

Unskilled and optimistic: Overconfident predictions despite calibrated knowledge of relative skill

Daniel J. Simons

Published online: 24 January 2013
© Psychonomic Society, Inc. 2013

Abstract Those who are less skilled tend to overestimate their abilities more than do those who are more skilled—the so-called *Dunning–Kruger* effect. Less-skilled performers presumably have less of the knowledge needed to make informed guesses about their relative performance. If so, the Dunning–Kruger effect should vanish when participants do have access to information about their relative ability and performance. Competitive bridge players predicted their results for bridge sessions before playing and received feedback about their actual performance following each session. Despite knowing their own relative skill and showing unbiased memory for their performance, they made overconfident predictions consistent with a Dunning–Kruger effect. This bias persisted even though players received accurate feedback about their predictions after each session. The finding of a Dunning–Kruger effect despite knowledge of relative ability suggests that differential self-knowledge is not a necessary precondition for the Dunning–Kruger effect. At least in some cases, the effect might reflect a different form of irrational optimism.

Keywords Social cognition · Decision making · Belief updating · Overconfidence

In domains ranging from driving to intelligence to sense of humor, well over 50 % of people believe themselves to be above-median performers (see Dunning, Heath & Suls, 2004). Moreover, those with less skill tend to overestimate their abilities more than do those with more skill—the so-called *Dunning–Kruger* effect (Kruger & Dunning, 1999). Why are

people so poorly calibrated in their self-assessments, and why is this bias greater for poor performers?

The traditional account (Kruger & Dunning, 1999) attributes the difference to metacognition: Due to gaps or distortions in their knowledge base, the unskilled inaccurately assess their own relative abilities. Because they tend not to realize how unskilled they actually are, they inflate their estimates of their own abilities. In contrast, skilled performers can better assess their own skill, so they are less overconfident in their relative self-assessments. Other factors, including regression to the mean (e.g., Krueger & Mueller, 2002) and task difficulty (e.g., Burson, Larrick & Klayman, 2006), can contribute to differences between good and poor performers, but the tendency for people as a group, and for the unskilled in particular, to exaggerate their own abilities persists (Dunning, 2011; Ehrlinger, Johnson, Banner, Dunning & Kruger, 2008).

In the vast majority of studies documenting the Dunning–Kruger effect, participants have judged how well they were performing without an objective metric of their own performance or of that of their peers (for an exception, see Park & Santos-Pinto, 2010). The use of subjective ability domains (e.g., sense of humor; Kruger & Dunning, 1999) and the lack of information about relative standings might make these domains more subject to biases in self-assessment. Even for more objective domains (e.g., math ability), participants often lack access to information about themselves or their peers when making their relative judgments.

In most studies, participants rate their performance on a single, just-completed measure, typically without feedback about relative performance. To my knowledge, only one published study has explored predictions about future performance by people who are fully aware of their own skill level (Park & Santos-Pinto, 2010). In that study, chess players (and poker players) predicted their final scores in a tournament before play began, and their expected results

D. J. Simons (✉)
Department of Psychology, University of Illinois,
603 E. Daniel Street, Champaign, IL 61820, USA
e-mail: dsimons@illinois.edu

exceeded the actual ones. Moreover, weaker players were more overconfident than better players (see also Chabris & Simons, 2010).

Studies of subjective judgments in the absence of feedback conflate several reasons why people might overestimate their abilities. Accurate self-assessments require both memory/knowledge of their own skill and knowledge of the skill levels in the comparison group. Should either of these forms of information be lacking, those with a positive self-bias will tend to give inflated self-assessments. For example, people might be biased to remember more of their own successes than failures, even if they were calibrated in their memory for the successes and failures of their peers (see Helzer & Dunning, 2012, for a related argument). Alternatively, someone who accurately represents their own skill level might underestimate those of their peers.

