Some Evidence for Unconscious Lie Detection

Journal:	*Psychological Science*
Manuscript ID:	PSCI-13-1054.R2
Manuscript Type:	Research article
Date Submitted by the Author:	n/a
Complete List of Authors:	ten Brinke, Leanne; University of California Berkeley, Haas School of Business Stimson, Dayna; University of California Berkeley, Carney, Dana; University of California Berkeley,
Keywords:	Social Cognition, Decision Making, Consciousness

Some Evidence for Unconscious Lie Detection

Leanne ten Brinke

Dayna Stimson

Dana R. Carney

University of California, Berkeley

Abstract

To maximize survival and reproductive success, alongside the tendency to tell lies, primates evolved the ability to accurately detect them. Despite the obvious advantage of detecting lies accurately, conscious judgments of veracity are only slightly more accurate than chance. However, findings in forensic psychology, neuroscience and primatology suggest that lies can be accurately detected when less conscious mental processes are utilized. We predicted that observing someone tell a lie would automatically activate cognitive concepts associated with deception and observing someone tell the truth would activate concepts associated with truth. Across two experiments results demonstrated that indirect measures of detecting deception are significantly more accurate than direct measures. Findings provide a new lens through which to reconsider old and approach new investigations of human lie detection. Further, we make available videotaped stimuli of 6 liars and 6 truth-tellers, and data detailing their nonverbal behaviors, affective self-reports, and cortisol reactivity.

Some Evidence for Unconscious Lie Detection

Human beings lie all day long and about all manner of things. Sometimes lies are small, prosocial, and without negative consequence. Other times, lies destroy precious, hard-earned value in personal, professional, and civic life (Ekman, 1992). Because deception is ubiquitous, our livelihood can depend on the ability to accurately detect it. However, when asked to make a simple decision about whether a person is lying or telling the truth, humans perform poorly. Individual studies consistently find human judgments of veracity to be no more accurate than the flip of a coin (Ekman & O'Sullivan, 1991; Porter, Woodworth, & Birt, 2000), and while a recent meta-analysis found average accuracy to be statistically greater than chance, that percentage was only 54% (Bond & DePaulo, 2006).

This general deception-detection incompetence is inconsistent with evolutionary theory, which suggests that the accurate detection of deception is critical to our survival (Krebs & Dawkins, 1974). Evolutionary theory suggests that for survival and reproduction, the ability to accurately detect deception must have evolved alongside the tendency to lie in a co-evolutionary arms race—after all, the acquisition of survival-related resources and attraction of quality mates may be enhanced by both successful deception and keen detection of the same (Bond, Kahler, & Paolicelli, 1985). In this dance of skill, deceivers are ever adapting to avoid detection, while targets of deception follow close behind in their counter-deception strategies, allowing only a few costly lies to evade detection before they become wise to the deceiver's new strategy (Dawkins & Krebs, 1978; von Hippel & Trivers, 2011). As the effectiveness of a deceiver's strategy decreases, as a result of increasing observer accuracy, the liar must adopt a new strategy to outwit his victim—and the cycle begins anew. Thus, while liars possess the basic architecture to continually transform those signals that betray their lies, lie detectors must have some mental architecture in place to support the commensurate shift in signal-detection. Yet, studies of explicit lie detection accuracy fall short of supporting this notion; even at 54% accuracy, this result provides little protection from manipulation by others, especially given that this above-chance finding is driven by the accurate detection of truths ($M_{accuracy}$ = 61%), not lies ($M_{accuracy}$ 47%; Bond & DePaulo, 2006). von Hippel and Trivers (2011) point to several shortcomings in the research to date which may systematically underestimate lie detection ability, and we offer one more—that the mental architecture promoting this skill may be unconscious, not captured by explicit assessments (Reinhard, Greifeneder, & Scharmach, 2013).

While humans are poor lie detectors, evidence from primatology and neuroscience suggests that without conscious awareness, parts of the human brain can automatically detect deception—and so can nonhuman primates (e.g., Grezes et al., 2004; Wheeler, 2010). Taking these findings together, we predicted that that indirect measures of deception detection—measures capable of accessing less conscious parts of the mind—would yield more deception detection accuracy than direct/conscious measures.

