Evolutionary Psychology

www.epjournal.net – 2014. 12(3): 635-654

Original Article

Deceiving Yourself to Better Deceive High-Status Compared to Equal-Status Others

Hui Jing Lu, Department of Applied Social Sciences, The Hong Kong Polytechnic University, Hong Kong SAR, China.

Lei Chang, Department of Educational Psychology, The Chinese University of Hong Kong, Hong Kong SAR, China. Email: leichang@cuhk.edu.hk (Corresponding author).

Abstract: The arms race between deception and detection is likely to have played out between individuals in different status hierarchies, with low-status individuals more likely to be deceivers and high-status individuals more likely to be detectors than the other way around. Memory and its distortion may be temporarily employed first to keep truthful information away from both self and others and later to retrieve accurate information to benefit the self. Using a dual-retrieval paradigm, we tested the hypothesis that people are likely to deceive themselves to better deceive high- rather than equal-status others. College student participants were explicitly instructed (Study 1 and 2) or induced (Study 3) to deceive either a high-status teacher or an equal-status fellow student. When interacting with the high- but not equal-status target, participants in three studies genuinely remembered fewer previously studied items than they did on a second memory test alone without the deceiving target. The results support the view that self-deception responds to status hierarchy that registers probabilities of deception detection such that people are more likely to self-deceive high- rather than equal-status others.

Keywords: self-deception, deception, memory, status

Introduction

Trivers (1976) introduced his theory of self-deception over three decades ago. According to his theory, individuals deceive themselves to better deceive others by placing truthful information in the unconscious while consciously presenting false information to others as well as the self without leaving cues to be detected of deception. This evolutionary theory has since been elaborated (Alexander, 1987; Trivers, 1985, 2000) and expanded (von Hippel and Trivers, 2011) but has not been empirically tested, in part because the theory does not specify a mechanism by which to place information in the unconscious or ways to falsify that such placement occurs. It is also unclear about the adaptive value of keeping truthful information from the self (Leeuwen, 2007), and it is difficult to separate deception from self-deception when inferred mainly from the self-serving outcome but not other conditions of deception. The purpose of the present study was to provide the first empirical test of Trivers' theory by addressing the aforementioned issues.

Specifically, we propose a dual-retrieval system to operationalize and empirically test self-deception within Trivers' theoretical framework. The dual-retrieval system that allows an individual to retrieve truthful information twice—once under the motivation to deceive and once after deception is achieved—resolves the potential controversy about the adaptiveness of keeping truthful information from the self (Leeuwen, 2007). We propose that, by keeping fitness-enhancing information away from both self and others, self-deception as an adaptation must cease to operate in most instances once the goal of deception has been achieved. Truthful information that once has been kept from both self and others will then be retrieved to benefit the self. Such information manipulation makes memory likely to have been co-opted to execute self-deception. Memory research thus provides a good paradigm within which to test hypotheses about self-deception. Although encoding may also be used (von Hippel and Trivers, 2011), especially in intra-individual self-deception, in this paper we focus on retrieval in inter-individual or social interactional contexts.

We also set conditions under which self-deception is deemed more likely than deception. Because the more costly self-deception evolved to escape detection, it must respond to social conditions that register the probability of deception detection such that people should self-deceive when they sense high rather than low probabilities of detection. The social status of the deceived, relative to that of the deceiver, represents one of the detection-registering conditions that should affect the likelihood and activation of self-deception. We argue that the evolutionary "arms race" between deception and deception detection is more likely to have played out between the low-status individuals as deceivers and high-status individuals as detectors rather than the other way around and is likely to lead to people deceiving themselves to deceive high- rather than low- or equal-status others. Whereas we focus on status as an example of increased detection probability, other variables such as the number of targets to be deceived also affect the probability of detection. Because it is generally easier to deceive one rather than multiple targets, high-status individuals are more likely to self-deceive the public or groups of individuals than single individuals. In this paper, we focus on social status, but not number of deceiving targets or others, as a condition in individual but not group deception.

The arms race between low-status deceivers and high-status detectors

Humans and other primates live in hierarchical social groups where status influences resource distribution (Boehm, 1999; Wright, 1994). High-status individuals who attained their position either by force or social intelligence (Ellis, 1995; Nettle, 2003) have more resources than low-status individuals and have the power to punish the latter for rule violations (Cummins, 1999, 2005). Low-status individuals who are under constant surveillance often attempt to hide resources from high-status individuals (Whiten and Byrne, 1988). In this case, low-status individuals should be more motivated to deceive, whereas high-status individuals should be more motivated to detect deception. High-status individuals have more honest means (through force or by changing the rules) to acquire resources than do low-status individuals. High-status individuals also have more resources—including information leading to the deception detection—and the means to punish deceivers. In contrast, low-status individuals are more limited detectors who may face revenge for detecting deception. Detecting deception does not enhance fitness if the detector is unable to punish but may be retaliated by the deceiver. There is thus more pressure to perfect deception when one has the need to deceive but also faces increased chances to be caught and punished. The same pressure to better deceive is much reduced when one expects little punishment from the detector if caught deceiving or has other honest means to pursue the same fitness gains. Social status may therefore shift selection pressure to favor low-status individuals over high-status individuals as fearful deceivers and the high-status individuals over low-status individuals as vigilant detectors. Thus, Trivers' arms race between deception and detection is likely to have played out between low-status deceivers and high-status detectors, leading to people deceiving themselves to better deceive high- rather than low- or equal-status others.

Animal research shows that subordinates spend time hiding desirable resources or rule violations from dominant peers who police the group (Byrne and Whiten, 1992; Mitchell, 1999). The young European male kestrel grows female-like plumage to avoid aggressive dominant older males and increase their chance of mating during the breeding season (Hakkarainen, Korpimäki, Huhta, and Palokangas, 1993). Ravens (Bugnyar and Kotrschal, 2004) and pigs (Held, Mendl, Devereux, and Byrne, 2002) take detours from food sources to mislead high-status conspecifics. Subordinate monkeys escorted their superiors far away from the food source, demonstrating a conspicuous and elaborate deceptive ploy (Amici, Call, and Aureli, 2009; Ducoing and Thierry, 2003). Low-status chimpanzees only retrieved their food when it would be unseen by a high-status competitor (Hare, Call, Agnetta, and Tomasello, 2000) or a human experimenter (Hare, Call, and Tomasello, 2006), whereas high-status chimpanzees took food openly (Hare et al., 2006). Field observations demonstrate social status-regulated deception and detection in more detail. In one case, a subordinate chimpanzee misled a dominant member away from where it had hidden a banana. Wary of deception, the high-status chimp did not go far. Later, when the subordinate retrieved the banana from its hiding place, the dominant chimp attacked (Whiten and Byrne, 1988). Relying on status, prestige, or violence in lieu of deception, high-status animals are able to disperse, drive away, or attack subordinates to obtain what they want (Wright, 1994). They also detect, prevent, and punish rule-breaking behaviors of low-ranking deceivers (Clutton-Brock and Parker, 1995) who, with no option of aggression or group alliance, are constantly pressured to improve their deceptive ploys.

