

Temporal view of the costs and benefits of self-deception

Zoë Chance^{a,1}, Michael I. Norton^a, Francesca Gino^b, and Dan Ariely^c

^aDepartment of Marketing and ^bDepartment of Negotiations, Organizations, and Markets, Harvard Business School, Boston, MA 02163; and ^cDepartment of Marketing, Duke University, Durham, NC 27708

Edited by Donald W. Pfaff, The Rockefeller University, New York, NY, and approved January 28, 2011 (received for review August 16, 2010)

Researchers have documented many cases in which individuals rationalize their regrettable actions. Four experiments examine situations in which people go beyond merely explaining away their misconduct to actively deceiving themselves. We find that those who exploit opportunities to cheat on tests are likely to engage in self-deception, inferring that their elevated performance is a sign of intelligence. This short-term psychological benefit of self-deception, however, can come with longer-term costs: when predicting future performance, participants expect to perform equally well—a lack of awareness that persists even when these inflated expectations prove costly. We show that although people expect to cheat, they do not foresee self-deception, and that factors that reinforce the benefits of cheating enhance self-deception. More broadly, the findings of these experiments offer evidence that debates about the relative costs and benefits of self-deception are informed by adopting a temporal view that assesses the cumulative impact of self-deception over time.

hindsight bias | lying | motivated reasoning | self-enhancement

People often rationalize their questionable behavior in an effort to maintain a positive view of themselves. We show that, beyond merely sweeping transgressions under the psychological rug, people can use the positive outcomes resulting from negative behavior to enhance their opinions of themselves—a mistake that can prove costly in the long run. We capture this form of self-deception in a series of laboratory experiments in which we give some people the opportunity to perform well on an initial test by allowing them access to the answers. We then examine whether the participants accurately attribute their inflated scores to having seen the answers, or whether they deceive themselves into believing that their high scores reflect new-found intelligence, and therefore expect to perform similarly well on future tests without the answer key.

Previous theorists have modeled self-deception after interpersonal deception, proposing that self-deception—one part of the self deceiving another part of the self—evolved in the service of deceiving others, since a lie can be harder to detect if the liar believes it to be true (1, 2). This interpersonal account reflects the calculated nature of lying; the liar is assumed to balance the immediate advantages of deceit against the risk of subsequent exposure. For example, people frequently lie in matchmaking contexts by exaggerating their own physical attributes, and though such deception might initially prove beneficial in convincing an attractive prospect to meet for coffee, the ensuing disenchantment during that rendezvous demonstrates the risks (3, 4). Thus, the benefits of deceiving others (e.g., getting a date, getting a job) often accrue in the short term, and the costs of deception (e.g., rejection, punishment) accrue over time.

The relative costs and benefits of self-deception, however, are less clear, and have spurred a theoretical debate across disciplines (5–10). Perhaps due to the inherent challenges of documenting self-deception, previous inquiries have tended to focus more broadly on the costs and benefits of people's general tendency to view themselves in an overly positive light (11–16). In

line with previous theorizing, we define self-deception as a positive belief about the self that persists despite specific evidence to the contrary (17). Consider a classic demonstration of self-deception (18, 19). After hearing that they had failed or succeeded on a test, participants were asked to distinguish between recordings of their own and others' voices. Those who believed they had failed the test were then more likely to deny hearing their own voice. Although participants' tendency to reject their identity at a moment when they wished to distance themselves from it provides strong evidence of self-deception, it remains unclear whether that denial is on balance adaptive or maladaptive.

We suggest that the debate on the relative costs and benefits of self-deception can be informed by adopting the same temporal view of the costs and benefits of deception of others. Most previous investigations of self-deceptive phenomena, like the one described above, are backward-looking, examining how people in the present cope with their past behavior. In contrast, we introduce a forward-looking paradigm to examine how self-deception influences predictions about the future. This temporal perspective also allows us to simultaneously investigate the costs and benefits of self-deception at different time points: we suggest, and our results demonstrate, that, like lying, although self-deception can be beneficial in the short term, basing decisions on erroneous beliefs can prove costly in the longer term.

