Running head: DIFFERENTIAL INFORMATION

When it's Rational for the Majority to Believe that They are Better Than Average

Don Moore
Carnegie Mellon University

Deborah Small
University of Pennsylvania

To appear in J. I. Krueger (Ed.), *Rationality and social responsibility: Essays in honor of Robyn M. Dawes*. Mahwah, NJ: Erlbaum.

The authors benefited from the support of the staff and facilities of the Center for Behavioral Decision Research at Carnegie Mellon University. This material is based upon work supported by the National Science Foundation under Grant No. 0451736. Thanks to Robyn Dawes and Joachim Krueger for helpful comments on the manuscript. Please address correspondence via email to: don.moore@alumni.carleton.edu.

When it's Rational for the Majority to Believe that They are Better Than Average

Behavioral decision research draws much of its energy from the fruitful tension between what people do and what people ought to do. Robyn Dawes brought this energy to bear on the false consensus effect when he asked how exactly people *ought* to use their own behavior to help them predict the behavior of others (Dawes, 1989, 1990; Dawes & Mulford, 1996). Dawes's insight came from the application of Bayes's Rule, a fundamental principle of normative decision making that routinely clashes with human intuition.

Bayes's Rule specifies how one ought to update one's beliefs based on evidence. Say, for instance, I am asked to estimate the probability that someone like me would consent to wearing a big sign reading "REPENT" around campus for an hour. This request is sufficiently unusual that I probably do not have strong prior beliefs about it—I do not begin this problem with much useful information. But I do have access to one useful data point—myself. Would I wear the "REPENT" sign? In their study, Ross, Greene, and House (1977) found a substantial difference between the beliefs of Stanford students who agreed to wear the sign and those who refused. Those who agreed believed that a greater proportion of other Stanford students would likewise agree than did those who refused.

Dawes (1989; 1990) suggested that this so-called "false consensus" effect might not be so false after all. Given that a majority of people are always in the majority on any given issue, people can improve the quality of their predictions if they assume that others will behave the way they do. Just so, Bayesian logic suggests that, if people use themselves as a useful data point in making inferences about others, people will expect others to be more similar to them than they actually are. This insight clarified exactly how much people ought to rely on what they know about themselves when making inferences about others. This represented real theoretical progress, and opened the door for researchers to ask whether there is a truly false consensus effect that is greater than what Bayesian logic would predict.

Better-Than-Average Effects

This paper extends the logic of Dawes's Bayesian critique to a different phenomenon: the so-called "better-than-average" (BTA) effect. Better-than-average effects describe the tendency for people to believe that they are different from—and better than—others. In the words of Peterson (2000): "Apparently, in our minds, we are all children of Lake Wobegon, all of whom are above average" (p. 45). Indeed, a number of important theories have begun with the assumption that human judgment is biased in this way, and have sought to explain it (Baumeister, 1998; Benabou & Tirole, 2002; Brown, 1998; Daniel, Hirshleifer, & Sabrahmanyam, 1998; Dunning, 1993; Epstein, 1990; Greenwald, 1980; Steele, 1988; Taylor & Brown, 1988). We will see that there are circumstances under which better-than-average beliefs may be justified, by applying the same Bayesian logic that Dawes used in his critique of the false consensus effect.

Evidence across many domains has demonstrated the important implications of BTA biases. If stock market investors believe they are better than other investors at identifying promising investments, that might help explain why there is so much more trading activity in most financial markets than would be predicted on the basis of traditional economic theory (Odean, 1998). If CEOs believe they are better managers than are other CEOs, that could help explain why there are so many more corporate acquisitions than there ought to be (Malmendier & Tate, 2005). If disputants believe that their claims are more justified than are those of others, it could explain the prevalence of costly and inefficient conflicts like labor strikes and lawsuits going to trial (Babcock & Loewenstein, 1997; Neale & Bazerman, 1985). And if nations believe that their armies are stronger than those of other nations, it could help explain their willingness to go to war (Howard, 1983).

As a rule, scholars have pronounced better-than-average beliefs as biases produced by self-serving cognitions that lead people to see themselves in an unrealistically positive light. In this paper we explore the normative issues surrounding the question of how people ought to infer the abilities and performances of others. In the process, we call into doubt the degree to which BTA effects are really biases. Our theory implies that there are circumstances under which a majority of rational Bayesians ought to believe that they are better than others. This theory presents a number of testable hypotheses. We test some of the key hypotheses in a series of three experiments.

Normative Explanations

Under what circumstances could more than 50% of rational people believe that they are better than average? There are at least two situations in which this can be the case. First, more than 50% of people will be above average in a negatively skewed distribution. For instance, most people have more legs than average.[1] If 1% of the population is missing a leg, then the average number of legs per person is 1.99, and 99% of the population has more legs than average. Negatively skewed distributions are especially likely whenever ceiling effects limit performance, such as the fact that a surplus of legs (over two) is highly improbable.

The second instance in which rational people will believe, on average, that they are better than average, is when they have imperfect information about performance and outcomes are better than expected (Moore & Small, in press). Consider an illustrative example. I attempt to accomplish some feat that I do not expect everyone to be able to accomplish, such as passing my driver's test. Let's say, coming in to it, I expect 70% of people to pass. I pass. Therefore, I am above average. But maybe part of the reason for my success is that the task is easy, so I should raise my estimate of the percentage of others who pass. But as long as I believe the average pass rate to be below 100%, I will be above average. If everyone expects a 70% pass rate but the test turns out to be easier than that, then the average test-taker will believe that he or she is above average.

This explanation holds similarly for tasks on which performance is not limited to 100% (success) and 0% (failure). To be precise, when average performance exceeds average expectations, then people will be normatively justified in believing that they are better than average (and better than the median, for that matter). We will refer to this as the differential information explanation, because it posits that differential information leads estimates of self and others to be differentially regressive to the prior. Where do priors come from? Often, people have some relevant experience, but evidence suggests that when people are unsure about how outcomes will be distributed, their predictions gravitate towards a uniform distribution in which all outcomes are predicted to be equally likely (Bruine de Bruin, Fischhoff, Millstein, & Halpern-Felsher, 2000; Fox & Rottenstreich, 2003).

We should note that three sensible assumptions are essential to obtaining this differential information effect. First, people have better information about themselves than they do about others. Second, people believe some portion of their own performances is due to their own idiosyncratic abilities or to luck. Without this assumption, rational people would expect that others to perform identically to themselves and there would be no BTA effects. However, the empirical evidence clearly supports the plausibility of this assumption: People use themselves as a useful, if imperfect, starting point for estimating others (Krueger, Acevedo, & Robbins, 2005). Third, when estimating self and others, people have to rely on both their prior expectations and the evidence at their disposal. If they ignore their prior expectations because they believe they are worthless, and instead believe that everyone will perform exactly as they have, there will be no BTA effects.

To clarify the differential information explanation, it might be useful to describe how it could explain prior BTA findings. For instance, a number of studies have found that people rate positive personality attributes as more descriptive of themselves than of others (Alicke, 1985; Brown, 1986). People generally express positive traits more often than negative ones. Furthermore, they are more familiar with their own intentions and behaviors than with those of others. For the majority of people who generally try to be friendly, cooperative, and dependable, they can reasonably infer that there are many ways for others to embody these traits *less* than they do. There are fewer ways in which people could embody these traits *more* than they do. It follows that rational people could believe that they are more friendly, cooperative, and dependable than are others (Fiedler, 1996). On the other hand, for negative behaviors such as being dishonest, phony, or rude, people will know that they personally express them rarely, but they cannot be as sure about others. It follows that they ought to infer that these negative traits (and rare behaviors) are less descriptive of themselves than of others (Fiedler, 2000).

Perhaps the most frequently cited BTA effect is Svenson's (1981) finding that people believe that they are better drivers than others. Undeniably, people have more information about their own driving habits than those of others. As a rule, people obey traffic laws and successfully avoid accidents; most of the information they have about themselves is positive. Other drivers might be just as competent, but it is hard to say with equal confidence since people have less information about others. It is easy to see the many ways in which others could be worse drivers than ones self. Estimates of others, then, should be more moderate than for self. It follows, then, that people will believe they are better drivers than are others.

It should be obvious at this point that both of the explanations we have offered (skewed distributions and differential information) also hypothesize circumstances under which people will believe themselves to be worse than others. In positively skewed distributions, the majority of people will be below average. This is the case, for instance, with income. Income is generally limited to zero at the bottom end. As long as there are a few CEOs raking in a few hundred million dollars, the vast majority of people will have incomes that are below average. One simple alternative is to ask participants to report their percentile ranks instead of comparing themselves with the average. It is statistically impossible for more than half of people to be above the 50^{th} percentile.

The differential information explanation would predict that people will believe themselves to be worse than others (and worse than the median) whenever their outcomes are worse or the probability of success is lower than expected. Recent empirical results bear this prediction out. For instance, people report that they are less likely than others to experience rare events, such as finding a $20 bill on the ground in the next two weeks (Chambers, Windschitl, & Suls, 2003). People believe that they are worse than others at difficult tasks such as coping with the death of a loved one (Blanton, Axsom, McClive, & Price, 2001; Kruger, 1999). And people believe that they will lose difficult competitions, such as trivia contests on indigenous vegetation of the Amazon (Moore & Cain, in press; Moore, Oesch, & Zietsma, in press; Windschitl, Kruger, & Simms, 2003).

Other researchers have provided explanations for these effects based on egocentric biases (for a review, see Chambers & Windschitl, 2004). However, in this paper, we focus on the differential information account because it is a more parsimonious explanation for these effects, because it offers broader explanatory power (Moore & Small, in press), and because it takes into account the normative question of how people ought to make comparative self-other comparative judgments. What are some of the testable implications of the differential information theory? We test three. First, estimates of others should be more regressive than are estimates of self. Second, to the extent that people have imperfect information about their own performances, self-estimates should also be regressive, but less so than estimates of others. The hypothesized patterns are illustrated in Figure 1. These first two implications are tested in all three experiments. The first experiment examines beliefs about the likelihood of future life events. The second experiment examines beliefs about performance on a task.

