

Available online at www.sciencedirect.com

Journal of Experimental Social Psychology xxx (2008) xxx–xxx

Journal of
Experimental
Social Psychologywww.elsevier.com/locate/jesp

Do people really believe they are above average? ☆

Elanor F. Williams *, Thomas Gilovich

Department of Psychology, Cornell University, 211 Uris Hall, Ithaca, NY 148503, USA

Received 31 August 2007; revised 20 December 2007

Abstract

A question that has plagued self-enhancement research is whether participants truly believe the overly positive self-assessments they report, or whether better-than-average effects reflect mere hopes or self-presentation. In a test of people's belief in the accuracy of their self-enhancing trait ratings, participants made a series of bets, each time choosing between betting that they had scored at least as high on a personality test as a random other participant, or betting on a random drawing in which the probability of success was matched to their self-assigned percentile rank on the test. They also reported the point at which they would switch their bet from their self-rating to the drawing, or vice versa. Participants were indifferent between betting on themselves or on the drawing, and it took only a slight change in the drawing's probability for them to switch their bet, indicating that people truly believe their self-enhancing self-assessments.

© 2008 Published by Elsevier Inc.

Keywords: Self-enhancement; Calibration; Above average effect

18 Q4 People are unlikely to know precisely where they stand
19 among their peers on such traits as warmth, wisdom, or
20 wastefulness, or on such abilities as wit, whistling, or
21 woodworking. That has not stopped psychologists from
22 asking people precisely where they think they stand on such
23 dimensions. Numerous investigators, including both
24 authors of this paper, have asked participants to assign
25 themselves a single percentile value specifying their stand-
26 ing among their peers (e.g., Dunning, Meyerowitz, & Holz-
27 berg, 1989; Gilovich, Kerr, & Medvec, 1993; Kruger, 1999;
28 Kruger & Dunning, 1999; Williams & Gilovich, in prepara-
29 tion). But what are we to make of these ratings? Imagine
30 someone who is uncertain of her standing on the trait of
31 intelligence, thinking she might rank as high as the 70th
32 percentile and no lower than the 40th. Is she likely to report
33 the midpoint of these high and low values (55th) as her per-

centile estimate? Or is she likely to give herself the benefit
of the doubt and report an estimate close to her subjective
ceiling?

Implicit in these questions is the broader issue of
whether people stand by their percentile estimates. The
issue is important in light of the frequently-observed ten-
dency for people's percentile estimates to yield an "above
average effect" (e.g., Dunning, Heath, & Suls, 2004). On
most positive traits and abilities, the average response is
well above average—roughly at the 60th or 65th percentile
across positive traits (e.g., Dunning et al., 1989; Kruger &
Dunning, 1999). Do people really mean it when they say
they are, on average, above average?

Surprisingly, this question has not been addressed in
previous research. We are unaware of any study that has
assessed whether people are truly committed to their self-
assessments, treating them the same way they treat, say,
the likelihood of various aleatory events. Does someone
who states that he is at the 75th percentile on the dimension
of friendliness think that his chances of being friendlier
than a randomly-selected peer are the same as rolling a four
or less with a six-sided die? Or, when push comes to shove
and people are asked to anchor their assessments in objec-

☆ This research was supported by National Science Foundation Research Grants SES0542486 and SES0241638. We thank David Dunning for his helpful comments on an earlier draft, and Michelle Jernstedt for her help in data collection.

* Corresponding author. Fax: +1 6072558433.

E-mail address: efw7@cornell.edu (E.F. Williams).

tive terms, do people back off their estimates, acknowledging that their estimates were inflated? The research we report here was designed to examine this question.