Evidence in humans is found in hierarchical social relations. The supervisor-subordinate relationship is an example. Sales staff in retail stores and supermarkets did not report customer complaints to their supervisors and restaurant servers hid incorrect orders from their managers (Harris and Ogbonna, 2010; Payne, 2008). On the other hand, workplace superiors were vigilant toward potential subordinates deceit (Wintrobe and

Breton, 1986). In an experiment using the Wason selection task, participants assigned a supervisor role were better able to detect cheaters than participants who played subordinates (Manktelow and Over, 1991). Once within a status hierarchy, such as the employer-employee setting, human beings, like other animals, may be driven by the additional motivation consistent with status-regulated deception and detection. Various effective management programs (e.g., self-managed work groups; Manz and Sims, 1993) aim to break away from this mindset by reducing hierarchical elements of the work environment to improve labor relations and productivity (Kirkman and Rosen, 1999). The state-citizen relationship provides another example in which the higher-level state employs law enforcement (i.e., detection) and correctional systems (i.e., punishment) to control deception by its citizens (e.g., committing and hiding crimes, including cheating on taxes). Corruption provides another example that can happen when lower-ranking individuals or institutions choose not to report the misdeeds of the higher-ups for fear of retribution. In each of these hierarchical relationships, the higher level of the hierarchy is vigilant about detecting deception from the lower level, thereby exerting pressure on subordinates to improve deception strategies, meeting Trivers's hypotheses of self-deception.

Memory helps to execute self-deception

According to Trivers (2000), a blatant deceiver keeps both true and false information in the conscious mind but presents only falsehoods to others. In doing so, the deceiver may leave clues about the truth due to its conscious access. A self-deceiver keeps only false information in consciousness. Lying to others and to the self at the same time, the self-deceiver thus leaves no clues about the truth retained in the unconscious mind (Trivers, 2000). This proposed mechanism has raised questions about the adaptiveness of self-deception because, by concealing information from others, the deceiver cannot access information to benefit him or herself (Leeuwen, 2007). To address this concern, we propose a dual-retrieval memory process in which the self-deceiver first fails to retrieve information when motivated to deceive others and, after deception has been achieved, retrieves information to benefit him or herself. The difference between the two retrievals in the self-serving direction suggests self-deception.

Specifically, the dual-retrieval paradigm offers a self-deceiver two chances to retrieve (or distort) encoded information assumed to be fitness enhancing given the motivation to keep it from the deceived. The first retrieval normally takes place in the presence of the deceived. The second retrieval often (but not always) takes place without the deceived after a successful deception. For example, in front of a high-status food seeker, a low-status food hider forgets completely or partially where he/she hides food. After the food seeker leaves, the food hider remembers where the food is. In the first retrieval, the deceiver honestly distorts, partially retrieves, or fails to retrieve encoded information due to the motivation to deceive and to escape detection. In the second retrieval (when motivation changes from hiding to seeking food), the deceiver retrieves the true or full information to achieve fitness gains. Thus, the dual-retrieval paradigm addresses the issue of adaptiveness and offers a way to verify self-deception, which is discrepancy between the two retrievals.

Memory serves well to execute self-deception because its performance is subject to the influence of goals and motivations. When motivated to survive, people memorize items encoded in survival-related situations better than other situations (Nairne, Pandeirada, and Thompson, 2008). In directed forgetting, participants show better recall of the items they are instructed to remember compared to items they are asked to forget (Bjork and Bjork, 1996); however, even though some of the forgetting items cannot be consciously retrieved, they have been encoded and can be probed by implicit memory tests (David and Brown, 2003). Thus, the mind suppresses unneeded information from memory. In hindsight bias, people unintentionally modify their memories of past experience to be consistent with current life goals (Leary, 1981; Pieters, Baumgartner, and Bagozzi, 2006). Such memory modifications or distortions are due to processes of retrieval rather than encoding or storage (Hasher, Attig, and Alba, 1981). These evidences suggest that retrieval of encoded materials respond to the context defined by the specific goals and motivations that influence memory at the time. The dual-retrieval paradigm, therefore, can operationalize Trivers's self-deception by comparing two memories driven by two different motivations.

Experimental operations and the present study

The dual-memory system provides a general approach by which to conduct research on self-deception. The central focus of this approach is to obtain self-serving inconsistencies between two memories, one of which is motivated by deception and the other of which is free from deceptive motives. The present paper reports four studies that all used the same dual-retrieval paradigm. Participants were assigned to one of two conditions: deceiving a high-status target or deceiving an equal-status target. Our hypothesis was that the status differential would make the first group self-deceive the high-status target and would make the second group (openly) deceive the equal-status target. The dual-retrieval paradigm consists of encoding and the first and second retrieval. In the encoding stage, participants learned items for an upcoming vocabulary or knowledge contest. In the first retrieval stage when participants were instructed or induced to deceive either a high-status (e.g., a teacher) or an equal-status target (e.g., a fellow student), they were asked to recognize or recall the items they had previously seen. In the second retrieval, which was conducted purely as a memory test for monetary incentives, participants remembered the items alone without the deceiving target. The first retrieval was motivated by deception. The second retrieval was free from deception. The two retrievals were compared. Self-deception was verified by more distorted memories in the first compared to the second retrieval that was self-serving.

Study 1

In the first two studies, participants were instructed to lie to a high- or equal-status target. This creates constant deception motivation for all participants, eliminating the suspicion that some participants may and some may not try to deceive. As perfect deception, self-deception is unconscious or is conscious of being honest. If memory co-opts by distorting itself to enable such perfect deception to escape detection, participants under the high-status condition were expected to show more memory loss or distortion in the first recognition when they were instructed to lie compared to the second recognition when they were not asked to lie, whereas participants lying to an equal-status target should experience little or no memory difference between the first and second recognition.

