
Actions, Intentions, and Self-Assessment: The Road to Self-Enhancement Is Paved With Good Intentions

Justin Kruger

University of Illinois at Urbana-Champaign

Thomas Gilovich

Cornell University

Actions and intentions do not always align. Individuals often have good intentions that they fail to fulfill. The studies presented here suggest that actors and observers differ in the weight they assign to intentions when deciding whether an individual possesses a desirable trait. Participants were more likely to give themselves credit for their intentions than they were to give others credit for theirs (Studies 1 and 2). This caused individuals to evaluate themselves more favorably than they evaluated others (Studies 3-5). Discussion focuses on the motivational and information-processing roots of this actor-observer difference in the weight assigned to intentions as well as the implications of this tendency for everyday judgment and decision making.

Keywords: *intentions; self-enhancement; above-average effect; bias; actor-observer difference*

I'm just a soul whose intentions are good,
Oh Lord, please don't let me be misunderstood.
—The Animals (1964)

It has been said that actions speak louder than words. But often, actions do not say enough. To understand behavior, one must understand not only what was done but what the actor intended to do—a fact emphasized by Jones and Davis (1965) but largely overlooked in subsequent attribution theories (but see Malle & Knobe, 1997, 2001, for an exception). Thus, a person is “cruel” only when he desires harm, “selfish” only when she wants for herself, and a “liar” only if a falsehood is knowingly told.

But as Jones and Nisbett (1971) also noted, “typically, the actor has more, and more precise, information than the observer about his own . . . intentions” (p. 85). To be sure, there are occasions when an actor may be unaware of his or her true intentions (and perhaps even occasions

when an observer may know more about an actor's intentions than the actor himself or herself), but “it is perhaps only with the self that one is familiar with the nuances of ‘unfulfilled potential’ and of the ‘excuses’ which explain why personal excellence has not been translated into tangible accomplishments” (James, 1890, pp. 326-327).

This implies that people are better positioned to take intentions into account when assessing their own traits, behaviors, and likely future outcomes than when assessing the traits, behaviors, and outcomes of others. Whereas others may be deemed considerate if they display considerate behavior (volunteering at the local soup kitchen, sending a note of appreciation, tutoring a student in statistics), judgments of the self are less stringent. People give themselves credit for their intentions as well—for *wanting* to help the needy, *planning* to send a note of appreciation, or *making oneself available* for a student befuddled by statistics. To be sure, intentions play some role in the judgments of others, as many legal scholars and moral philosophers have noted. But intentions play a greater role, we argue, in judgments of the self.

Authors' Note: This research was supported by Research Grant 1-2-69853 from the University of Illinois and SBR9319558 from the National Science Foundation. We thank Adi Bar-Lev, Lisa Libby, Bryn Lovejoy-Grinnell, John Lubric, and Nev Vujosevic for their help in collecting and coding the data and Bill Altermatt and Matt Evans commenting on an earlier draft of this article. Correspondence concerning the article can be directed to Justin Kruger, Department of Psychology, University of Illinois at Urbana-Champaign, 603 East Daniel Street, Champaign, IL 61820; e-mail: jkruger@uiuc.edu or tgd1@cornell.edu.

PSPB, Vol. 30 No. 3, March 2004 328-339

DOI: 10.1177/0146167203259932

© 2004 by the Society for Personality and Social Psychology, Inc.

The greater availability of one's own intentions noted by Jones and Nisbett (1971) is not the only reason intentions may loom larger in self-assessments than in the assessments of others. People also may be motivated to give themselves more credit for their intentions than they give others for theirs (at least when they are positive). Numerous studies have documented the powerful impact that the motive to think well of oneself has on people's judgments of themselves and others (e.g., Alicke, 1985; Brown, 1986, 1990; Campbell, 1986; Dunning, Leuenberger, & Sherman, 1995; Kunda, 1990; Taylor & Brown, 1988). It stands to reason, then, that one reason people fail to give others the same benefit of their good intentions that they give themselves is to look good in comparison.

People also may weight their own intentions more heavily than the intentions of others because they are cynical about the likelihood that someone else's good intentions and plans for the future will translate into effective action. Consider New Year's resolutions. Although people recognize that some of their own resolutions may pass unfulfilled, they are likely to expect even fewer resolutions to be fulfilled by others (as illustrated by health clubs that sell more memberships than they can accommodate, secure in the knowledge that only a fraction of their new "customers" will make it to the gym). Indeed, research on the "planning fallacy" has found that although people tend to be overly optimistic about how long it will take them to complete important projects, they are much less overconfident about the likelihood that others will meet deadlines (Buehler, Griffin, & Ross, 1994, 2002).

Finally, people may be reluctant to give much credence to other people's intentions because they doubt the strength and sincerity with which they are held. What we know about others' intentions, after all, often comes from what they tell us—in the excuses offered when a person's actions do not match their goals or expectations. The failure to translate an intention into effective action is often taken as a sign that the intention was weak or nonexistent. Our own intentions, in contrast, are experienced and felt rather than heard and evaluated and, hence, treated as sovereign. The net result is that we are more likely to take intentions into account when making self-assessments than when making assessments of others.

Regardless of the cause, the tendency to give oneself more credit for good intentions than one gives others is likely to contribute to a number of biases in judgments of the self and how the self stacks up against others. As the recent literature on the subject indicates, self-assessment is often associated with self-enhancement. People are overconfident in their ability to predict their own futures (Dunning, Griffin, Milojkovic, & Ross, 1990) and in the

likelihood that they will meet important deadlines (Buehler et al., 1994, 2002). They believe that their own preferences and responses enjoy more consensus than the preferences and responses of others (Gilovich, 1990; Marks & Miller, 1987; Ross, Green, & House, 1977), and they think that they are more likely than others to experience the good things in life and avoid the bad (Klein, 1999; Klein & Weinstein, 1997; Weinstein, 1980). Finally, people tend to rate their own traits and abilities more favorably than they rate the traits and abilities of others (Alicke, 1985; Dunning, Meyerowitz, & Holzberg, 1989; Kruger & Dunning, 1999; Larwood, 1978; Messick, Bloom, Boldizar, & Samuelson, 1985; although see Heine & Lehman, 1997; Kruger, 1999, for exceptions).

We believe the tendency to give oneself more credit for intentions than one gives others can further our understanding of these common biases in self-assessment. Previous theorists have written about the differential access to intentions on the part of actors and observers (S. M. Anderson, 1984; S. M. Anderson & Ross, 1984; Jones & Davis, 1965; Jones & Nisbett, 1971; Malle & Knobe, 1997, 2001; Pronin, Gilovich, & Ross, in press), and a voluminous literature supports the existence of numerous biases in self-assessments (Alicke, 1985; Brown, 1990; Buehler et al., 1994, 2002; Dunning et al., 1989; Ross et al., 1977; Taylor & Brown, 1988). The purpose of the present research is to connect these two literatures. Our focus was the "above-average effect," the tendency of individuals to rate themselves more favorably than they rate others. We propose that asymmetries in access to and assumed credibility of one's own intentions versus those of others lead people to apply a more lenient standard when evaluating the self than when evaluating others. As a consequence, people come to evaluate themselves more favorably than they evaluate others.