Even if people are perfectly calibrated in their knowledge of their own performance and of that of their peers, they might still be overconfident in evaluating their recent performance or in predicting their future performance: They might predict that they will “do better this time.” If levels of optimism varied with skill levels, differential self-knowledge would not be a necessary precondition for a Dunning–Kruger effect. Some forms of optimistic forecasting appear resistant to feedback. For example, football fans show persistent overconfidence when predicting their favored team’s outcome each week, despite receiving feedback about the accuracy of their predictions each time the team wins or loses (Massey, Simmons & Armor, 2011). Moreover, people tend to rely more on their aspirations than on their past performance when predicting their own future performance (e.g., Helzer & Dunning, 2012).

Determining whether the Dunning–Kruger effect emerges only when the less skilled have relatively impoverished knowledge requires a task that would be subject to all of these possible sources of overconfidence. Ideally, it should be a task in which participants generate multiple judgments, so that it would be possible to separate distorted memory for past experiences from distorted assessments of current or future performance. If biased self-assessments arise from poor calibration, then providing people with performance feedback should lead to better calibration and less-biased self-assessments. Thus, a reduction in the relationship between skill and overconfidence following feedback would be consistent with the differential-self-knowledge explanation. However, if the bias results from overly optimistic predictions (Massey et al., 2011), it should persist even when people do know their own skill level; they would consistently expect to “do better this time” than they have done on average.

To distinguish among these alternatives, I explored predictions made by competitive bridge players in a local

duplicate bridge club over a period of 2 months. The “matchpoints” score in a bridge session is normed to the performance of the rest of the pairs of players, and after each session, players receive objective feedback about their relative performance (a percentile ranking). Duplicate bridge players are more skilled than more casual players, and they play regularly against the same competitors for years. Consequently, if weaker players are more overconfident than stronger players, in this context the difference could not be solely due to differential knowledge of relative performance levels.

By obtaining repeated predictions, it is possible to determine whether accurate feedback about relative performance leads to better-calibrated predictions. If so, then differential self-knowledge might well underlie the Dunning–Kruger effect, even with experienced participants. This approach also can distinguish between mechanisms that might contribute to overconfidence. For example, players might be overconfident because they are biased to discount their poor results and to remember their positive results; overconfident predictions would follow from this memory distortion. Alternatively, players might be overconfident in their predictions even if they do accurately remember their past performance. If so, their overconfidence would not be based on a lack of knowledge of their relative skill. Finally, if strong and weak players do not differ in their levels of overconfidence, then the Dunning–Kruger effect might be limited to cases in which people make judgments about subjective domains or lack objective feedback about their relative performance.

Method

Participants

Players at the Champaign, Illinois, bridge club participated voluntarily. The player who made the most accurate predictions over the course of the study and the player who made the most predictions (one per session played) each won a token prize of a free entry to one session at the club (valued at either \$4 or \$5, depending on the session). Otherwise, participation was uncompensated.

The club hosts two 3-h-long sessions on most weekdays. A total of 165 players competed in at least one session during the course of the study, and 111 of them made at least one prediction. Most of the analyses focused on the 50 players who made at least five predictions. This subset consisted of players who tended to play at least once each week on average, and most had been playing in the club for some years, so they were familiar with the club, their opponents, and their own performance. All of the participants were experienced players, although they varied substantially

in skill (note, though, that the weakest club players tend to be better than the vast majority of recreational bridge players).

The study was conducted with the approval of the University of Illinois Institutional Review Board and in cooperation with those bridge club directors interested in helping with the study. Copies of a detailed information page were provided to participants and posted in the club.

Procedure

At the beginning of each of 41 sessions occurring between June 9 and August 31 of 2009, each player received a pre-printed card on which they wrote the session date and time, their name, their partner's name, and their current master-points (a rough measure of skill/experience in duplicate bridge that is confounded with number of sessions played). Finally, the card instructed them to "Please predict your final percentage score for this session. It doesn't matter if you go over or under—just try to predict your final score as accurately as possible. You can use up to 2 decimal places."