Attempts to Explain Deception Detection Incompetence

Theories from social, forensic, and evolutionary psychology have attempted to explain consistently poor deception detection accuracy; however, most accounts fall short in explaining the full suite of findings. For example, some blame the lack of accuracy on the "absence of a Pinocchio's nose" because deceptive behavior is subtle and variable across time and person (e.g. DePaulo et al., 2003; Hartwig & Bond, 2011). Instead, liars (versus truth-tellers) emit a complex array of nonverbal cues, and research suggests that—even in the presence of many deception cues—perceivers have inaccurate beliefs about which to rely on (The Global Deception Research Team, 2006). Further, deceivers do not necessarily feel or behave in line with predominant stereotypes; for example, the commonly held expectation that liars avert their gaze and fidget is false (Bond & DePaulo, 2006). Other theoretical frameworks point to the fact that humans often live in conditions of such abundance and safety that we lack the motivation and suspicion necessary to detect deception (Vrij, Granhag, & Porter, 2010), consistent with research suggesting that under conditions of environmental scarcity, deception detection accuracy increases (Carney et al., 2013).

Deception *Can* be Accurately Detected: Evidence from Primatology and Neuroscience

The evolutionary argument that humans should be accurate at detecting deception finds some traction in primate work, which suggests that *non-human* primates can both successfully produce and detect lies (Byrne & Corp, 2004; Menzel, 1974). Jane Goodall (1986) and others have documented sophisticated and accurate deception detection by chimpanzees, allowing them to find (and subsequently steal) food hidden by a dishonest counterpart. Similarly, capuchin monkeys accurately detect deception, choosing to ignore false 'alarm calls' aimed at luring feeding monkeys away from their meal (Menzel, 1974; Wheeler, 2010). Accurate deception detection to acquire (or maintain) access to resources is precisely the mechanism thought to promote this ability in humans, too (Krebs & Dawkins, 1984).

Recent brain imaging work suggests that three brain regions are activated when deceptive acts are correctly detected: the orbitofrontal cortex (involved in understanding others' mental states), anterior cingulate cortex (associated with monitoring inconsistencies), and the amygdala (associated with threat detection) (Grezes et al., 2004; Lissek et al., 2008). Abnormal functioning in these regions is associated with deficits in basic social cognition more generally and impaired deception detection accuracy (e.g., among autistics; Sodian & Firth, 1992). In contrast, aphasics—individuals sustaining left cerebral hemisphere damage (particularly the left orbitofrontal cortex), who cannot comprehend spoken sentences and must rely on nonverbal cues to detect deception—are *more* accurate than healthy observers (Etcoff, Ekman, Magee, & Frank, 2000). Together, these findings suggest that when conscious thought is impaired or stripped away, deception detection accuracy is enhanced. In fact, Albrechtsen et al. (2009), Ask et al. (2013), and Hartwig and Bond (2011), have all hinted at the possibility that *the ability to accurately detect deception may linger below the reaches of conscious introspection.*

The Unconscious Mind is Better Equipped to Accurately Detect Deception

From a dual-process perspective, it is possible that less conscious parts of the mind are equipped with the architecture for accurate deception detection but that conscious reasoning compromises accuracy by imposing attribution biases and incorrect stereotypes about how liars behave during deception (Evans & Stanovich, 2013; Vrij, Granhag, & Porter, 2011). Evidence for this notion comes from data showing that imposing cognitive load or interrupting conscious deliberation about a target's veracity *increases* explicit deception detection accuracy by up to 15% (Albrechtsen et al., 2009; Reinhard et al., 2013). These results suggest a tension between explicit and implicit deception detection processes and a consolidative or "corrective" mental design in which bottom-up accuracy by the unconscious is dampened by the extent to which cognitive resources are available to provide top-down interference (Gilbert, 1999). Indirect lie detection strategies too increase accuracy; rating the extent to which a potential liar appears to be ambivalent or thinking hard can more accurately predict veracity than direct truth vs. lie decisions (Vrij, Edward, & Bull, 2011; Sporer & Masip, 2012).

To shift explicit decisions toward accuracy, researchers have provided lie detectors with detailed information about deceptive behavior. Modest success of training programs has been observed ranging from gains of 4% (Frank & Feeley, 2003) to over 30% (Shaw, Porter, & ten Brinke, 2012). While such effortful and cerebral approaches can increase deception detection accuracy, evidence points overwhelmingly to the idea that somewhere below the reaches of introspective access, deception-detection accuracy already exists. Here, we seek direct empirical support for that proposition.

The Current Research

In the current report, we hypothesized that indirect measures of deception detection—capable of accessing less conscious parts of the mind—would yield greater accuracy than direct/conscious measures. Experiment 1 employed a high-stakes mock crime paradigm to produce videos of people ($N_{stimuli}$ = 12) who were either lying or telling the truth about stealing $100. We used these videotaped stimuli to test our unconscious lie detection hypothesis with the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998; Nosek, Greenwald & Banaji, 2007). Experiment 2 further tested the hypothesis by using subliminally presented images and a semantic classification task (Draine & Greenwald, 1998). In both experiments naïve participants also made direct, self-report judgments of whether each person genuine or deceptive.