To test this hypothesis, our first two studies examined whether people believe that intentions are more relevant when assessing their own characteristics than when assessing the characteristics of others. The remaining three studies investigated the implications of this differential weighting for comparative assessments of traits and abilities.

STUDY 1

Participants in Study 1 were asked to write down several aspects of either themselves or someone they knew that are indicative of a particular trait and then to rate how much these aspects capture the extent to which the person possesses that trait. In addition to this self-manipulation versus other manipulation, half of the participants were told to confine what they wrote to overt behaviors and observable actions (no-intention condition) and half were given no such restriction (control condition).

We reasoned that if people give themselves more credit for their intentions than they give others, then preventing participants from listing intentions ought to influence their ratings in predictable ways. Specifically, participants prohibited from listing intentions should rate their examples as less descriptive of “who they really are” than nonrestricted participants. Critically, this difference should be greater among participants rating themselves than among participants rating someone else.

Method

Participants. One hundred Cornell University students (73 women, 27 men) earned course credit for their participation.

Procedure. Participants were randomly assigned to one of four conditions in a 2 (target: self vs. other) \times 2 (condition: no-intention vs. control) between-subject design. Participants in the self/no-intention condition were asked to write down five overt behaviors or observable actions in which they have engaged that are indicative of a particular trait and then to rate the extent to which the five examples capture how much they possess the trait on a scale from 1 (*doesn't capture me at all*) to 7 (*captures me very well*). Participants considered a total of four traits: fair, considerate, cooperative, and generous. The specific instructions were as follows:

For each of the four traits that follow, we'd like you to write down 5 things that you have done that are reflective of the trait. For example, for the trait *thoughtful*, you would write down 5 thoughtful behaviors in which you have engaged. For instance, you might write that you recently gave money to a charity, that you helped a roommate with homework, or that you listened to a friend when she really needed to talk. There are no right or wrong answers: simply write down 5 *overt behaviors* or *observable actions* that you have engaged in that are indicative of the particular trait. Finally, we'd like you to rate the extent to which the 5 behaviors you wrote down capture the amount of the trait that you possess. For instance, for the trait *thoughtful*, you would rate how much the 5 thoughtful things you wrote down capture how thoughtful you are. That is, if a stranger knew *only* about the 5 behaviors you wrote down, to what extent would he or she have an accurate view of how thoughtful you are? (emphases in original)

Participants in the other three conditions followed a similar procedure, except that those in the control conditions were not given any restrictions regarding what kind of examples to write down (i.e., they were simply told to write down five aspects of the person in question that are indicative of the particular trait) and those in the other conditions were asked to rate, not themselves, but

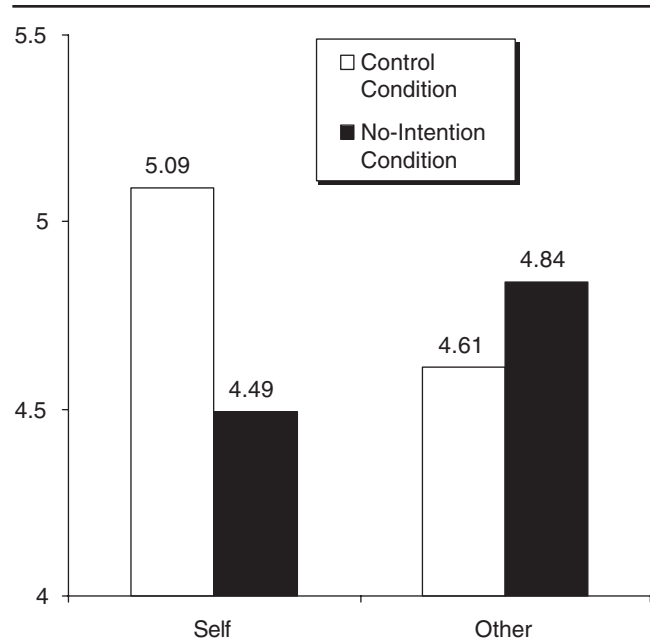


Figure 1 Participants' mean ratings (on a 1 to 7 scale) of how much the examples they listed captured the extent to which the target possessed the trait, by condition (Study 1).

a specific person they knew. We asked participants to exclude close friends because we were interested in self-other differences and the self-other border can be a bit less distinct among close friends, particularly intimates (Aron & Aron, 2001). Participants wrote the person's initials on each page of the questionnaire, which we hoped would help keep the particular person in mind.

Results and Discussion

Gender did not alter the pattern of results in this or any of the studies presented in this article and is not discussed further.

Participants' ratings of how much the examples listed captured the extent to which the target possessed the trait in question were averaged across the four traits to create a single measure for each participant.¹ These data were analyzed in a 2 (target: self vs. other) \times 2 (condition: no-intention vs. control) between-subject ANOVA that yielded the predicted interaction, $F(1, 96) = 3.87, p = .05, \eta^2 = .04$. Neither main effect was significant, $F_s < 1$. As Figure 1 illustrates, when rating themselves, participants thought the examples they wrote down captured the extent to which they possessed the trait in question less well when they were prohibited from listing intentions than when they were not, $t(50) = 2.38, p < .03, \eta^2 = .10$. Participants rating someone else, in contrast, were unaffected by the restriction. If anything, they thought the examples captured the person's attributes *more* effectively when they were prohibited from listing intentions, although this difference was not significant, $t(46) < 1$.

These results suggest that intentions are a more integral part of people's self-representations than their representations of others. Participants who listed examples from their own lives thought the examples captured the extent to which they possessed a given trait less well when they were prohibited from listing intentions than when they were not so restricted. This restriction had no effect on participants who rated someone else. Participants seemed to believe that the full extent of their cooperativeness, for example, could not be understood unless their cooperative intentions were taken into account, whereas someone else's cooperative behavior was a relatively good measure of his or her cooperativeness.

STUDY 2

Study 2 was designed to examine more directly the differential weight actors and observers place on intentions. Participants were presented with a series of traits and asked to rate the extent to which a person's intentions would need to be considered to get an accurate picture of the extent to which he or she possessed the trait. For half of the participants, the person in question was someone they knew. For the other half, the person in question was the participant himself or herself.

We predicted that participants would find intentions more essential for gauging the extent to which they possessed a specific trait than when gauging the extent to which someone else possessed that trait.

Method

Participants. Seventy Cornell University students (53 women, 17 men) earned course credit for their participation.