In our paradigm, participants take tests assessing their general knowledge and IQ. Some are given the opportunity to view an answer key while taking an initial test, whereas those in a control group have no such opportunity. We compare these two groups on two primary measures: performance on the first test (to assess the impact of having the answers) and predictions of future performance on a similar second test lacking an answer key (to assess self-deception). We predict that participants given access to the answers will outperform the control group on the first test, using the answer key to their advantage. In the absence of self-deception, there should be no difference between groups in predicted performance on the second test. Access to the answer key, however, makes initial performance a noisy signal of ability: a high score on the first test will be due to some combination of ability and the presence of the answers. Whether participants deliberately look at the answers to cheat or, having seen an answer to a question, come to believe that they “knew it all along” and fall victim to hindsight bias (20–23), they lack clear cues as to which factor—their ability or the presence of the answers—accounts for more of the variance. We suggest that participants given the answers to the first test will overweight their ability and

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Quantification of Behavior,” held June 11–13, 2010, at the AAAS Building in Washington, DC. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/quantification.

Author contributions: Z.C., M.I.N., F.G., and D.A. designed research; Z.C., M.I.N., and F.G. performed research; Z.C., M.I.N., and F.G. analyzed data; and Z.C., M.I.N., F.G., and D.A. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: zchance@hbs.edu.

underweight the presence of the answers, and therefore expect continued superior performance. In short, we predict that having access to the answers will both enhance performance on the first test and trigger self-deception—holding positive beliefs about the self (“I am a good test-taker”) despite negative information to the contrary (“I saw the answers”).

We first demonstrate that people self-deceive by failing to account for the impact of having the answers to an initial test when predicting performance on a subsequent test, and then examine whether people can foresee this self-deception. In experiment 1, some participants were given the answer key for an initial test (answers condition), and others were not (control condition). After completing and scoring this test, participants in both groups predicted their performance on a hypothetical, longer test that did not have the answers. We observed score inflation on the first test: those in the answers condition reported solving more problems correctly than those in the control condition (mean answers = 6.45, SD = 1.18; mean control = 5.58, SD = 1.41), $t(74) = 2.92$, $P < 0.01$. More importantly, those who had the answer key to the first test also expected to perform better on the second (mean answers = 81.4, SD = 15.3; mean control = 72.7, SD = 19.8), $t(73) = 2.14$, $P < 0.04$. Thus, participants in the answers condition, as predicted, failed to correct for the effect of the answers on their performance, instead deceiving themselves into believing that their strong performance was a reflection of their ability.

To test whether people are aware of this tendency to self-deceive, we asked a separate group of participants to estimate both their performance on a first test and their subsequent predictions for a second test in a similar experiment in which they were assigned to either the answers or control condition. Participants who imagined having the answers did expect to achieve higher scores on the first test (mean = 89.2, SD = 12.9) than those in the control condition (mean = 69.2, SD = 30.6), $t(34) = 2.60$, $P < 0.02$. However, they did not foresee self-deception: in contrast to participants in experiment 1 who actually took the test, those who had merely imagined having the answers expected that they would predict worse performance when the answer key was absent (mean = 77.9, SD = 17.0), paired $t(18) = 3.29$, $P < 0.01$, suggesting that they were unaware that they would attribute their improved performance to the presence of the answers.

In experiment 2, we sought additional evidence for the presence of self-deception in two ways. First, whereas in experiment 1, participants predicted only how they would perform on the second test, in experiment 2, participants actually took a second test, allowing us to confirm that the estimates for the second test were inflated, and therefore self-deceptive. Second, we included a measure of dispositional self-deception (24, 25) to show that these overestimations tracked with participants' chronic inclination to self-deceive. Because high self-deceivers are more likely than others to ignore evidence of their failures (26, 27), we predicted that dispositional self-deception would moderate their biased performance predictions. We expected that high self-deceivers would be likely to provide inflated predictions both with and without the answers; we also expected the increased ambiguity in the answers condition to enhance the opportunity for self-deception (28), amplifying the difference in predicted scores between high and low self-deceivers in this condition.

In experiment 2, participants completed a test of general knowledge either with or without the answers at the bottom of the page, looked over a second test without answers, then predicted their scores on and completed the second test. As in experiment 1, those in the answers condition solved more problems correctly on the first test (mean answers = 8.97, SD = 2.08) than those in the control group (mean control = 6.29, SD = 2.52), $t(129) = 6.65$, $P < 0.001$, and again expected to outperform the control group on the second test (mean answers = 7.73, SD =

2.20; mean control = 6.32, SD = 1.75), $t(129) = 4.04$, $P < 0.001$. As predicted, however, there was no significant difference between the two groups in actual performance on the second test (mean answers = 5.18, SD = 2.13; mean control = 4.86, SD = 1.35), $t(129) = 1.03$, $P = 0.31$, demonstrating that participants with the answers deceived themselves into believing they were smarter than their results on the second test proved them to be (Fig. 1).