Experiment 1

The goals of Experiment 1 were to: (1) create a set of videotaped targets insisting they did not commit a crime—half of whom were genuinely pleading their innocence and half of whom were lying, and (2) to use these pleaders to test the hypothesis that indirect measures of deception detection would yield more accuracy than traditional self-reports. These videotaped stimuli are freely available for research use from the authors.

Method

Participants

Seventy-two (51 female) undergraduate students completed the study for $16 in compensation. Participants were 20.7 years old, on average (*SD* = 3.1). Participants observed pleaders—both liars and truth-tellers—in the videos described next.

Materials

High-Stakes Mock Crime: Development of Truthful and Deceptive Pleader Videos

A "high-stakes mock-crime paradigm" was borrowed from the field of criminal justice (Kircher et al., 1988) and social psychology (e.g., Frank & Ekman, 2004; see DePaulo et al., 2003). An experimenter sat with each suspect and explained that they would have an opportunity to earn $100 by convincing the experimenter that they did not steal a $100 bill hidden in the testing room.

Suspects were randomly assigned to steal the $100 or not and all were equally incentivized to do their best to convince the experimenter they did not steal the money (regardless of whether or not they did); all had the chance to earn the $100, and be entered into a lottery to win an additional $500. Failure to convince the experimenter resulted in a loss of the $100 prize.

After instructions were given, the experimenter left the room and suspects received one of two instruction sets (steal condition):

> "STEAL THE MONEY OUT OF THE ENVELOPE!!! Be very quiet. Put the envelope and books back exactly as you found them. Put the money ON YOU somewhere—pocket, sock, wherever, but make sure the experimenter can't see it (obviously). When you are done STEALING the money come back to the computer and click 'continue'."

Or, (no-steal condition):

> DO NOT steal the money in the envelope. Leave the money in the envelope and put it back where you found it. Be very quiet. Put the envelope and books back exactly as you found them. When you are done putting the money and envelope back in the books, come back to the computer and click 'continue'."

After the possible theft, the experimenter re-entered, turned on a video camera, and began the interrogation. Suspects were asked a series of questions in an affectively neutral, firm manner. Ten questions were asked, including: "baseline questions" (i.e., neutral questions not pertaining to the mock theft about verifiable facts) and "pleading questions" (i.e., questions about the possible theft; Kircher et al., 1988). Baseline questions included: "what are you wearing today?"

and "what is the weather like outside?" The "pleading questions" were adapted from Frank and Ekman (2004) and included, "did you steal the money?", "why should I believe you?" and "are you lying to me now?" Using these procedures, *N* = 12 mock crime videos were created for use in the present research, 6 were genuine (6 deceptive) and 5 were male (7 female). On average, interrogations lasted 97 seconds (*SD* = 21.62 seconds) and captured a frontal view of the participant from the shoulders up.

Three variables were measured to assess evidence of deception: emotional distress (self-reported on 4 items: afraid, frightened, scared, jittery), physiological stress (salivary cortisol reactivity: time 1 taken 10 minutes after arrival and time 2 taken ~27 minutes after the manipulation), and nonverbal tells (8 reliably coded deception variables taken primarily from DePaulo et al., 2003: less speaking time, faster speech rate, more nervous, more lip presses, less cooperativeness, more vocal uncertainty, more one-sided shoulder shrugs). The 3 variables were *z*-scored and combined with principle components analysis to produce a deception-stress composite variable. As expected, deceptive pleaders showed more evidence of deception-related stress ($M = .66$; $SD = .99$) than truth-tellers ($M = -.66$; $SD = .42$), $F(1,11) = 9.05$, $p < .013$; $d = 1.81$ (these data accompany the freely available set of videotaped stimuli).

Indirect Measure of Automatic Deception Detection. The IAT was used as a measure of indirect deception detection (see Ask, Granhag, Juhlin, & Vrij, 2013 for use of the same paradigm to reveal deceptive intentions). "Accurate deception detection" was operationalized as a mental association between the liar or truth-teller and congruent deception-related concepts. The principle underlying the IAT is that stimuli sharing conceptual features are more mentally associated. In this context, we were interested in whether observing someone tell a lie would, outside of awareness, activate mental concepts associated with deception. After viewing videos of genuine and deceptive pleaders, we tested whether participants were more likely to conceptually link deceptive pleaders' faces with deception-related concepts (untruthful, dishonest, invalid, deceitful) relative to truth-related concepts (truthful, honest, valid, genuine); and truthful pleaders' faces more with truth-related concepts (relative to deception-related concepts). Specifically, participants watched pairs of pleader videos (one liar, one truth-teller) and completed an IAT juxtaposing the two targets. In the IATs, pleaders' pseudonyms were displayed in the upper left and upper right hand corner of the screen along with the words TRUTH and LIE. Still photographs from each video and words associated with lies and truths were presented in the middle of the screen and classified into the right- or left-hand category. A 5-block IAT format with counterbalancing and scoring procedures by Greenwald, Nosek, and Banaji (2003) was used. Each IAT provided an effect size (*d* score) representing strength of association between liars and deception concepts, and truth-tellers and truth concepts (relative to incongruent pairings). Indirect deception detection was measured as the average *d* score across the 6 IATs with higher values (i.e., those above zero) indicating greater accuracy.