In duplicate bridge, each player has the same partner throughout a session and competes against other pairs of players. The session design includes counterbalancing to eliminate much of the randomness resulting from individual deals of the cards; each pair of players plays several bridge deals against each of the other pairs sitting in the opposite direction (i.e., a pair would sit "north/south" or "east/west" throughout the session), and each deal of the hands (known as a "board") is played once by each pair. After playing a board, the cards are not intermixed. Rather, they are placed back into a special holder and passed to the next table. That is, the deals of the cards are "duplicated." Results are tallied across an entire 3-h session, during which each pair of players plays between 15 and 30 bridge deals. Each pair's final result for the session is based on how they did relative to everyone else who played the same hands during that session, expressed as a percentage. Scores above 50 % represent above-average performance for that session, and most scores fall between 40 % and 60 %.

Each player independently predicted the pair's final percentage result without discussion, and the prediction cards were placed in a closed box before play began for that session. The actual score results from each session were matched to each predicted score (all results are posted publicly).

In order to determine whether they accurately remembered their performance in the prediction phase of the study, in sessions following the final prediction session, players in the club were asked to estimate their average result during the study period and over the previous year. The form read: "Please estimate your average percentage score across games you've played at the Champaign club since the start of the prediction study on June 9, 2009" and "Please

estimate your average percentage score across games you've played at the Champaign club for all of the past year (since August 1, 2008)."

Limits in scope and generalizability

The results from this study should generalize to experienced bridge players at other duplicate bridge clubs, as well as to other domains in which players regularly compete against the same group of players in games of skill, provided that the outcome of any individual match or session is not determined entirely by skill.

Results and discussion

Across all sessions in which predictions were collected (including players who did not submit predictions), the mean percentage score was 49.99 % ($SD = 7.74$ %, median = 50.00 %, range = 24.70 to 72.22, $N = 2,108$ player scores). Note that these scores are not independent; each member of a pair receives the same score for a given session, some players contributed more scores because they played in more sessions, and scores within a session are calculated on the basis of a comparison to the performance of other players in that session.

A total of 165 players participated in at least one session during the study period, and the mean of the average player scores during that period was 49.85 ($SD = 5.69$). Of those 165 players, 111 made at least one prediction, and 54 made no predictions (nonparticipants). The average score of the players participating in the study ($M = 50.04$, $SD = 4.53$) did not differ from that of nonparticipants ($M = 49.46$, $SD = 7.56$), $t(163) = 0.613$, $p = .541$. Moreover, players who made five or more predictions ($N = 50$; mean score = 49.94, $SD = 4.29$) also performed no better or worse than nonparticipants, $t(102) = 0.391$, $p = .696$. Consequently, any systematic deviations between predicted and actual scores are not attributable to systematic differences in skill between those who chose to participate in the study and those who did not. The remaining analyses focused on the 50 participants who made five or more predictions ($M = 10.46$ predictions, $SD = 5.91$, range = 5 to 30): They were among the most frequent and experienced players at the club, playing an average of 87 club sessions (not including tournaments) over the 12 months up to and including the study period ($SD = 52.16$, range = 15 to 226). They typically had played at the club for years and spanned a range of ability levels (the average scores during the prediction period ranged from 37.64 to 56.34). Given their extensive playing experience and the feedback that they received about their relative performance after each session, they had ample opportunity to calibrate their predictions.

From each player's predicted score, I subtracted the actual score to determine whether the predictions were overconfident. On average, players predicted that they would score 52.13 %, which was significantly higher than their actual average score of 50.13 % for those sessions, paired $t(49) = 5.67, p < .001$. In other words, players were 2 % overconfident in their predictions. Although this difference seems small, it is not—with the distribution of scores in the sample, a 2 % overestimate translates to 0.41 standard deviations above their actual performance, meaning that an average player expected to outperform approximately 66 % of his or her peers. Most of the players (39 out of 50) were overconfident in their predictions.