A third implication of the differential information explanation is that differential regressiveness should be exaggerated when people expect themselves to be different from others, and should be reduced when people expect themselves to be similar to others. This third implication is tested in the third experiment.

EXPERIMENT 1: RELATIVE PROBABILITIES

Experiment 1 applies the differential information explanation to account for relative probability judgments. For some time, researchers have believed that people were motivationally biased to believe that they were more likely than others to experience positive life events and less likely than others to experience negative life events (Weinstein, 1980). Two recent studies have shown that while people do believe themselves to be above-average in their likelihood of experiencing common events (such as living past 70), they often rate themselves as below average in their likelihood of experiencing rare life events (such as living past 100)

(Chambers et al., 2003; Kruger & Burrus, 2004). Experiment 1 replicates these effects and also elicits participants' beliefs regarding the actual probabilities of experiencing common and rare events for both self and others.

Method

Participants. We recruited 158 participants from the student body at Carnegie Mellon University in Pittsburgh, Pennsylvania. Half of the participants (79 individuals) received payment of $2 for their participation; the other 79 participants received credit towards a course requirement. The participants had an average age of 21 years and 57 percent of them were male.

Procedure. Participants saw a list of 24 life events (the list of events appears in Table 1). These 24 events were chosen because they varied with respect to both their frequency and their valence. Half of these events are positive and half are negative; half of the events are likely to occur and half are rare. For each event, we asked participants to estimate the probability that the event would happen to them at some point in their lives. A separate page asked participants to estimate, for each event, the probability that it would occur at some point in the life of the typical participant in the experiment. A third page asked participants to estimate their comparative likelihood of experiencing the event. Participants were asked to estimate their percentile rankings relative to all other participants in the experiment:

> "If you think you are more likely to than anyone else in this experiment to experience the event, enter '100' as your percentile. If you think that you are the least likely person to experience the event, enter '0' as your percentile. If you think your chances are exactly in the middle, enter '50' as your percentile. All numbers between 0 and 100 are acceptable responses."

The experimental materials included two different order manipulations to rule out idiosyncratic effects of order not relevant to the present hypothesis. The first order manipulation varied the sequence of the three estimation tasks (estimating for self, for others, and comparatively) using a Latin-squares design. The three orders were: (1) self-other-comparison; (2) other-comparison-self; (3) comparison-self-other. The second order manipulation varied the order in which the 24 events appeared on each of the three pages in the experimental materials. Some participants saw the 24 events in one randomly determined order; the other participants saw the same events in the reverse order.

Results and Discussion

We dropped five participants from the analysis who failed to complete the questionnaire.

Results for the 24 life events appear in Table 1. First, we sought to verify replication of BTA and WTA effects. In order to do this, we computed average percentile rankings for each participant for each of the four categories of events (positive/frequent, positive/rare, negative/frequent, and negative/rare). We then subjected these measures to a 2 (valence) x 2 (frequency) repeated-measures ANOVA. Recall that participants were all asked to rate themselves relative to other participants in the experiment, so if participants were accurate in their assessments, the average percentile rank for all participants would have to be 50. The results are consistent with prior results: People rated their relative likelihood of experiencing events as higher when the events were common (Mean percentile = 64.7, *SD* = 15.27) than when they were rare (Mean percentile = 30.8, *SD* = 15.19), F (1, 150) = 359, $p < .001$, $\eta^2 = .71$. There was also a significant main effect of valence: Participants rated themselves as more likely to experience events that were positive (Mean percentile = 53.1, *SD* = 11.3) than those that were negative (Mean percentile = 42.4, *SD* = 12.2), F (1, 150) = 93, $p < .001$, $\eta^2 = .38$. The interaction is not significant, F (1, 150) = 1.25, $p = .265$, $\eta^2 = .01$.

To test for the statistical significance of the hypothesized differences in participants' estimates of absolute probabilities, we computed means for each participant for each of the four categories of events (positive/frequent, positive/rare, negative/frequent, and negative/rare). We subjected these means to a 2 (valence) x 2 (frequency) x 2 (target: self vs. other) repeated-measures ANOVA. The hypothesized target x frequency interaction is significant, F (1, 152) = 88.9, $p < .001$, $\eta^2 = .37$. Participants rated themselves more likely (*M* = 74%, *SD* = 15%) than others (*M* = 70%, *SD* = 15%) to experience common events but less likely (*M* = 22%, *SD* = 14%) than others (*M* = 24%, *SD* = 16%) to experience rare events.

This two-way interaction is qualified by a significant three-way target x frequency x valence interaction, F (1, 152) = 7.79, $p = .006$, $\eta^2 = .05$. This three-way interaction is illustrated in Figure 2. It describes the fact that the hypothesized target x frequency interaction is strongest where it is consistent with motivational bias, and smallest where motivation is acting against it. The effect is strongest for rare negative events and common positive events. This same effect did not show up in the direct comparative judgments (percentile rankings), since the frequency x valence interaction is not significant there. There are two reasons why this apparent inconsistency is not troublesome. First, it is not troublesome for our theory because our theory makes no claims about motivational influences on comparative judgment. Second, it does not imply inconsistency between direct and indirect comparative judgments because the way we measured them did not allow us to infer their percentile ranks from their absolute assessments of self and others.

We acknowledge that the absolute differences in probability estimates for self and other are not large. More important than the absolute size of these differences is their direction. Of nearly 3672 comparisons (24 events x 153 participants), only 26 percent go in the opposite direction from that predicted by our theory (i.e., participants estimating that they are more likely to experience rare events and less likely to experience common events than is the typical participant). It is the consistency of this directional difference that produces such a reliable target x frequency interaction.

Not every result of the second experiment is perfectly consistent with our expectations. For our participants (ambitious students at a selective university), contrary to our expectations, amassing a million dollars in savings was not seen as a rare event. Yet the overall tests are clearly consistent with our hypotheses. Furthermore, the regression argument would predict that the consistency in expected differences between self and other should decrease as the task domain moves away from the extremes in performance or probability. The anomalies we do observe tend to occur where they are least problematic for the differential information theory.

We predicted that, in addition to obtaining BTA effects for common events and WTA effects for rare events, we would find that participants also underestimated their chances of experiencing common events and overestimated their chance of experiencing rare events. While we cannot say for certain how likely our participants will be to experience each of the 24 events during the courses of their lives, the results shown in Table 1 appear to be consistent with our hypotheses. While our participants estimated that their chances of being struck by lightening at 20%, the average American's probability is closer to .009% (U.S. Census Bureau, 2002). And while our participants estimated that they stood a 14% chance of winning over $10 million in the lottery, at least in Pennsylvania where they reside, the probability of winning the Powerball jackpot (the prize most likely to exceed $10 million) is closer to 0.0000008% (Pennsylvania Lottery, 2004).[2] As for common events, while participants estimated that they stood only an 87% chance of obtaining a starting salary over $25,000 per year, the records of the Carnegie Mellon Career Center indicate that approximately 97.5% of Carnegie Mellon's class of 2003 started at jobs with salaries over $25,000 (Carnegie Mellon Career Center, 2003). And while our participants estimated that they only had a 71% chance of cutting themselves shaving, virtually everyone will cut themselves shaving at some point.

The findings of Experiment 1 and the differential information explanation can help make sense of a set of inconsistent findings in research on risk perception. Research indicates that people tend to overestimate, sometimes radically, the probability that they will experience rare events. For example, Lerner, Gonzalez, Small, and Fischhoff (2003) reported that after September 11th, 2001, Americans estimated their probability of being injured in a terrorist attack as 20%. Other examples come from perceptions of health risks. In one study, smokers reported a 37% chance that they will get cancer due to smoking (Viscusi, 1990). The actual risk that smokers will fall ill with lung cancer is around 5% to 10%. However, when smokers are asked whether they are at more or less risk than other smokers, the frequently report believing that their risk is below average (Slovic, 2000; Weinstein, 1984). Similarly, women have been found to overestimate the probability that they will fall ill with breast cancer, often by as much as eight times the true probability (Lipkus, Biradavolu, Fenn, Keller, & Rimer, 2001). At the same time, people often believe that their risk for experiencing these rare events is below average (Woloshin, Schwartz, Black, & Welch, 1999). The regression explanation can account for both these findings: people's estimates of their own chances are regressive, but their estimates of others' chances are even more regressive.

EXPERIMENT 2: THE TRIVIA QUIZ

The first experiment presents evidence consistent with the regression explanation. However, it is limited because it does not include measures of actual performance. This limitation prevented us from assessing the accuracy of judgment. Experiment 2 solves this problem. Experiment 2 also introduces a new test of the differential information explanation. The differential information explanation applies to people comparing themselves with others. When people are comparing other individuals to each other and self-knowledge is not relevant or useful, then the regression account would not predict biases in comparative judgments. In order to test this implication, Experiment 2 includes a condition in which participants are asked to compare two randomly chosen individuals.

Method

Participants. We recruited participants after classes at Carnegie Mellon University. An experimenter invited students in each of six classes to remain for 10 minutes after class and complete an experiment for cash payment. Two-hundred fifteen individuals agreed to participate.

Procedure. Each participant was given a packet of questionnaires, beginning with a 10-item trivia quiz and an 11th tiebreaker question. Half the participants received a simple quiz including questions like “How many inches are there in a foot?” whereas others received a difficult quiz including questions like, “What is Avogadro’s number?” The tiebreaker question (“How many people live in Pennsylvania?”) was scored based on participants’ distance from the correct answer, and virtually eliminated the possibility of a tied score.

Next, instructions informed participants that they had earned $3 and that they could wager any amount of it on the trivia competition. If they bet and won, the amount they bet would be doubled. Half the participants bet on whether they would beat a randomly chosen opponent (the tiebreaker question resolved tied scores). The other half of participants bet on whether a randomly selected (anonymous) protagonist would beat a randomly chosen opponent. For those betting on this random protagonist, the random nature of the selection of the protagonist was driven home by asking them to draw a number out of a hat in order to determine on whom they would be betting.[3]

The experiment, therefore, employed a 2 (difficulty) x 2 (protagonist) between-subjects design.