Direct / Self-Report Measures of Deception-Detection. Participants indicated in a forced-choice format whether each person was lying or telling the truth (Bond & DePaulo, 2006).

Procedure

Participants viewed pairs of pleader videos presented in the middle of a computer screen, approximately 4" x 4" in size. A unique pseudonym was displayed across the top of the screen (above the video) for each pleader (e.g., "John"; pseudonyms were balanced for length and commonality within and across pairs). Participants then saw pleaders' images on the screen and completed a direct, self-report judgment of whether each person was: "Lying?" or "Telling the truth?". Following the self-report judgment task, participants completed an IAT. This was repeated 6 times—once for each pair of truthful and deceptive pleader videos, for a total of 12 videos[1].

Results & Discussion

Direct (Self-Report) Deception Detection Accuracy was Poor

Explicit accuracy in discriminating liars from truth-tellers was poor ($M = 46.83\%$; $SD = 13.54$). This accuracy rate was marginally below chance (50%), $t(71) = -1.97$, $p = .053$, $d = -.23$, but falls well within the range of accuracy outcomes included in Bond and DePaulo's (2006) meta-analysis. Detection of lies ($M = 43.75\%$, $SD = 14.92$) specifically was below chance ($t(71) = -3.56$, $p = .001$), while accuracy for truthful statements did not differ from chance ($M = 48.61\%$, $SD = 19.53$), $t(71) = -.60$, $p = .55$. No participant gender differences were evident, $ps > .05$. Findings support the claim that consciously considered self-report judgments lead to poor deception-detection accuracy.

Indirect (Unconscious) Deception-Detection Accuracy was Better

Mean *d* scores for each participant ($M = .06$, $SD = .19$) were significantly greater than zero, $t(71) = 2.63$, $p = .011$, $d = .32$, indicating that in less consciously accessible and controllable parts of the mind, viewing a liar (truth-teller) automatically activates concepts associated with deception (truth) (see Figure 1). Female participants' less-conscious cognitions were significantly more "accurate" ($M = .10$; $SD = .17$) than male participants' ($M = -.03$; $SD = .20$), $t(70) = -2.74$, $p < .01$, $d = -.65$. Consistent with this finding, sex differences in person-perception accuracy tend to show a female advantage around more intuitive judgments (i.e., "gut" or "less conscious" judgments; Hall, 1978).

Experiment 2

[1] In both experiments, participants completed the TIPI (Gosling, Rentfrow & Swann, 2003) and a demographic questionnaire after completing all judgments. Only gender and age were analyzed. No additional conditions were run nor critical DVs measured.

Results of Experiment 1 suggested that when less conscious processing is allowed to render a decision about whether a person is lying or telling the truth, accuracy is revealed. However, Experiment 1 contained methodological limitations inherent to the IAT. First, videos were watched two-at-a-time after which an IAT was taken juxtaposing the two videos—a sequence that was repeated six times. Importantly, each pair included one liar and one truth-teller. The contrast in pleader sincerity could have artificially increased accuracy. Second, the images of liars and truth-tellers presented in the context of the IAT were supraliminal. A stricter test of the unconscious deception detection hypothesis would be to use subliminally presented images to test for unconscious and automatic concept activation. Thus, Experiment 2 used a semantic classification task as the indirect measure of deception detection accuracy in which images of liars and truth-tellers were presented subliminally.

Method

Participants

Sixty-six undergraduates (42 female) completed the study for $16 in compensation. Participants were, on average, 20.33 years old (*SD* = 1.82).

Materials

Observing the Pleader Videos. The same set of twelve pleaders (6 genuine, 6 deceptive) were presented in randomized pairs in Experiment 2. Importantly, videos were randomly linked which resulted in two different (truth-lie) and four same (two truth-truth; two lie-lie) veracity pairings.

