general are more tainted by bias than a specific set of forecasts, $F(1, 95) = 8.78$, $p < .005$.

To test the significance of our predicted pattern of ratings, we subjected participants' bias ratings to a planned contrast with the self-specific condition assigned a weight of -3 and the other conditions assigned weights of $+1$. This contrast revealed that the lowest attributions of bias were made when participants were most likely to introspect—for specific predictions that the participants themselves had just made, $F(1, 95) = 14.29$, $p < .001$.

As a test of whether the tendency to consult abstract theories of bias predicts the extent to which forecasts are seen as biased, we utilized the ratings made by our additional group of participants. Recall that these participants indicated how they expected the typical Cornell student to respond when asked to rate the likelihood of experiencing the six target events. We used these ratings to compute a measure of how much each prediction in the main sample (in the self-specific and other-specific conditions) deviated from expectation by computing the difference between that rating and the average expected rating for that item. We then averaged the difference scores across the six items to create a measure of how much each person's predictions deviated from what might be expected. As hypothesized, perceptions of bias in another's predictions were highly correlated with the degree to which those predictions deviated from the expected norm, $r = .45$, $z(25) = 2.27$, $p < .05$. When determining whether the predictions they had just made themselves were influenced by bias, however, participants paid no attention to how much their estimates deviated from the norm, $r = -.01$, $z(26) = -.03$, $p = ns$. Thus, people appear to consult abstract theories of bias when assessing the extent to which others' assessments are biased but introspect when considering their own.

DISCUSSION

Study 1 provided three noteworthy results. First, it provided additional evidence of the previously documented tendency for people to see bias more in the judgments of others than in their own judgments (Pronin et al., 2002). Second, it provided additional support for the contention that this self-other difference arises in part from the different strategies people use to ascertain whether a judgment is biased (Pronin et al., 2004). That is, participants were shown to rely on information about how much a set of predictions strays from what is expected only when they cannot introspect—when assessing bias in the predictions made by another. Third, this study provided evidence for an important implication of the tendency for people to ascertain whether they have been biased by looking inward to find evidence of bias. Because one's own *particular* judgments are the ones most likely to prompt introspection, and because introspection is unlikely to turn up evidence of bias, one's own specific judgments are least likely to be thought of as biased—precisely the pattern of results we observed in this study.

Although we obtained results consistent with our predictions, one might be concerned that we did not directly demonstrate that individuals rely heavily on introspection to assess bias in their own judgments but rely more on abstract theories to assess bias in others. Study 2 was designed to deal with this objection.

STUDY 2

By definition, one can introspect to try to find evidence of bias only when it comes to judgments made by the self. What is less certain is whether introspection is the preferred strategy even in such cases. Thus, in Study 2 we examined more directly the degree to which individuals rely on introspection versus abstract theories of bias when deciding whether a judgment has been unduly influenced by goals or motives. We asked participants to report the degree to which they introspected and consulted abstract theories to determine whether particular assessments might be overly influenced by a desire to see oneself in a favorable light. We predicted that participants would rely more on introspection and less on abstract theories when assessing bias in the self than in others and when assessing bias in specific assessments than assessments in general, particularly for one's own judgments. As a result, we also predicted that as in Study 1, participants would most disavow the possibility of bias when considering their own specific judgments they recently rendered, seeing more bias in the judgments of others and their own judgments broadly and abstractly construed.

Participants were asked to assess the likelihood of bias in various trait ratings (either their own or the average Cornell student's). People's assessments of their own traits and abilities are often unrealistically favorable: On average, individuals tend to rate themselves above average on positive traits (Alicke et al., 1995). The mechanisms underlying this above-average effect tend to work outside of awareness and thus are unlikely to leave a trace that would be detected through introspection (Gilovich, Epley, & Hanks, in press). For example, to the extent that a trait can be defined in multiple ways, people tend to construe it in a self-serving fashion—namely, in terms of characteristics they possess when it comes to positive traits and characteristics they lack when it comes to negative traits (Dunning et al., 1989). People also appear to selectively recall their past in ways that sustain positive views about themselves (Sanitioso, Kunda, & Fong, 1990). Both of these tendencies make it more likely that individuals will rate themselves as above average but do so without it being obvious that these self-assessments might not be justified.

METHOD

Participants. One hundred twenty-seven Cornell undergraduates were recruited from psychology courses and given extra credit for their participation.

Procedure. Participants estimated the degree to which one of four sets of judgments might be influenced by bias. Those in the self-specific condition first rated the degree to which they were more or less intelligent, ethi-

cal, realistic, and likable than the average Cornell student on a 9-point scale. Participants in the other-specific condition were each yoked to a participant in the self-specific condition and shown that student's trait ratings. Those in the self-general and other-general conditions did not make these trait ratings nor were they exposed to trait ratings made by anyone else. Rather, they were simply asked to consider the trait ratings that generally might be made by themselves (self-general condition) or by another student (other-general condition).