Procedure. Participants in the self-rating condition ($n = 37$) were given a list of 40 desirable traits and asked to indicate "how much weight should be placed on your intentions in order to get an accurate picture of the extent to which you possess the trait" on a scale from 1 (*none*) to 7 (*a lot*). The traits were selected from N. H. Anderson's (1968) list of 555 traits and varied in terms of ambiguity and intentionality. To ensure that the issue of intentionality was clear, participants read the following example:

Suppose the trait was *thoughtful*. Should a person judging how thoughtful you are consider only your overt behavior (such as throwing a birthday party or helping a roommate with homework) or should he or she also consider your intentions (such as *planning* a birthday party or *trying* to help a roommate with homework)? For each trait listed on the following page, we'd like you to indicate how much weight should be placed on your intentions on a scale from 1 (*none*) to 7 (*a lot*).

After providing their ratings, participants were given the same list of traits again and asked to rate "the extent to which the trait describes you" on a scale from 1 (*doesn't describe me at all*) to 7 (*describes me very well*).

Participants in the other-rating condition ($n = 33$) followed a similar procedure except that they were asked to make their ratings with another person in mind. Specifically, they were asked to think of a specific person they knew but not someone they would consider a close friend. They then rated "how much weight should be placed on this person's intentions in order to get an accurate picture of the extent to which he or she possesses the trait" and "the extent to which the trait describes this person" using the same scale used by participants in the self-rating condition. As in the previous study, participants were asked to write the person's initials on each page of the questionnaire to ensure that they kept this person in mind when making their ratings.

Results and Discussion

Our first prediction was that participants would rate themselves more favorably than they rated the other person. This hypothesis was tested in two ways. First, we conducted a subject-level analysis by computing, for each subject, the average trait rating across the 40 traits and comparing these ratings in a two-sample t test. Consistent with prior work, participants evaluating themselves made higher ratings ($M = 5.11$) than did participants rating someone else ($M = 4.49$), $t(68) = 4.41$, $p < .001$, $\eta^2 = .22$. This pattern also was evident at the level of the trait, which was tested by computing, for each trait, the average ratings of the participants in the two conditions and then comparing these means in a paired t test, $t(39) = 6.25$, $p < .001$, $\eta^2 = .55$.

Our primary prediction, however, was that participants would stress the importance of intentions more when rating themselves than when rating someone else. As Table 1 depicts, this was true for 36 of the 40 traits (90%), with a mean self-other difference of 0.37. This difference corresponded to a marginally significant effect at the subject level of analysis, $t(68) = 1.71$, $p = .092$, $\eta^2 = .04$, and a large and highly significant effect at the trait level of analysis, $t(39) = 6.41$, $p < .001$, $\eta^2 = .51$.

Not surprisingly, despite the self-other discrepancy in the weight placed on intentions, there was a high correlation between the two ratings, $r(38) = .91$, $p < .001$, suggesting that some traits are seen as defined more by a person's intentions (e.g., considerate, generous, and kind-hearted) than others (e.g., lucky, popular, and intelligent). Of key interest, the degree of intentionality inherent in a given trait was a strong and reliable predictor of the size of the above-average effect: The correlation across the 40 traits between the mean intentionality rating (averaging across participants in both the self-

TABLE 1: Mean Ratings of the Importance of Intentions for Assessing the Self Versus Someone Else on 40 Traits, Study 2

<i>Trait</i>	<i>Self</i>	<i>Other</i>	<i>Difference</i>
Alert	3.54	3.52	0.02
Appreciative	5.54	4.48	1.06
Articulate	3.62	3.36	0.26
Artistic	4.05	3.21	0.84
Bright	4.03	3.85	0.18
Brilliant	3.22	3.21	0.01
Capable	4.39	3.79	0.60
Clean	3.73	3.52	0.21
Competent	4.11	3.79	0.32
Confident	4.25	4.15	0.10
Conscientious	5.05	4.76	0.29
Considerate	5.51	5.15	0.36
Constructive	4.67	3.97	0.70
Cooperative	5.24	4.36	0.88
Decent	4.97	4.47	0.50
Educated	4.03	3.88	0.15
Energetic	3.65	3.79	-0.14
Fair	5.14	4.58	0.56
Friendly	5.16	4.94	0.22
Generous	5.41	4.91	0.50
Gracious	4.54	4.73	-0.19
Grateful	5.38	4.30	1.08
Helpful	5.43	5.36	0.07
Honorable	4.83	4.52	0.31
Humble	4.43	3.91	0.52
Intelligent	3.84	3.42	0.42
Kind-hearted	5.53	5.45	0.08
Lucky	2.22	2.21	0.01
Moral	4.97	4.39	0.58
Nice	5.03	4.70	0.33
Open-minded	5.11	4.66	0.45
Popular	2.81	3.15	-0.34
Positive	5.11	4.67	0.44
Quick-witted	3.54	2.63	0.91
Refined	3.84	3.61	0.23
Sharp-witted	3.89	3.09	0.80
Skillful	3.62	3.24	0.38
Sympathetic	5.41	4.91	0.50
Talented	3.68	3.33	0.35
Wise	3.38	3.79	-0.41
<i>Grand mean</i>	<i>4.40</i>	<i>4.03</i>	<i>0.37</i>

NOTE: Higher numbers reflect greater perceived importance of intentions.

condition and other condition) and the size of the above-average effect (the difference between the average rating of participants in the self-condition and other condition) was .69 ($p < .001$). This provides the first hint that the oft-documented above-average effect may be due in part to the tendency to give oneself more credit for one's intentions than one is willing to give others for theirs. The next three studies were designed to examine this possibility directly.

STUDY 3

The results of Study 2 suggest that the more heavily intentions are weighed in the assessment of a given trait,

the greater the tendency to rate oneself higher on that trait than someone else. But note that the "someone else" participants evaluated in Study 2 actually may have been lower on these traits—that is, participants may have selected comparison others who were in fact less talented, accomplished, and so forth, than they were themselves. As a result, the findings do not establish that participants' assessments were biased per se. In addition, although the traits that were high in intentionality tended to yield the greatest above-average effects, the traits may have differed along one or more other criteria (such as desirability or ambiguity) that are correlated with intentionality, and it may have been one of these other criteria that produced these effects.

Study 3 was designed to overcome these limitations. Participants rated themselves along the same 40 traits used in Study 2. This time, however, participants compared themselves with the average student from their course. They then rated the "intentionality" of the trait, that is, the extent to which intentions are important for knowing whether a person possesses the trait. Finally, participants rated each trait on three other dimensions that prior research has shown to influence self-assessments: desirability, ambiguity, and observability (Alicke, 1985; Dunning et al., 1989; Felson, 1981; Funder, 1980; John & Robins, 1993; Paunonen, 1989). If the tendency to give oneself but not others credit for intentions is partly responsible for the above-average effect, then participants should see themselves as more above average as the relevance of intentions increases—independent of trait desirability, ambiguity, and observability.