Consistent with the pattern of a Dunning–Kruger effect, the higher the average score (a measure of skill),¹ the less overconfident participants were ($r = -.480$; see Fig. 1). The bottom third of the players were significantly more overconfident (mean overconfidence = 3.06 %, or 0.62 standard deviations) than the top third (mean overconfidence = 0.25 %), $t(27) = 2.806, p = .0092$. Even the more skilled players showed some overconfidence, although they were better calibrated on average. Consequently, the relationship between skill and overconfidence cannot result from regression to the mean: Even the most skilled players were slightly overconfident.

The traditional explanation for the Dunning–Kruger effect is that less-skilled players have gaps in their knowledge that contribute to flawed judgments of their relative skill; they are unaware of their relative lack of skill, leading them to rely on less-informed beliefs about their performance to make their judgments. The negative correlation between overconfidence and skill in these data demands a different explanation; the bridge players in the study were familiar with their relative skill (they had played regularly with the same partners and against the same opponents) and received regular feedback about their performance and prediction accuracy, so they did have the information needed to make calibrated predictions.

Overconfident predictions could result from a memory failure: If people remember having performed better than they actually did, they would generate inflated predictions. For example, a player who averaged 50 % but remembered averaging 52 % would sensibly predict 52 % performance

for the current session. Such a memory distortion could result from a bias to remember positive results and to discount or forget negative ones. And if weaker players are more prone to memory inflation, that might explain the Dunning–Kruger pattern in the data.

Following the prediction phase of the study, players in the club were asked to estimate their average performance for the period of the study and for the year prior to the start of the study. A total of 31 of the 50 players who made five or more predictions were available to complete the remembered-performance phase, and their remembered average performance ($M = 49.95, SD = 4.59$) did not differ from their actual performance during the study ($M = 49.98, SD = 3.98, t(30) = 0.073, p = .942$). (The performance of the 19 who did not report their memories, $M = 49.86, SD = 4.86$, was comparable to that of the 31 who did.) Players misremembered their past performance by an average absolute deviation of 1.73 % ($SD = 1.30$). The size of these memory errors was unrelated to skill level ($r = -.009$), showing that players of different skills did not differ systematically in their ability to remember their actual performance.

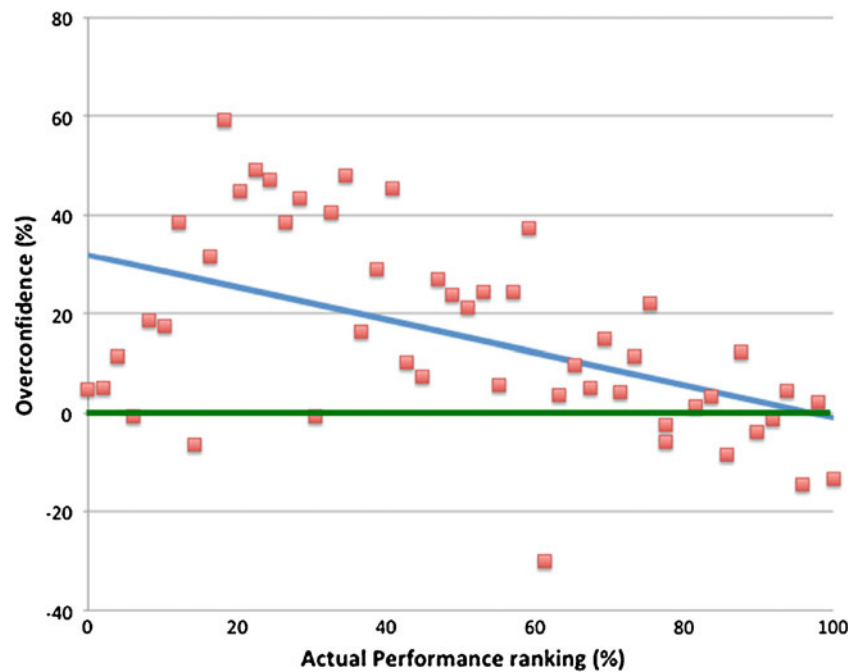
Crucially, remembered performance levels were equally distributed above and below actual performance. That is, as a group, participants were calibrated in their memory for their past performance; although their *predictions* were biased in the direction of overconfidence, their *memories* were not (mean of the signed errors = 0.028 %, $SD = 2.19$ %). The mean signed error—the extent to which remembered performance exceeded actual performance—also was uncorrelated with average performance during the study period ($r = .027$), showing no relationship between memory bias and skill.² Consequently, biased or distorted memory for past performance cannot account for the relationship between overconfidence and skill with respect to predicted performance.