Materials and Methods

Participants

Forty-four college students (Mean age = 20.20, *SD* = 1.69; 22 males) from a university in southern China were recruited to take part in an ostensible Chinese vocabulary test for monetary awards.

Design and procedure

The design was a 2 (deceiving target: high- vs. equal-status) × 2 (recognition: first vs. second) mixed design. Participants were randomly assigned into one of two experimental groups: high-status deception group (n = 22), where participants were asked to lie to a teacher, and equal-status deception group (n = 22), where participants were asked to lie to a fellow student. Participants took two recognition tests after an encoding stage where they learned Chinese words. In the first recognition test, participants were instructed to recognize words they had seen and, at the same time, to report the opposite to the high- or equal-status target. Specifically, if participants remembered to have seen or not to have seen a word, they would need to indicate so and, at the same time, need to lie to the target by communicating the opposite of the word recognition result. In the second recognition, participants were asked to recognize the words alone for monetary incentives. Moderate monetary awards were promised and later given to participants for lying to the target.

Encoding. Participants were told to compete in a Chinese vocabulary contest by making a sentence for each word presented to them and, before the test, they had a chance to review practice words similar to those on the real test. A total of 60 Chinese words were presented successively on a computer screen for 3 s each. These are later referred to as studied or old items. Another batch of 60 words, matched in usage frequency, was also prepared. These are referred to as new items. Two old and two new items were used in practice trials prior to the first recognition, and 30 old and 30 new items were used in the first recognition. The other 30 old and 30 new items were used in the second recognition. The order of these four types of items was counterbalanced among participants. After the encoding, participants were asked to count backwards for 30 s.

Target status. Each participant was informed that the practice words he or she had just seen were from the actual test due to a mistake made by the experimenter. The participant was told that he or she must conceal the exposed test words from a detector who was going to check with each participant about the test exposure. The participant was instructed further that, once the detector found out about the concealment, the participant's score on the vocabulary test would be deducted accordingly, thus, affecting the participant's chance to win the monetary prices for the vocabulary test. Thus, the participant must pretend to be as honest as possible while concealing the test words. The participants were also promised (and later paid a constant amount) moderate payment for successfully hiding the test exposure from the detector. Under the high-status deception condition, the detector was introduced as a teacher, who was matched with the participant's gender, looked 10 years older than the participant, and was formally dressed. Under the equal-status condition, the person was introduced as a fellow student, who was of similar age and same gender as the participant, and wore causal wear.

First recognition. A participant was told that, as each word appeared on the screen, he or she must identify, as accurately as possible, whether he or she had seen the word before by pressing one of two particular keys on the keyboard and that the participant must lie to the detector by saying the opposite of his or her true memory. The participant was also made aware of the fact that the detector was only watching the participant's face but was not able to see the participant's hands which were hidden behind the computer. After these instructions which were given to the participant without the detector's presence, the detector showed up and was introduced to the participant either as a teacher or as a fellow student. The participant and the detector were then seated facing each other across a computer table and the participant proceeded with the recognition and deception task.

Second recognition. Each participant was debriefed and asked to identify as accurately as possible the words the participant had seen for monetary incentives.

Results

Participants' performance on the recognition test was represented by three indicators: Hit, False Alarm, and Sensitivity *d'* (see Table 1). Hit is the percentage of old items correctly identified as old. False alarm is the percentage of new items incorrectly identified as old. Sensitivity *d'* represents participants' discrimination between old and new items, with larger numbers representing better discrimination as well as better memory.

Table 1. First and the second recognition results

	Deceive High-Status Target		Deceiving Equal-Status Target	
	1^{st} Recognition While Deceiving	2^{nd} Recognition While Alone	1^{st} Recognition While Deceiving	2^{nd} Recognition While Alone
Hit	.65 (.17)	.77 (.12)	.76 (.16)	.69 (.18)
False Alarm	.35 (.23)	.32 (.14)	.35 (.17)	.33 (.20)
Sensitivity *d'*	.86 (.81)	1.28 (.38)	1.22 (.50)	1.12 (.58)

A 2 (deceiving target: high- vs. equal-status) × 2 (recognition: first vs. second) mixed model ANOVA was conducted separately using each of the three indicators as the dependent variable. For Hit, there was a significant interaction effect, $F(1,42) = 27.06$, $p < .001$, $\eta^2 = .39$. Under the high-status deception condition, participants correctly identified fewer old items the first time while lying ($M = 19.50$, $SD = 5.14$) than they did the second time while being alone ($M = 23.01$, $SD = 3.64$), $t(21) = -4.43$, $p < .001$, whereas participants under the equal-status condition showed a reversed pattern (first $M = 22.67$, $SD = 4.66$; second $M = 20.76$, $SD = 5.52$), $t(21) = 2.82$, $p = .01$. For False Alarm, neither main nor interaction effects were significant. For sensitivity *d'* there was a significant interaction, $F(1,40) = 6.14$, $p = .02$, $\eta^2 = .13$. Participants under the high-status condition showed worse discrimination while lying (M = .86, SD = .81) than they did while being alone (M = 1.28, SD = .38), $t(20) = -2.47$, $p = .02$, whereas participants under the equal-status condition showed similar discrimination at both times (First M = 1.22, SD = .50; Second M = 1.12, SD = .58), $t(20) = .82$, $p = .42$.

Given the different results between Hit and False Alarm, we further conducted a 2 (deceiving target: high- vs. equal-status) × 2 (recognition: first vs. second) × 2 (items: old vs. new) mixed model ANOVA to examine how participants correctly identified old item as old and new items as new. Results showed a significant three-way interaction, $F(1,42) = 6.65$, $p = .01$, $\eta^2 = .14$. Further analysis showed that, for correctly identifying old items, an interaction between target and recognition existed, $F(1,42) = 27.06$, $p < .001$, $\eta^2 = .39$, whereas, for correctly rejecting new items (i.e., 1 – false alarm), no interaction was found, $F(1,42) = .05$, $p = .82$, $\eta^2 = .00$. These results suggest that participants had worse memory when they deceived a high- but not equal-status target. The worse memory was due to incorrectly discriminating old items they had seen but not new items they had not seen.