Dependent measures. Participants made both comparative and absolute judgments of performance. Comparative measures included participants' bets, estimates of their probability of winning, and their responses to the question "How do you expect that you [the person whose number you drew] will score relative to all other people taking the same test?" The response scale ran from 1 to 7, with labels at 1 (well below average), 4 (average), and 7 (well above average).

Participants' scores on the actual quiz served as measures of absolute performance. In addition, participants were asked to estimate scores for (a) the protagonist (self or the randomly chosen protagonist), (b) the randomly chosen opponent, and (c) the average person. We assessed confidence in these estimates by asking participants to specify (for both the protagonist and opponent) scores above and below their guesses such that they were 90 percent sure the true score fell within that range.

Results and Discussion

Manipulation check. Scores on the simple test were in fact higher (M = 8.51 out of 10, SD = 1.48) than were scores on the difficult test (M = 1.66 out of 10, SD = 1.26), t (213) = 36.54, p < .001.

Comparative judgments. To assess the effects of the manipulations on comparative judgments, we subjected participants' bets to a 2 (difficulty) x 2 (protagonist) ANOVA. The results reveal a main effect for difficulty: Those taking the simple quiz bet more (M = \$1.84, SD = \$1.01) than did those taking the difficult quiz (M = \$1.38, SD = \$1.14), F (1, 210) = 11.34, p = .001, η^2 = .05. These results parallel results of asking participants to estimate the protagonist's percentile rank. Again, the same 2 x 2 ANOVA reveals a main effect for difficulty, such that the protagonist's percentile rank is estimated to be higher in the simple quiz condition (M = 57, SD = 17) but than in the difficult quiz condition (M = 41, SD = 17), F (1, 209) = 67.23, p < .001, η^2 = .24.[4] However, this main effect is qualified by a significant difficulty x protagonist interaction, F (1, 209) = 23.54, p < .001, η^2 = .10. The effect of difficulty on comparative evaluation was greater for those betting on self than for those betting on a randomly selected person (see Table 2).

Absolute judgments. To test the effect of the experimental manipulations on participants' estimates of absolute performance for self and opponent, we conducted a 2 (difficulty) x 2 (protagonist) x 2 (target: protagonist vs. opponent) mixed ANOVA with repeated measures on target. Naturally, the results reveal a significant between-subjects main effect of difficulty: Participants taking the simple test predicted higher scores (M = 7.5, SD = 1.4) than did participants taking the difficult test (M = 3.0, SD = 1.5), F (1, 173) = 443.36, p < .001, η^2 = .72. This main effect is qualified by a significant target x difficulty interaction, F (1, 173) = 25.36, p

< .001, η^2 = .13. Participants estimated that the easy test would be easier for the protagonist (*M* = 7.9, *SD* = 1.5) than for the opponent (*M* = 7.4, *SD* = 1.5), whereas the difficult test would be more difficult for the protagonist (*M* = 2.9, *SD* = 1.6) than for the opponent *(M* = 3.3, *SD* = 1.5). This two-way interaction is qualified by the expected three-way target x difficulty x protagonist interaction, *F* (2, 173) = 8.72, *p* = .004, η^2 = .05. This three-way interaction reveals that, consistent with the differential information explanation, differences between scores predicted for protagonist and opponent are greater for those betting on themselves than for those betting on a randomly selected person. For those betting on a randomly selected person, the estimated score for that protagonist is not significantly different from the score estimated for their randomly chosen opponent for either the easy quiz, *t* (57) = 1.84, *p* = .071, or the difficult quiz, *t* (58) = -1.00, *p* = .32.

Measures of regressiveness. Participants were imperfect estimators of performance. As shown in Figure 3, participants' estimates look regressive, because error leads estimates to be less extreme than the actual scores. Consistent with the assumption that individuals have more information about themselves than about others, participants estimated their own scores with greater accuracy than others' scores. Nevertheless, people estimated their own scores with some error. When performance was low, they overestimated it; those taking the difficult test only got an average of 1.66 (*SD* = 1.26) correct, yet estimated that they got 2.22 (*SD* = 1.34) correct. When performance was high, they underestimated it; those taking the simple test got an average of 8.51 (*SD* = 1.59) correct, yet estimated that they got 8.37 (*SD* = 1.46) correct. This effect is statistically significant, as shown by the difficulty x measure interaction, *F* (1, 211) = 71.72, *p* < .001, η^2 = .25, in a 2 (difficulty) x 2 (measure: actual vs. estimated score) ANOVA with repeated measures on the second factor. These results are consistent with the well-documented hard/easy effect in judgments of confidence (Erev, Wallsten, & Budescu, 1994).

Each participant estimated a 90 percent confidence interval for both protagonist and opponent. Participants knew that their estimates of their own scores were not perfectly accurate; they established confidence intervals around their answers that were, on average, 2.97 points in width (SD = 1.40). Their confidence intervals are even wider when estimating others' performances (M = 4.39, SD = 2.08). In a 2 (difficulty) x 2 (protagonist) x 2 (target: protagonist vs. opponent) mixed ANOVA with repeated measures on target, the main between-subjects effect of protagonist, F (1, 173) = 11.45, p < .01, and within-subject effect of target, F (1, 173) = 29.91, p < .001, are significant, but they are qualified by a significant target x protagonist interaction, F (1, 173) = 37.75, p < .001, reflecting greater confidence when predicting their own scores than when predicting others' scores (Table 2).

We also expected that participants would use the self as a basis for estimating others, which would also contribute to greater regressiveness in estimates of others. In other words, for those betting on themselves, actual performance would serve as the basis for predicting the performance of the protagonist, which would in turn serve as the basis for predicting the performance of the opponent. In order to examine this hypothesis, we conducted a mediational analysis (Baron & Kenny, 1986) in which actual performance and self estimate were used to predict opponent estimate. The results appear in Figure 4A, and indicate that estimates of self mediate the effect of actual performance on estimates of opponent. The significance of the indirect effect was tested using Sobel's (1988) equation (z = 2.00, p < .05). The converse, however, is not true (see Figure 4B); opponent estimates do not fully account for the effect of test performance on estimates of self.

Both greater error in estimating others' scores and the use of self as a basis for estimating others contribute to more regressive estimates of others than of self. After just having overestimated their own performances, those taking the difficult quiz overestimated the performances of others to an even greater extent. Similarly, those who had taken the simple quiz underestimated their performances, and underestimated the performances of their opponents even more.

Mediational tests of comparative judgment. If it is the more regressive predictions of others that lead to both BTA and WTA effects, then the estimated differences between one's own score and those of others should mediate the effect of test difficulty on comparative judgments for those who were betting on themselves. In order to test this hypothesis, we conducted three sets of regressions, shown in Figure 4C. First, comparative judgments were regressed on a dummy variable for test difficulty. Indeed, those who took the difficult test rated themselves as worse than those who took the simple test. In the second regression, test difficulty had a dramatic influence on the predicted difference between scores for self and for opponent. When both test difficulty and predicted score difference are used to predict comparative judgments, both remain significant. Differences between predicted absolute scores for self and the group account for 62.2% of the variance in comparative self-evaluation. Nevertheless, these predictions do not fully mediate the relationship between difficulty and relative evaluation because test difficulty remains a significant predictor of comparative evaluation after controlling for differences in predicted absolute scores (Sobel test: $z = -.49$, $p = .63$).[5]

While it is comforting that differences between participants' absolute estimates of self and other influence self-other comparisons, our theory makes a bolder prediction regarding the confidence with which participants make these estimates of absolute performance. Specifically, we would predict that when self is predicted with much greater confidence than other, both BTA and WTA effects should be stronger, since estimates made with lower confidence will naturally be more regressive. We do in fact have measures of the confidence with which participants made their estimates of absolute performance: their 90 percent confidence intervals. We constructed a measure of differential confidence by subtracting the size of the confidence interval for protagonist from the size of the confidence interval for opponent. We then correlated this measure of differential confidence with participants' comparative judgments. Consistent with the regression explanation, that correlation was positive and significant among those who took the simple quiz, r (88) = .26, $p < .05$: The greater the difference between the confidence with which they estimated self and other, the more likely they were to rate themselves above average. Also consistent with the regression explanation, that correlation reversed itself among those taking the difficult quiz, r (88) = -.20, $p = .058$: The greater the difference in confidence between self and other, the more likely participants were to rate themselves below average. The correlation difference test reveals the difference between these two correlations to be statistically significant, $Z = 20.06$, $p < .001$.

How much variance does this theory account for? The results clearly replicate prior results showing that task difficulty affects beliefs about relative performance. The results also show that differential regressiveness is indeed at work. But how much of the effect of task difficulty can be accounted for by differential regression? In order to answer this question we first have to assess the size of the effect of difficulty. When difficulty is used as the sole independent variable in a regression predicting self-reported percentile rank for those betting on themselves, the resulting R-squared value indicates that difficulty accounts for 39.4% of the variance. Differential regressiveness in participants' absolute estimates of performance by self and other, as measured by their indirect comparative judgments (estimated score for self minus estimated score for opponent), on its own, accounts for 62.2% of the variance in estimated percentile rank. When both difficulty and the indirect comparison are included as independent variables, the resulting R-squared value indicates that combined they account for 66.9% of the variance in estimated percentile rank. What this means is that task difficulty accounts for 4.7% of the variance in comparative judgments, over and above the differential regressiveness in people's absolute judgments. This 4.7% represents just 12% of the total effect of difficulty (39.4%) on direct comparative judgments. The implication is that differential regressiveness accounts for 88% of the effect of difficulty on comparative judgments, at least for self-estimates, leaving 12% of the variance for other explanations, such as differential weighting (Klar & Giladi, 1999; Kruger, Windschitl, Burrus, Fessel, & Chambers, 2006).