There is some evidence in the literature on self-assessment that suggests that people's typical estimates are purposefully (and hence knowingly) inflated. For example, people's confidence in their ability to accomplish various tasks reliably diminishes as the moment to perform the task draws near (Carroll, Sweeny, & Shepperd, 2006; Gilovich et al., 1993; Shepperd, Ouellette, & Fernandez, 1996). There are many possible interpretations of this result, not all of them involving deliberate distortion. But one plausible interpretation is that people feel free to give knowingly inflated estimates of their likely success far from the moment of truth and there is no possibility—or no imminent possibility—of their predictions being disconfirmed. As the moment of truth draws near, however, people may feel accountable to what is about to happen and thus rein in their inflated estimates accordingly. This work can be taken to imply, in other words, that people would tend to stand behind their proximate estimates but not their distant estimates.

More recent research reinforces this possibility. Armor and Sackett (2006) found that for hypothetical tasks, people make overconfident predictions about their performance, predictions that are uncorrelated with their actual performance. When these same tasks are presented as real ones the participants are about to undertake, however, people's predictions are much less overconfident, and correlate quite strongly with their actual performance. The authors suggest that predictions about hypothetical or non-imminent tasks are construed at a more abstract level, and thus are more ambiguous, allowing for a more optimistic interpretation of what the tasks will be like and how one is likely to perform (Dunning et al., 1989). In support of this contention, specifying the conditions under which participants are to perform the tasks constrains their ability to construe them in a self-serving manner, much like how, in the context of trait-based self-assessments, people who make comparisons to specific, individuated others are less able to self-enhance (Alicke, Klotz, Breitenbecher, Yurak, & Vredenburg, 1995).

Further evidence that people may be less than fully committed to the self-assessments they typically provide in psychology experiments comes from research on accountability (Lerner & Tetlock, 1999; Sedikides, Herbst, Hardin, & Dardis, 2002). That is, accountability appears to reduce overconfidence and self-enhancement. In one study, participants wrote and graded their own essays. Those who were told they would have to justify their grades to a graduate student in English Composition assigned themselves lower grades than those who were led to believe that their self-assigned grades were anonymous (Sedikides et al., 2002). Furthermore, accountability may not just diminish inflated self-assessments; it may also increase calibration, such that the correlation between accuracy and confidence is higher for accountable participants than for unaccountable participants (Lerner & Tetlock, 1999).

Although these findings hint that people may not really mean what they say when they provide percentile estimates of their abilities, studies in the confidence and calibration literature suggests that they may be truly committed to their estimates. For example, Dunning, Griffin, Milojkovic, and Ross (1990) had participants make predictions about another person's behavior and state how confident they were that their predictions were correct. Then, inspired by de Finetti's exchangeability hypothesis (Edwards, Lindman, & Savage, 1963), they offered the participants a series of gambles, one for each prediction. For each gamble, they were given a choice between betting that their prediction of the target person's behavior was correct (for which, if they were properly calibrated, the odds of success were equivalent to their stated confidence) or betting on a random drawing, with the chance of winning varying between 55% and 95%. Thus, for some of the bets, the random drawing had a lower chance of paying off than their prediction about the target's behavior; for others it had a higher chance of paying off; and for others it had an equivalent chance of paying off. Consistent with the idea that participants stood behind their estimates, they tended to choose the bet with the higher payoff, regardless of whether it involved betting on their prediction or betting on the drawing, and they were indifferent between the two when their chance of success was equal. If the participants were aware that they were overconfident, one would expect to see significantly more bets being placed on the drawing than on their predictions, but this was not the case (see also Fischhoff, Slovic, & Lichtenstein, 1977).

Similar results were obtained by Heath and Tversky (1991), who gave participants a choice between betting that their judgments were correct or betting on a random lottery, for which the chances of winning were equal to their stated chances that their judgments were correct. When the domain of judgment was one about which they did not feel highly knowledgeable, and thus they did not feel confident about their judgments, they were more likely to bet on the lottery. However, when the domain was one in which they were knowledgeable and they were confident in their judgments, they were more likely to bet on their judgments than the random drawing, and were willing to pay a premium to do so. Likewise, Camerer and Lovallo (1999) obtained similar results in their study of excess entry into financial markets. They found that investors were willing to bet on a random market at rates almost exactly commensurate with the risk involved, but were willing to bet on a market based on their own competence at rates unwarranted by their probability of success.