Indirect Measure of Deception Detection. A semantic classification task following Draine and Greenwald (1998) was used. Each trial began with a 500ms fixation point (+) in the center of the screen, followed by the subliminally-presented stimulus (< 17ms or, 1 screen refresh) which was forward- and backward-masked by abstract faces taken from Cunningham et al. (2004)—each presented for 200ms (see Figure 2 for trial sequence). Stimuli were one of the four still images extracted from each of the previously watched pleas and subliminal presentation ensured that any target-sincerity spreading activation was unconscious.[2] One of eight target words (truthful, honest, valid, genuine, untruthful, dishonest, invalid, deceitful) then appeared in the center of the screen until the participant sorted the word into the TRUTH or LIE category (appearing in the upper right and left corner of the screen, counterbalanced). All test blocks included 64 trials. A practice block was 8 trials long and did not include subliminal primes; instead, a black screen appeared between the fixation point and each target word. These trails familiarized participants with the task using a red-X-incorrect-feedback procedure.

Procedure

Participants viewed two pleaders (in random pairs and presented in random order) and then completed a semantic classification task in which subliminal stimuli were faces of the two previously seen pleaders. Participants then saw pleaders' images on the screen and completed explicit self-report judgments of sincerity. This procedure was repeated for six video-pairs.

Results & Discussion

Direct (Self-Report) Deception-Detection Accuracy was Poor

Consistent with our hypothesis and evidence from traditional deception detection paradigms, participants performed at chance when veracity judgments were harvested from consciousness (M = 49.62%; SD = 11.36), $t(65) = -.27$, $p = .79$, $d = -.01$. Truthful statements were accurately detected at a rate greater than chance (M = 62.63%; SD = 22.66), $t(65) = 4.53$, $p < .001$. The detection of lies (M = 36.62%; SD = 17.59), however, was significantly below chance, $t(65) = -6.18$, $p < .001$.

[2] At the end of the experiment, participants completed a subliminal threshold sensitivity task. Participants were unable to discriminate male and female faces above chance (M = 48.35%, SD = 15.14) suggesting that faces were presented below conscious perception, $t(65) = -.88$, $p = .38$.

Indirect (Unconscious) Deception-Detection Accuracy was Better

Automatic deception-detection, as represented by a *d* score for each participant ($M = .03$, $SD = .11$) was significantly greater than zero, $t(65) = 2.26$, $p = .027$, $d = .27$, demonstrating that subliminally-presented faces of liars and truth-tellers *activated* and *facilitated* congruent concepts.[3]

Are Indirect Deception-Detection Measures *More Accurate* than Direct Measures?

To directly compare direct and indirect measures of deception-detection accuracy, we conducted a mini meta-analysis of Experiments 1 and 2, comparing effect sizes (see Figure 1) with a *z* test. The average effect size for indirect (unconscious) measures was $r = .28$ and $r = -.11$ for direct (self-report) measures. As expected, automatic associations were significantly more accurate than controlled, deliberate decisions, $z = -3.32$, $p < .001$. These findings support our hypothesis that: (1) viewing a liar automatically, and unconsciously, activates deception-related concepts (and viewing a truth-teller activates truth concepts), and (2) indirect measures of deception-detection yield more accuracy than historically used, direct self-reports.

General Discussion

Across 2 experiments, indirect measures of accuracy in deception detection were superior to traditional, direct measures, providing strong evidence for the idea that, although humans can't consciously report who is lying and who is telling the truth, somewhere on a less conscious level we do actually have a sense of when someone is lying. The current results are consistent with research on primates who lack self-awareness and yet demonstrate the ability to detect deception (Wheeler, 2010). It is also consistent with evidence that both cognitively-taxed humans and certain parts of the brain can discern liars from truth-tellers (Albrechtsen, et al., 2009; Etcoff et

[3] Unlike Experiment 1, no gender differences were evident, $t(64) = 1.67$, $p = .10$.

al., 2000). Findings provide long sought-for support for the evolutionary perspective that accurate deception detection is adaptive and should be favored (Krebs & Dawkins, 1984).

Characterizing human deception detection as an error-fraught process, no more accurate than chance, is a misleading summary of scientific insight on the topic, given interdisciplinary findings and the results presented here. But how does consciousness interfere with our natural ability to detect deception? From a dual process perspective, our results—in combination with insights from Albrechtsen et al. (2009) and Hartwig and Bond (2011)—suggest that the unconscious can make efficient and effective use of cues to deception, but that resulting accurate unconscious assessments are made inaccurate either by consolidation with, or "correction" by, conscious biases and incorrect decision-rules (Gilbert, 1999).