However, among those who had little useful information about the people whose percentile ranks they were estimating, this same analysis suggests that differential regressiveness accounts for less than 9% of the effect of difficulty on comparative judgments. Part of the issue here is that the effect of difficulty is so small among this group (R-squared = 6.8%). However, our theory would not predict differential regressiveness among this group, given that they do not have better information about the protagonist than the opponent.

Reconciliation with prior results. The disappearance of BTA and WTA effects among people betting on a randomly selected individual stands in contrast to the results of Moore and Kim (2003) who replicated BTA and WTA effects even when comparing two other individuals. The key difference between Experiment 3 and their approach is probably related to the degree to which participants focused on the other person on whom they were betting. Moore and Kim (Experiment 4) showed that focusing on the other person was key to the effect. When people took the perspective of the person on whom they were betting, they made predictions about that person in the same way they made predictions about themselves. Sanbonmatsu et al. (1987) have shown that focusing on a particular target can be enough to lead people to make more regressive estimates of other individuals who are not in focus. However, in the present study participants drew that person's number from a hat, highlighting the random and anonymous nature of the other.

Skewed distributions? As noted earlier, the majority of people *are* in fact above average in a negatively skewed distribution. Of course, one cause of negatively skewed distributions is a performance ceiling. If participants knew that the distributions were skewed and responded to our questions regarding randomly chosen opponents as if we were asking about the average other, it could account for our results. This raises the question of the degree to which BTA effects depend on the presence of a ceiling effect, and the degree to which WTA effects depend on a floor. It is noteworthy that most prior findings of BTA and WTA effects occur in domains where there are ceilings and floors in performance, or at least on the scales used to measure performance. Our theory does not depend on ceilings or floors for the production of BTA and WTA effects, and so predicts that we will replicate these biases in comparative judgment even when no ceilings or floors are present. Experiment 3 tests this prediction.

Experiment 3 also tests a second prediction of our theory: that perceived similarity between the target and the referent will moderate BTA and WTA effects. When the two are assumed to be similar, then information about the target's performance will generalize to the referent (Krueger, 2000; Mussweiler, 2003), and one should not expect the target to be much better or worse than the referent. However, when the two are assumed to be different, our theory would predict that BTA and WTA effects would be stronger.

Experiment 3 allows us to examine the process of updating from prior expectations because it includes measures of expected performance. This is important because our theory posits an important role for expectations. That is, differential regression will only produce BTA when the task is easier than expected and will only produce WTA effects when the task is more difficult than expected. This is a reasonable supposition for the tasks used in Experiments 1 and 2 our case would be more convincing if we measured these beliefs. We do so in the third experiment.

EXPERIMENT 3: THE DATES TEST

Method

Participants. Participants were 187 volunteers from Carnegie Mellon University who participated in exchange for pay.

Task. When participants arrived they were told that they would be taking "The Dates Test": "In this experiment, you will take a test that measures how well informed you are about the world in which we live. Each question on the test asks you to identify the calendar year in which some event occurred." They were to answer as many of these questions as they could, and as accurately as they could, in the space of two minutes. They had to answer every question in sequence and could not skip questions.

For every question they answered they could earn up to 100 points if their answer was exactly right. If they were not exactly right, one point would be deducted for each year off their answer was from the correct answer. If their answer was more than 100 years off, they could lose points, up to a maximum loss of -100 points for an answer that was 200 or more years off.

Procedure. After reading the instructions for the task, participants were asked to predict their own scores and the average score on the upcoming test. Then they actually took the test.

Each participant received a set of either easy or difficult questions. The first item on the easy quiz was: "*At the start of what year, known as 'Y2K,' were computers expected to crash due to the 'millennium bug'?*" Answer: 2000. The first item on the difficult quiz was: "*In what year was Pope Pius the first selected as Pope?*" Answer: 142.

After having taken the test, participants again estimated their own scores. Then participants read: "*Next we are going to ask you to estimate the score of the average participant. But before you do that, please spend about 5 minutes writing down all the reasons why the average participant is likely to have a score that is similar to [different from] your own*." Below these instructions were several blank lines on which they then wrote. Participants then estimated the average quiz score.

Participants were then asked to estimate precisely how much they expected their score to exceed that of the average person. They were told to use negative numbers to indicate the extent to which their score would be below that of the average person. This was the direct comparative judgment.

Design. The experimental design then, was a 2 (difficulty) x 2 (similarity priming) between subjects design. The experiment also included measures of participants' beliefs about scores, both their own and the average score, both before and after having taken the actual quiz.

Results and Discussion

Pre-test predictions. We submitted pre-test predicted scores for self and other to a 2 (difficulty) x 2 (similarity) x (2) (target: self vs. other) mixed ANOVA with repeated measures on the third factor. The effect of target emerges as significant, F (1, 180) = 4.93, p = .03, η^2 = .03. Before taking the test, participants predicted that they would score higher (M = 751, SD = 631) than the average (M = 669, SD = 601). No other effects are significant.

Test performance. Average scores were higher on the easy (M = 1318, SD = 568) than on the difficult quiz (M = -649, SD = 389), t (185) = 27.58, $p < .001$. Participants answered more questions on the easy (M = 13.9, SD = 5.42) than the difficult (M = 12.1, SD = 4.57) quiz, t (185) = 2.54, p = .012. The distributions of test scores were roughly normal, as indicated by modest levels of skewness for the easy (.18, SE = .25) and the difficult (-.51, SE = .25) tests. Note that the directions of skew here allow for a conservative test of our hypotheses, in the sense that given positive skew on the easy test, the majority of people (60%) are in fact be below average; and given negative skew on the difficult test, the majority of people (56%) are in fact be above average.

Post-test beliefs. In order to examine participants' beliefs about their performance relative to others, we submitted their estimated scores for self and for the average other to a 2 (difficulty) x 2 (similarity) x (2) (target: own estimated score vs. estimated average score) mixed ANOVA with repeated measures on the third factor. Of course, the effect of difficulty is significant because estimated scores are higher on the easy than the difficult test, F (1, 182) = 228, $p < .001$, $\eta^2 = .56$. However, the within-subjects effect of target did not emerge as significant, F (1, 182) = .54, p = .47, $\eta^2 = .003$. This reflects the fact that, after taking the test, participants did not believe they (M = 499, SD = 946) scored any better than did others (M = 516, SD = 738).

The results do reveal the expected target x difficulty interaction effect, F (1, 182) = 39.86, $p < .0001$, $\eta^2 = .18$. Participants reported that on the simple quiz, they did better (M = 1213, SD = 455) than average (M = 1002, SD = 378), t (44) = 3.91, $p < .001$, but that on the difficult quiz they did worse (M = -206, SD = 829) than average (M = -57, SD = 733), t (45) = -2.21, p = .032.

Furthermore, the predicted 3-way target x difficulty x similarity interaction effect emerges as marginally significant, *F* (1, 182) = 3.71, *p* = .056, η^2 = .02. This effect, which is illustrated in Figure 5, shows that the target x difficulty effect described above is weakened when participants were primed to think about how the average is similar to them, whereas it was exacerbated when they were primed to think about how the average is different from them. No other effects in this mixed ANOVA emerge as significant.

Finally, and also consistent with our theory, participants' estimates of their own scores (*M* = 1196, *SD* = 553) were lower than their actual scores (*M* = 1318, *SD* = 568) on the easy test, *t* (93) = 4.23, *p* <.001. But on the difficult test, participants' estimates were higher (*M* = -211, *SD* = 705) than were their actual scores (*M* = -646, *SD* = 390), *t* (91) = -5.72, *p* < .001. Participants made regressive estimates of themselves, but their estimates of others were even more regressive. See Figure 6.

Updating from priors. The differential information theory is based on the Bayesian notion that people begin with some prior expectation and then update that belief when they get new evidence. To clarify how their prior expectations influenced participants' evaluations of self and other, we conducted two regressions using post-test estimated scores for self and other as dependent variables. Our differential information theory would predict that people use their own quiz performances to update beliefs about their own scores more so than beliefs about others' scores. Self-knowledge is more useful for predicting the self than it is for predicting others.

The regression predicting participants' post-test beliefs about their own performance employs three independent variables: (1) the participant's pre-test estimated score for self; (2) the participant's own actual score on the quiz; and (3) the difficulty of the quiz the participant received. The results of this regression appear in the first two columns of Table 3. The second regression predicting participants' post-test beliefs about others' performance employed the participant's pre-test estimated score for the other and the same last two variables. These results appear in the right-hand columns of Table 3.

Consistent with our expectations, participants' own quiz performances exert a weaker influence on post-test estimates of others (B = .32, SE = .07, t = 4.46, $p = 1.44 \times 10^{-5}$) than on post-test estimates of self (B = .61, SE = .07, t = 8.26, $p = 3.01 \times 10^{-14}$). When participants were estimating their own scores, they had excellent information. They relied heavily on their own scores, but their pre-test priors were also a significant influence. When estimating the scores of others, the regression results suggest participants relied less on their own experiences and instead tried to account for the ease or difficulty of the task—hence the significance of quiz difficulty.

The significant effects of pre-test priors for estimations of both self and others suggest, consistent with our theory, that people's priors affect their subsequent judgments. They updated from these priors using new information, and since the information they had (their own quiz performances) was more useful for estimating self than others, this information was weighted more heavily when estimating their own scores than when estimating others' scores.

Moreover, comparing those who were primed with similarity vs. difference clarifies how this manipulation had its effect: The degree to which participants projected information about themselves onto others depended upon similarity priming. When estimating others' scores, those who expected to be similar to others weighted their own scores more heavily (B = .42, SE = .10, t = 4.14, $p = 7.60 \times 10^{-5}$) than did those who expected others to be different (B = .27, SE = .11, t = 2.53, p = .013).