Next, all participants read a passage explaining that individuals might be affected by competing desires when rating themselves on various traits—people generally want to be accurate and honest but are also influenced by a desire to feel good about themselves. Participants were asked to indicate the degree to which the ratings they considered might be influenced by these competing motives by placing an X on a 14.5-cm line anchored by “affected primarily by a desire to give an honest appraisal” and “affected primarily by a desire to see oneself in a favorable light.”

Finally, all participants were asked to rate the degree to which they had used different strategies to determine whether the trait ratings under consideration might be influenced by bias. All ratings were made on a 9-point scale ranging from 1 (*not at all*) to 9 (*quite a bit*). Two questions assessed the degree to which the participants consulted abstract theories to assess possible bias. Participants rated the degree to which they (a) considered how well the tendency to make inflated ratings fits the way people tend to behave and (b) considered the extent to which the context of rating the self would tempt individuals to be biased by what they want to believe. An additional question assessed the extent to which participants relied on introspection. Specifically, participants rated the degree to which they tried to “get inside their [the average Cornell student's] head” to find evidence of self-serving motives.² The wording of each question was altered slightly to make it appropriate for considering bias in the self or another.

RESULTS

Perceptions of bias. We expected a person's own trait ratings to be seen as less biased than the trait ratings made by another and specific trait ratings to be seen as less biased than trait ratings considered in general. The results of a 2 (self vs. other) \times 2 (abstract vs. specific ratings) ANOVA on participants' assessments of bias supported these predictions (see Figure 2). There was a significant main effect of the source of the trait ratings such that participants thought their own trait ratings were less influenced by wishful thinking than the trait ratings of others, $F(1, 120) = 34.45, p < .0001$.³ There was also a significant main effect of the nature of the trait ratings such that




Figure 2 Perceived bias in self trait ratings, Study 2.

participants thought that a specific set of trait ratings that had actually been made were less likely to be tainted by bias than trait ratings that might be made, $F(1, 120) = 7.58, p < .001$.

As in Study 1, we subjected participants' bias ratings to a planned contrast with the self-specific condition assigned a weight of -3 and the other conditions assigned weights of $+1$. This contrast revealed that the lowest attributions of bias were made when participants were most likely to introspect—for specific predictions that participants themselves had just made, $F(1, 120) = 14.01, p < .001$.

Reported strategy use. We predicted that participants would more often report consulting theories of bias and less often report a strategy of introspection when assessing bias in the judgments of another. To test these predictions, we examined participants' responses to the question tapping the degree to which they reported searching within the person (self or other) for evidence of having been biased and their average responses to the two questions representing the degree to which they reported consulting theories of bias ($r = .50$). A 2 (self vs. other) $\times 2$ (abstract vs. specific predictions) ANOVA was conducted on each measure, and each yielded a significant self-other main effect. Participants asked about the possibility of being biased themselves reported using a strategy of introspection more ($M = 5.38$) than those asked about the possibility of bias in someone else ($M = 4.39$), $F(1, 123) = 8.15, p = .005$. The opposite result was obtained on the measure of how much participants consulted abstract theories of bias. Participants asked about the possibility of being biased themselves reported consulting abstract theories less ($M = 4.60$) than those asked about the possibility of bias in someone else ($M = 5.74$), $F(1, 123) = 12.70, p < .0005$.⁴

We also expected participants to rely on introspection more often to assess bias in specific judgments than when considering judgments in the abstract. This prediction was largely disconfirmed, as we did not obtain a significant main effect of specific versus abstract judgments on either the use of introspection measure, $F(1, 123) = 2.36, p > .10$, or the use of abstract theories measure, $F < 1.0$. Note, however, that this prediction applies most readily to one's own judgments rather than judgments made by another. One can no better introspect about a specific judgment made by another than about the sort of judgments others generally make. Indeed, the prediction fared a bit better when confined to the responses of participants who assessed bias in their own judgments. A planned comparison revealed a significant tendency for participants to rely more on introspection when considering a specific set of their own trait-ratings ($M = 5.90$) than when considering the sort of ratings they might make in the abstract ($M = 4.88$), $t(62) = 2.20, p < .05$. A similar planned comparison on the measure of the extent to which participants relied on abstract theories, however, did not yield even a marginally significant effect, $F < 1.0$.

To parallel our analyses of the bias ratings in both this study and Study 1, we performed two planned contrasts in which the self-specific condition was assigned a weight of -3 and the other conditions were assigned weights of $+1$. These contrasts revealed that participants relied on introspection most, $F(1, 125) = 4.09, p < .01$, and abstract theories least, $F(1, 125) = 4.60, p < .005$, when assessing bias in their own specific judgments.