Future Directions & Limitations

While these findings provide little in the way of practical implications at present, they provide a new lens through which to examine future questions and shed light on a process through which accurate lie detection may occur. Future investigations should replicate and extend this effect to different forms of deception. As evidenced by the variability in accuracy illustrated by Bond and DePaulo (2006), new stimuli may lead to more or less accurate explicit responses than reported here; regardless, we expect that implicit measures would always outperform explicit judgments of the same, reflecting our proposition that the unconscious mind identifies and processes cues to deception (to the extent that they are available) more efficiently and effectively than the conscious mind. Further, an important question is whether implicit accuracy is associated with the enhanced ability to detect lies, truths, or both. The evolutionary arms race framework would predict increased accuracy in the detection of lies, in particular. Our focus on overall discrimination accuracy, while the norm in deception detection research, is a limitation of the current work, and the methodologies we chose to measure implicit responses. Accuracies reported by Reinhard et al. (2013), however, appear to support the arms race prediction; preventing conscious deliberation about credibility improved lie detection more than truth detection. Future research should utilize alternative measures of implicit thought that allow for the direct examination of implicit truth and lie accuracies.

Conclusion

In short, while the detection of lies is of great importance in personal, professional, and civic domains, past research has concluded that conscious determinations of deception are error ridden—a dismal conclusion that contradicts evolutionary theory. Our findings suggest that accurate lie-detection is, indeed, a capacity of the human mind, potentially directing survival and reproduction-enhancing behavior from below introspective access.

References

Albrechtsen, J. S., Meissner, C. A., & Susa, K. J. (2009). Can intuition improve deception detection performance? *Journal of Experimental Social Psychology*, *45*(4), 1052-1055.

Ask, K., Granhag, P. A., Juhlin, F., & Vrij, A. (2013, in press). Intending or pretending? Automatic evaluations of goal cues discriminate true and false intentions. *Applied Cognitive Psychology*.

Bond, C. F., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review*, *10*(3), 214-234.

Bond C. F., Kahler, K. N., & Paolicelli, L. M. (1985). The miscommunication of deception: An adaptive perspective. *Journal of Experimental Social Psychology*, *21*, 331-345.

Byrne, R. W., & Corp, N. (2004). Neocortex size predicts deception rate in primates. *Proceedings of the Royal Society B: Biological Sciences, 271*(1549), 1693.

Carney, D. R., Dubois, D., Nichiporuk, N., ten Brinke, L., Rucker, D. D., & Galinsky, A. D. (2013). The deception equilibrium: The powerful are better liars but the powerless are better lie-detectors. *Manuscript submitted for publication.*

Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of black and white faces. *Psychological Science*, *15*(12), 806-813.

Dawkins R., & Krebs J. R. (1978). Animal signals: information or manipulation? In J. R. Krebs & N. B. Davies (Eds.) *Behavioural Ecology* (pp. 282-309). Oxford, England: Blackwell Scientific Publications.

DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, *129*(1), 74-118.

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology-General*, *127*(3), 286-303.

Ekman, P. (1992). *Telling lies: Clues to deceit in the marketplace, politics, and marriage.* New York: W. W. Norton.

Ekman, P., & O'Sullivan, M. (1991). Who can catch a liar? *American Psychologist, 46*(9), 913-920.

Etcoff, N. L., Ekman, P., Magee, J. J., & Frank, M. G. (2000). Lie detection and language comprehension. *Nature*, *405*(6783), 139.

Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223-241.

Frank, M. G., & Ekman, P. (2004). Appearing truthful generalizes across different deception situations. *Journal of Personality and Social Psychology*, *86*(3), 486.

Frank, M. G., & Feeley, T. H. (2003). To catch a liar: Challenges for research in lie detection training. *Journal of Applied Communication Research*, *31*(1), 58-75.

Gilbert, D. (1999). What the mind's not. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 3-11). New York: Guilford.

Goodall, J. (1986). *The chimpanzees of Gombe: Patterns of behavior.* Cambridge: Harvard University Press.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit social cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464-1480.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology, 85*, 197-216.

Grèzes, J., Frith, C., & Passingham, R. E. (2004). Brain mechanisms for inferring deceit in the actions of others. *The Journal of Neuroscience*, *24*(24), 5500-5505.

Kircher, J. C., Horowitz, S. W., & Raskin, D. C. (1988). Meta-analysis of mock crime studies of the control question polygraph technique. *Law and Human Behavior*, *12*(1), 79-90.

Hall, J. A. (1978). Gender effects in decoding nonverbal cues. *Psychological Bulletin*, *85*(4), 845-857.

Hartwig, M., & Bond, C. F. (2011). Why do lie-catchers fail? A lens model meta-analysis of human lie judgments. *Psychological Bulletin*, *137*(4), 643.

Krebs, J. R., & Dawkins, R. (1984). Animal signals: Mind-reading and manipulation. *Behavioural Ecology: An Evolutionary Approach, 2*, 380-402.

Lissek, S., Peters, S., Fuchs, N., Witthaus, H., Nicolas, V., Tegenthoff, M., Juckel, G., & Brune, M. (2008). Cooperation and deception recruit different subsets of the theory-of-mind network. *PLoS One, 3*(4), e2023.