Method

Participants. Eighty-three University of Illinois students (57 women, 26 men) earned course credit for their participation.

Procedure. Participants were given the same list of 40 traits used in Study 2. For each trait, participants compared themselves with the average student from their course on a scale from 1 (*I'm much worse than average*) to 4 (*I'm average*) to 7 (*I'm much better than average*). They also provided separate estimates of the extent to which they possessed the trait and the extent to which the average person from their course possessed the trait, again on 1 to 7 scales. Participants then rated the extent to which a person's intentions can provide information about whether one possesses the trait. To ensure that what we meant by "intentions" was clear to participants, the following instructions were provided:

For each of the following traits, we'd like you to indicate the extent to which one's intentions can provide information about whether one possesses the trait. Some traits, such as thoughtful, are defined not only by one's

overt behavior (such as throwing a birthday party or helping a friend with homework) but also by one’s intentions (such as *planning* a birthday party or *intending* to offer help with homework). After all, if you plan a birthday party or try to offer a roommate help with homework but are unsuccessful, you are still thoughtful. Other traits, such as quick, are less defined by one’s intentions. Trying or intending to be quick just doesn’t make it so. Please rate the following traits on a scale from 1 (*intentions don’t count toward the trait*) to 7 (*intentions count toward the trait*). (emphases in original)

Finally, participants rated each of the 40 traits along three other dimensions: desirability, ambiguity, and observability. Each of these ratings was made on separate 1 to 7 scales with endpoints labeled *undesirable–desirable*, *concrete/has only one meaning–ambiguous/has many meanings*, and *unobservable/it is difficult or impossible to tell if a person has this trait by watching his or her behavior–observable/it is easy to tell whether a person has this trait by watching how he or she behaves*, respectively.

Results

On average, participants rated themselves as above average, consistent with prior research. This was true regardless of whether these effects were calculated directly, by comparing participants’ responses ($M = 4.60$) to the question “How do you compare with your peers?” to the midpoint of the scale (4), $t(39) = 10.47, p < .001$, or indirectly, by comparing participants’ ratings of the extent to which they possess the trait ($M = 5.43$) with their ratings of the extent to which the average person possesses the trait ($M = 4.94$), $t(39) = 6.55, p < .001$, $t(39) = 6.41, p < .001, \eta^2 = .52$. This was also true regardless of whether the data were analyzed at the level of the trait (just reported) or the level of the subject. With the subject as the unit of analysis, the comparative ratings again exceeded the midpoint of the scale, $t(82) = 10.09, p < .001$, and on average, self-ratings exceeded peer ratings, $t(82) = 6.61, p < .001, \eta^2 = .35$.

As the first two rows of Table 2 indicate, however, the extent to which participants rated themselves above average depended on the perceived relevance of intentions. As predicted, the more intentions were deemed relevant to the assessment of a given trait, the greater the above-average effect. (Tellingly, the trait participants rated as most intention-relevant, *considerate*, also yielded the greatest above-average effect.) Of key importance, these effects remained strong after holding constant differences in trait desirability, ambiguity, observability, or all three (all $r_s > .60$, all $p_s < .001$).² In fact, when the size of the above-average effect was predicted by all four variables (intentionality, desirability, ambiguity, and observability) simultaneously in a multiple regression, intentionality was by far the strongest predictor.

TABLE 2: Simple and Partial Correlations Between Trait Intentionality and Trait Endorsements, Study 3

Correlation	Controlling for				All Three
	Raw r	Ambiguity	Desirability	Observability	
Intentionality and direct comparison	.76**	.76**	.65**	.76**	.64**
Intentionality and indirect comparison	.67**	.65**	.60**	.69**	.63**
Intentionality and self-rating	.62**	.41**	.41**	.63**	.42**
Intentionality and rating of other	.15	.13	-.29†	.14	-.30*

† $p < .10$. * $p < .05$. ** $p < .01$.

The last two rows of Table 2 summarize the relationship between the relevance of intentions and participants’ self-rating and other rating separately. Consistent with the earlier analyses, as the perceived relevance of intentions increased, so did the tendency of participants to rate themselves positively. This was not the case, however, when participants rated someone else. If anything, as the perceived relevance of intentions increased, participants’ ratings of the extent to which others possessed a given trait *decreased*.

Because participants rated all aspects of every trait, it is possible to perform these analyses at the level of the individual by computing, for each participant, the correlation between his or her ratings of intentionality across the 40 traits and the various self-trait and other trait ratings provided by the participant across the 40 traits (as well as the partial correlation between these variables holding constant the participant’s own ratings of desirability, ambiguity, and observability). After Fisher’s r -to- z transformations, these correlations and partial correlations can then be compared against a null of zero.

As expected, the results of this subject-level analysis revealed the same pattern as the trait-level analysis. The average correlation between trait intentionality and comparative trait assessments (i.e., “how do you compare with your peers”) was .26, which a one-sample t test revealed was significantly greater than zero, $t(82) = 9.97, p < .001$.³ This was also true when controlling for desirability, mean $r = .23, t(82) = 9.68, p < .001$; ambiguity, mean $r = .24, t(82) = 9.69, p < .001$; or observability, mean $r = .20, t(82) = 8.15, p < .001$. A similar pattern emerged when the above-average effects were calculated indirectly by subtracting, for each trait, participants’ rating of the average person from their rating of themselves. Here, too, the mean correlation exceeded zero, mean $r = .21, t(82) = 8.56, p < .001$, even after controlling for desirability, mean $r = .17, t(82) = 6.61, p < .001$; ambiguity,

mean $r = .20$, $t(82) = 8.15$, $p < .001$; or observability, mean $r = .20$, $t(82) = 7.80$, $p < .001$.

This subject-level analysis also mirrors the trait-level finding in that the more important intentionality was deemed to be, the more self-ratings increased relative to other ratings. The average correlation between trait intentionality and ratings of how well the trait described the self was $.29$, $t(82) = 12.36$, $p < .001$, and was largely unchanged after controlling for desirability (mean $r = .23$), ambiguity (mean $r = .29$), or observability (mean $r = .28$). The average correlation between trait intentionality and the ratings of how well the trait described the average person, in contrast, yielded mean correlations between $.04$ and $.08$, which, although statistically significant ($ps < .06$), were substantially smaller than the correlations between trait intentionality and self-ratings.

Discussion

The results of Study 3 were as predicted. The traits highest in intentionality yielded the greatest above-average effects. This substantial and very reliable relationship was found at both the trait and subject level of analysis and, of crucial importance, remained robust after controlling for other factors known to influence self-assessments, such as the desirability of the trait in question. Taken together, these results are consistent with our thesis that the tendency for people to give themselves more credit for intentions than they give others contributes to people's inclination to believe they are better than average.