Finally, the degree to which participants indicated that they consulted theories of bias more than introspection (as represented by a difference score between the two self-report indices) predicted greater attributions of bias ($r = .22, z = 1.68, p < .10$).

Indirect evidence of strategy use. As in Study 1, we conducted an additional test of the extent to which participants relied on introspection versus abstract theories by examining the correspondence between their ratings of bias and the degree to which the judgment in question strayed from what might be expected. Our measure of the trait ratings participants were likely to have expected were taken directly from the responses of participants in the other-general condition who were asked this very question. A measure of the extent to which a given set of trait ratings deviated from expectations was computed by subtracting from each trait rating the mean expectation rating for that trait and then averaging across all trait ratings made by a given individual. As predicted, participants appeared to pay attention to the extent to which trait ratings strayed from what was expected only when estimating the degree to which another's trait rat-

ings might be biased, $r = .39$, $z(31) = 2.20$, $p < .05$. In contrast, participants paid no attention to the degree to which their own ratings deviated from expectations, $r = .09$, $z < 1.0$.

DISCUSSION

Study 2 provides further evidence that individuals see greater bias in others' judgments than in their own, and acknowledge more bias in their own judgments in general than in specific judgments they have just made. We also replicated the finding that others' judgments, but not their own, are seen as biased to the extent that those judgments deviate from the perceived norm. The new findings in this study involve the apparent weight given to lay theories versus introspections and the tendency for people to treat their own introspections—but not other's introspections—as probative. These tendencies, we argue, account for the asymmetries in assessments of bias that we have documented. That is, people deny bias in their own judgments because introspection provides little evidence of bias. They acknowledge the possibility of bias in their past judgments and their judgments in general because they have no introspective evidence to rely on and therefore rely on their general theories about bias. And they infer bias in others because they do not have direct access to the introspection of others and doubt the candor or insightfulness of the self-reported introspections that others offer.

It is noteworthy that participants did not claim to be completely unbiased, even in their immediate judgments. The mean bias ratings for the self were simply low in absolute terms, and participants thought their own judgments were less influenced by bias than the judgments made by others. This finding reflects both that people give theory (e.g., theory about motivation influences) at least some weight, even in their self-assessments, and that there are some occasions in which introspection does yield evidence of bias, or at least evidence of the motives that prompt bias. (A Little League umpire is apt both to realize and recollect how hard it was to utter "Strike three!" when it was his or her own child in the batter's box.)

Our contention is simply that introspection will more often than not yield a verdict of not guilty. Indeed, introspection is apt to yield the impression that one acted in spite of one's preferences, not because of them. One's conscious efforts to have avoided bias, in fact to have "bent over backwards" to do so, are likely to be highly salient. An interesting test case is provided by instances in which one has a group identity or some other personal connection to an issue that constitutes a potential source of bias. To what extent do we recognize the biasing influence of such connections on our own judgments? We propose that people may not only fail to acknowledge

the biasing influence in question, they may be led by their introspections to claim that their connections constituted a source of enlightenment.

Anecdotal evidence for this proposition abounds. In advocating funding for stem cell research, Congressman Dick Gephardt spoke of his experience caring for his elderly mother and insisted that "unless you've gone through something, you really don't understand it" (Solomon, 2001). Similarly, Senator Strom Thurmond claimed that "as a father of a daughter with juvenile diabetes, I know first-hand the devastating nature of this disease"—and thus justified his long-time support of medical research (Pear, 2001). Both politicians not only denied bias, they claimed unique enlightenment.

Undeniably, experiences of this sort do provide useful information—and appreciation of the significance of that information—that those unconnected to the issue are unlikely to share. But at the same time, such connections give rise to vested interests and policy preferences that can distort the individual's evaluation of relevant evidence and arguments (Ditto & Lopez, 1992; Lord et al., 1979). We contend that when individuals with a vested interest introspect, they will find lots of evidence of enlightening experiences and little evidence of conscious bias in assessing information. By contrast, when an observer considers these same individuals, the congruence between their positions and their self-interest will be taken as evidence of bias in accord with the observer's theories about the potent effects of self-interest in distorting judgment. Studies 3 and 4 were designed to explore this asymmetry in perceived bias versus enlightenment.

STUDY 3

In October 2000, an Israeli-Palestinian truce collapsed, leading to an eruption of violence that killed more than 100 people in the first few weeks and that continues to this day. The events became one of the most talked-about issues on campus and across the nation. Not surprisingly, Israeli and Jewish individuals held different views of the conflict than Arab and Muslim individuals. In the aftermath of this eruption of violence, we approached Arab, Muslim, and Jewish students at rallies and debates on campus and asked them to consider the degree to which a personal connection to the Middle East might have influenced either their own views regarding the crisis or the views of someone on the "other side." We expected participants to think that such a connection would serve as a source of bias in the views of the other side but as a source of enlightenment in their own views.