Unlike in most studies of the Dunning–Kruger effect, participants in this study made multiple predictions. If judgments become more calibrated with repeated predictions and feedback, the number of predictions that a player made would be negatively correlated with the extent of their overconfidence. However, these factors were slightly *positively* correlated ($r = .162$). In other words, feedback about prediction accuracy did not reduce overconfidence in predictions (see also Massey et al., 2011). Consistent with this conclusion, the negative correlation between actual performance and overconfidence ($r = -.480$) was effectively the same after controlling for the

¹ I am using actual past performance as a measure of skill rather than trying to infer an underlying latent “skill” construct from independent measures of knowledge or ability. Given that the player judgments were predictions about performance and not judgments about relative knowledge, past performance was the optimal measure of skill. For domains with regular “head-to-head” competition (including bridge, chess, or tennis), past performance represents the most direct measure of actual playing ability or skill, even if it does not fully capture all of the knowledge that underlies that playing ability. Were participants required to make relative judgments about knowledge (e.g., understanding of bidding conventions), past performance might not be the optimal operational definition for skill.

² Overall, the difference between remembered and actual performance was correlated with the level of overconfidence in predictions ($r = .28$); players who had relatively greater memory inflation (or less memory deflation) had relatively greater overconfidence as well. But this pattern cannot account for the relationship between overconfidence and skill in the predictions, because skill is subtracted from both memory estimates and predictions to determine the extent of inflation or overconfidence. This correlation shows that an individual difference other than skill leads to consistently larger or smaller estimates *across* predictions and recollections.

Fig. 1 Overconfidence as a function of actual performance. Each square represents an individual player. The x -axis reflects the percentile ranking for that player's average score for all sessions in which he or she made a prediction. The y -axis was calculated by determining the percentile rank of that player's average predicted score and subtracting the ranking for his or her actual performance. The percentile ranking for the predicted scores was calculated with respect to the distribution of actual scores rather than the distribution of predictions in order to reflect the extent of overconfidence in predicted performance. Note that overconfidence decreased as actual performance increased



number of predictions (partial correlation $r = -.471$). The extent of overconfidence also was positively correlated with the total number of sessions played over the year preceding the study, $r = .174$, so experience at the club did not diminish overconfidence.

To determine whether individual players became better calibrated with successive predictions, for each participant I calculated the slope of the regression line predicting overconfidence from the prediction number. If participants became less overconfident with repeated predictions, the average slope should be negative, but in fact it was slightly positive (mean slope = 0.039, $SD = 0.673$). Even after excluding the 11 participants who showed underconfident predictions, the slope was essentially flat (mean slope = 0.011, $SD = 0.731$). These findings again suggest that overconfident predictions result from factors other than self-awareness. Players were informed that their previous predictions were overconfident, but they remained just as overconfident in making new predictions (see also Massey et al., 2011).

Given that overconfident predictions persisted in the face of repeated, accurate performance feedback and that memory for past performance was not inflated and was equally accurate across players of varied skill, the overconfident predictions in this study must result, at least in part, from a source other than differential knowledge of relative ability.