One week before experiment 2, all participants had completed a self-deception scale (24, 25). Dispositional self-deception was related to predicted scores across both conditions, $\beta = 0.30$, $P < 0.001$, controlling for performance on test 1. Most importantly, and consistent with our prediction, there was a significant interaction between dispositional self-deception and our manipulation in predicting anticipated performance on test 2 ($\beta = 0.11$, $P = 0.05$; Fig. 2), controlling for scores on the first test. In the control condition, dispositional self-deception was correlated with predicted performance (average $\beta = 0.20$, $P = 0.01$), but having the answers magnified this relationship substantially (average $\beta = 0.41$, $P < 0.001$), suggesting that dispositional self-deceivers were particularly prone to taking credit for their answers-aided performance (Fig. 2). The nature and significance of the results did not change when we conducted additional regression analyses in which we controlled for other personality traits, and these traits did not significantly predict the gap between actual and anticipated performance on test 2 (all β 's ranged from -0.07 to 0.10 , all P 's > 0.17), suggesting a unique role for dispositional self-deception in the inflation of predicted performance.

To examine the temporal dimension of self-deception (short-term gains with longer-term costs), we explored whether monetary incentives could temper inflated predictions of performance. In experiment 3, after completing an initial test with or without answers, participants learned they could earn up to \$20 on the second test, depending on both their test performance and the accuracy of their prediction of that performance. If participants were aware of the impact of having the answers, then the monetary bonus for prediction accuracy should lower their inflated estimates for the second test; we expected, however, that because people are unaware of self-deception, incentives would fail to serve as a corrective force, such that participants would pay the price when they underperformed relative to their predictions.

As before, we observed cheating on the first test (mean answers = 7.61, SD = 2.20; mean control = 4.58, SD = 2.43), $t(76) = 5.77$, $P < 0.001$. Despite the monetary incentive to predict their scores as accurately as possible, participants in the answers condition persisted in predicting superior performance




Fig. 1. Predicted vs. actual performance on test 2 (experiment 2). The group with the previous advantage of the answer key again anticipated superior performance. However, we observed no difference in actual performance between the two groups.




Fig. 2. Regression model interaction with dispositional self-deception (experiment 2). Predictions of performance on the second test are moderated by dispositional self-deception, with high self-deceivers (+1 SD above the mean) showing greater inflation in the answers condition than low self-deceivers (−1 SD below the mean).

on the second test (mean answers = 7.24, SD = 1.79; mean control = 4.98, SD = 1.41), $t(76) = 6.22, P < 0.001$. Because both groups again performed equally well (mean answers = 4.47, SD = 1.47; mean control = 4.45, SD = 1.85), $t < 1$, participants in the answers condition made larger prediction errors (mean answers = 2.76, SD = 2.16; mean control = 1.13, SD = 0.97), $t(76) = 4.36, P < 0.001$, and therefore earned less money on the second test (mean answers = \$14.47, SD = 4.32; mean control = \$17.75, SD = 1.93), $t(76) = 4.36, P < 0.001$ (Table 1). Thus, the self-deception resulting from the short-term benefit of cheating on the first test led to longer-term (monetary) costs on the second test. The results of experiment 3 demonstrate that, unlike some biases which are mitigated by monetary incentives, such as conformity (29), self-deception is not reduced by financial costs, at least at the incentive levels used here. We note that though this experiment was designed to demonstrate that self-deception can have longer-term costs, these results do not suggest that self-deception is invariably costly; there are likely some situations in which it proves beneficial. The fact that people are unable to undo or correct for their self-deception even when it is costly, however, offers further strong evidence of a lack of awareness of this process.