METHOD

Participants. We approached individuals attending campus events designed to foster discussion of the Middle East peace crisis and asked them to complete a questionnaire on the subject. Sixty-eight agreed.

Procedure. Participants provided demographic information and then read a passage stating that by virtue of their ethnic or religious background, Jews, Arabs, and Muslims tend to feel personally connected to the events in the Middle East. Participants who reported that they were Jewish, Muslim, or Arabic were randomly assigned to consider either their own viewpoint or that of someone with the opposite personal connection. They then read the following:

Your status as member of a group with a lot at stake in this issue might influence how you view the history of this conflict, how you view the events going on right now, and what you think is the best course for the future. A person's stake in a given issue might give them a perspective . . . that is particularly illuminating (i.e., that gives them an understanding of the issues that someone without such a stake simply cannot have). Alternatively, a person's stake in a given issue might give them a biased perspective on the issue (i.e., that makes it hard for them to see the issues fairly).

For those considering the viewpoint of someone on the opposite side of the issue, "you" was replaced by "this person." Participants were then asked to indicate the extent to which they thought the view in question (their own or that of someone on the opposite side) might be influenced in one of these ways versus the other. Specifically, participants circled a number on a 9-point scale anchored by "is likely to give me [this student] an *illuminating* perspective" and "is likely to give me [this student] a *biased* perspective."

RESULTS AND DISCUSSION

We expected participants to think that their own views were less biased and more enlightened by personal connections to the Middle East than the views of those on the opposite side. A 2 (Jewish vs. Arab/Muslim participant) \times 2 (Jewish vs. Arab/Muslim target) ANOVA yielded the predicted significant interaction between participant and target background, $F(1, 64) = 26.42, p < .0001$. As Figure 3 indicates, Jewish participants were more likely than Arab/Muslim participants to think Arab/Muslim targets were unduly biased by their background, $t(64) = -3.96, p < .0005$, and less likely to think that Jewish participants were biased by theirs, $t(64) = 3.31, p < .005$. A one-sample t test indicated that Arabic/Muslim participants rated their own views as more illuminated than biased by their personal connection to the




Figure 3 Perceptions of the degree to which one's ethnic or religious status offers an enlightened as opposed to a biased perspective on issues related to the conflict in the Middle East, Study 3.

NOTE: Zero represents the midpoint of the scale, with positive numbers representing a belief that the target individual's view is likely to be more enlightened than biased by a personal connection and negative numbers representing a belief that the target individual's view is likely to be more biased than enlightened.

issue, $t(23) = 2.17, p < .05$. A similar test on the responses of Jewish participants indicated that although they did not think their personal connections to the Middle East biased their views of the pertinent issues, they also did not think that their personal connections were a significant source of enlightenment, $t < 1$.

Thus, participants with ethnic or religious connections to the conflict in the Middle East viewed their own personal connection as more enlightening and less biasing than the personal connections of those on the opposing side. Arab and Muslim participants, furthermore, viewed their connections to the issue as a significant source of enlightenment. Jewish participants (most of whom were born in the United States and were thus perhaps less personally involved in the crisis than our Arab/Muslim participants) did not consider their own connection to the issue to be particularly enlightening, but they also did not see it as a source of bias. We have no way of knowing whether this difference between Jewish and Arab/Muslim students was due to the amount of passion they brought to the issue, the amount of knowledge they brought, or both.

In this study, we demonstrated that individuals' identity-based connections to a controversial issue are seen as less biasing—indeed, as enlightening to some—than the personal, identity-based connections of others. Someone else's connection to a given issue is necessarily rather abstract, bereft of the personal experiences and detailed knowledge that come with one's own connection. This raises the question of whether the effects observed here would still be obtained when the pronounced difference in the abstract versus concrete quality of the personal connection was reduced. Would the

same effects be observed when considering a personal connection on the part not of oneself and a member of the other side but someone else on one's own side and an opponent? One cannot introspect about the influence of a personal connection on someone else's thinking, even when that someone else is from one's own group and shares one's own values and ideological leanings. Nevertheless, it is certainly possible to introspect and use one's own internal experience as a proxy for the experience of someone with the same connection to the issue at hand. Participants in Study 2, for example, reported that they tried to "get into the head" of another student to decide if he or she may have been guilty of bias—but not nearly as much as they tried to get into their own heads to decide if they themselves had been biased. The possibility of introspection by proxy implies that a personal connection to a controversial issue will be seen as less biasing and more enlightening when it applies to someone on one's own side of the issue than when it applies to someone on the other side.

STUDY 4

This study featured two issues for which respondents' personal connections might be seen as either biasing or enlightening. Specifically, we examined assessments of minority and majority students' input on university affirmative action policies and assessments of varsity and intramural athletes' input on access to the most desirable weight training facility on campus.