Of course, these results are correlational and so the usual cautions apply. To establish a causal link between the above-average effect and the self-other asymmetry in the weighting of intentions, it is necessary to experimentally manipulate whether people consider intentions in making their assessments and see whether this influences the tendency to rate oneself more favorably than others. Study 4 was designed to do precisely that.

STUDY 4

Participants in Study 4 were given the same list of traits used in Studies 2 and 3. As in Study 2, half of the participants were asked to rate themselves and half to rate someone they knew.

Unlike Study 2, however, half of the participants were asked to base their ratings exclusively on overt behavior and were told explicitly to ignore other components of the trait such as "good intentions, plans for the future, or unfulfilled potential." If the tendency to give oneself rather than others the benefit of good intentions underlies inflated trait ratings, this restriction should lower participants' ratings of themselves more than their ratings of someone else.

Method

Participants. One hundred Cornell University students (70 women, 30 men) earned course credit for their participation.

Procedure. Participants were assigned to one of four conditions in a 2 (target: self vs. other) \times 2 (ratings: no-intentions vs. control) between-subject design. Participants in the control conditions were asked to rate either themselves (self condition) or someone they knew (other condition) on the 40 traits depicted in Table 1 on a scale from 1 (*doesn't describe me/person at all*) to 7 (*describes me/person very well*) and then to compare that person with the average person from their course on each trait on a scale from 1 (*much worse than average*) to 4 (*average*) to 7 (*much better than average*). As in the previous studies, participants who were to rate someone else were asked to think of a specific person they knew, to answer all questions in reference to that person, and to write his or her initials on each page of the questionnaire.

Participants in the no-intentions conditions followed a similar procedure except that they were instructed to base their ratings exclusively on "overt behavior, observable actions, or tangible accomplishments," ignoring other components of the trait such as "good intentions, plans for the future, or unfulfilled potential." To ensure that participants understood this distinction, we fleshed out several examples similar to the one described in Study 3.

Results

We predicted that preventing participants from considering intentions would have a different effect on ratings of the self than on ratings of someone else. Our predictions were confirmed. On average, restricting participants from considering intentions lowered self-ratings, $M_{\text{control}} = 5.33$ vs. $M_{\text{no-intention}} = 5.13$, $t(39) = -3.87$, $p < .001$, but increased other ratings, $M_{\text{control}} = 4.18$ vs. $M_{\text{no-intention}} = 4.48$, $t(39) = 5.77$, $p < .001$. In addition, restricting participants from considering intentions lowered their comparative ratings of themselves (i.e., How do you compare with the average person?), $M_{\text{control}} = 4.91$ vs. $M_{\text{no-intention}} = 4.74$, $t(39) = -3.65$, $p < .001$, but increased participants' comparative ratings of someone else, $M_{\text{control}} = 4.18$ vs. $M_{\text{no-intention}} = 4.45$, $t(39) = 5.06$, $p < .001$. The interaction was significant for both sets of ratings with the trait as the unit of analysis ($ps < .001$). As in Study 2, supplemental analyses with the subject as the unit of analysis yielded similar, albeit more modest, results ($ps < .15$).

Discussion

Study 4 provides experimental evidence that the tendency to give oneself more than others the benefit of good intentions gives rise to inflated self-assessments.

Participants who were restricted from considering their good intentions provided lower self-ratings than those who were not so restricted.

Of interest, restricting participants from considering someone else's intentions had just as strong an effect, but in the opposite direction. Participants restricted from considering someone else's intentions provided higher ratings than participants not so restricted. This mirrors the correlational finding observed in Study 3 (see the bottom-right cell of Table 2). Note that this finding helps rule out a demand characteristic interpretation of our results. It could be argued that the reason participants provided lower self-ratings when told to restrict their ratings is that they thought that was what the experimenter expected. Note, however, that this same demand would presumably hold for participants rating someone else, and yet we observed an increase in other ratings.

Why did participants' ratings of another person increase when they were instructed to consider only overt behavior? One (admittedly speculative) possibility involves a phenomenon we discussed earlier: One is often uncertain about whether another person's professed intentions are genuine. In this study, an unconstrained consideration of another's cleanliness, conscientiousness, or generosity may have brought up instances in which the individual in question stated an intention to tidy up, exercise great care, or give to others—an intention that the rater may have been unwilling to accept. Any doubts aroused in this fashion would lower participants' ratings, making the unconstrained ratings lower than those in the "no-intention" condition.

STUDY 5

Seldom is there a bigger rift between actions and intentions than when it comes to helping others. Even the most altruistic among us do not attend to our children, care for our parents or grandparents, or give to the needy as much as we might like. There is always one more bedtime story we could read, one more visit we could make, one more dime we could spare.

This discrepancy between altruistic acts and intentions, together with the results presented thus far, suggest that as disappointing as our failures to meet our altruistic intentions may be to us, they are likely to be judged even more negatively by others. Whereas another person's failure to give to charity may seem a sure sign of selfishness, the fact that we would *like* to give more to charity—and indeed, fully intend to do so when our bank account more readily allows it—may leave us convinced that we are altruistic even when our actions suggest otherwise. Indeed, although altruism was not directly assessed in the previous studies, the traits closest to it—generous, considerate, and sympathetic—tended to yield the largest above-average effects.

Study 5 was designed to explore more directly this self-other discrepancy in the weight placed on altruistic intentions. Participants underwent a painful cold-pressor task to earn money for charity and then rated: (a) their actions, (b) their altruism during the study, and (c) their altruistic intent. Yoked observers provided similar ratings after watching a video recording of an individual participant's performance.

The cold-pressor task required participants to submerge their hand in a bucket of ice water for as long as they could and was selected for two reasons. First, it provided a compelling laboratory analogue to everyday altruistic behavior: We donated 50 cents for every minute the subject kept her hand submerged and allowed subjects to select the charity themselves. Thus, the greater participants' self-sacrifice, the more their selected beneficiary was rewarded.

Our second reason for selecting the cold-pressor task was that it made a rift between actions and intentions likely: Pretesting revealed that most people have a difficult time keeping their hand submerged for very long but that most people try to do precisely that. If people give themselves (but not others) credit for their altruistic intentions, then participants' ratings of their altruism should match, not their actual performance, on the cold-pressor task but their altruistic intentions. Participants rating the altruism of someone else, in contrast, should instead base their ratings on the target's cold-pressor performance.

Method

Participants. Fifty-four Cornell University women earned course credit as well as \$0.50 to \$5.00 for their favorite charity in exchange for their participation. We recruited only one gender to reduce variability in pain tolerance (Hellstrom & Lundberg, 2000).