Thus far we have considered self-deception as an intrapersonal phenomenon, but people's private acts of self-deception often become public. In several documented cases, for example, individuals who posed as war heroes—acquiring fame, money, and political office—appear to have come to believe their own lies (30). In experiment 4, we examined the reinforcing influence of social feedback on self-deception, by adding an additional manipulation to our standard answers/no answers paradigm. After completing the first test, but before predicting their scores on the second, some participants were randomly assigned to receive a certificate of recognition. The experimenter told each of these participants that certificates were given to everyone who scored above average on the test, and wrote the participant's name and score on the certificate. As before, participants in the answers conditions reported higher scores on the first test (mean answers = 7.85, SD = 2.18;

Table 1. Self-deception proves costly: Experiment 3

	Control	Answers
Test 1 score	4.58	7.61***
Test 2 prediction	4.98	7.24***
Test 2 score	4.45	4.47
Earnings	\$17.75	\$14.47***

Participants given the answer key for the first test performed better on that test than those in the control condition, and erroneously predicted higher scores on the second test, resulting in lower earnings due to their similar performance. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$.

mean control = 4.01, SD = 2.11), $t(134) = 10.44, P < 0.001$, and also predicted higher scores on the subsequent longer test that lacked answers (mean answers = 72.93, SD = 17.05; mean control = 50.59; SD = 20.41), $F(1, 132) = 52.87, P < 0.001$. Receiving a certificate also increased performance predictions overall (mean certificate = 66.93, SD = 21.11; mean no certificate = 56.62, SD = 21.48), $F(1, 132) = 11.23, P < 0.01$.

Most importantly, we observed a significant interaction, $F(1, 132) = 4.05, P < 0.05$, such that the certificate's enhancement of self-deception was restricted to the answers condition (mean certificate, answers = 81.2, SD = 14.9; mean no certificate, answers = 64.7, SD = 15.1), $t(66) = 4.54, P < 0.001$, with no effect on the control group (mean certificate, control = 52.7, SD = 16.2 vs. mean no certificate, control = 48.5, SD = 24.0), $t(66) = 0.83, P = 0.41$. As we had expected, social recognition exacerbated self-deception: those who were commended for their answers-aided performance were even more likely to inflate their beliefs about their subsequent performance. The fact that social recognition, which so often accompanies self-deception in the real world, enhances self-deception has troubling implications for the prevalence and magnitude of self-deception in everyday life. As the eventual disgrace of those individuals caught exaggerating their military service demonstrates, however, those who benefit from self-deception may eventually pay a high price when those beliefs are publicly repudiated.

Our experiments demonstrate that people who use an answer key to perform well on a test interpret their resultant high scores as evidence of superior intelligence. As a result, when asked to predict their performance on a future task, they fail to account for the impact of having had the answers, even when inflated predictions prove costly. In addition, once that initial behavior has occurred, debiasing the ensuing self-deception (e.g., with monetary incentives) proves difficult. Finally, we show that self-deception that occurs at the level of the individual can be intensified in a social context, when the rewards that accrue as the result of self-deception are reinforced by others.

This research offers several contributions to the larger understanding of self-deceptive processes. First, whereas much of the previous literature on self-deception has been more theoretical than empirical (with few exceptions) (18, 19, 31, 32), we introduce a paradigm that reliably elicits self-deception, allowing for closer empirical examination of the phenomenon. Second, whereas previous research views self-deception as a filtering process by which negative information is excluded from consciousness to preserve a positive self-view, our experiments suggest that negative behavior can in fact be the source of people's inflated opinions of themselves. Third, although the construct of self-deception has a long history in psychology, the nature of its underlying mechanism is still subject to debate (33). Our research contributes to this literature by distinguishing people's awareness of poor behavior from their understanding of that behavior's role in self-deception. We show that people are aware that they occasionally engage in questionable behavior, but fail to predict the aftermath of having engaged in that behavior; people understand they will deceive, but fail to perceive the processes by which that deception leads to self-deception. Finally, and more broadly, our findings inform the larger debate on the adaptive or maladaptive nature of self-deception, by demonstrating that costs and benefits depend on temporal (short- or long-term) and contextual (private or public) variables.

We have focused on one particular instantiation of self-deception: the impact of cheating on people's beliefs about their test-taking ability. Given the difficulty of behaving consistently with one's ideals when they conflict with one's wishes, coupled with the importance of positive self-regard, self-deception is likely common. Sadly, though people are willing to use a single ambiguous incident to make globally negative judgments about others (34), our findings show that people not only fail to judge

themselves harshly for unethical behavior, but can even use the positive results of such behavior to see themselves as better than ever.