METHOD

Participants. Seventy-nine Cornell University students participated in the affirmative-action version of the experiment and received extra credit in upper level psychology courses for their efforts. The sample included 41 Caucasian participants and 38 from various ethnic or racial groups (15 African American, 9 Hispanic, 6 biracial, and 8 international students, 4 of whom hailed from African countries, 2 from India, and 1 each from Indonesia and Jamaica).⁵ Thirty-eight varsity athletes and 37 intramural athletes participated (on a purely voluntary basis) in the athletic-facilities version of the experiment.

Procedure. Participants in the affirmative-action version of the experiment arrived at the laboratory and were run in individual cubicles. They were given a questionnaire that asked them to imagine that Cornell was considering changing its admissions policy. Under current policy, students read, the admissions committee takes into account a student's ethnic background in addition to such factors as high school grades and standardized test scores when rendering admissions decisions. Under the new policy being considered, a student's ethnic background would not be considered.

Participants were asked to imagine that the administration had assembled a student panel to provide input on the proposed change. Participants were randomly assigned to consider how the opinion of either a minority or a Caucasian student on the panel might be influenced by his or her ethnicity. Participants indicated how they thought this student would be influenced by his or her ethnicity by placing an X on a 16-cm line anchored by "likely to detract from his or her ability to see the issues clearly" and "likely to enhance his or her ability to see the issues clearly."

Participants in the athletic-facilities version of the study were approached at locations around campus and asked to fill out a questionnaire. The questionnaire described the weight training equipment housed in a newly built facility on campus and reminded participants that use of this facility was limited to varsity athletes. The questionnaire asked participants to imagine that in response to student criticism the administration had formed a student committee to consider whether the university should open the varsity weight room to intramural athletes. Participants who were themselves varsity or intramural athletes were asked to imagine either a varsity or intramural athlete on this committee. They indicated the degree to which they thought this athlete's status either as a varsity or as intramural athlete would influence his or her ability to see the issue clearly by placing an X on a 16-cm line anchored by "likely to detract from his or her ability to see the issues clearly" and "likely to enhance his or her ability to see the issues clearly."

RESULTS AND DISCUSSION

Affirmative action. We predicted that each group would see the other's ethnicity as a source of bias but their own ethnicity as a source of insight. Confirming this prediction, a 2 (Caucasian vs. minority participant) \times 2 (Caucasian vs. minority target) ANOVA yielded the predicted interaction (see Figure 4a), $F(1, 75) = 19.20$, $p < .0001$. Caucasian students thought that a panel member's ethnicity would serve as more of a biasing influence when that student was a member of an ethnic or racial minority ($M = -1.89$) than when that student was a fellow Caucasian ($M = .53$), $t(75) = 1.99$, $p = .05$. The opposite was true of minority students, who thought a student's ethnicity would constitute more of a source of bias when that student was Caucasian ($M = -2.40$) than when the student was a fellow minority ($M = 2.86$), $t(75) = -4.17$, $p < .0005$. Indeed, ethnic minority students asked about a fellow ethnic minority student on the university panel thought that such a student's ethnicity would serve as a significant source of enlightenment rather than bias, one-sample $t(17) = 2.69$, $p < .05$. A complementary effect emerged for Caucasian students, who thought the ethnicity of a fellow Caucasian would serve as more of a

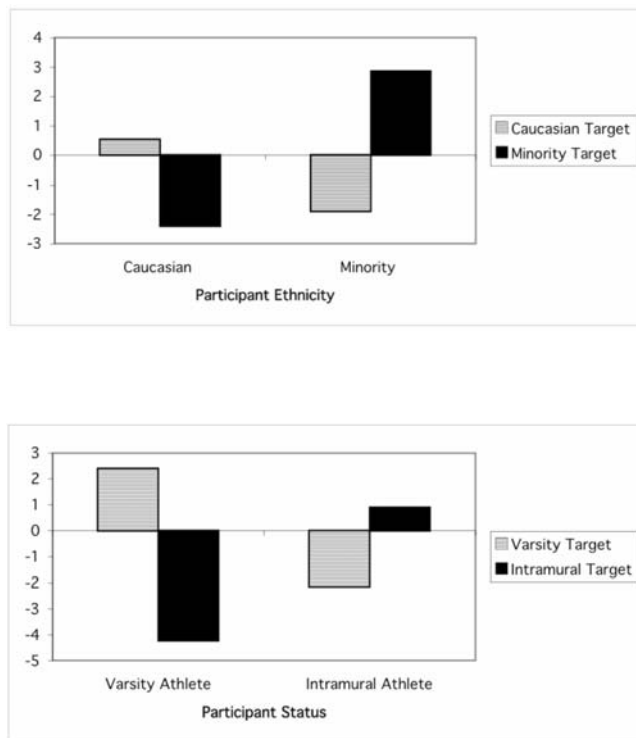


Figure 4 (a) Perceptions of the degree to which one's ethnicity offers an enlightened as opposed to biased perspective on attitudes about affirmative action, Study 4.
(b) Perceptions of the degree to which one's athletic status offers an enlightened as opposed to biased perspective on allocation of athletic space, Study 4.