Procedure: Cold-pressors. Participants were recruited individually for a study of "charitable behavior" in which they would have an opportunity to earn money for a favorite charity. On arrival at the lab, participants in the cold-pressor condition (henceforth referred to as "cold-pressors") were given a consent form explaining that they would be asked to submerge their hand in a bucket of cold water for as long as they could and that for each minute they kept their hand submerged, a local charity of their choice would receive 50 cents. They were warned that although the task is harmless, it is quite painful: the water, they were told (correctly), was between 32 and 35 degrees Fahrenheit and was very cold to the touch. They also were told that they could stop the experiment at any point and still receive course credit. All 20 participants in the cold-pressor condition agreed to participate.

Participants were then asked to select a charity to which they would like to donate their earnings from a list

of 17 popular charities (e.g., Planned Parenthood, Society for the Prevention of Cruelty to Animals, Zero Population Growth). They also were asked to indicate how long they intended to keep their hand submerged.

Next, participants were asked to roll up the sleeve of their nondominant arm and to submerge it in a bucket of ice water. Prior to doing so, they were reminded that their charity would earn 50 cents for each minute they kept their hand submerged and that they could remove their hand at any time. They also were told that, with their permission, the entire experiment would be videotape-recorded. None objected.

After completing the cold-pressor task, participants were given a towel and a questionnaire. The latter asked them to compare their performance on the cold-pressor task to that of the other study participants. Specifically, they were asked to compare themselves to others, in terms of (a) the “time you kept your hand submerged,” (b) their “altruism” during the study,⁴ and (c) their altruistic intent (“how much you wanted to help your selected charity” and “your willingness to perform the cold-pressor task for your charity”). Participants were told to answer each question on a scale from 0 (*less than everyone else*) to 50 (*exactly average*) to 100 (*more than everyone else*). By prior decision, the experimenter stopped participants who kept their hand submerged for more than 10 minutes. Three participants were excluded on this basis.

Procedure: Yoked observers. Observers were recruited from the same subject pool as the cold-pressors. After completing a consent form that described the cold-pressor task, observers were told that instead of completing the cold-pressor task themselves, they would watch a videotape of someone else doing so. Observers were shown the instructions provided to the cold-pressors and then watched a videotape of a single cold-pressor’s performance. Finally, they rated the target on the same dimensions as the cold-pressors themselves.

Each target was yoked to two observers who followed the same procedure with one exception: One observer was asked, prior to watching the videotape, to submerge her own hand in a bucket of ice water. We ran this “empathic” observer to ensure that any self-other differences we might obtain were not due to observers simply not realizing that the water was everything we said it was (i.e., really, really cold). As it turned out, the two observer conditions did not differ, and so we averaged their responses in the analyses presented next.

Results

We predicted that participants who completed the cold-pressor task would give themselves credit for their intentions when rating their altruism but that observers watching them would not. We tested this hypothesis in two ways. First, we compared the ratings provided by

TABLE 3: Ratings of Cold-Pressor Submersion Time, Altruism, and Altruistic Intent by Self and Observers, Study 5

Rater	Percentile Estimate of		
	Submersion Time	Altruism	Altruistic Intent
Self	41	62*	63*
Observer	54	54	59*

NOTE: Values significantly greater than average (the 50th percentile) are denoted as follows: * $p < .05$.

each cold-pressor to the average ratings made by her two yoked observers in a 2 (rater: self vs. observer) \times 3 (rating: submersion time vs. altruism vs. altruistic intent) ANOVA. We predicted that the cold-pressors’ estimates of their altruism would be greater than their self-rated cold-pressor performance but indistinguishable from their estimates of their altruistic intentions. In contrast, we predicted that the yoked observers would not distinguish between the target’s behavior and the target’s altruism.

As the significant interaction suggests, $F(1, 15) = 11.50$, $p < .001$, $\eta^2 = .61$, and as Table 3 reinforces, all predictions were confirmed. Participants’ assessments of their own altruism, although considerably greater than their self-rated submersion time, $t(16) = 3.49$, $p < .01$, $\eta^2 = .43$, was indistinguishable from their ratings of their altruistic intentions, $t(16) < 1$. The yoked observers, in contrast, rated the target’s altruism only as favorably as they rated the target’s actual task performance. This was true despite the fact that the observers, similar to the targets, perceived a rift between the target’s intentions and actualized behavior, $t(16) = 2.38$, $p < .05$, $\eta^2 = .26$. Thus, both targets and observers saw a discrepancy between actions and intentions; where they differed was in how they interpreted that rift.⁵

A second test of participants’ tendency to give themselves (but not others) credit for their intentions examined the relationship between ratings of altruism, submersion time, and altruistic intent in two multiple regressions: one predicting cold-pressors’ ratings of their own altruism from their ratings of their actions and their intentions, and the other predicting the observers’ ratings of the target’s altruism from their ratings of the target’s actions and intentions. We predicted that whereas participants’ ratings of their own altruism would be predicted better by their altruistic intentions than by their actual behavior, the opposite would be true of participants’ ratings of the altruism of someone else. As Table 4 illustrates, these predictions were supported. Participants’ ratings of their intentions accounted for considerably more variance in their ratings of their own altruism than did their ratings of their actual behavior,

TABLE 4: Relative Weights Placed on Submersion Time and Altruistic Intent in Judgments of Target's Altruism by Targets and by Observers, Study 5

Rater	Rating	
	Submersion Time	Altruistic Intent
Target	-.13	.85*
Observer	.59†	.09

† $p < .10$. * $p < .05$.

whereas the opposite was true of the yoked observers' ratings of the target's altruism.

Discussion

The results of Study 5 support our claim that the tendency for people to give themselves (but not others) credit for their good intentions leads to favorable self-assessments. Participants who underwent a painful ice submersion task for charity rated their intentions to help the charity higher than their actual performance but not higher than their assessment of the altruism they displayed. Observers, in contrast, did not discriminate between the cold-pressor's actions and altruism—despite the fact that they, too, rated intentions more favorably than actions. In short, whereas participants judged others by what those individuals did, they judged themselves by what they tried to do.

GENERAL DISCUSSION

What I've been and done doesn't show my true worth. . . . There remains within me, unused and quite viable, a host of propensities, inclinations, possibilities, that one wouldn't guess from the mere series of things I've done. (Sartre, 1957, pp. 31-32)

It is a fact of life that actions and intentions do not always align. The present studies suggest that when they do not, actors and observers differ in the weight they place on intentions. People give themselves more credit for their own good intentions than they give others for theirs (Studies 1-2). Because actions often fall short of intentions, the use of these different criteria leads individuals to rate themselves more favorably than the average person and more favorably than they are rated by others (Studies 3-5).