How much variance does this theory account for? How much of the effect of task difficulty in this experiment can be accounted for by differential regression? We sought to answer this question as we did for Experiment 2. We first computed the percentage of the variance in direct comparative judgments accounted for by the difficulty manipulation. When difficulty is used as the sole independent variable in a regression predicting participants' estimates of the difference between their own score and the average score, the resulting R-squared value indicates that difficulty accounts for 14.3% of the variance. Differential regressiveness in participants' absolute estimates of performance by self and others, as measured by their indirect comparative judgments (estimated score for self minus estimated score for others), on its own, accounts for 29.7% of the variance in their direct comparative judgments. When both difficulty and the indirect comparison are included as independent variables, the resulting R-squared value indicates that combined they account for 26.2% of the variance in direct comparisons. What this means is that task difficulty accounts for 3.5% of the variance in comparative judgments, over and above the differential regressiveness in people's absolute judgments. This 3.5% represents just 24% of the total effect of difficulty (14.3%) on direct comparative judgments. The implication is that differential regressiveness accounts for 76% of the effect of difficulty on comparative judgments in Experiment 3, leaving 24% of the variance for other explanations.

General Discussion

The results of the three experiments we present are consistent with the differential information explanation for both BTA and WTA effects. Estimates of performance were regressive, especially when people made estimates about others about whom they had poorer information. As a consequence, when performance was high, people believed they were better than others. When performance was low, people believed that they were worse than others.

There is a great deal of evidence that shows that people prefer flattering information about themselves (Baumeister, 1998; Greenwald, 1980; Taylor, 1989). They seek it out and they accept it uncritically (Gilovich, 1991). It is often comforting or flattering to believe that one has done well or that one is better than others. But these theories have trouble accounting for WTA effects. The differential information explanation can account for both BTA and WTA effects. In contrast to the previous theories, it does not suggest an egocentric or self-enhancing bias. Rather, it is a rational, Bayesian explanation for why people might exhibit biases in comparative judgments.

Our theory is useful for reconciling BTA and WTA effects with a set of results that appear to contradict them. The so-called hard/easy effect on overconfidence has documented the fact that people overestimate performance on difficult tasks and tend to underestimate performance on easy tasks (Burson, Larrick, & Klayman, 2006; Krueger & Mueller, 2002; Kruger & Dunning, 1999). We should note that hard/easy effect involves measures of absolute performance whereas BTA and WTA effects involve measures of relative performance. Regressiveness in self-estimates is sufficient to produce the hard/easy effect (Erev et al., 1994). Yet this pattern is even more pronounced for other-estimates, resulting in the co-occurrence of hard/easy effects and WTA/BTA effects.

The Normative Question

We have examined BTA and WTA effects in order to clarify what the appropriate normative judgment is. We have argued, using evidence from three experiments, that people ought to believe that they are better than others when their performance is better than expected. When people learn about their own likely performance or their own probability of experiencing some event, that knowledge is often useful for updating their beliefs about others as well.

But there is a problem. Real people do not obey Bayes's Rule all that well. Sometimes, people appear to neglect priors (such as base rates), overweighting recent evidence (Grether, 1980, 1990). Other times, people appear too conservative, overweighting priors and neglecting useful new evidence (Edwards, 1968; McKelvey & Page, 1990). Which of these errors people commit depends on the order and form in which they acquire information (Hogarth & Einhorn, 1992; Wells, 1992). What is important for our differential information explanation, however, is that although people are imperfect Bayesians, they rarely abandon Bayesian logic completely. Indeed, under some circumstances human judgment is impressively close to Bayesian prescription (Griffiths & Tenenbaum, in press). However, all that is necessary for our explanation to hold is that people's estimates of others lie between their priors and their beliefs about themselves. Both the results from our experiments and from other experiments bear this assumption out (Krueger et al., 2005).

If people neglected their priors (base rates) and assumed that others behaved exactly as they did, they would, in effect, commit the false consensus effect in grand form. They would predict that others' scores would be the same as theirs, and that others would experience the identical probability of various possible outcomes. If people did the opposite and assumed that their own outcomes were irrelevant for determining the outcomes of others, then learning about their own performances would not shift predictions of others off of baseline. Instead, what we observe is a combination of the two. Across all three experiments, people's estimations of themselves are highly correlated with their estimates of others, but estimates of others are less extreme (McFarland & Miller, 1990; Miller & McFarland, 1987). We have come full circle, then, back to the false consensus effect.

False consensus?

Dawes's critique of the false consensus effect did not attempt to argue that there was no such thing as the false consensus effect, only that there was a rational explanation that would hypothesize an effect that looked like the false consensus effect. This useful clarification set the stage for a further refinement: Was there, in fact, a false consensus effect that was stronger than that predicted by normative theory? The answer, as shown by Krueger and Clement (1994) is yes. There really is a false consensus effect, and it is smaller than researchers had previously assumed before Dawes clarified the matter.

Can we accomplish the same refinement of theories of bias in comparative judgment? We believe we can. Our evidence suggests that differential regression caused by differential information can indeed account for a good deal of the variance in comparative judgments. However, it cannot account for 100% of such variance. WTA and BTA effects are stronger than one would predict based solely on differential regressiveness in estimates of absolute performance by self and other (Moore, in press). Clearly, there are other causes of BTA and WTA effects.

One such cause is likely to be differential weighting. A number of researchers have shown that people tend to overweight information about the self when making comparative judgments (Klar & Giladi, 1997, 1999; Kruger, 1999; Windschitl et al., 2003). Kruger and his colleagues (2006) have shown that this, too, can have a rational basis. When people possess high quality information it makes sense to give that information greater weight than they give to mere speculation about the performances of others.

We wish to conclude by pointing out that it is not our intention to claim that people are perfectly rational or that their judgments are unbiased. Evidence of human irrationality is abundant and undeniable (Bazerman, 2002). Indeed, it is precisely because we take seriously the evidence of imperfections in human judgment that we seek to understand it as best we can. We want to hold research on judgment and decision making to the high standard that Robyn Dawes would set: Before we accuse people of being irrational, we ought to have an excellent idea of precisely what it means to make rational judgments. This approach has two clear benefits. First, it helps us more clearly understand the normative benchmarks that we can use to provide advice to others or make better judgments ourselves. Second, we gain a richer understanding of human judgment by better understanding exactly when it complies with rational prescriptions and when it deviates from them. And finally, we must note that by offering a rational explanation for BTA and WTA effects, we are not saying that they are not real. On the contrary, if there is a rational basis for these effects we ought to expect them to be particularly durable because we cannot expect to be able to correct or to debias them.

References

Alicke, M. D. (1985). Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology, 49*(6), 1621-1630.

Babcock, L., & Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives, 11*(1), 109-126.

Baumeister, R. F. (1998). The self. In D. T. Gilbert, S. T. Fiske & G. Lindzey (Eds.), *The handbook of social psychology* (Vol. 3rd, pp. 680-740). Boston: McGraw-Hill.

Bazerman, M. H. (2002). *Judgment in managerial decision making* (5th ed.). New York: Wiley.

Benabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics, 117*(3), 871-915.

Blanton, H., Axsom, D., McClive, K. P., & Price, S. (2001). Pessimistic bias in comparative evaluations: A case of perceived vulnerability to the effects of negative life events. *Personality and Social Psychology Bulletin, 27*(12), 1627-1636.

Brown, J. D. (1986). Evaluations of self and others: Self-enhancement biases in social judgments. *Social Cognition, 4*(4), 353-376.

Brown, J. D. (1998). *The self*. Boston: McGraw-Hill.

Bruine de Bruin, W., Fischhoff, B., Millstein, S. G., & Halpern-Felsher, B. L. (2000). Verbal and numerical expressions of probability: 'It's a fifty-fifty chance'. *Organizational Behavior and Human Decision Processes, 81*(1), 115-131.

Burson, K. A., Larrick, R. P., & Klayman, J. (2006). Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *Journal of Personality and Social Psychology, 90*(1), 60-77.

Carnegie Mellon Career Center. (2003, September 29). *2003 post-graduation survey results*. Retrieved March 17, 2004, from http://www.studentaffairs.cmu.edu/career/student/student.html

Chambers, J. R., & Windschitl, P. D. (2004). Biases in social comparative judgments: The role of nonmotivational factors in above-average and comparative-optimism effects. *Psychological Bulletin, 130*(5).

Chambers, J. R., Windschitl, P. D., & Suls, J. (2003). Egocentrism, event frequency, and comparative optimism: When what happens frequently is "more likely to happen to me". *Personality and Social Psychology Bulletin, 29*(11), 1343-1356.

Daniel, K. D., Hirshleifer, D. A., & Sabrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. *Journal of Finance, 53*(6), 1839-1885.

Dawes, R. M. (1989). Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology, 25*(1), 1-17.

Dawes, R. M. (1990). The potential nonfalsity of the false consensus effect. In R. M. Hogarth (Ed.), *Insights in decision making: A tribute to Hillel J. Einhorn.* (pp. 179-199). Chicago, IL, USA: University of Chicago Press.

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes, 65*(3), 201-211.

Dunning, D. (1993). Words to live by: The self and definitions of social concepts and categories. In J. M. Suls (Ed.), *Psychological perspectives on the self* (Vol. 4, pp. 99-126). Hillsdale, NJ: Erlbaum.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal representation of human judgment* (pp. 17-52). New York: Wiley.

Epstein, S. (1990). Cognitive-experiential self-theory. In L. A. Pervin (Ed.), *Handbook of personality: Theory and research* (pp. 165-192). New York, NY, US: Guilford Press.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519-527.

Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review, 103*(1), 193-214.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*(4), 659-676.

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science, 14*(3), 195-200.

Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.

Greenwald, A. G. (1980). The totalitarian ego: Fabrication and revision of personal history. *American Psychologist*, *35*(7), 603-618.

Grether, D. M. (1980). Bayes' rule as a descriptive model: The representative heuristic. *Quarterly Journal of Economics*, *95*, 537-557.

Grether, D. M. (1990). Testing Bayes rule and the represetativeness heuristic: Some experimental evidence. *Journal of Economic Behavior and Organization, 17,* 31-57.

Griffiths, T. L., & Tenenbaum, J. B. (in press). Optimal predictions in everyday cognition. *Psychological Science*.