What, then, is one to make of the self-assessments documented in the self-enhancement literature? Are they the result of people knowingly giving themselves the benefit of considerable doubt—and thus represent judgments to which they have little commitment? Research showing that people back away from their estimates about their future performance when the tasks they face become more real, more imminent, and more public suggests that they may

indeed. Or, are people's assessments of their standing on various trait dimensions like the confidence estimates they provide in calibration paradigms—overly optimistic but fully endorsed? The present research used an exchangeability procedure to find out. The question of interest was whether people would be committed to their assessments—treating, for example, the statement that they are at the 67th percentile in intelligence the same as a lottery bet with a two-thirds chance of success—or whether they would back away from their estimates when push came to shove, knowing that their estimates had been shaded in a flattering direction.

Overview

Participants were given a bogus personality test and then asked to predict how they would score on it relative to other Cornell students. They were asked to make predictions about their scores on four traits known to yield above average effects. After making their ratings, participants were told that their actual scores on the test would be compared to those of a randomly-selected participant. With this in mind, they were to make a series of gambles, one for each of the four traits. For each gamble, they could bet on whether they would score higher on the test than the other participant, or they could bet on a random drawing with a probability of success equal to the percentile ranking they had assigned themselves on the trait in question. After they chose their bet, they were also asked at which point, after we altered the probability of winning the random draw, they would switch their bet to the other option. If participants' self-assessments are knowingly inflated, as work on accountability and the time course of self-enhancement suggests they may be, they should prefer to bet on the random drawing. They would know that the random drawing offered them a better chance of winning. If, however, participants are committed to their percentile estimates, as previous work on the calibration of confidence judgments has indicated, they should be indifferent between betting on their standing on each of the traits and betting on the random draw.

Methods

Participants

Forty-nine Cornell undergraduates participated for extra-credit in various psychology courses and an opportunity to win up to \$4.

Procedure

Participants arrived individually and were administered a bogus personality test consisting of 58 yes-or-no statements, such as "You often do jobs in a hurry," and "You are consistent in your habits." The test was designed to have a mix of transparent and opaque questions to give

it the feel of a genuine personality inventory. After completing the questionnaire, participants were asked to rate how they believed had they scored, relative to other Cornell students, on four traits—intelligence, creativity, maturity, and positivity. More specifically, participants were asked to assign themselves a percentile value, representing the percentage of their Cornell peers they would score higher than, on each of the traits.

Participants were then told that they were to make a series of four bets, each with a chance to win \$1. On each bet, they were given the option of betting on (a) drawing a numbered token from an urn, or (b) the prospect of scoring higher than a randomly selected Cornell student on the trait in question. Participants were shown the urn containing the numbered tokens and were told they had been randomly yoked to an earlier participant, someone whose identity they would never know and who would never know theirs. They were given the following detailed instructions about the bets:

In one, you would bet on whether you have scored higher on a particular trait, as measured by the test you just took, than the person with whom you were randomly paired. Or, you can bet on whether you select, without looking, a piece from this jar here that contains a number with a value equal to or below some cutoff number. The cutoff number will be set up so that the probability of drawing a value that equals or falls below it will be the same as the percentile ranking you gave yourself on the trait in question.

Let's say, for example, that you ranked yourself at the 60th percentile on intelligence. This means that you believe that of 100 random Cornell students, you will rank higher than 60 of them, and lower than 40 of them. This would mean that you have a 60% chance of winning the bet if you chose to bet on your standing on intelligence. If, however, you were to choose the matched-probability bet, you would choose, without looking, a piece from this jar here. Each of the pieces has a number on it, from 1 to 100. To win, the piece you chose from the jar would have to be numbered 60 or below, giving you a 60% chance of winning.