Mann, S., Ewans, S., Shaw, D., Vrij, A., Leal, S., & Hillman, J. (2013). Lying eyes: Why liars seek deliberate eye contact. *Psychiatry, Psychology and Law, 20*(3), 452-461.

Menzel, E.W. Jr. (1974). A group of young chimpanzees in a one-acre field: Leadership and communication. In A.M. Schrier & F. Stollnitz (Eds.), *Behavior of nonhuman primates* (pp. 83−153). Academic Press.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The implicit association test at age 7: A methodological and conceptual review (pp.265-292). In J. A. Bargh (Ed.), *Automatic processes in social thinking and behavior.* Psychology Press.

O'Sullivan (2003). The fundamental attribution error in detecting deception: The boy-who-cried-wolf effect. *Personality and Social Psychology Bulletin, 29,* 1316-1327.

Reinhard, M. A., Greifeneder, R., & Scharmach, M. (2013). Unconscious processes improve lie detection. *Journal of Personality and Social Psychology, 105*, 721-739.

Shaw, J., Porter, S., & ten Brinke, L. (2013). Catching liars: training mental health and legal professionals to detect high-stakes lies. *Journal of Forensic Psychiatry & Psychology*, *Online First.*

Sodian, B., & Frith, U. (1992). Deception and sabotage in autistic, retarded and normal children. *Journal of Child Psychology and Psychiatry*, *33*(3), 591-605.

Sporer. S. L., & Masip, J. (March 14, 2012). Indirect vs. direct measures to detect deception: A cross-national comparison. *Paper presented at the annual meeting of the American Psychology-Law Society, San Juan, Puerto Rico.*

The Global Deception Research Team (2006). A world of lies. *Journal of Cross-Cultural Psychology*, *37*(1), 60-74.

von Hippel, W., & Trivers, R. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences, 34,* 1-56.

Vrij, A., Edward, K., & Bull, R. (2001). Police officers' ability to detect deceit: The benefit of indirect deception detection measures. *Legal and Criminological Psychology, 6,* 185-196.

Vrij, A., Granhag, P. A., & Porter, S. (2010). Pitfalls and opportunities in nonverbal and verbal lie detection. *Psychological Science in the Public Interest*, *11*(3), 89-121.

Wheeler, B. C. (2010). Decrease in alarm call response among tufted capuchins in competitive feeding contexts: Possible evidence for counter-deception. *International Journal of Primatology*, *31*(4), 665-675.

Wiseman, R., Watt, C., ten Brinke, L., Porter, S., Couper, S. L., & Rankin, C. (2012). The eyes don't have it: Lie detection and neuro-linguistic programming. *PloS One*, *7*(7), e40259.