Discussion

The results support our hypothesis that memory is temporarily worsened when one is deceiving a high- but not equal-status target. The temporary memory impairment driven by fear for detection thus helps to execute self-deception or perfect deception without detection. Fear for detection is correlated with the social status of the deceived relative to the deceiver (Cummins, 1999; Manktelow and Over, 1991) among other social conditions contributing to the probabilities of deception detection. We further speculate that self-deception may be automated by the mere presence of such social conditions as the social status of the deceived. Such social conditions have been so frequently tied to deception detection and retribution in the evolutionary past that no conscious calculation of probabilities of detection and punishment may be necessary or ever carried out. Participants under the high-status deception condition failed to recognize items they had actually seen. We speculate that the fear of detection was so strong and was so closely tied to social status that, at the sight of a teacher, a student participant would at times automatically engage in correcting the deception that he or she was instructed to commit by misremembering an item. We speculate that, within a "forced to lie" context such as the present experiment, fear for detection potentially imbedded in a recurrent hierarchical social relationship (e.g., student-teacher relationship) would facilitate self-deception through impaired memory without having to experimentally induce deceiving motivations.

Our speculation of a selected association between fear for detection and self-deception is in part also supported by the finding that, when lying to the high-status target, participants only mis-remembered by not remembering items they had seen (i.e., reduced hit rates), but did not mis-remember by remembering items they had not seen (i.e., no difference in false alarm). The first but not the second type of mis-remembering matches with one common form of hierarchical relationship among primates, including humans: Low-status individuals hide from the high-status peers resources they have in possession or mistakes they have made. In one form or another, this type of hierarchical social interaction should exert selection pressure to favor hiding and covering strategies among the low-status persons and finding and detecting strategies among the high-status persons. The subsequent arms race (Trivers, 1976) is likely to be between status hierarchies, leading to the low-status individuals self-deceiving high-status individuals by not remembering what they have rather than fabricating what they do not have. The high-status individuals, on the other hand, may be more likely to deceive by fabricating what they do not have or do not have enough of (e.g., abilities and resources or good will and altruism). More empirical work is needed to test these speculations. There are also alternative explanations, especially given our fallible design, which we discuss in the end.

Study 2

An alternative interpretation of the results of Study 1 may be that participants' memories became worse simply by interacting with a high-status target. That is, even without being motivated or instructed to deceive the target, memory would become worse when interacting with a high- compared to low-status person. This is particularly true given the unequal teacher-student relationship. We addressed this potential problem in Study 2 by asking participants to lie about some questions, which were test questions in Study 2, and to tell the truth about other questions, which were called practice questions. If an individual's memory became worse simply by interacting with a high-status person but not by trying to deceive the person, the participant's memory on both kinds of material would be equally worse when compared to interacting with an equal-status person. If memory impairment occurred only on the test but not practice questions, then we had additional confidence beyond that of Study 1 to support the hypothesis that people may use memory to execute self-deception especially when low-status individuals interact with high-status individuals. We also included exit interview questions to measure how worried and afraid participants felt about being caught lying to the target to see if there were differences in these emotions responding to the two kinds of targets. The design and procedure of the experiment in Study 2 were identical to those of Study 1, but the encoding material was different and was distinguished between test and practice questions.

Materials and Methods

Participants

Fifty-two college students (Mean age = 19.88, SD = 1.22; 30 males) from a university in southern China took part in an ostensible knowledge contest for monetary awards.

Procedure and design

Participants were randomly assigned into one of two experimental groups: high-status interaction group (n = 28), where the detector was a teacher, and equal-status interaction group (n = 24), where the detector was a student. Participants were told to participate in a history and geography knowledge contest. They were first given certain practice questions to study (encoding stage) and were then asked to lie about seeing some of these questions and to tell the truth about seeing other questions to a detector who was either a teacher or a fellow student. While lying or telling the truth, participants were at the same time asked to indicate the true memory about seeing each question (first recognition). After debriefing, participants were asked to remember the questions they had seen for monetary incentives (second recognition). The design was a 2 (interacting target: high- vs. equal-status) × 2 (recognition: first vs. second) × 2 (questions: lie vs. truth) mixed ANOVA design.

Encoding. Participants were given 20 history questions (e.g., “In which dynasty was papermaking technology developed?”) and 20 geography questions (e.g., “Which is the longest river in the world?”). By random assignment, half of the participants were told that the 20 history questions were practice questions which they were supposed to study and that the 20 geography questions were actual test questions which were given to them by mistake and they were not supposed to have read. For the other participants, the order of these two types of questions was reversed. These 40 questions were randomly presented one by one on a computer in a speed of 5 s per question. Each question was labeled either as “history practice question,” “geography test question,” “geography practice question,” or “history test question.” After encoding, participants were asked to count backwards for 30 s.

Two parallel sets of 10 history and 10 geography questions that participants had not read were prepared as new items to be included in each of the two subsequent recognition tests. Thus, 10 history and 10 geography questions presented in the encoding stage were used as old items and 10 history and 10 geography questions not presented were used as new items in the first recognition, and the same combination of the other 10 old and 10 new questions was used in the second recognition. The assignment of 40 old or new items was counterbalanced among participants.

First recognition. Similar to Study 1, a participant was seated behind a computer table across from a person who was introduced either as a teacher (high status) or as a fellow student (equal status). Not in the presence of the detector, a participant was told that the detector came to check which questions the participant had read, and that the participant should lie about seeing test questions. Specifically, if a participant had previously seen a test question (old item), the participant should say to the detector that he/she had not and, if the participant had not seen a test question (new item), the participant should say he/she had seen the question. As each question appeared on the screen, the detector would ask, for example, “Did you read geography (history) question: which is the longest river in the world? (“in which dynasty was papermaking technique developed?”). A participant should identify, as accurately as possible, whether he/she had previously seen the question by pressing one of two particular keys on the keyboard, which was hidden from the detector. As the participant keyed in the truth about each question privately, he/she either lied to the person about what he/she remembered as a test question or told the truth about what he/she remembered as a practice question.

Second recognition. Each participant was then debriefed and asked to identify as accurately as possible the words he/she had seen for monetary incentives.

In the exit interview, a participant was asked three questions: “Were you worried about being caught?” “Did you feel nervous while lying?” and “What is the probability, from 0 to 100%, that you might be caught by the detector?”

Results and Discussion

Participants reported to have felt more nervous while lying to the high-status target (M = 2.96, SD = .99) than lying to the equal-status target (M = 2.25, SD = 1.15), $t(50)$ = 2.40, p = .02, and the participants thought that the chance of being caught by the high-status detector was higher (M = .56, SD = .19) than the equal-status detector (M = .40, SD = .22), $t(50)$ = 2.73, p = .009. Participants showed no difference in worrying about being caught by the high-status (M = 2.71, SD = 1.15) or equal-status detector (M = 2.29, SD = 1.52), $t(50)$ = 1.14, p = .26.