We are hardly the first to point out that actors and observers differ in their access to one another's intentions (James, 1890; Jones & Nisbett, 1971; Pronin et al., 2003) or that people treat private information such as intentions as more diagnostic of the self than of others (S. M. Anderson, 1984; S. M. Anderson & Ross, 1984; Malle & Knobe, 1997, 2001). Nor are we the first to suggest that people view their own traits and abilities more

favorably than the traits and abilities of others (e.g., Alicke, 1985; Brown, 1986, 1990; Campbell, 1986; Dunning et al., 1989; Kruger & Dunning, 1999; Kunda, 1990; Taylor & Brown, 1988). The present results, however, represent the first demonstration that the greater weight assigned to one's own intentions relative to those of others is an important determinant of self-enhancement.

Should People Give Themselves Credit for Their Intentions?

One question the present findings leave unanswered is whether this extra credit people give themselves for their good intentions reflects an unwarranted bias or a perfectly defensible judgmental tendency. In some cases, the inclination to base self-assessments not only on what one has done but also on what one has tried to do is entirely justified. Webster's (1999) definition of *considerate*, for instance, is "having or showing regard for others and their feelings" (p. 311). One may be considerate, then—or, for that matter, appreciative, thoughtful, or sympathetic—regardless of whether one's intentions are actualized. What is less clear, however, is whether the failure to give others the same benefit is similarly justifiable. The robust finding across the five studies reported here was that participants not only gave themselves credit for their intentions but they failed to give others credit for theirs.

Note also that although good intentions may make one considerate, they do not, say, make one healthy. Nevertheless, other evidence from our lab suggests that one reason people often underestimate their likelihood of experiencing common health problems such as liver disease or lung cancer (Weinstein, 1980) is that they base their predictions not on the health-related behaviors in which they actually engage but on the health-related behaviors in which they *intend* to engage (Kruger, Gilovich, & Staggs, 2003). Because actions often fall short of intentions, people are not as immune to health problems, relative to their peers, as they believe.

The Role of Intentions in Other Judgmental Biases

Although we have focused on inflated trait judgments, the tendency to weight one's intentions more heavily than those of others may play a role in a number of other judgmental biases. For instance, people tend to underestimate their task completion times, a fact apparent to anyone who has ever reviewed a manuscript for *Personality and Social Psychology Bulletin* or submitted a manuscript for some (usually) anonymous "other" to review. Of interest, although the planning fallacy is common in people's predictions of their own completion times, people's predictions of the completion times of others are typically more accurate (Buehler et al., 1994, 2002). The differential weight that actors and observers

place on intentions may help explain why. Whereas people may give considerable weight to the details of how they intend to complete a given task, they are less likely to do the same for others. As a consequence, people are likely to provide more optimistic assessments of their own task completion times than the completion times of others—precisely the pattern typically observed.

The differential weight assigned to intentions on the part of actors and observers also may play a role in relationship discord. In relationships, as elsewhere, actions and intentions often diverge. A spouse may try to pick up dinner on the way home only to get stuck in traffic, or a roommate may plan to do the dishes but forget. When actions fall short in these ways, individuals are likely to differ in the standing granted to intentions. Whereas good intentions often exonerate one's own faulty behavior, we suspect that they figure less prominently in one's appraisals of the faulty behavior of others.

Judgments of intention appear to play an especially large role in one particular aspect of relationships, namely, teasing. Teasers and their targets tend to differ in their evaluations of teasing: Although it is often seen as innocent and playful by the teaser, it tends to be seen as more malicious by the target (Kowalski, 2000; Kruger, Gordan, & Kuban, 2003). One reason for this discrepancy stems from the fact that teasing is often delivered with the best of intentions, often as playful ribbing designed to inform the target that he or she is well liked (Keltner, Young, Heerey, Oemig, & Monarch, 1998). This intent, however, turns out to be less apparent—and less relevant—to the target of the tease than to the person doing the teasing (Kruger, Gordan, & Kuban, 2003).

Our discussion thus far has focused on intentions to do good that are not actualized in behavior. But there are doubtless pronounced actor/observer differences in the weight assigned to intentions surrounding bad behavior as well. In particular, moral transgressions are likely to be seen as more egregious to observers than to the actors themselves because only the actors know the limits they intended, all along, to place on their wayward actions: the prankster who only intended to frighten, not injure, by tampering with a mountain biker's handbrakes; the flirt who only intended to spice up an evening, not damage another's relationship. In such cases, the actors may cling to their intentions in an attempt to contextualize and circumscribe their misdeeds—pleas that may fall on the deaf ears of observers, who are less inclined to accept the sincerity of the intentions offered. That said, we cannot help but wonder whether there are some occasions in which individuals give themselves less credit for their intentions than they are given by others, such as when the end result is particularly tragic. A mother who accidentally drowns her child, for instance,

may blame herself more than she is blamed by others despite the fact that her actions were unintended.

Why Do People Give Themselves Credit for "The Best of Intentions"?

A final question is the familiar one of whether the phenomenon examined here is cognitive or motivational in nature (or both). People give little weight to the intentions of others because others' intentions, unlike one's own, are largely unknown. But people also may give themselves more credit for their good intentions than they give others to boost their self-esteem. As Sartre (1957) put it, "Often the only way [people] can bear their wretchedness is to think, 'circumstances have been against me—what I've been and done doesn't show my true worth'" (pp. 31-32).

We suspect that both mechanisms operate in tandem. On one hand, participants in Study 5, despite knowing that the participants they had viewed had the best of intentions (see Table 3), failed to consider these intentions when making their overall ratings, consistent with a motivated reasoning interpretation. On the other hand, unpublished data from our lab suggests that people have an easier time thinking of their own intentions than the intentions of others, consistent with an availability interpretation.

Regardless of the cause, the outcome is the same: an inflated view of the self relative to others. It is often said (in reference to others, at least) that "the road to Hell is paved with good intentions." It appears that the road to self-enhancement is similarly constructed.

NOTES

1. The number of examples listed per trait did not differ by condition.
2. These ratings also enable a test of whether the relationship between intentionality and the above-average effect observed in Study 2 is similarly independent of desirability, ambiguity, and observability. It did: the correlation ($r = .46$) remains strong after holding constant desirability ($r = .43, p < .01$), ambiguity ($r = .45, p < .01$), observability ($r = .40, p < .01$), or all three ($r = .40, p < .05$).
3. The means reported in the text are based on raw correlations, whereas the tests are based on Fisher's r to z transformed values.
4. We had originally included a second question to assess altruism. In addition to asking about altruism directly, we also asked participants to rate how "self-sacrificing" they thought the target had been. We dropped this question because it failed to correlate significantly with the altruism measure, $r = .43, p = .084$, and in fact correlated more highly with the behavior measure, $r = .54, p < .025$.
5. Surprisingly, although our analyses revealed that targets gave themselves more credit for their altruistic intentions than they were given by yoked observers, this tendency did not yield a reliable self-other difference in participants' ratings of the target's altruism, $t(16) = 1.31, p = .210, \eta^2 = .10$. One reason may be that our experiment, although designed to create a rift between actions and intentions for all participants, did so only incompletely: nine of the cold-pressors met or exceeded their stated expectations. Excluding these participants, the self-other difference was more pronounced: $M_s = 71$ st percentile (self) vs. 50th percentile (other), $t(7) = 2.36, p < .05, \eta^2 = .44$.