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24*, 1-55.

Howard, M. E. (1983). *The causes of war and other essays*. Cambridge, MA: Harvard University Press.

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*(5), 885-901.

Klar, Y., & Giladi, E. E. (1999). Are most people happier than their peers, or are they just happy? *Personality and Social Psychology Bulletin, 25*(5), 585-594.

Krueger, J. I. (2000). The projective perception of the social world: A building block of social comparison processes. In J. Suls & L. Wheeler (Eds.), *Handbook of social comparison: Theory and research* (pp. 323-351). Dordrecht, Netherlands: Kluwer Academic Publishers.

Krueger, J. I., Acevedo, M., & Robbins, J. M. (2005). Self as sample. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 353-377). New York: Cambridge University Press.

Krueger, J. I., & Clement, R. W. (1994). The truly false consensus effect: An ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology, 67*(4), 596-610.

Krueger, J. I., & Mueller, R. A. (2002). Unskilled, unaware, or both? The better-than-average heuristic and statistical regression predict errors in estimates of own performance. *Journal of Personality and Social Psychology, 82*(2), 180-188.

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*(2), 221-232.

Kruger, J., & Burrus, J. (2004). Egocentrism and focalism in unrealistic optimism (and pessimism). *Journal of Experimental Social Psychology, 40*(3), 332-340.

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*(6), 1121-1134.

Kruger, J., Windschitl, P. D., Burrus, J., Fessel, F., & Chambers, J. R. (2006). *The rational side of egocentrism in social comparisons*. Unpublished manuscript.

Lerner, J. S., Gonzalez, R. M., Small, D. A., & Fischhoff, B. (2003). Effects of fear and anger on perceived risks of terrorism: A national field experiment. *Psychological Science, 14*(2), 144-150.

Lipkus, I. M., Biradavolu, M., Fenn, K., Keller, P., & Rimer, B. K. (2001). Informing women about their breast cancer risks: Truth and consequences. *Health Communication, 13*(2), 205-226.

Malmendier, U., & Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance*(60), 6.

McFarland, C., & Miller, D. T. (1990). Judgments of self-other similarity: Just like other people, only more so. *Personality and Social Psychology Bulletin*, *16*(3), 475-484.

McKelvey, R. D., & Page, T. (1990). Public and private information: An experimental study of information pooling. *Econometrica, 58*, 1321-1339.

Miller, D. T., & McFarland, C. (1987). Pluralistic ignorance: When similarity is interpreted as dissimilarity. *Journal of Personality and Social Psychology, 53*(2), 298-305.

Moore, D. A. (in press). Not so above average after all: When people believe they are worse than average and its implications for theories of bias in social comparison. *Organizational Behavior and Human Decision Processes*.

Moore, D. A., & Cain, D. M. (in press). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior & Human Decision Processes*.

Moore, D. A., Oesch, J. M., & Zietsma, C. (in press). What competition? Myopic self-focus in market entry decisions. *Organization Science*.

Moore, D. A., & Small, D. A. (in press). Error and bias in comparative social judgment: On being both better and worse than we think we are. *Journal of Personality and Social Psychology*.

Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review, 110*(3), 472-489.

Neale, M. A., & Bazerman, M. H. (1985). The effects of framing and negotiator overconfidence on bargaining behaviors and outcomes. *Academy of Management Journal, 28*(1), 34-49.

Odean, T. (1998). Volume, volatility, price, and profit when all traders are above average. *Journal of Finance, 53*(6), 1887-1934.

Pennsylvania Lottery. (2004, Feb 2). *Powerball game information*. Retrieved March 17, 2003, from http://www.palottery.com/lottery/cwp/view.asp?a=3&q=457089&lotteryNav=|29736|

Peterson, C. (2000). The future of optimism. *American Psychologist, 55*(1), 44-55.

Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology, 13*(3), 279-301.

Sanbonmatsu, D. M., Shavitt, S., Sherman, S. J., & Roskos-Ewoldsen, D. R. (1987). Illusory correlation in the perception of performance by self or a salient other. *Journal of Experimental Social Psychology, 23*(6), 518-543.

Slovic, P. (2000). Rejoinder: the perils of Viscusi's analyses of smoking risk perceptions. *Journal of Behavioral Decision Making, 13*(2), 273-276.

Steele, C. M. (1988). The psychology of self-affirmation: Sustaining the integrity of the self. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 21, pp. 261-302). New York: Academic.

Svenson, O. (1981). Are we less risky and more skillful than our fellow drivers? *Acta Psychologica, 47*, 143-151.

Taylor, S. E. (1989). *Positive illusions: Creative self-deception and the healthy mind*. New York, NY, US: Basic Books, Inc.

Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin, 103*(2), 193-210.

U.S. Census Bureau. (2002). *Statistical abstract of the United States* (122nd ed.). Washington, DC: Library of Congress.

Viscusi, W. K. (1990). Do smokers underestimate risks? *Journal of Political Economy, 98*(6), 1253-1269.

Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology, 39*(5), 806-820.

Weinstein, N. D. (1984). Why it won't happen to me: perceptions of risk factors and susceptibility. *Health Psychology*, *3*(5), 431-457.

Wells, G. L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology, 62*(5), 739-752.

Windschitl, P. D., Kruger, J., & Simms, E. (2003). The influence of egocentrism and focalism on people's optimism in competitions: When what affects us equally affects me more. *Journal of Personality and Social Psychology, 85*(3), 389-408.

Woloshin, S., Schwartz, L. M., Black, W. C., & Welch, H. G. (1999). Women's perceptions of breast cancer risk: How you ask matters. *Medical Decision Making*, *19*(3), 221-229.

Footnotes

[1] Thanks to Shane Frederick for suggesting this helpful example.

[2] This number represents the probability of a single lottery ticket winning, but people routinely buy more than one ticket. However, most people do not buy 17.5 million lottery tickets, which is how many you would need to buy to raise your chances of winning to the 14% participants estimated.

[3] In order to rule out timing effects of this procedure, half of these participants drew their random number before they bet, half after they bet. Because this manipulation did not yield any main or interaction effects with variables of interest, these participants are grouped together in the analyses reported.

[4] Degrees of freedom fluctuate slightly between tests, due to missing data for some participants.

[5] Similar patterns hold for bets and probability of winning as measures of comparative judgment. Bets, however, are influenced by a number of other motivations, including risk aversion and feelings about gambling.

Table 1

Estimates of experiencing each of 24 life events for (1) comparative likelihood expressed as a percentile rank; and absolute probabilities both for (2) self and for (3) the average participant. Standard deviations in parentheses, Experiment 2.

Events	Percentile	Self	Other
Positive events			
Rare events			
Record your own music CD	33.5 (33.5)	29.1 (37.4)	31.0 (32.9)
Win over $10 million in the lottery	22.9 (25.9)	14.1 (26.3)	15.6 (26.2)
Inherit over $10,000 from a distant relative	28.4 (27.9)	23.9 (29.0)	35.1 (28.2)
Write a book	40.0 (30.0)	35.5 (32.0)	35.3 (28.5)
Live past the age of 100	32.5 (27.2)	25.2 (29.9)	23.6 (29.8)
Have over $1 million in savings	61.5 (26.3)	61.5 (28.2)	49.7 (27.5)
Common Events			
Have a stranger spontaneously complement your appearance	59.3 (25.8)	64.8 (31.4)	54.6 (25.9)
Obtain a starting salary over $25,000 per year	76.3 (21.0)	87.1 (18.6)	80.2 (18.0)
Live past the age of 70	60.3 (23.5)	66.1 (24.0)	62.4 (22.5)
Be elected an officer of an organization	68.3 (24.4)	72.7 (27.3)	60.0 (26.1)
Own your own home	76.5 (20.4)	89.1 (16.8)	79.0 (17.6)
Own your own car	79.1 (20.9)	92.0 (17.7)	84.6 (19.0)
Negative events			
Rare events			
Be killed in a terrorist attack	31.5 (25.7)	15.9 (20.2)	14.5 (19.4)
Be struck by lightening	32.9 (26.5)	20.1 (31.8)	20.2 (30.6)
Become addicted to crack cocaine	13.1 (21.2)	6.7 (15.4)	16.1 (21.7)
Go to jail	19.9 (22.6)	9.3 (14.4)	17.0 (20.2)
Be shot with a gun	31.2 (26.0)	15.4 (19.8)	15.0 (16.9)
Get AIDS	19.3 (21.0)	10.5 (16.9)	15.2 (17.5)
Common events			
Be interrupted during dinner by calls from telemarketers	54.7 (26.9)	68.3 (35.7)	72.0 (32.0)
Slip and fall on the ice	67.6 (25.2)	79.2 (26.7)	78.2 (21.7)
Be in an automobile accident	50.0 (28.0)	53.9 (32.6)	53.6 (29.5)
Attend the funeral of a loved one	62.9 (26.9)	79.0 (32.4)	76.8 (32.6)
Get lost driving in an unfamiliar city	64.1 (29.5)	75.6 (29.3)	69.3 (27.8)
Cut yourself shaving	59.0 (32.8)	71.3 (36.3)	69.6 (32.0)

Table 2

Predicted and actual scores by experimental condition, standard deviations in parentheses, Experiment 3. Figures on the same row with different superscripts are significantly different from each other ($p < .05$).