After it was clear that participants understood this part of the task, they were also asked how much of a change in probabilities it would take for them to switch their bet from the jar to their standing on the trait in question, or vice versa. They were asked whether they would change their bet if the probability of a winning draw from the jar increased or decreased by 1 point, 2 points, and so on until 5 points above or 5 points below the original matched bet. Although no participants changed their bet in the "wrong" direction (for example, changing their bet from the self to the jar if the jar's probability of winning were to decrease), they were still asked if they would switch at any of the five "wrong" points to ensure that they understood the nature of the process. In addition, if none of the five new probabilities was a large enough change to prompt the partici-

pants to switch their bet, they were asked outright how much the jar's probability of paying off would have to increase or decrease to get them to switch. Because this process would most likely be unfamiliar to participants, the spoken instructions were accompanied by a visual aid depicting two yardsticks with moving markers (see Fig. 1), which demonstrated how a bet on one's self-assessment and a bet on the drawing would compare, and also how changing the probability of a winning draw from the jar would compare to the original bet they had placed.

After this process was completed for all four traits, participants were carefully debriefed, being told that: (1) the test was constructed specifically for this experiment and cannot, in fact, predict or assess their personality, and (2) that their scores would never be compared to other participants' scores. Because they were offered the opportunity to win cash in addition to the extra credit they were to receive, they were given the opportunity to draw a piece from the jar, with the color of the piece determining how much they won. There were equal numbers of five different colors in the jar, and each color corresponded to a dollar amount from \$0 to \$4.

Results

Self-enhancement

To test whether people are committed to their beliefs that they are above average on positive traits, it is necessary to establish that participants did, in fact, self-enhance. A composite of the four relevant trait ratings revealed that participants did indeed make self-enhancing estimates on the four traits, with the average rating well above the 50th percentile ($M = 60.96, SD = 13.33$), $t(48) = 5.76$, $p < .001$, $d = .82$. The four individual traits, *intelligent*

($M = 55.51, SD = 16.75, t(48) = 2.30, p < .05, d = .33$), *creative* ($M = 56.22, SD = 18.94, t(48) = 2.30, p < .05, d = .33$), *mature* ($M = 67.86, SD = 15.44, t(48) = 8.09, p < .001, d = 1.16$), and *positive* ($M = 64.24, SD = 21.11, t(48) = 4.73, p < .001, d = .67$), each revealed a significant above average effect as well.

Commitment

The distribution of participants' bets on their percentile ratings versus a draw from the jar should reveal whether they truly believe the self-assessments they provided. If they were to bet disproportionately on their percentile ratings, one would have to infer that participants thought they had understated their true standing and that their chance of outscoring their yoked peer was actually higher than their stated percentile standing. If they were to bet disproportionately on the random draw, one would have to infer that participants thought they had overstated their true standing and that their chance of outscoring their yoked peer was lower than their stated percentile standing. If, however, participants stand behind their self-assessments, they should show no preference between betting on the random draw and betting on their self-ratings.

Across all 196 bets, participants were in fact indifferent between betting on their chances of scoring higher on the personality test than their yoked peer and betting on the probability of winning the random draw. Participants placed 95 bets on the random draw and 101 bets on the chance that they would score higher than their randomly-selected peer, $\chi^2(1, N = 196) = .18, ns$. This indifference was apparent on three of the four individual traits, (*creative*, *mature*, and *positive*; all $\chi^2 < .51, ns$), with participants exhibiting a marginally significant preference to bet on the




Fig. 1. A graphical depiction of the visual aids used to help participants understand the nature of the series of bets they were to make. Fig. 1a represents the initial choice between bets. The top yardstick represents the participant's self-rating and thus their supposed chance of scoring higher than the other person—in this case 55%. The lower yardstick represents the matched-probability of drawing a winning token from the urn. Fig. 1b represents a subsequent choice, assuming the participant originally chose to bet on her self-ratings. The probability of drawing a winning token from the jar has increased by 3% to 58%. The participant then decides whether that change in probability is enough for her to switch her bet from her self-ratings to the random drawing.

jar rather than their relative standing when it came to the trait of *intelligence*, $\chi^2(1, N = 49) = 3.45, p = .06, \phi = .27$.