Participants' performance on the recognition test is shown in Table 2. A 2 (interaction target: high- vs. equal-status) × 2 (recognition: first vs. second) × 2 (question: test vs. practice) mixed model ANOVA was conducted separately using each of the three memory indicators as the dependent variable. For Hit, there was a main effect of recognition, $F(1,50) = 11.14$, $p = .002$, $\eta^2 = .18$, with Hit in the first recognition (M = .86, SD = .12) being lower than that in the second (M = .91, SD = .11), and there was an interaction between recognition and interaction target, $F(1,50) = 7.48$, $p = .009$, $\eta^2 = .13$, an interaction between recognition and question type, $F(1,50) = 5.89$, $p = .019$, $\eta^2 = .11$, an interaction between interaction target and question type, $F(1,50) = 7.21$, $p = .010$, $\eta^2 = .13$, and a three-way interaction, $F(1,50) = 11.36$, $p = .001$, $\eta^2 = .19$. For test questions that participants should lie about having read, a 2 (target) × 2 (recognition) mixed model ANOVA showed a main effect of recognition, $F(1,50) = 16.47$, $p < .001$, $\eta^2 = .25$, and the interaction, $F(1,50) = 21.14$, $p < .001$, $\eta^2 = .30$. Under the high-status target condition, Hit in the first recognition was lower than that in the second recognition, $t(27) = -7.30$, $p <$ .001, whereas under the equal-status target condition, Hit in the two recognitions was similar, $t(23) = .40$, $p = .70$. For practice questions that participants told the truth about reading, a 2 (target) × 2 (recognition) mixed model ANOVA showed no significant main effect or interaction effect.

Table 2. First and the second recognition results of test questions and practice questions in Study 2

	Deceive High-Status Target		Deceiving Equal-Status Target	
	1st Recognition While Deceiving	2nd Recognition While Alone	1st Recognition While Deceiving	2nd Recognition While Alone
Test Questions				
Hit	.79 (.12)	.95 (.08)	.90 (.09)	.89 (.10)
False Alarm	.13 (.12)	.12 (.11)	.10 (.11)	.12 (.11)
Sensitivity d'	2.01 (.65)	2.66 (.56)	2.55 (.52)	2.43 (.63)
Practice Questions				
Hit	.91 (.10)	.92 (.11)	.87 (.12)	.90 (.14)
False Alarm	.14 (.10)	.13 (.11)	.11 (.13)	.15 (.11)
Sensitivity d'	2.44 (.58)	2.49 (.57)	2.37 (.65)	2.28 (.65)

For False Alarm, a 2 (interaction) × 2 (recognition) × 2 (questions) mixed model ANOVA showed neither main effects nor interaction effects. For sensitivity *d'*, there was an interaction between recognition and interaction target, $F(1,50) = 7.04$, $p = .011$, $\eta^2 = .12$, an interaction between question type and interaction target, $F(1,50) = 4.89$, $p = .032$, $\eta^2 = .09$, and a three-way interaction, $F(1,50) = 4.59$, $p = .037$, $\eta^2 = .08$. For test questions that participants should lie about having read, a 2 (target) × 2 (recognition) mixed model ANOVA showed a main effect of recognition, $F(1,50) = 4.77$, $p = .034$, $\eta^2 = .09$, and an interaction, $F(1,50) = 9.90$, $p = .003$, $\eta^2 = .17$. Under the high-status target condition, sensitivity *d'* in the first recognition was lower than that in the second recognition, $t(27) = -3.81$, $p = .001$, whereas under the equal-status target condition, sensitivity *d'* in the first recognition was similar to that in the second recognition, $t(23) = .68$, $p = .50$. For practice questions that participants told the truth about reading, a 2 (target) × 2 (recognition) mixed model ANOVA showed neither main effects nor an interaction effect.

These results suggest that participants had worse memory on material about which they had to lie to a high- compared to equal-status target, whereas for material about which participants did not have to lie, their memory did not differ when interacting with a high- or an equal-status target. The worsened memory shown only when lying to a high-status target was due to the motivation of deception rather than omnibus emotions, because memory for honestly presented material was unimpaired. Such temporarily worsened memory helps low-status individuals execute self-deception when lying to high-status others to whom low-status individuals dare not lie for fear of detection.

Study 3

In the previous two studies we explicitly instructed participants to lie in order to ensure the motivation to deceive, because otherwise a deceptive motivation is, by definition, consciously inaccessible to individuals engaging in self-deception. However, it is more natural for the participants to choose to deceive on their own rather than being told to do so. In the present study, we created the situation such that it was inherently advantageous for the participants to lie. The procedures, material, and instructions were identical to those of Study 2 except that, instead of being instructed to lie to the detector as in Study 2, each participant was told that a detector would come to check on what questions the participant had seen and that certain marks would be deducted from his/her contest score in the end for each test question the participant reported having seen. Like Study 2, because of the experimenter's mistake, participants studied certain test questions that they were not supposed to see. These test questions would help participants' performance on the contest. Thus, participants were motivated to hide these test questions from the detector. We also increased the monetary rewards for winning different prizes of the knowledge contest to increase participants' motivation to deceive. Like in Study 2, participants were asked in the exit interview whether they lied to the detector and how afraid they were to lie to the detector.

Materials and Methods

Participants and designs

Forty-eight college students (Mean age = 20.27, *SD* = 1.31; 25 males) from a university in southern China were recruited to take part in an ostensible contest of general knowledge for monetary awards. They were randomly assigned into one of two experimental groups: high-status interaction (n = 24), where the test detector was a teacher, and equal-status interaction (n = 24), where the detector was a student. The design was a 2 (interaction target: high- vs. equal-status) × 2 (recognition: first vs. second) mixed design.

Results and Discussion

Participants' performance on the two recognition tests is shown in Table 3. A 2 (interaction target: high- vs. equal-status) × 2 (recognition: first vs. second) mixed model ANOVA was conducted separately using each of the three memory indicators as the dependent variable. For Hit, there was a main effect of recognition, $F(1,46) = 13.59$, $p = .001$, $\eta^2 = .23$, with Hit in the first recognition ($M = .80$, $SD = .14$) being lower than that in the second recognition ($M = .87$, $SD = .11$), and there was an interaction between recognition and interaction target, $F(1,46) = 10.31$, $p = .002$, $\eta^2 = .18$. Under the high-status target condition, Hit in the first recognition was lower than in the second recognition, $t(23) = -4.55$, $p < .001$, whereas under the equal-status target condition, there was no difference in Hit between the two recognitions, $t(23) = .37$, $p = .72$.