	Betting on self		Betting on a random person	
	Simple	Difficult	Simple	Difficult
Actual score (out of 10)	8.58[a] (1.59)	1.76[b] (1.34)	8.48[a] (1.43)	1.60[b] (1.22)
Bet (up to $3)	$2.13[a] ($.88)	$1.41[b] ($1.23)	$1.69[a,b] ($1.05)	$1.37[b] ($1.15)
Probability of winning	64%[a] (20%)	37%[c] (24%)	51%[b] (21%)	39%[c] (21%)
Percentile rank	61.3[a] (17.9)	30.6[d] (20.5)	54.3[b] (16.7)	46.4[c] (12.4)
Protagonist’s score (estimated)	8.37[a] (1.46)	2.22[c] (1.34)	7.66[a] (1.48)	3.32[b] (1.62)
Opponent’s score (estimated)	7.68[a] (1.02)	3.06[b] (1.41)	7.32[a] (1.64)	3.37[b] (1.57)
Average score (estimated)	7.12[a] (1.11)	3.07[b] (1.16)	7.05[a] (1.48)	3.29[b] (1.33)
Size of 90% conf. interval—protagonist	2.79[a] (1.43)	3.09[a] (1.31)	4.33[b] (2.05)	4.69[b] (1.94)
Size of 90% conf. interval—opponent	3.92[a] (1.75)	4.39[a] (2.15)	4.13[a] (2.09)	4.89[a] (2.14)

Table 3

Regressions predicting post-test score estimates for self and other, Experiment 3. (Standard errors in parentheses.)

Model 1 predicting post-test beliefs about own performance		Model 2 predicting post-test beliefs about other's performance	
Independent variable	Unstandardized *B* coefficient	Independent variable	Unstandardized *B* coefficient
Pre-test estimated score for self	.37*** (.05)	Pre-test estimated score for other	.23*** (.06)
Own actual score	.61*** (.06)	Own actual score	.33*** (.07)
Difficult quiz dummy	-244 (164)	Difficult quiz dummy	-427** (163)
R^2	.75***	R^2	.58***

** $p < .01$, *** $p < .001$

Figure 1. Estimated scores for self and other (hypothesized).

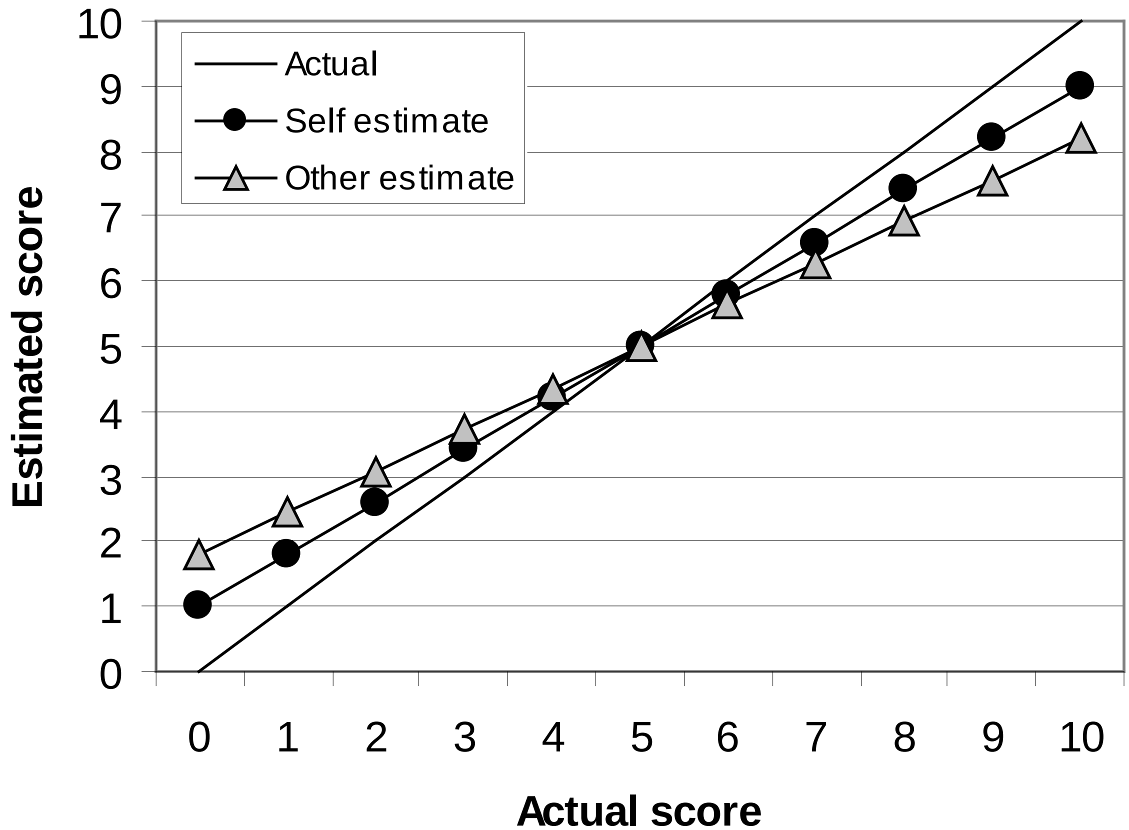

Figure 2. Estimated absolute probabilities of experiencing negative (Panel A) and positive (Panel B) events, Experiment 1. Error bars show standard errors.

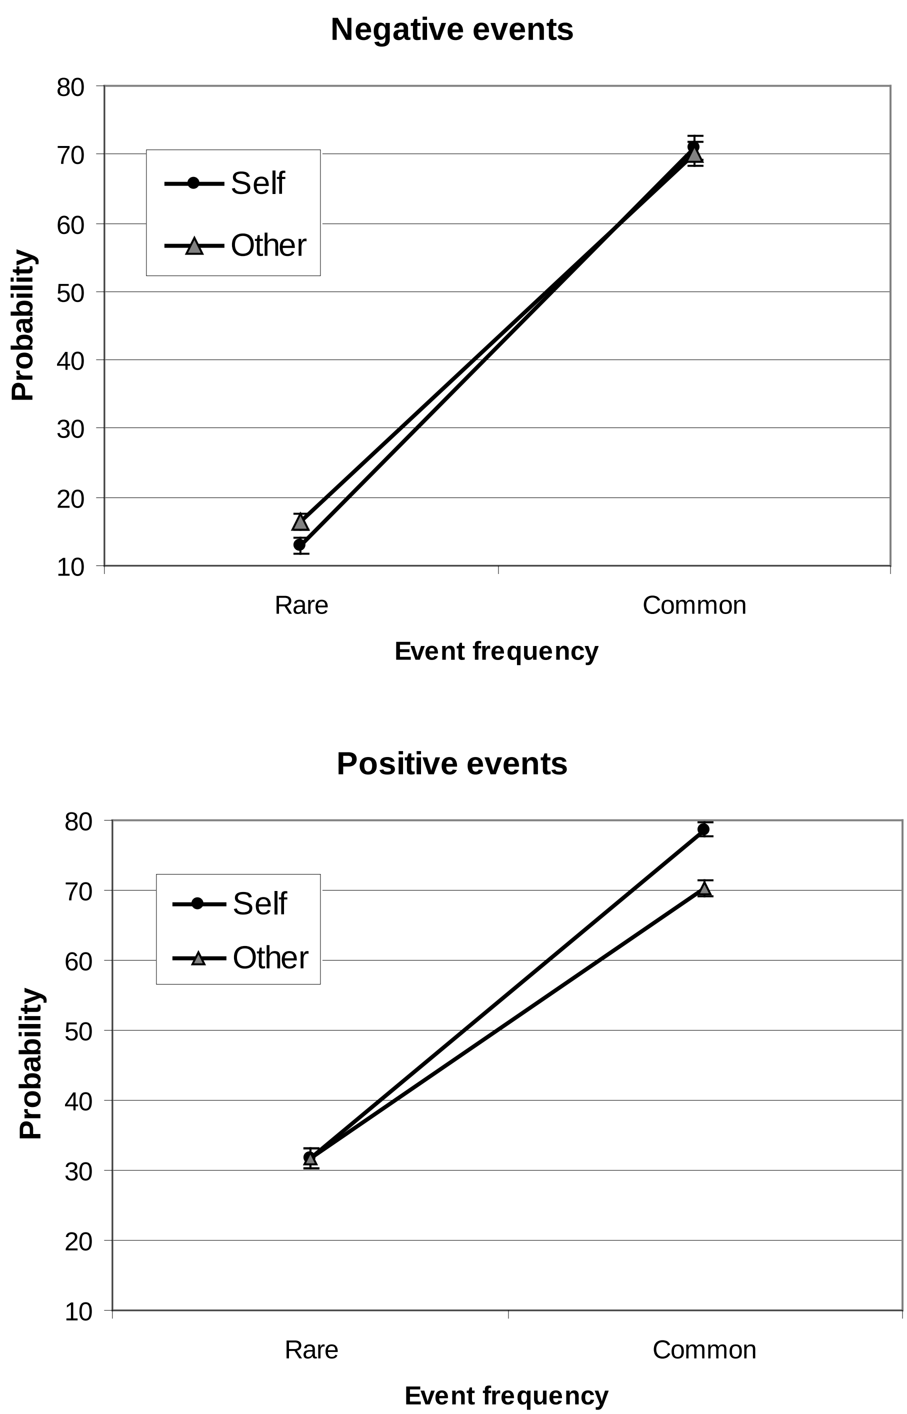

Figure 3. Estimated scores for self and opponent by those betting on the self, Experiment 2. There is more noise and variability in the middle of the scale on Figure 3 because there are so few participants with scores of 5 and 6. Most of the scores lie at the extremes.