General discussion

Experienced bridge players who were aware of their relative skill level predicted that they would do better than they actually did, and this overconfidence persisted despite

regular performance feedback (see Massey et al., 2011, for a similar pattern). The extent of overconfidence was less for more skilled players, consistent with the pattern of a Dunning–Kruger effect (Kruger & Dunning, 1999).

Most of the evidence for the Dunning–Kruger effect has focused on judgments about subjective ability dimensions (e.g., sense of humor) or used tasks rarely performed by the participants. In many cases, participants also made just one evaluation of their relative performance, without any direct feedback about their own performance or that of their peers. Under such conditions, their only sources of information were their own performance expectations and their ability to evaluate how difficult they found the task, both of which are influenced by their lack of ability (Dunning, Johnson, Ehrlinger & Kruger, 2003)

In contrast, the bridge players in this study received feedback about their relative standing after every session and accurately recalled their own past performance. Moreover, they were knowledgeable about the subject domain, demonstrating that true incompetence is not necessary for excessive confidence; a Dunning–Kruger pattern can emerge from *relative* differences in skill level. Taken together, these findings show that the Dunning–Kruger effect for predictions of bridge performance does not result solely from a lack of knowledge of relative skill or from memory distortions favoring positive outcomes. Note, though, that the extent of overconfidence in this study (0.41 standard deviation units, or about a 16 % inflation in relative skill) is somewhat smaller than those in many other studies that have shown a relationship between overconfidence and ability. Differences in self-knowledge might well have inflated the extent of overconfidence in those studies, even if such

differences are not strictly necessary to produce a pattern consistent with the Dunning–Kruger effect.

Why, if players have a good sense of their own relative skill and past performance, would they remain overconfident? And why would weaker players be more overconfident?

One possibility is that players might have underappreciated the skill of their opponents in each session, leading them to predict a higher final score. Most players in the club are familiar with the skills of the other regular players at the club. Consequently, their biases were not likely due to a misevaluation of individual player abilities. However, the combinations of players and partners change from session to session, depending on who attends on a particular day. When scanning the room and making their judgments before the start of a session, players might tend to focus more on the weaker players, leading to underestimates of their overall competition. That possibility is consistent with evidence that, when making judgments of relative health risks, people tend to anchor their judgments to those in worse condition, leading to overly optimistic judgments (Weinstein & Klein, 1995). Such biased sampling and anchoring might not be extinguished by feedback because of the variability in the participants across sessions.

Alternatively, rather than all players focusing on weaker players, if players focused more attention on those players who were like them, then weak players would pay more attention to other less-skilled players, and strong players would devote more attention to stronger players. Assuming that judgments of the average strength of the players in a session are biased by such differential attention, then weak players would underestimate the average playing strength of their opponents to a much greater extent than would strong players. Thus, a Dunning–Kruger pattern could emerge even if weaker players were fully aware of their relative lack of skill. This possibility is consistent with evidence that, despite full knowledge of their relative ranking, weaker chess players expect to perform better in tournaments than they actually do (Park & Santos-Pinto, 2010; see also Chabris & Simons, 2010). Note, though, that these accounts differ from the typical explanation for overconfidence, in that they focus on overly negative assessments of the competition rather than overly positive self-assessments. Future studies could assess judgments of the relative skill of all other players in a room to test this account.

Another possibility is that players lacked a good sense of how strong their partners were. If they thought that their partners were better than they actually were, they would predict better-than-expected performance for that session. In most cases, though, players at the club play with the same small group of partners many times over a period of years, so they should have been at least as familiar with that player's ability as they were with the skills of others at the club.