Table 3. First and the Second Recognition Results in Study 3

	Deceive High-Status Target		Deceiving Equal-Status Target	
	1st Recognition While Deceiving	2nd Recognition While Alone	1st Recognition While Deceiving	2nd Recognition While Alone
Hit	.75 (.16)	.87 (.10)	.86 (.10)	.87 (.11)
False Alarm	.08 (.07)	.10 (.07)	.07 (.05)	.07 (.07)
Sensitivity *d'*	2.12 (.65)	2.56 (.54)	2.60 (.51)	2.64 (.59)

For False Alarm, a 2 (interaction) × 2 (recognition) mixed model ANOVA showed neither main effects nor an interaction. For sensitivity *d'*, there was a main effect of recognition, $F(1,46) = 8.25$, $p = .006$, $\eta^2 = .15$, with *d*' of the first recognition ($M = 2.36$, $SD = .58$) being lower that that of the second ($M = 2.60$, $SD = .57$), and there was an interaction between recognition and interaction target, $F(1,46) = 5.70$, $p = .021$, $\eta^2 = .11$. Under the high-status target condition, sensitivity *d*' in the first recognition was lower than that in the second recognition, $t(23) = -4.02$, $p = .001$, whereas under the equal-status target condition, *d*' in the first recognition was not different than that of the second recognition $t(23) = .32$, $p = .75$.

The exit interview question showed that, among participants under the high-status condition, 17 reported no deception and 7 reported deception, whereas the corresponding numbers were 10 and 14 for participants under the equal-status condition, $\chi^2 = 4.15$, $p = .042$. Participants were more nervous while lying to the high-status target ($M = 2.92$, $SD = 1.10$) than lying to the equal-status target ($M = 2.17$, $SD = 1.13$), $t(46) = 2.33$, $p = .02$, and they thought the chance of being caught by the high-status target was higher ($M = .55$, $SD = .25$) than that with the equal-status target ($M = .41$, $SD = .22$), $t(46) = 2.14$, $p = .04$. Participants were more worried about being caught by the high-status target ($M = 2.99$, $SD = 1.38$) compared to the equal-status target ($M = 2.13$, $SD = 1.15$), $t\ (46) = 2.38$ $p = .02$.

These results, derived from a paradigm where participants had the motivation to deceive on their own, are consistent with those of Studies 1 and 2 where participants were instructed to deceive. Together, these results support the hypothesis of temporary memory impairment when voluntarily lying or being told to lie, but not when telling the truth, to a high- compared to equal-status target. Through temporarily worsened memory, individuals deceived themselves to better deceive high- compared to equal-status others.

General Discussion

This study is among the first to empirically test Trivers's theory of self-deception (Trivers, 1985, 2000). It also provides a methodological framework within which to conduct research on self-deception in general. This framework is based on memory research and approaches self-deception as inconsistencies between two memory tasks. This approach demonstrates how self-deception is adaptive when one self-deceives by being unaware of fitness enhancing information that is meant only to be kept away from others (Leeuwen, 2007). We show that self-deception is adaptive if true information is kept from the self only during deception and is accessible to the self after the goal of deception has been achieved. Memory and its temporary loss or distortion aid this mental state of self-deception. True information is not retrieved when the self-deceiver deceives others. When deception ceases, the lost memory is recovered so that the self-deceiver benefits from accurate information. By suppressing truthful information (i.e., concealing to deceive) or by distorting encoded material (i.e., fabricating to deceive), memory helps one to "honestly" offer null or false information to others. This aspect of self-deception does not derive net fitness gains. By later retrieving the true information, memory helps to complete self-deception to achieve net fitness. Memory is a direct target of selection (Nairne, Pandeirada, and Thompson, 2008). It also may be co-opted to perform self-deception by switching between the two states of consciousness. That almost all deceptive ploys use memory as an honest or dishonest excuse may not be a coincidence.

Although we suggest that memory or the decoding stage of information processing executes self-deception, encoding may also achieve self-deceptive goals. Motivational biases in information searching and screening influence attention, perception, and interpretation of incoming materials (Balcetis, 2008; Mele, 1997). These encoding biases of information processing contribute to a type of self-deception widely studied as an intrapersonally-oriented information manipulation process that achieves mental wellbeing. Well-known examples of intrapersonal self-deception include moderate optimism (Taylor and Brown, 1988), self-services (Greenwald, 1988), and competency and morality self-enhancement (Paulhus and John, 1998). In addition to being involved in intrapersonal self-deception, encoding may also be involved in interpersonal self-deception. People may differentially take in and decipher information with a motivation to deceive others. Through biased encoding, the self-deceiver may somehow know in advance that the information so encoded will enhance fitness when retrieved to misinform others. Self-deception may be achieved in the encoding stage where information is already distorted before storage (Von Hippel and Trivers, 2011). This is in contrast to our decoding approach, which assumes distorting information when being recalled. However, to be adaptive, biased encoding has to enhance fitness all of the time or the biased encoder must re-encode the information to stop self-deceiving after others have been deceived. Empirical work is needed to test the encoding mechanism in interpersonal self-deception.

The present study also extends Trivers' theory by showing that the arms race between deception and deception detection is more likely to have been played out with low-status people as deceivers and self-deceivers and high-status people as detectors rather than the other way around. In the hierarchical groups of primates and ancestral (as well as modern) humans, social status provides protection from punishment and retribution when conflicts arise among individuals or between individuals and the group (Wright, 1994). When caught deceiving in a hierarchical group setting, high-status individuals are less likely to be punished than low-status individuals who often face more severe consequences. Status differentials should pressure low-status individuals to be careful and fearful deceivers, whereas high-status deceivers should become unscrupulous and bold. Research on humans and other primates confirm this speculation. An alpha chimpanzee can freely exaggerate (i.e., deceive) his display of dominance and aggressiveness without being challenged (Boehm, 1999). This is deemed similar to human leaders bragging about their efforts or abilities. High-status humans are calmer and show fewer changes in heart rate and skin conductance than their low-status counterparts when deceiving (Carney, 2010). Low-status deceivers are also more likely to be caught than high-status deceivers who have more resources, including information leading to detection. Facing a higher likelihood of detection with a higher certainty of punishment, low-status deceivers yet have a higher need to deceive than high-status individuals because, being low in status, they have fewer resources and fewer honest means (e.g., either through force or alliance) to acquire them. High-status individuals who have more resources to lose, on the other hand, have a greater need to detect deception. The selection pressures to evade detection and detect deception affect low-status and high-status individuals more than their counterparts, respectively. Our empirical findings imply that self-deception, especially when carried out by low-status individuals, may have been selected to better deceive high- rather than equal-status others.