Methods

Experiment 1. Seventy-six participants were approached in the Massachusetts Institute of Technology student center and asked to take a test that would measure their math IQ; each was paid \$3. We provided two sample questions followed by an eight-item test of math IQ (e.g., "If a man weighs 75% of his own weight, plus 42 pounds, how much does he weigh?"). Tests for those in the answers condition had an answer key printed at the bottom of the page, but were otherwise identical to the control condition. Participants scored their own tests. Finally, they predicted how many questions they could answer correctly if asked to complete another 100 similar questions.

Prediction Experiment. Thirty-six Harvard undergraduates volunteered to complete the experiment as they waited for a lecture to begin. They were asked to imagine they would be taking a difficult, 100-question test of their IQ, and were provided with two sample questions. Those in the answers condition were told to imagine that as they took the test, they would be able to refer to an answer key at the bottom of each page; those in the control condition were told to imagine taking the test without the answers. All participants then estimated how many questions they would answer correctly. Finally, they were asked to estimate how well they would predict they would do on an additional 100-question test for which they would not have the answers.

Experiment 2. One hundred thirty-one University of North Carolina undergraduate and graduate students completed the experiment on paper and were paid based on their performance on the task (\$1 per correct answer on each of the two tests) in addition to a \$5 flat fee for completing an online questionnaire. All participants completed the online questionnaire 1 wk before the experiment. Upon reporting to the laboratory, participants took a 10-question general knowledge test of medium difficulty (e.g., "How many US states border Mexico?" and "In which US state is Mount Rushmore located?"). Half of the participants were given the answers at the bottom of the page; half were not. All participants then glanced over a second test (and could see that this test had no answer key), predicted their score on this test, and completed it.

The online questionnaire included measures of dispositional self-deception and personality traits. To measure trait self-deception, we used the Balanced Inventory of Desirable Responding (BIDR), consisting in its complete form of three 20-item subscales: Self-Deceptive Enhancement, Impression Management, and Self-Deceptive Denial (24, 25). We focused on Self-Deceptive Denial because it highly correlates with Impression Management

and Self-Deceptive Enhancement, thus making the two other subscales somewhat redundant (26). To test for the specificity of the relationship between measures of self-deception and our paradigm, we also measured the Big Five personality factors (neuroticism, extraversion, openness to experience, agreeableness, and conscientiousness) using the Ten Item Personality Inventory (35).

Experiment 3. Seventy-eight University of North Carolina undergraduate students completed the same test of general knowledge used in experiment 2. As in our other experiments, participants in the answers condition had the answers printed at the bottom of each page, whereas those in the control condition did not. After scoring and receiving payment of \$0.50 per question for the first test, participants learned they would receive \$2 for each question on the next test if they guessed their score on that second test correctly; if their guess were off by one question, they would get \$1.80 for each question; if off by two questions, they would get \$1.60, and so on. They were given 1 min to look over the second test before completing it.

Experiment 4. One hundred thirty-six University of North Carolina undergraduate students completed the experiment on paper and were paid based on their performance on the task (\$1 per correct answer) in addition to a \$2 show-up fee. They were presented with the same 10 general-knowledge questions used in experiment 2. Each participant completed the session alone. The experiment used two between-subjects manipulations. First, as before, we manipulated the presence of the answers on the first test. Second, we manipulated whether participants received a certificate of recognition. In the certificate condition, upon completion of the first test, participants were asked to walk to the experimenter's desk to pick up the materials for the second part of the experiment, which consisted of the prediction task used in experiment 2. Along with this material, the experimenter gave the participant a certificate printed on a thick piece of paper. The experimenter told participants that the certificates were being given to those who scored above average on the test as recognition of their good performance, and recorded the participant's name and the number of correct answers on the first test on the certificate. In the no-certificate condition, participants only received the materials for the second part of the experiment. Participants returned to their seats and completed the prediction task.

ACKNOWLEDGMENTS. The authors thank Hayley Barna, Daniel Bravo, Jennifer Fink, Jennifer Ford, Leonard Lee, Daniel Mochon, Katie Offer, Carrie Sun, and Leslie Talbott for their assistance with data collection; and Shane Frederick, Don Moore, and the members of our laboratory group for their comments on a previous draft.