NOTE: Zero represents the midpoint of the scale, with positive numbers representing a belief that the target individual's view is likely to be more enlightened than biased by a personal connection and negative numbers representing a belief that the target individual's view is likely to be more biased than enlightened.

source of enlightenment than bias, although this effect was not statistically significant, one sample $t < 1$.

Athletic facilities. Here, too, we predicted that each group would see the other's connection to the issue at hand as a source of bias but their own connection as a source of insight. Confirming this prediction, a 2 (varsity vs. intramural participant) \times 2 (varsity vs. intramural target) ANOVA yielded the predicted significant interaction, $F(1, 71) = 31.53, p < .0001$ (see Figure 4b). Varsity athletes thought that a panel member's athletic status would serve as more of a biasing influence when that student was an intramural athlete ($M = -4.22$) than when that student was a fellow varsity athlete ($M = 2.41$), $t(71) = -5.47, p < .0001$. The opposite was true of intramural athletes, who thought a student's athletic status would constitute more of a source of bias when the student was a varsity athlete ($M = -2.14$) than when the student was a

fellow intramural athlete ($M = .90$), $t(71) = 2.48, p < .05$. Indeed, each group asked about a fellow athlete of their own kind thought that such a student's athletic status would serve as a source of enlightenment rather than bias, $t(17) = 2.14, p < .05$ for varsity athletes and $t(18) = 1.74, p < .10$ for intramural athletes.

These data demonstrate that individuals more readily see a personal connection to a controversial issue as a source of bias among those holding an opposing viewpoint than among those who share their own connection to the issue. Furthermore, participants often viewed their own side's personal connection as more of a source of enlightenment than bias.

General Discussion

Determining whether a given judgment has been tainted by bias is not easy. Past research has shown that people are more apt to detect bias in the judgments of others than in their own judgments (Pronin et al., 2002), an asymmetry that appears to result from the different strategies people use to detect bias in their own judgments versus the judgments of others (Pronin et al., 2004). People tend to introspect to determine whether their own judgments are tainted by bias but to consult abstract theories to determine whether others' judgments are biased. Because many biases work below the surface and leave no trace of their operation, an introspective search for evidence of bias often turns up empty. Having conducted an internal search that produced little evidence of bias, people feel justified in believing their own judgments to be untainted by bias. One cannot conduct the same internal search to determine the causes of others' judgments, and so observers rely on their abstract theories about the types of judgments that are likely to be biased. Those theories in turn lead observers to infer bias whenever the actor's judgments seem self-serving.

Consistent with this analysis, our first two studies indicate that people believe judgments to be least biased when they are most likely to introspect to find evidence of bias. That is, people suspect the judgments of others to be more biased than their own, and they view their own judgments in the abstract as more prone to bias than specific judgments they have recently rendered. Furthermore, the degree to which a given judgment departs from the norm has a significant influence on perceived bias in others' judgments but not on one's own. Finally, participants in Study 2 directly reported using different strategies to assess the presence of bias in ways that mirrored the degree to which those judgments were thought to be tainted by wishful thinking. Participants reported looking inward to detect traces of bias more for their own judgments than for the judgments of others. They also reported looking inward to assess bias

in judgments they had just made than when considering whether they are generally prone to a particular type of bias. Thus, people are least likely to think a judgment is biased when they are most likely to introspect—when assessing whether their own specific judgments they just rendered are likely to be biased.

This does not imply that people never acknowledge that their own judgments are biased. People are often willing to concede, for example, that they are guilty of bias in their assessments of their friends or children. In such cases, the motivation to be seen as unbiased is not as great—or is balanced by a countervailing motive to be a stand-up friend or a protective parent—so it is easier to admit to the possibility of bias. Note that even here, however, one is more likely to confess to being biased in general than in any specific instance.

People will sometimes own up to being biased even in specific instances if and when their introspections yield some trace of their judgments being pulled in a favored direction. This can occur when one experiences conflict between an outcome and one's wish for a better result. In a relevant experiment by Ditto and Lopez (1992), for example, participants were led to believe that they could be assured that they were not predisposed to a troubling medical condition if a piece of litmus paper turned color when placed in a test medium. When for some of the participants the paper did not change color, they kept putting it back in the medium in an apparent effort to prompt the desired change. Most participants—and this was the main thrust of the Ditto and Lopez paper—surely saw no bias in what they did, interpreting their actions as simply being sufficiently thorough to allow the test to work properly. But we suspect a small number of participants may have been aware of their aversion to the initial result (“This can't be!”) and how it led them to do something—test further—that they would not have done otherwise. A more familiar instance of such awareness of bias occurs when the loser of a coin flip proposes “two out of three,” fully recognizing they would not propose such a change in criterion if the initial outcome had been in their favor.