Although this paper emphasizes status differences, we do not deny the possibility that the arms race between deception and detection may have been advanced by the intrapersonal experience of being both a deceiver and deceit detector (von Hippel and Trivers, 2011). Having deceived makes one aware of the need to detect possible deception; in turn, the detecting experience makes one cautious not to be detected. Such intrapersonal experience helps to advance both ends of the arms race. It is also apparent that all humans and many other animals have the potential to cheat and an equal urge to eliminate cheating, independent of social ranking. There is clearly no division of labor between deceiving and detecting. However, like any adaptation, environmental conditions shift selection power to favor certain individuals over others as beneficiaries of the adaptation. We speculate that the hierarchical context of primate groups (including humans) should favor high-status members as detectors and low-status members as deceivers more than the other way around. Both intrapersonal and interpersonal experience advances the arms race between deceiving and detecting and should lead, as Trivers (2000) first theorized, to self-deception. Our findings support this speculation.

Furthermore, our findings suggest that social status is so closely related to getting caught and facing consequences that it may motivate self-deception without conscious calculation of the odds of detection. Study 1 showed that when participants were instructed to lie free from punishment, they still showed comparatively worse first-round memory when lying to high-status compared to low-status targets. However, social status is not the direct cause of self-deception; rather, it is a conditional variable that regulates the likelihood of detection and punishment, and the latter motivates self-deception to avoid detection. The number of targets to be deceived also moderates the probability of detection, with a higher number increasing the likelihood of detection (Lu and Chang, 2011a, 2011b). In other words, escaping detection from one person is easier than from a group because each group member contributes to detective pressure and raises the likelihood of detection. In this context, self-deception is adaptive because by deceiving one target (i.e., the self) the person successfully deceives multiple targets (i.e., the group). Like social status, the number of deceived may moderate self-deception such that it is more likely to occur when one has to deceive a group of individuals rather than a single individual. We speculate that, especially when high-status people engage in self-deception, their targets are more likely to be groups than individuals. Religious and political leaders, in relation to their congregations and constituencies, are good examples of self-deception serving to deceive groups. Intrapersonal self-deception, such as self-enhancement, may represent a mental state that is under constant pressure to group-deceive everyone around the deceiver (i.e., the general public). Other moderators may include perceived determination of the target to detect and avenge. Based on our findings regarding social status, we expect all of these social conditions to enhance the likelihood of self-deception. That is, self-deception is more likely to happen when low-status people try to deceive high-status people, when a single individual tries to deceive a group, or when the deceiver perceives a clear determination from the target to detect and avenge.

There are several limitations of this study, many of which are imbedded in the operationalization of the unconscious process of self-deception. In deception research, a lie can be verified either by post-experiment interviews or by readings of the participants' heart rate and skin conductance. Because self-deception is partially unconscious, it cannot be verified even physiologically, although some researchers believe that skin conductance may show signs contradicting the conscious belief (Gur and Sackeim, 1979). Self-deception is inferred by establishing the motivation to deceive and by obtaining inconsistencies between two memories which must also be distinguished from blatant deception. In Study 1 and 2, we instructed participants to lie. This ensures the motivation to deceive. In Study 1, 2, and 3, we asked participants to remember which items they had seen before reporting them to the target. This separates out the direct deceivers who would remember and report different items. However, these design features reduce the ecological validity of our studies because in real deception, a deceiver or a self-deceiver only lies or reports to the target without first recording his/her memory.

We used factual items in all studies to examine how memory executes self-deception. This kind of memory offers insufficient room for distortion. Memory for episodes or personal experience may provide more room for distortion to better facilitate self-deception. Future studies are also needed to test self-deception on participants from other non-Chinese cultural backgrounds. Even though the issue under discussion represents pan-cultural evolutionary predictions, culture and the social settings from which culture evolved may have particularly relevant bearing on social interactive strategies. More deeply rooted in social compared to individual learning (Chang et al., 2011), Asians who also seem to pay special attention to social hierarchy may be particularly prone to self-deception compared to straight deception. More replications using more diverse cultural groups will help to confirm and illuminate evolutionary predictions.

Finally, whereas self-deception is originally construed mainly as an interpersonal strategy to facilitate the deception of others (Trivers, 1976, 1985), it has also been studied as an intrapersonal variable that represents chronic mis-representation or selective representation of the self without immediate reference to others (Kurzban and Aktipis, 2007; Lu and Chang, 2011b; Paulhus and Reid, 1991; von Hippel and Trivers, 2011). The present study focuses on the interpersonal process of self-deception. Future research may examine both interpersonal and intrapersonal self-deception together to better understand this social adaptation. Despite these limitations, this is among the first studies to test an important evolutionary theory about self-deception. Methodologically, this study demonstrates the use of dual-retrieval as a useful paradigm within which to conduct future studies on self-deception. Theoretically, the study demonstrates how the memory system may have been co-opted to execute self-deception, and how social hierarchy may have shifted selection power to favor low- over high-status individuals as deceivers and self-deceivers.

Received 27 November 2013; Revision submitted 23 April 2014; Accepted 25 April 2014

References

Alexander, R. D. (1987). *The biology of moral systems*. New York: Aldine.

Amici, F., Call, J., and Aureli, F. (2009). Variation in withholding of information in three monkey species. *Proceedings of the Royal Society B: Biological Sciences, 276*, 3311-3318.

Balcetis, E. (2008). Where the motivation resides and self-deception hides: How motivated cognition accomplishes self-deception. *Social and Personality Psychology Compass, 2*, 361-381.

Bjork, E. L., and Bjork, R. A. (1996). Continuing influences of to-be-forgotten information. *Consciousness and Cognition, 5*, 176-196.

Boehm, C. (1999). *Hierarchy in the forest: The evolution of egalitarian behavior*. Cambridge, MA: Harvard University Press.

Bugnyar, T., and Kotrschal, K. (2004). Leading a conspecific away from food in ravens (*Corvus corax*)? *Animal Cognition, 7*, 69-76.

