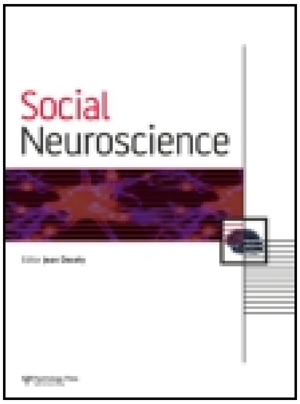

Social Neuroscience
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/psns20

Types of deception revealed by individual differences in cognitive abilities

Charity J. Morgan [a], Julia B. LeSage [a] & Stephen M. Kosslyn [a]
[a] Harvard University, Cambridge, Massachusetts, USA

Available online: 25 Sep 2009

To cite this article: Charity J. Morgan, Julia B. LeSage & Stephen M. Kosslyn (2009): Types of deception revealed by individual differences in cognitive abilities, Social Neuroscience, 4:6, 554-569

To link to this article: http://dx.doi.org/10.1080/17470910802299987

SOCIAL NEUROSCIENCE, 2009, 4 (6), 554–569

Psychology Press
Taylor & Francis Group

Types of deception revealed by individual differences in cognitive abilities

Charity J. Morgan, Julia B. LeSage, and Stephen M. Kosslyn
Harvard University, Cambridge, Massachusetts, USA

The two studies reported in this article are an extension of the neuroimaging study by Ganis et al. (2003), which provided evidence that different types of lies arise from different cognitive processes. We examined the initial response times (IRTs) to questions answered both deceptively and truthfully. We considered four types of deceptive responses: a coherent set of rehearsed, memorized lies about a life experience; a coherent set of lies spontaneously created about a life experience; a set of isolated lies involving self-knowledge; and a set of isolated lies involving knowledge of another person. We assessed the difference between truthful and deceptive IRTs. Scores from cognitive tasks included in the MiniCog Rapid Assessment Battery (MRAB) were significant predictors of IRT differences. Each type of lie was predicted by a distinct set of MRAB scores. These results provide further evidence that deception is a multifaceted process and that different kinds of lies arise from the operation of different cognitive processes.

INTRODUCTION

Deception involves intentionally misleading another person, either by omission or commission. For both theoretical and practical reasons, many researchers have studied how to detect deception, from numerous vantage points, over the course of many years (Burgoon & Buller, 1994; Ekman, 1992, 2001; Horvath, Jayne, & Buckley, 1994; Mehrabian, 1971; Vrij, 1994; Zuckerman, DeFrank, Hall, Larrance, & Rosenthal, 1979). Nevertheless, to date no method of detecting deception has proven satisfactory (e.g., National Research Council, 2003). In particular, the standard approach, based on measuring signals at the periphery (such as skin conductance, respiration rate, blood pressure, and heart rate), has been highly criticized. Polygraph tests rely on the assumption that deception creates arousing emotions such as guilt or fear, and such tests are designed to detect the physiological changes that occur during deception. Such changes are measured at the periphery, and typically rely on assessing heart rate, skin conductance, and blood pressure. However, not only has no unique physiological pattern been found to be associated with deception (Iacono, 2000), but also many researchers even question the validity of the assumptions on which the polygraph is based (Iacono & Lykken, 1997). Furthermore, the results of laboratory studies on the validity of the polygraph may not be generalizable to the real world—and the accuracy of polygraph results in the field is difficult to determine because

Correspondence should be addressed to: Charity Morgan, Department of Statistics, One Oxford Street, Cambridge, MA 02138, USA. Email: cmorgan@stat.harvard.edu

Julia B. LeSage is now in Macon, Georgia. This material is based partially on work supported by the National Science Foundation under Grant BCS-0322611. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors wish to thank Jennifer M. Shepard, William L. Thompson, and Giorgio Ganis for their helpful advice and constructive criticism.

www.psypress.com/socialneuroscience DOI:10.1080/17470910802299987

ground truth cannot always be established (Vrij, 2008).