The coin-flip example notwithstanding, we suspect that such instances of recognizing one's own biases in specific instances are rare, a contention reinforced by the results of Studies 3 and 4. In those studies, even a highly salient potential source of bias—a vested interest in a controversial issue—was not seen as having a distorting influence on participants' own views of the issue in question. In contexts ranging from the trivial (e.g., campus athletic privileges) to the geopolitically significant (e.g., the second intifada), participants reported that a personal connection to the issue in question did not bias either their own judgments (Study 3) or the judgments of those who share their connection (Study 4). In fact, in

both studies participants claimed that such a personal connection served as an important source of enlightenment. No such charitable interpretation was granted to those holding an opposing viewpoint.

Failing to detect bias in one's own judgment can sometimes be perfectly harmless and sometimes consequential. Dismissing a supervisor's unflattering performance appraisal as inaccurate without questioning whether it is one's own view that is in error can prevent a person from learning valuable information that can foster more rewarding future employment. Deciding that a division of common resources beneficial to the self is fair can lead to bitterness on the part of those who believe another division to be more equitable (Ross & Sicoly, 1979).

The latter example highlights the fact that asymmetric assessments of bias are likely to have the greatest implications in situations of conflict. Conflict is exacerbated—and resolution harder to achieve—when the two sides differ not only in what they deem fair and reasonable but when each takes the other's position as the product of unintentional bias or willful distortion of the relevant facts and entitlements. The tendency to believe that one's own experiences afford a particularly enlightened perspective on the conflict, and that one thereby sees the relevant issues more clearly than anyone else, can lead individuals to reject proposed solutions not only when offered by members of the opposing side but also when suggested by neutral third parties.

The role of the bias blind spot in interpersonal and intergroup conflict calls attention to how its untoward effects might most readily be overcome or dampened. On one hand, the task is formidable. The mental operations that give rise to the blind spot—looking inward to detect bias in oneself but consulting abstract theories to assess bias in others—are perfectly natural and perhaps even automatic processes that can seem beyond reproach. The impact of these mental operations is hard to overcome, and there is little in everyday experience that is likely to call them into question. On the other hand, here, as in so many other instances of research inspired by the cognitive revolution in psychology, knowledge itself is surely liberating. Recognizing that the divergent views expressed by others are apt to be sincerely held rather than a motivated strategy to seek advantage can help to defuse the spiral of conflict. Recognizing also that the biases that influence others result from normal psychological processes—processes to which we too are susceptible—rather than some essential characteristic of the other group, should further be helpful. We hope that the present research will serve to foster such insight.

NOTES

1. Because participants in the other-specific condition rendered their judgments about ratings provided by a specific participant in the self-specific condition, the ratings of these two groups are not independent, potentially contaminating the reported self-other difference. To overcome this problem, we computed separate *t* tests for the self-other comparisons in both the specific and general conditions. A paired *t* test revealed that the self-other difference in the specific conditions was significant, $t(24) = 2.59, p < .01$, and an independent-samples *t* test revealed that the corresponding difference in the general conditions was also significant, $t(46) = 3.42, p < .005$. One extra participant was run in the self-specific condition (i.e., one participant in this condition was not yoked to a participant in the self-general condition), which explains why the degrees of freedom in these two *t* tests do not match what would be expected from the ANOVAs reported in the main text.

2. Participants also answered two other questions that we originally designed as measures of the extent to which participants utilized a strategy of introspection: (a) whether they (the average Cornell student) felt (would feel) a temptation to give themselves the benefit of the doubt and (b) the extent to which they bent over backwards (would bend over backwards) not to give in to temptation. As was pointed out to us, however, it is not clear whether these questions really do tap into the use of introspection rather than abstract theories (indeed, the average correlation between responses to these three questions was only .32), and so they were dropped from the analysis. Note that when they are included in the analysis, the results are the same as with the single-item measure reported in the text.

3. There are 120 degrees of freedom because 3 participants did not complete this measure. Because, as in Study 1, participants in the other-specific condition rendered their judgments about ratings provided by a specific participant in the self-specific condition, the ratings of these two groups are not independent, potentially contaminating the reported self-other difference. To overcome this problem, we computed separate *t* tests for the self-other comparisons in both the specific and general conditions. A paired *t* test revealed that the self-other difference in the specific conditions was significant, $t(30) = 4.02, p < .0001$, and an independent-samples *t* test revealed that the corresponding difference in the general conditions was also significant, $t(60) = 4.12, p < .0001$.