The point at which participants would choose to switch their bets also speaks to how committed they were to their self-assessments. If they believed their estimates to be accurate, and hence were indifferent to betting on their own standing or the random draw, they should be willing to switch their bet to the rejected option when the probability of winning with the random draw is altered only slightly. Fig. 2 presents the percentage of participants who switched their bets when the probability of winning with the random draw was increased (for those who initially bet on themselves) or decreased (for those who initially bet on the draw) in 1% increments. Consistent with the idea that participants were truly committed to their original estimates of their own percentile standing, they tended to be quite responsive to changes in the probability of winning with the random draw. When the chances of winning with the random draw were increased by a single percentage point, participants who initially bet on themselves stated that they would switch to a bet on the draw 14.7% of the time, and that they would do so 60.4% of the time when the probability of winning with the random drawing was increased by 5 percentage points. Similar results were obtained from participants who originally bet on the random draw. When the chances of winning with the random draw were decreased by a single percentage point, they stated that they would switch to bet on themselves 14.8% of the time, and would do so 60.0% of the time when the probability of winning with the random draw was decreased by 5 percentage points.

For the 77 bets that participants said they would not switch even when the chances of winning with the random draw was altered by 5%, the median change in the probability of winning with the draw that would prompt them to alter their original bet was just 10 percentage points. Clearly, it did not take much of a change in the probability

of winning the random draw to alter participants' preference for betting on their own standing versus the outcome of a random draw. And participants were not more reluctant to switch one type of bet over the other, requiring, for example, a much larger change in probability to switch from the jar to the self-ratings. Those bets that would require larger changes in probability to be switched were evenly balanced between initial bets on self-ratings and initial bets on the jar. The mean change in probability that would prompt a switched bet, across all bets as well as only those that require a change in probability greater than 5%, is not significantly different from 0, ($M_{\text{overall}} = -1.38, t(195) = -1.41, n; M_{>5\%} = -3.38, t(76) = -1.4, ns$). It thus appears that participants took their percentile estimates seriously, viewing an estimate that, say, they fall at the 67th percentile on the trait of creativity much like they would the probability of drawing from an urn with a two-thirds probability of winning.

Discussion

Psychologists have known for at least thirty years that people are prone to unrealistically positive assessments of their traits and abilities—and they probably suspected it much longer than that. Researchers have offered two interpretations of this effect, both viable: (1) people knowingly inflate their standing on positive traits, and thus do not truly believe them and would modify them downward if held responsible for their estimates; and (2) people fully believe their positive self-assessments, and are willing to take actions based on them that are commensurate with the ratings they have provided. We sought to test whether people do, in fact, truly believe the self-enhancing trait ratings to which they are prone.

Participants in our study were indifferent between betting on the percentile rankings they assigned themselves and a matched-chance random drawing, indicating that they believed the two numerically equivalent probabilities—the probability that they would score higher than a random person on a personality test and the probability that they would win a random lottery—were truly equal. Their indifference between the two bets indicates that they believe they were neither overestimating nor underestimating their standing among their peers on the traits in question. In addition, participants did not hesitate to change their bets when the relative probabilities were altered only slightly, indicating that they truly believed the two initial bets gave them an equal chance of winning, and therefore that a small change in one should prompt a switch to the other.