Byrne, R. W., and Whiten, A. (1992). Cognitive evolution in primates: Evidence from tactical deception. *Man, New Series, 27*, 609-627.

Carney, D. (2010). Powerful people are better liars. *Harvard Business Review, 88*, 32-33.

Chang, L., Mak, M. C. K., Li, T., Wu, B. P., Lu, H. J., and Chen, B.B. (2011). Cultural adaptations to environmental variability: An evolutionary account of East-West differences. *Educational Psychology Review, 23*, 99-129.

Clutton-Brock, T. H., and Parker, G. A. (1995). Punishment in animal societies. *Nature, 373*, 209-216.

Cummins, D. D. (1999). Cheater detection is modified by social rank: The impact of dominance on the evolution of cognitive functions. *Evolution and Human Behavior, 20*, 229-248.

Cummins, D. D. (2005). Dominance, status, and social hierarchies. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 676-697). Hoboken, NJ: John Wiley & Sons.

David, D., and Brown, R. (2003). The impact of different directed forgetting instructions on implicit and explicit memory: New evidence from a modified process dissociation procedure. *Quarterly Journal of Experimental Psychology: Section A, 56*, 211-231.

Ducoing, A., and Thierry, B. (2003). Withholding information in semifree-ranging Tonkean macaques (*Macaca tonkeana*). *Journal of Comparative Psychology, 117*, 67-74.

Ellis, L. (1995). Dominance and reproductive success among nonhuman animals: A cross-species comparison. *Ethology and Sociobiology*, *16*, 257-333.

Greenwald, A. G. (1988). Self-knowledge and self-deception. In J. S. Lockard and D. L. Paulhus (Eds.), *Self-deception: An adaptive mechanism?* (pp. 113-131). Englewood Cliffs, NJ: Prentice-Hall.

Gur, R., and Sackeim, H. A. (1979). Self-deception: A concept in search of a phenomenon. *Journal of Personality and Social Psychology, 37*, 147-169.

Hakkarainen, H., Korpimäki, E., Huhta, E., and Palokangas, P. (1993). Delayed maturation in plumage colour: Evidence for the female-mimicry hypothesis in the kestrel. *Behavioral Ecology and Sociobiology, 33*, 247-251.

Hare, B., Call, J., Agnetta, B., and Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour*, *59*, 771-785.

Hare, B., Call, J., and Tomasello, M. (2006). Chimpanzees deceive a human competitor by hiding. *Cognition, 101*, 495-514.

Harris, L., and Ogbonna, E. (2010). Hiding customer complaints: Studying the motivations and forms of service employees' complaint concealment behaviours. *British Journal of Management, 21*, 262-279.

Hasher, L., Attig, M. S., and Alba, J. W. (1981). I knew it all along: Or did I? *Journal of Verbal Learning and Verbal Behavior, 20*, 86-96.

Held, S., Mendl, M., Devereux, C., and Byrne, R. W. (2002). Foraging pigs alter their behaviour in response to exploitation. *Animal Behaviour, 64*, 157-166.

Kirkman, B. L., and Rosen, B. (1999). Beyond self-management: Antecedents and consequences of team empowerment. *Academy of Management Journal, 42,* 58-74.

Kurzban, R., and Aktipis, C. A. (2007). Modularity and the social mind. *Personality and Social Psychology Review, 11*, 131-149.

Leary, M. R. (1981). The distorted nature of hindsight. *The Journal of Social Psychology, 115*, 25-29.

Leeuwen, D. S. N. V. (2007). The spandrels of self-deception: Prospects for a biological theory of a mental phenomenon. *Philosophical Psychology, 20*, 329-348.

Lu, H. J., and Chang, L. (2011a). Self-deceive to countermine detection. *Behavioral and Brain Sciences*, *34*, 33.

Lu, H. J., and Chang, L. (2011b). The association between self-deception and moral self-concept as functions of self-consciousness. *Personality and Individual Differences, 51*, 845-849.

Manktelow, K. I., and Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition, 39*, 85-105.

Manz, C. C., and Sims, H. P., Jr. (1993). *Business without bosses: How self-managing teams are building high-performance companies.* New York: Wiley.

Mele, A. R. (1997). Real self-deception. *Behavioral and Brain Sciences, 20*, 91-102.

Mitchell, R. W. (1999). Deception and concealment as strategic script violation in great apes and humans. In S. T. Parker and R. W. Mitchell (Eds.), *The mentalities of gorillas and orangutans: Comparative perspectives* (pp. 295-315). New York: Cambridge University Press.

Nairne, J. S., Pandeirada, J. N. S., and Thompson, S. R. (2008). Adaptive memory: The comparative value of survival processing. *Psychological Science, 19*, 176-180.

Nettle, D. (2003). Intelligence and class mobility in the British population. *British Journal of Psychology, 94*, 551-561.

Paulhus, D. L., and John, O. P. (1998). Egoistic and moralistic biases in self-perception: The interplay of self-deceptive styles with basic traits and motives. *Journal of Personality, 66*, 1025-1060.

Paulhus, D. L., and Reid, D. B. (1991). Enhancement and denial in socially desirable responding. *Journal of Personality and Social Psychology, 60*, 307-317.

Payne, H. J. (2008). Targets, strategies, and topics of deception among part-time workers. *Employee Relations*, *30,* 251-263.

Pieters, R., Baumgartner, H., and Bagozzi, R. (2006). Biased memory for prior decision making: Evidence from a longitudinal field study. *Organizational Behavior and Human Decision Processes, 99*, 34-48.

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

Trivers, R. (1976). Preface. In R. Dawkins, *The selfish gene*. Oxford: Oxford University Press.

Trivers, R. (1985). *Social evolution*. Menlo Park, CA: Benjamin/Cummings.

Trivers, R. L. (2000). The elements of a scientific theory of self-deception. *Annals of the New York Academy of Sciences, 907*, 114-131.

Von Hippel, W., and Trivers, R. (2011). The evolution and psychology of self-deception.

Behavioral and Brain Sciences, 34, 1-16.

Whiten, A., and Byrne, R.W. (1988). The manipulation of attention in primate tactical deception. In R. W. Byrne and A. Whiten (Eds.), *Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes, and humans* (pp. 211-224). New York: Oxford University Press.

Wintrobe, R., and Breton, A. (1986). Organizational structure and productivity. *The American Economic*, *76*, 530-538.

Wright, R. (1994). *The moral animal: Evolutionary psychology and everyday life*. New York: Pantheon Books.