4. Once again, to deal with the nonindependence of the self-specific and other-specific conditions, we conducted separate *t* tests for the self-other comparisons in both the specific and general conditions. For the introspection measure, a paired *t* test revealed that the self-other difference in the specific conditions was significant, $t(30) = 2.28, p < .05$, but an independent-samples *t* test revealed that the corresponding difference in the general conditions was not significant, $t(63) = 1.37, p < .20$. For the abstract-theories measure, a paired *t* test revealed that the self-other difference in the specific conditions was significant, $t(30) = 3.90, p < .001$, and an independent-samples *t* test revealed that the corresponding difference in the general conditions was marginally significant, $t(63) = 1.78, p < .10$.

5. Asian students were not included because although their daily experience on campus is in many ways that of a racial minority, their relationship to issues of affirmative action is not.

REFERENCES

- Alicke, M. D., Klotz, M. L., Breitenbecher, D. L., Yurak, T. J., & Vredenburg, D. S. (1995). Personal contact, individuation, and the better-than-average effect. *Journal of Personality and Social Psychology, 68*, 804-825.
- Armor, D. A. (1999). The illusion of objectivity: A bias in the perception of freedom from bias. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 59*, 5163.
- Brandstatter, E., & Schwarzenberger, H. (2001). Beyond the gambling paradigm: Internal controllability in decision-making. *Psychological Reports, 89*, 259-266.
- Buehler, R., Griffin, D., & Ross, M. (1994). Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology, 67*, 366-381.
- Canedy, D. (2001, July 18). Official defends her conduct in the vote counting in Florida. *The New York Times*, p. A20.
- Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology, 63*, 568-584.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology, 57*, 1082-1090.
- Eiser, J. R., Pahl, S., & Prins, Y. R. A. (2001). Optimism, pessimism, and the direction of self-other comparisons. *Journal of Experimental Social Psychology, 37*, 77-84.
- Epley, N., & Dunning, D. (2000). Feeling "holier than thou": Are self-serving assessments produced by errors in self- or social prediction? *Journal of Personality and Social Psychology, 79*, 861-875.
- Friedrich, J. (1996). On seeing oneself as less self-serving than others: The ultimate self-serving bias? *Teaching of Psychology, 23*, 107-109.
- Gallup. (2000, December 22). *The Florida recount controversy from the public's perspective: 25 insights. Poll Releases*. Retrieved December 17, 2003 from <http://www.gallup.com/poll/releases/pr001222bii.asp>.
- Gilovich, T. (1991). *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*. New York: Free Press.
- Gilovich, T., Epley, N., & Hanko, K. (in press). Shallow thoughts about the self: The automatic components of self-assessment. In M. Alicke, D. Dunning, & J. Krueger (Eds.), *The self in social perception*. New York: Psychology Press.
- Kruger, J., & Gilovich, T. (1999). "Naive cynicism" in everyday theories of responsibility assessment: On biased assumptions of bias. *Journal of Personality and Social Psychology, 76*, 743-753.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*, 480-498.
- Lord, C. G., Ross, L., & Lepper, M. R., (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*, 2098-2109.
- Miller, D. T., & Ratner, R. K. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology, 74*, 53-62.
- Miller, D. T., & Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction? *Psychological Bulletin, 82*, 213-225.
- Murray, S. L., & Holmes, J. G. (1997). A leap of faith? Positive illusions in romantic relationships. *Personality and Social Psychology Bulletin, 23*, 586-604.
- Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice Hall.
- Pear, R. (2001, June 19). Several G.O.P. Senators back money for stem cell research. *The New York Times*, p. 18.
- Price, P. C., (2001). A group size effect on personal risk judgments: Implications for unrealistic optimism. *Memory and Cognition, 29*, 578-586.
- Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others. *Psychological Review, 111*, 781-799.
- Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin, 28*, 369-381.
- Ross, M., & Sicoly, F. (1979). Egocentric biases in availability and attribution. *Journal of Personality and Social Psychology, 37*, 322-336.
- Ross, L., & Ward, A. (1996). Naive realism in everyday life: Implications for social conflict and misunderstanding. In T. Brown, E. Reed, & E. Turiel (Eds.), *Values and knowledge* (pp. 103-135). Hillsdale, NJ: Lawrence Erlbaum.
- Sanitioso, R., Kunda, Z., & Fong, G. T. (1990). Motivated recruitment of autobiographical memories. *Journal of Personality and Social Psychology, 59*, 229-241.
- Solomon, J. (2001, August 15). Family crisis shifts politics. *USA Today*, p. 15A.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology, 38*, 806-820.
- Whitley, B. E., & Frieze, I. H. (1985). Children's causal attributions for success and failure in achievement settings: A meta-analysis. *Journal of Educational Psychology, 77*, 608-616.

Wilson, T. D., Centerbar, D. B., & Brekke, N. (2002). Mental contamination and the debiasing problem. In T. Gilovich, D. W. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 185-200). New York: Cambridge University Press.

Received February 23, 2004.

Revision accepted August 25, 2004.