A critic might argue that our results are an artifact of self-presentation—that participants gave the responses they did in order to convey an impression that they are the type of person whose word can be trusted. Two elements of our study suggest otherwise, however. First, real money was riding on the bets participants chose. Therefore, if participants were aware that they had over-




Fig. 2. The percentage of bets that was switched from an initial bet on the self to a bet on the random drawing when the probability of the drawing paying off increased by 5 percentage points or less (triangles), and the percentage of bets that were switched from an initial bet on the drawing to a bet on the self when the probability of the drawing paying off decreased by 5 percentage points or less (circles).

underestimated themselves but attempted to bet as if they had not, they did so at a real cost to themselves. Second, suppose participants were interested in conveying the impression that they were committed to their estimates. How would they do so? By betting on their standing relative to the other Cornell student or by betting on the random draw with the same stated probability? Neither bet does much to support or impugn the impression that one stands behind one's original percentile estimate. Indeed, the fact that participants bet almost exactly as often on the random draw as they did on their own relative standing indicates that no systematic self-presentational artifact was at play. Although some participants split their four bets evenly between their own standing and the random draw, many did not. Twenty-two of the 49 participants chose three bets of one type and one of the other and four participants chose all four bets of the same type.

Note that although our results indicate that people truly believe in their (self-enhancing) trait ratings, treating an estimated likelihood that they rank higher than 65% of their peers as equivalent to a random drawing with a 65% chance of success, we cannot specify just how confident participants are in their estimates. A person might be absolutely convinced that she ranks at the 65th percentile on the dimension of creativity, or she might merely think that that is her most likely standing, with little faith in the precise value. Distinguishing between these two possibilities and ascertaining exactly how confident people are in their trait ratings was not the purpose of this research. Instead, the aim was to ascertain whether the self-enhancing trait ratings frequently reported in the literature are the result of deliberate distortion or whether they are truly believed—at whatever level of confidence. Our results clearly support the latter. Participants took their estimates sufficiently seriously to guide their choices, both in terms of which bets to pick initially as well as when it was appropriate to switch to the other. They were sufficiently confident in them, in other word, to act on them.

Several processes have been offered to account for self-enhancement, including some form of wishful thinking (e.g., Sedikides, 1993), self-serving resolutions of ambiguity (e.g., Dunning et al., 1989), and an egocentric substitution error (e.g., Klar & Giladi, 1997; Kruger, 1999; Moore & Small, 2007). It is possible that any or all of these processes gave rise to the ratings we report here, and we make no claim about the source of the self-enhancement we observed. What is clear, however, is that participants, as a whole, self-enhanced. The average trait ratings were significantly above the 50th percentile on all four traits. Although we cannot say, of course, whether any given participant's ratings were overly optimistic (e.g., Armor & Taylor, 1998; Colvin & Block, 1994), the ratings as a whole did not comport with reality. As long as the comparison group is an appropriate one (in this case, Cornell students as a whole), and the participants are not a select sample from that group, the average estimated percentile ranking must equal the 50th. The mean ratings we observed were

significantly higher than that and hence constitute evidence of significant self-enhancement.

Our findings have implications for the longstanding controversy over the question of whether various biases in everyday judgment are “real,” or simply the product of artificial laboratory assessment procedures that encourage an excessively negative assessment of human capabilities (Gigerenzer, 1996; Gilovich, 1991; Gilovich & Griffin, 2002; Kahneman & Tversky, 1996; Krueger & Funder, 2004). Had participants in our study provided evidence that their initial assessments were inflated, the results would have to be taken as evidence that there is indeed something wrong with the usual assessment procedures used in this area of research and that the problems with the methods promote a misleadingly negative impression of people's capacity to evaluate themselves accurately. But our participants' readiness to exchange one type of bet with another when the stated probability of the aleatory event was slightly altered indicates that the self-enhancement observed in this study is anything but an artifact. Overly optimistic self-assessments thus appear to be a genuine feature of everyday experience. The main contribution of this research, then, is to demonstrate that these self-assessments are not an artifact of the methods that are used to assess them, but instead reflect beliefs to which people are genuinely committed.