In response to the drawbacks of detecting deception by registering signals at the periphery, researchers have turned to studying the brain events that underlie deception. Two broad approaches have emerged. On one hand, some researchers have sought to discover a "neural signature" of deception. For example, studies have shown that the anterior cingulate—a brain area thought to be involved in monitoring conflict between observed and expected events—is active when one engages in deception (Kozel et al., 2004; Langleben et al., 2002). On the other hand, Ganis, Kosslyn, Stose, Thompson, and Yurgelun-Todd (2003) suggested that there is no single neural signature of deception. Instead, these researchers argue that there are numerous different ways to lie, and different neural networks are evoked depending on the type of lie. In particular, Ganis et al. (2003) showed that different brain areas are activated when a person lies spontaneously than when a person lies on the basis of a previously memorized scenario.

The present research follows up the Ganis et al. (2003) functional magnetic resonance imaging (fMRI) study, using an entirely different logic to document that different processes can be used to tell lies. In these studies we exploit the fact that people differ in how effectively they can perform specific information processing, such as monitoring conflict or retrieving stored information. If different processes are evoked to produce different sorts of lies, then individual differences in the efficacy of different processes should predict performance when participants produce different types of lies. We use this logic to extend the inferences drawn by Ganis et al., and document an additional dimension that underlies the production of deception.

The present studies relied on initial response times (IRTs), which were defined as the length of time between the end of a question and the beginning of a response. Greene, O'Hair, Cody, and Yen (1985) characterized an increased latency between question and response as "indicative of heavy demands upon the central processing capacity" and reasoned that longer latencies when providing a deceptive response would be consistent with the theory that "lying generally requires more cognitive work than telling the truth." Cody, Marston, and Foster (1984) also consider the IRT to be a reliable indicator of "the cognitive effort required for successful deception [that] cause[s] communicators to exhibit longer response latencies when responding deceptively . . . than when responding truthfully" (as cited in Stiff, Corman, Krizek, & Snider, 1994). Seymour, Seifert, Shafto, and Mosmann (2000) found that IRT could reliably distinguish which participants in their study were in possession of "guilty knowledge." Similarly, Farrow et al. (2003) asked participants questions about themselves and found that participants required longer IRTs when lying than when telling the truth. Sporer and Schwandt's (2006) meta-analytic study concluded that people tend to take slightly longer to respond to questions when lying than when telling the truth.

However, telling a lie may not always be signaled by increased IRT. McCornack (1997) notes that the cognitive effort required to produce a deceptive message can vary. We propose that such variations could reflect differences in the type of lie being produced as well as individual differences in the deceiver's cognitive abilities. It comes as no surprise, then, that a number of studies have found that participants do not always have longer IRTs during deception. For example, Greene et al. (1985) found that people actually responded more quickly with a prepared lie than they did with the truth. Greene et al. did not find any group differences when the lies told were spontaneous. Vrij (2008) summarizes several studies that consider IRT and finds an inconsistent relationship between the length of IRT and deception.

One limitation of previous research on deception is that investigators often ignore or underestimate the importance of the ways in which lies can differ. For example, deTurck and Miller (1985) designed an experiment in which participants cheated on a task performed with a confederate. Participants who subsequently tried to conceal their cheating therefore had to lie not only about their own behavior but also about the behavior of the confederate. The fact that some of the participants' deceptive statements concerned their own behavior and others concerned the behavior of the confederate is a potentially confounding factor that was not examined nor acknowledged. In contrast, Greene and colleagues (1985) examined prepared lies and spontaneous lies separately, and hypothesized that creating a lie spontaneously would require more cognitive effort than reciting a lie that had been prepared beforehand. However, they did not investigate possible differences in cognitive demands that might underlie different sorts of lies.

Another limitation of previous deception research is the lack of focus on individual differences, particularly cognitive differences. For example, Ganis et al. (2003) and Greene et al. (1985) report group differences in response time and do not compare IRT within individuals. Although Farrow et al. (2003) do consider individual differences in IRT, they cast these differences solely in the context of sex and personality traits. In the studies reported here, we compared IRT for telling the truth versus lying for each individual, essentially using each participant as his or her own control. Furthermore, we examine the relationship between an individual's IRT differences and his or her cognitive abilities.