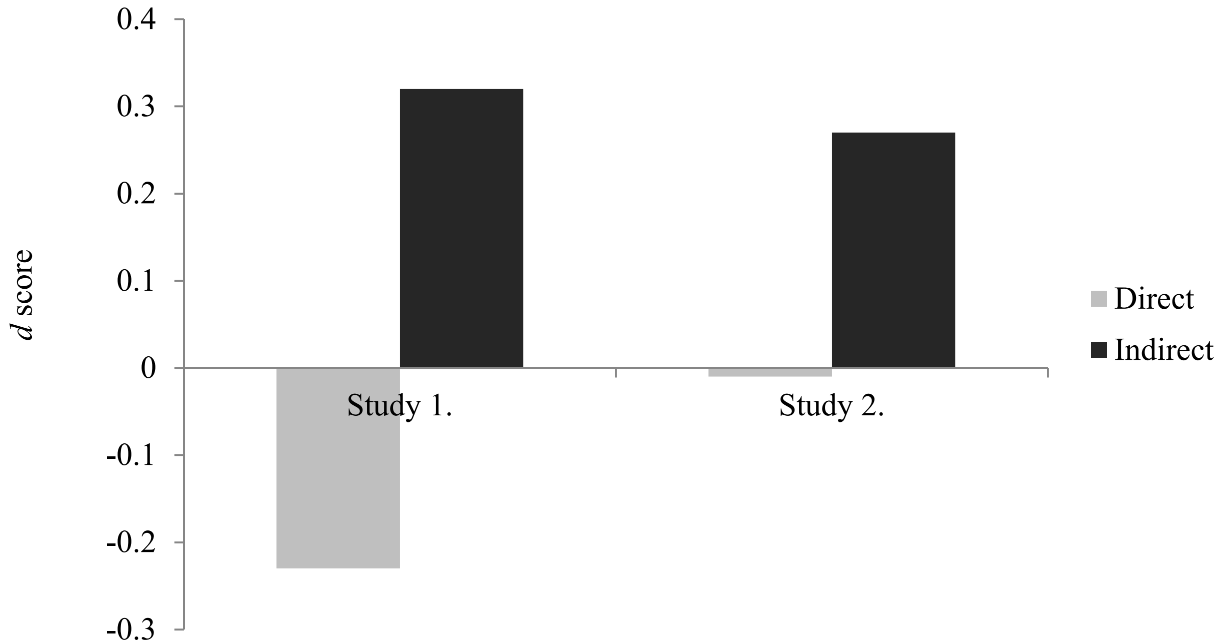

Figure 1. Effect sizes (*d* scores) of analyses testing for the presence of discrimination between liars and truth-tellers, using direct and indirect (IAT or SCT, in Experiments 1 and 2 respectively) tasks.

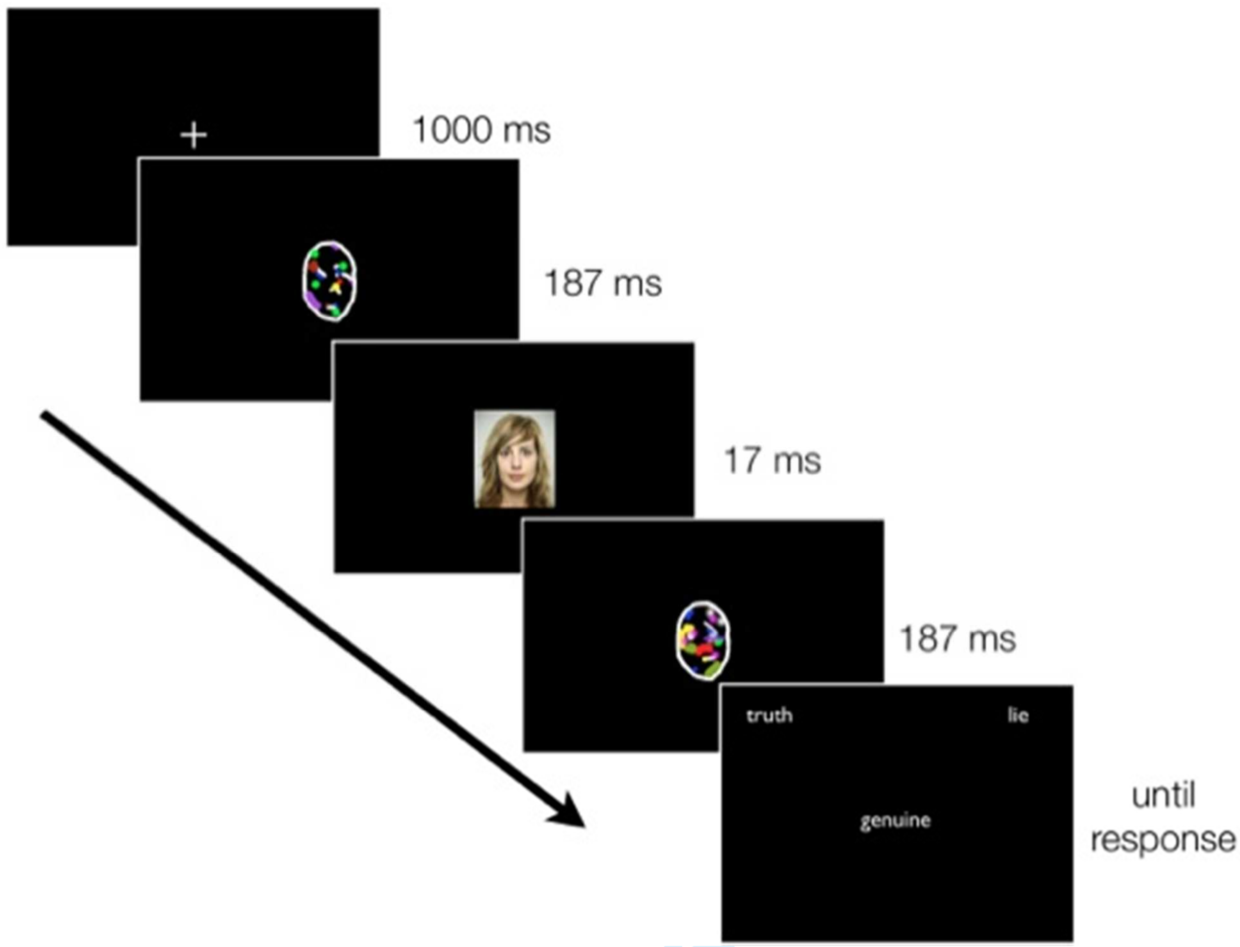

Figure 2. Sequence of stimuli presentation in each test trial of the semantic classification task.