The Dunning–Kruger effect in these data also might have resulted from some aspect of self-knowledge other than awareness of relative ability and past performance. Both weak and strong players remembered how they had performed in the past, suggesting that they understood their relative skill level. And, if they just used their past performance to predict their future performance, they would make calibrated predictions. But their predictions might draw on other factors that could lead to a greater bias for less skilled players. For example, when predicting future performance, people tend to rely more on evidence of past performance when judging the skill of others, but to rely more on their aspirations when predicting their own performance (Helzer & Dunning, 2012). If so, the fact that players are familiar with the skill of other players in a session might lead to accurate predictions for the performance of their opponents but inflated predictions for themselves. It is unclear, though, why that difference would interact with skill.

Another subtle factor that could contribute to a Dunning–Kruger pattern is that more-skilled players might be better able to assess the interaction of their skill level with that of their partner when predicting the results of a session. Mitigating this possibility is the fact that most players in the club play regularly with the same partners and do accurately recall their past performance with those partners. That familiarity and experience also dampens the possibility that weaker players could be more overconfident because they think that their skill has improved more rapidly than it actually has—many of the weaker players in the sample had been playing regularly for years, with relatively stable relative performance.

Perhaps the most plausible mechanism for the Dunning–Kruger effect in this context would be unmerited optimism (Massey et al., 2011). Weaker players know that they are less skilled, but they hold out hope for a good result in any given session. When thinking about their future, people are more likely to imagine positive than negative outcomes (Newby-Clark & Ross, 2003), and the same principle might apply to anticipated bridge performance. Moreover, the idea that weaker players might be even more optimistic about their chances is plausible, in that undue optimism might provide them with a source of motivation. Being irrationally optimistic about their future prospects might keep them playing regularly, even if they do not consistently generate top-level results. In fact, performing poorly in general while occasionally earning a high score in a session could reinforce that optimism, just as the occasional win at a slot machine reinforces people for playing slots. Their optimism might reflect a gambler's fallacy, a belief that they are *due* for a good night. More experienced players do not need to be overly optimistic in order to continue playing—they regularly achieve higher scores.

Future studies could explore whether weaker players are more likely to underestimate the average skill of the competition or whether they just have an overly rosy view of their future prospects.

Author note Thanks to Chris Fraley and Gary Dell for feedback on an earlier draft of the manuscript and to Gary Dell and Chris Chabris for helpful input during the process. Thanks especially to Karen Walker, John Brandeberry, Debbie Avery, and Martha Owen Leary for helping coordinate data collection at the Champaign–Urbana Community Bridge Club and to all of the players who participated.

References

- 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*, 60–77.
- Chabris, C. F., & Simons, D. J. (2010). *The invisible gorilla, and other ways our intuitions deceive us*. New York, NY: Crown.
- Dunning, D. (2011). The Dunning–Kruger effect: On being ignorant of one's own ignorance. *Advances in Experimental Social Psychology, 44*, 247–296.
- 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., Johnson, K., Ehrlinger, J., & Kruger, J. (2003). Why people fail to recognize their own incompetence. *Current Directions in Psychological Science, 12*, 83–87.
- Ehrlinger, J., Johnson, K., Banner, M., Dunning, D., & Kruger, J. (2008). Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent. *Organizational Behavior and Human Decision Processes, 105*, 98–121.
- Helzer, E. G., & Dunning, D. (2012). Why and when peer prediction is superior to self-prediction: The weight given to future aspirations versus past achievement. *Journal of Personality and Social Psychology, 103*, 38–53.
- Krueger, J., & 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*, 180–188.
- 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.
- Massey, C., Simmons, J. P., & Armor, D. A. (2011). Hope over experience: Desirability and the persistence of optimism. *Psychological Science, 22*, 274–281.
- Newby-Clark, I. R., & Ross, M. (2003). Conceiving the past and future. *Personality and Social Psychology Bulletin, 29*, 807–818.
- Park, Y. J., & Santos-Pinto, L. (2010). Overconfidence in tournaments: Evidence from the field. *Theory and Decision, 69*, 143–166.
- Weinstein, N. D., & Klein, W. M. (1995). Resistance of personal risk perceptions to debiasing interventions. *Health Psychology, 14*, 132–140.