REFERENCES

- Agnes, M. (Ed.). (1999). *Webster's New World College Dictionary* (4th ed.). New York: Macmillan.
- Alicke, M. D. (1985). Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology*, *49*, 1621-1630.
- Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, *9*, 272-279.
- Anderson, S. M. (1984). Self knowledge and social inference II: The diagnosticity of cognitive/ affective and behavioral data. *Journal of Personality and Social Psychology*, *46*, 294-307.
- Anderson, S. M., & Ross, L. (1984). Self knowledge and social inference I: The impact of cognitive/ affective and behavioral data. *Journal of Personality and Social Psychology*, *46*, 280-293.
- The Animals. (1964). *Don't let me be misunderstood* (Song). Written by Bennie Benjamin, Sol Marcus, and Gloria Caldwell. Miami, FL: WB Music Corp.
- Aron, A., & Aron, E. N. (2001). The self-expansion model: Motivation and including other in the self. In W. Ickes & S. Duck (Eds.), *The social psychology of personal relationships*. Chichester, UK: Wiley.
- Brown, J. D. (1986). Evaluations of self and others: Self-enhancement biases in social judgments. *Social Cognition*, *4*, 353-376.
- Brown, J. D. (1990). Evaluating one's abilities: Shortcuts and stumbling blocks on the road to self-knowledge. *Journal of Experimental Social Psychology*, *26*, 149-167.
- 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.
- Buehler, R., Griffin, D., & Ross, M. (2002). Inside the planning fallacy: An empirical exploration of optimistic time predictions. In T. Gilovich, D. W. Griffin, & D. Kahneman (Eds.), *The psychology of intuitive judgment: Heuristics and biases*. New York: Cambridge University Press.
- Campbell, J. D. (1986). Similarity and uniqueness: The effects of attribute type, relevance, and individual differences in self-esteem and depression. *Journal of Personality and Social Psychology*, *50*, 281-294.
- Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality & Social Psychology*, *58*, 568-581.
- Dunning, D., Leuenberger, A., & Sherman, D. A. (1995). A new look at motivated inference: Are self-serving theories of success a product of motivational forces? *Journal of Personality & Social Psychology*, *69*, 58-68.
- 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.
- Felson, R. B. (1981). Ambiguity and bias in the self-concept. *Social Psychology Quarterly*, *44*, 64-69.
- Funder, D. C. (1980). On seeing ourselves as others see us: Self-other agreement and discrepancy in personality ratings. *Journal of Personality*, *48*, 473-493.
- Gilovich, T. (1990). Differential construal and the false consensus effect. *Journal of Personality & Social Psychology*, *59*, 623-634.
- Heine, S. J., & Lehman, D. R. (1997). The cultural construction of self-enhancement: An examination of group-serving biases. *Journal of Personality and Social Psychology*, *72*, 1268-1283.
- Hellstrom, B., & Lundberg, U. (2000). Pain perception to the cold-pressor test during the menstrual cycle in relation to estrogen levels and comparison with men. *Integrative Physiological and Behavioral Science*, *35*, 132-141.
- James, W. (1890). *The principles of psychology*. New York: Henry Holt.
- John, O. P., & Robins, R. W. (1993). Determinants of interjudge agreement on personality traits: The Big Five domains, observability, evaluativeness, and the unique perspective of the self. *Journal of Personality*, *61*, 521-551.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 220-266). New York: Academic Press.
- Jones, E. E., & Nisbett, R. E. (1971). The actor and the observer: Divergent perceptions of the causes of behavior. In E. E. Jones, D. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 79-94). New York: General Learning Press.
- Keltner, D., Young, R. C., Heerey, E. A., Oemig, C., & Monarch, N. D. (1998). Teasing in hierarchical and intimate relations. *Journal of Personality and Social Psychology*, *75*, 1231-1247.
- Klein, W. M. (1999). Justifying optimistic predictions with minimally diagnostic information under conditions of outcome dependency. *Basic & Applied Social Psychology*, *21*, 177-188.
- Klein, W. M., & Weinstein, N. D. (1997). Social comparison and unrealistic optimism about personal risk. In P. B. Buunk & F. X. Gibbons (Eds.), *Health, coping, and well-being: Perspectives from social comparison theory* (pp. 25-61). Mahwah, NJ: Lawrence Erlbaum.
- Kowalski, R. M. (2000). "I was only kidding!": Victims' and perpetrators' perceptions of teasing. *Personality and Social Psychology Bulletin*, *26*, 231-241.
- Kruger, J. (1999). Lake Wobegon be gone! The "below-average effect" and the egocentric nature of comparative ability judgments. *Journal of Personality and Social Psychology*, *77*, 221-232.
- Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*, 1121-1134.
- Kruger, J., Gilovich, T., & Staggs, V. (2003). *Actions and intentions in unrealistic optimism*. Unpublished manuscript.
- Kruger, J., Gordan, C., & Kuban, J. (2003). *Intentions in teasing: When "just kidding" just isn't good enough*. Unpublished manuscript.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480-498.
- Larwood, L. (1978). Swine flu: A field study of self-serving biases. *Journal of Applied Psychology*, *18*, 283-289.
- Malle, B. F., & Knobe, J. (1997). Which behaviors do people explain? A basic actor-observer asymmetry. *Journal of Personality and Social Psychology*, *72*, 288-304.
- Malle, B. F., & Knobe, J. (2001). The distinction between desire and intention: A folk-conceptual analysis. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition*. Cambridge, MA: MIT Press.
- Marks, G., & Miller, N. (1987). Ten years of research on the false consensus effect: An empirical and theoretical review. *Psychological Bulletin*, *102*, 72-90.
- Messick, D. M., Bloom, S., Boldizar, J. P., & Samuelson, C. D. (1985). Why we are fairer than others. *Journal of Experimental Social Psychology*, *21*, 480-500.
- Paunonen, S. V. (1989). Consensus in personality judgments: Moderating effects of target-rater acquaintanceship and behavior observability. *Journal of Personality and Social Psychology*, *56*, 823-833.
- Pronin, E., Gilovich, T., & Ross, L. (in press). Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others. *Psychological Review*.
- Ross, L., Green, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attributional processes. *Journal of Experimental Social Psychology*, *13*, 279-301.
- Sartre, J. P. (1957). *Existentialism and human emotions*. New York: Philosophical Library.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin*, *103*, 193-210.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, *39*, 806-820.

Received November 1, 2002

Revision accepted April 16, 2003