- Goleman D (1996) *Vital Lies, Simple Truths* (Simon & Schuster, New York).
- Trivers R (2000) The elements of a scientific theory of self-deception. *Ann N Y Acad Sci* 907:114–131.
- Rowatt WC, Cunningham MR, Druen PB (1999) Lying to get a date: The effect of facial physical attractiveness on the willingness to deceive prospective dating partners. *J Soc Pers Relat* 16:209–223.
- Norton MI, Frost JH, Ariely D (2007) Less is more: The lure of ambiguity, or why familiarity breeds contempt. *J Pers Soc Psychol* 92:97–105.
- Baumeister RF, Campbell JD, Krueger JI, Vohs KD (2003) Does high self-esteem cause better performance, interpersonal success, happiness, or healthier lifestyles? *Psychol Sci Public Interest* 4:1–44.
- Colvin CR, Block J (1994) Do positive illusions foster mental health? An examination of the Taylor and Brown formulation. *Psychol Bull* 116:3–20.
- Rorty A (1972) Belief and self-deception. *Inquiry* 15:387–410.
- Schafer R (1992) *Retelling a Life: Narration and Dialogue in Psychoanalysis* (Basic Books, New York).
- Swann WB, Jr., Stein-Seroussi A, Giesler RB (1992) Why people self-verify. *J Pers Soc Psychol* 62:392–401.
- Taylor SE, Brown JD (1988) Illusion and well-being: A social psychological perspective on mental health. *Psychol Bull* 103:193–210.
- Chambers JR, Windschitl PD (2004) Biases in social comparative judgments: The role of nonmotivated factors in above-average and comparative-optimism effects. *Psychol Bull* 130:813–838.
- Dunning D, Leuenberger A, Sherman DA (1995) A new look at motivated inference: Are self-serving theories of success a product of motivational forces? *J Pers Soc Psychol* 69:58–68.
- Kunda Z (1987) Motivated inference: Self-serving generation and evaluation of causal theories. *J Pers Soc Psychol* 53:636–647.
- Moore DA, Kim TG (2003) Myopic social prediction and the solo comparison effect. *J Pers Soc Psychol* 85:1121–1135.
- Robins RW, Beer JS (2001) Positive illusions about the self: Short-term benefits and long-term costs. *J Pers Soc Psychol* 80:340–352.
- Sedikides C, Strube MJ (1997) Self-evaluation: To thine own self be good, to thine own self be sure, to thine own self be true, and to thine own self be better. *Advances in Experimental Social Psychology*, ed Zanna MP (Academic, San Diego), Vol 29, pp 209–269.
- Mitchell J (2000) Living a lie: Self-deception, habit, and social roles. *Hum Stud* 23: 145–156.
- Gur RC, Sackeim HA (1979) Self-deception: A concept in search of a phenomenon. *J Pers Soc Psychol* 37:147–169.
- Sackeim HA, Gur RC (1978) Self-deception, self-confrontation, and consciousness. *Consciousness and Self-Regulation, Advances in Research and Theory*, eds Schwartz GE, Shapiro D (Plenum, New York), Vol 2, pp 139–197.
- Arkes HR, Wortmann RL, Saville PD, Harkness AR (1981) Hindsight bias among physicians weighing the likelihood of diagnoses. *J Appl Psychol* 66:252–254.
- Fischhoff B (1975) Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Exp Psychol Hum Perception and Performance* 104:288–299.
- Nestler S, Blank H, Egloff B (2010) Hindsight ≠ hindsight: Experimentally induced dissociations between hindsight components. *J Exp Psychol Learn Mem Cogn* 36: 1399–1413.
- Wood G (1978) The knew-it-all-along effect. *Exp Psychol Hum Perception and Performance* 4:345–353.
- Paulhus DL (1986) Self-deception and impression management in test responses. *Personality Assessment via Questionnaire*, eds Angleitner A, Wiggins JS (Springer, New York), pp 143–165.
- Paulhus DL (1991) Measurement and control of response bias. *Measures of Personality and Social Psychological Attitudes*, eds Robinson JP, Shaver PR, Wrightsman LS (Academic, New York), pp 17–59.
- Paulhus DL (2002) Socially desirable responding: The evolution of a construct. *The Role Of Constructs in Psychological and Educational Measurement*, eds Braun HI, Jackson DN, Wiley DE (Erlbaum, Mahwah, NJ), pp 46–69.