The results of this study provide further evidence of a “bias blindspot” in everyday judgment (Ehrlinger, Gilovich, & Ross, 2005; Pronin, Gilovich, & Ross, 2004). Although people seem to be aware that others are prone to self-enhancing trait ratings (e.g., Krueger, 1998), they clearly treat their own ratings as unbiased summaries of their true standing. To the extent that the causes of unrealistically favorable self-assessments stem from purely cognitive processes, blindness to one's own self-enhancement is nearly inevitable. Those who exert great effort to drive carefully, for example, will quite naturally think of carefulness when evaluating driving ability, note their own efforts to be careful, and conclude that they are especially good drivers (e.g., Dunning et al., 1989; Svenson, 1981). And nothing in this process is likely to trigger any recognition that their resultant judgments are biased. Even many motivational sources of self-enhancement are unlikely to leave palpable evidence of bias. Accumulating evidence indicates that motivational biases are not heavy-handed and transparent, leading people simply to seize on desirable propositions or willfully adopt favorable self-assessments. Instead, they appear to operate more subtly, leading individuals to employ different—but, individually, entirely reasonable—standards for evaluating propositions they want to believe and those they want to reject (Ditto & Lopez, 1992; Gilovich, 1991; Lord, Ross, & Lepper, 1979). Thus, whether a given instance of self-enhancement is cognitive or motivational in origin, the result is the same: the assessment rings true to the one who made it. So true, in fact, that it is experienced as more of a fact about the world than a product of judgment (e.g., Pronin et al., 2004), and therefore as some-

thing exchangeable with a comparable fact about the world, such as an aleatory gamble with equal probability.

What is one to make of the finding that people who feel socially accountable for their judgments tend to make less flattering assessments—or that people's estimates of their likely performance are less favorable right before the time to perform is at hand than when it is far in the future? Do people view one set of assessments as more valid and "real" than another? Are they aware of the inconsistency? We suspect not. Research from across a broad range of areas in psychology indicates that many judgments and mental representations are constructed "on the fly" using information that is highly context-dependent and hence yield judgments that are quite variable and sensitive to the existing circumstances (e.g., Bem & McConnell, 1970; Kahneman & Tversky, 1982; Slovic & Lichtenstein, 1983; Tversky & Thaler, 1990; Yeh & Barsalou, 2006). In the present case, some circumstances lead to the recruitment of mental operations and subsets of existing knowledge that yield highly flatteringly self-assessments; other circumstances encourage the recruitment of different mental operations and subsets of information that yield more modest assessments. The judgments rendered in both circumstances are experienced as entirely valid—as appropriate conclusions derived from the pertinent facts. Neither judgment, one rendered from a distant temporal perspective and the other from a more proximate perspective, or one rendered under conditions of accountability and the other not, is experienced as more real or valid than the other. It thus appears that people are *flexible* (Seligman, 2006) or *situated* (Armor & Taylor, 1998) *optimists*—but genuine optimists nonetheless.

Uncited reference

Paulhus and Levitt (1987).

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., & Sackett, A. M. (2006). Accuracy, error, and bias in predictions for real and hypothetical events. *Journal of Personality and Social Psychology*, *91*, 583–600.
- Armor, D. A., & Taylor, S. E. (1998). Situated optimism: Specific outcome expectancies and self-regulation. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 30, pp. 309–379). New York: Academic Press.
- Bem, D. J., & McConnell, H. K. (1970). Testing the self-perception explanation of dissonance phenomena: On the salience of premanipulation attitudes. *Journal of Personality and Social Psychology*, *14*, 23–31.
- Camerer, C., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, *89*, 306–318.
- Carroll, P., Sweeny, K., & Shepperd, J. A. (2006). Forsaking optimism. *Review of General Psychology*, *10*, 56–73.
- Colvin, C. R., & Block, J. (1994). Do positive illusions foster mental health? An examination of the Taylor and Brown formulation. *Psychological Bulletin*, *116*, 3–20.