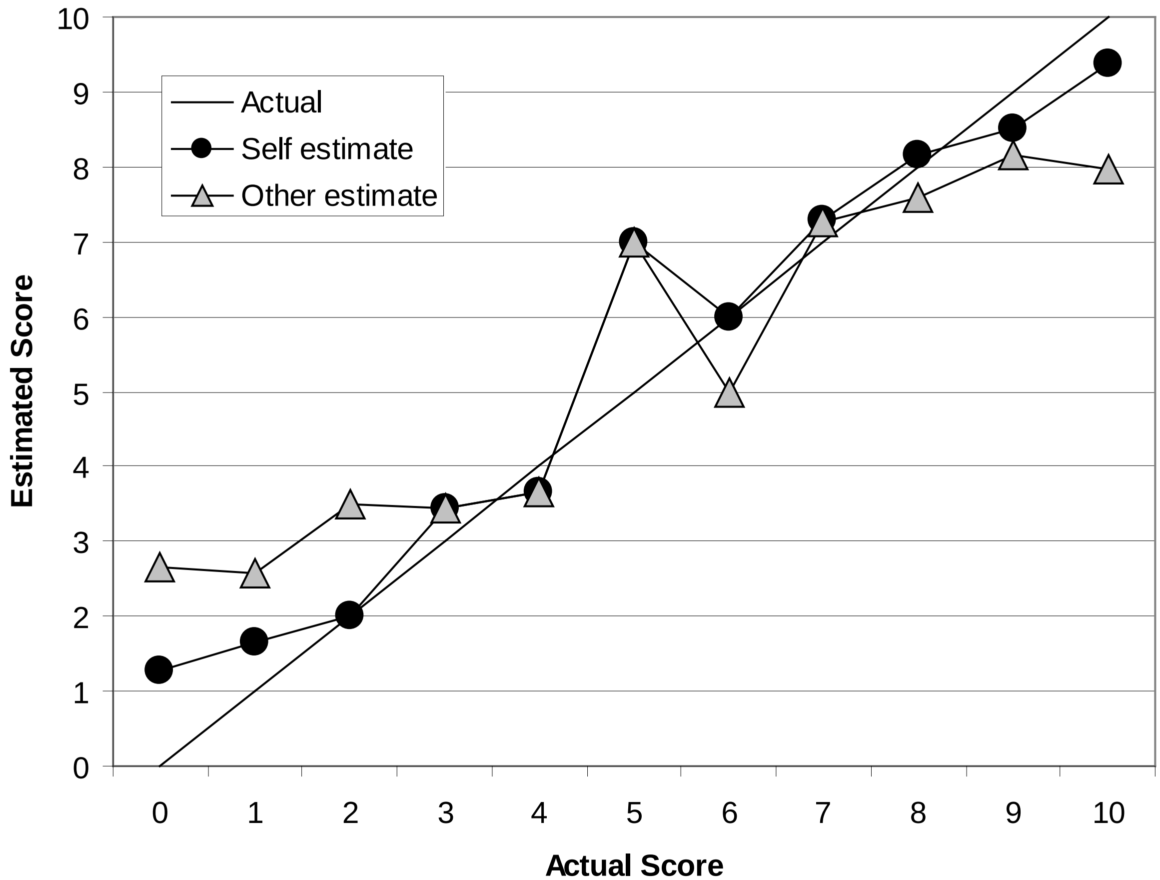

Figure 4. Mediation analyses (among those betting on themselves), Experiment 2.

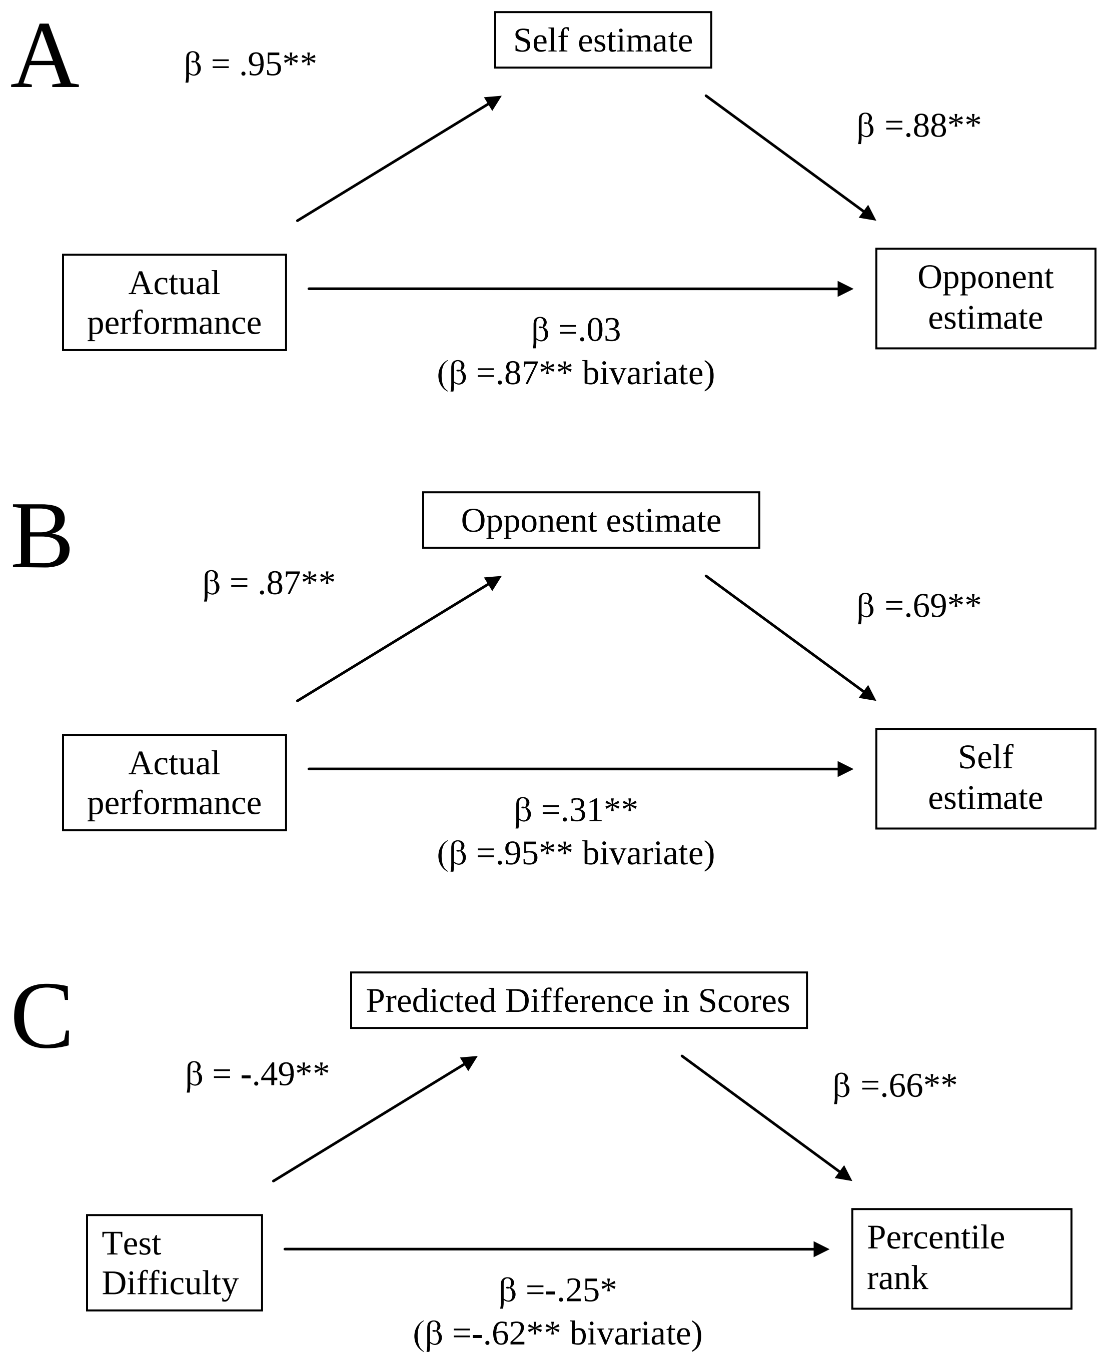

*p < .05, **p < .001

Figure 5. Estimated relative scores, Experiment 3.

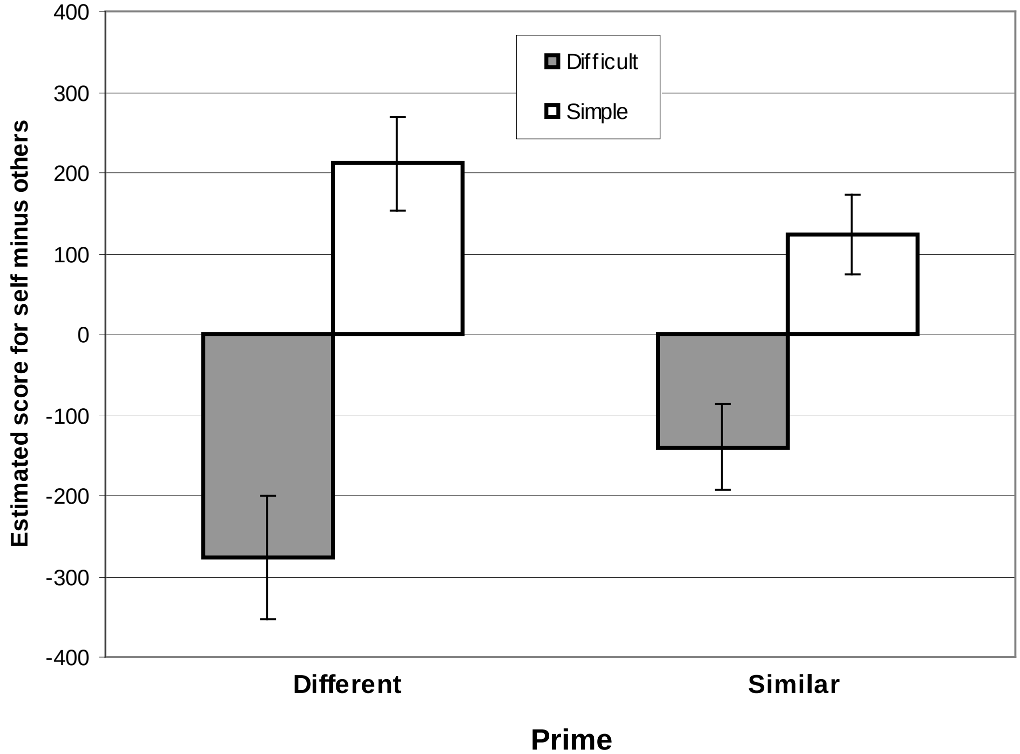

Figure 6. Estimated scores for self and opponent, both before and after taking the actual test, Experiment 3. Respondents are grouped by their own actual scores, by rounding their scores to the nearest 500 points.