- 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*, *65*, 568–584.
- Dunning, D., Griffin, D., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*, 568–581.
- Dunning, D., Heath, C., & Suls, J. M. (2004). Flawed self-assessment: Implications for health, education, and the workplace. *Psychological Science in the Public Interest*, *5*, 69–106.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of others. *Journal of Personality and Social Psychology*, *57*, 1082–1090.
- Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inferences for psychological research. *Psychological Review*, *70*, 193–242.
- Ehrlinger, J., Gilovich, T., & Ross, L. (2005). Peering into the bias blind spot: People's assessments of bias in themselves and others. *Personality and Social Psychology Bulletin*, *31*, 680–692.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 552–564.
- Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.
- Gilovich, T., & Griffin, D. (2002). Heuristics and biases: Then and now. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 1–18). Cambridge: Cambridge University Press.
- Gilovich, T., Kerr, M., & Medvec, V. H. (1993). Effect of temporal perspective on subjective confidence. *Journal of Personality and Social Psychology*, *64*, 552–560.
- Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, *4*.
- Kahneman, D., & Tversky, A. (1982). The psychology of preferences. *Scientific American*, *246*, 160–173.
- Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, *103*, 582–591.
- Klar, Y., & Giladi, E. E. (1997). No one in my group can be below the group's average: A robust positivity bias in favor of anonymous peers. *Journal of Personality and Social Psychology*, *73*, 885–901.
- Krueger, J. (1998). Enhancement bias in description of self and others. *Personality and Social Psychology Bulletin*, *24*, 505–516.
- Krueger, J., & Funder, D. C. (2004). Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral and Brain Sciences*, *27*, 313–327.
- Kruger, J. (1999). Lake Wobegone 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.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, *125*, 255–275.
- 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.
- Moore, D. A., & Small, D. A. (2007). Error and bias in comparative judgment: On being both better and worse than we think we are. *Journal of Personality and Social Psychology*, *92*, 972–989.
- Paulhus, D. L., & Levitt, K. (1987). Desirable responding triggered by affect: Automatic egotism? *Journal of Personality and Social Psychology*, *52*, 245–259.

- 671 Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the
672 beholder: Divergent perceptions of bias in self versus others. *Psycho-*
673 *logical Review*, *111*, 781–799. 684
- 674 Sedikides, C. (1993). Assessment, enhancement, and verification determi-
675 nants of the self-evaluation process. *Journal of Personality and Social*
676 *Psychology*, *65*, 317–338. 685
- 677 Sedikides, C., Herbst, K. C., Hardin, D. P., & Dardis, G. J. (2002).
678 Accountability as a deterrent to self-enhancement: The search for
679 mechanisms. *Journal of Personality and Social Psychology*, *83*, 592–605. 686
- 680 Seligman, M. E. P. (2006). *Learned optimism: How to change your mind*
681 *and your life*. New York: Vintage Books. 687
- 682 Shepperd, J. A., Ouellette, J. A., & Fernandez, J. K. (1996). Abandoning
683 unrealistic optimism: Performance estimates and the temporal prox- 688
- imity of self-relevant feedback. *Journal of Personality and Social*
Psychology, *70*, 844–855. 689
- Slovic, P., & Lichtenstein, S. (1983). Preference reversals: A broader
perspective. *American Economic Review*, *73*, 596–605. 690
- Svenson, O. (1981). Are we all less risky and more skillful than our fellow
drivers? *Acta Psychologica*, *47*, 143–148. 691
- Tversky, A., & Thaler, R. H. (1990). Anomalies: Preference reversals.
Journal of Economic Perspectives, *4*, 201–211. 692
- Williams, E.F., & Gilovich, T. (in preparation). *Take the best: The impact*
of typical and extreme performances in self-assessments and the
assessments of others, Cornell University. Q3 693
- Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts.
American Journal of Psychology, *119*, 349–384. 694
695
696
697

UNCORRECTED PROOF