IRT differences can also be useful in judging whether an individual is skilled at deception. An individual's attempt to deceive another can only truly be judged successful if the recipient of the deceptive message is indeed fooled by it. In the studies reported here, we focus only on the deceivers and therefore cannot judge the effectiveness of the participants' attempts at deception in this way. Instead we reason that in order for attempts at deception to be successful, a lie must be indistinguishable from the truth. One of several possible ways that a deceptive response can differ from a truthful one is in the amount of time required to produce it. We expect that some successful deceivers will require approximately the same amount of time to tell a lie as to tell the truth; that is, effective deception should result in IRT difference scores close to zero. However, the ability to deceive effectively should depend, at least in part, on how well a person can manifest relevant cognitive abilities. Thus, we expected that the difference in IRTs between truthful and deceptive responses should vary depending on the type of lie and on the relevant cognitive strengths and weaknesses of the participant.

To assess individual differences in cognitive abilities, we used the MiniCog Rapid Assessment Battery (MRAB, Shephard & Kosslyn, 2005). This battery consists of nine tasks, which rely upon response time and error rate measurements. These tasks assess three sorts of attention (selective, divided, and vigilance), two kinds of working memory (verbal and spatial), three kinds of reasoning (three-term series verbal deduction, spatial mental rotation, and set-switching), and simple perceptual-motor reaction time. A summary of these tasks is provided in the appendix. If different sets of individual differences predict IRTs for different types of lies, that is evidence that different sorts of processing underlie the production of different types of lies.

EXPERIMENT 1

This experiment essentially replicated the behavioral methods of the Ganis et al. (2003) fMRI study. We compared lies when participants lied spontaneously versus when they lied based on a previously memorized story. However, instead of assessing patterns of neural activity, we measured individual differences on the MRAB in advance, and then measured IRT and error rates (ERs) during the production of deception and truthful responses.

Based on the earlier fMRI results, we had one strong prediction: Individual differences on the MRAB variant of the classic Stroop task—a "filtering" task that required indicating how many repetitions of a digit were present, independent of what the digit actually was—should predict the results when participants lied spontaneously but not when they lied on the basis of a memorized story. Ganis et al. (2003) found consistent activation of the anterior cingulate only during the spontaneous condition, and this structure is strongly activated during variants of the Stroop task (Bush et al., 1998).

Method

Participants

Thirty-nine participants (20 females and 19 males) volunteered to take part in the study for pay. Five participants (four females and one male) were removed from the sample because they had unusually high error rates (13.64% or greater) on the computer-administered deception task. The final sample consisted of 16 females (mean age 23.5 years; range 19–31 years) and 18 males (mean age 23.44 years; range 18–35 years). Seven females and 11 males were Harvard students, and the remaining participants were from the community or other local colleges. The mean for the Bors and Stokes (1998) short form (12-item) of Raven's Advanced Progressive Matrices (APM) for the females was 8.06 (range 2–12) and for the males was 7.78 (range 3–12). The mean number of years of education for the females was 16.25 (range 12–22) and for the males was 15.44 (range 12–20); 14 females and 14 males were right-handed. All participants were US citizens or permanent residents, had lived in the USA for at least the past 10 years, spoke fluent English, did not stutter, and were physically and psychologically healthy. The fliers and studypool information used to recruit participants did not mention that the study was designed to investigate the nature of deception, but instead described it only as a "cognitive" study. The study was conducted with the informed consent of each participant and the approval of the university's Institutional Review Board.

Materials

A critical part of this study, and the most time-consuming, was the interaction between the investigator and the participant *before* the participant arrived in the lab. The investigator sent participants an email message, asking them to review their most memorable work experience and their most memorable vacation experience. The participant was asked to write short summaries of these events, and to send the descriptions to the investigator, along with a phone number, and a time the participant could come into the lab. Each of the two experiences was to be summarized in one page, which was to include a detailed description (who, what, when, where, and why) of the event. If the descriptions did not include enough detail, the investigator asked for more detail, but this had to be done only twice (both were cases in which the participants sent only half a page instead of a full page for each event).

The investigator generated questions from the descriptions, and called the participant the day before the scheduled session (which was an average of one week after the descriptions were received). During this phone call, the participant was instructed to invent an alternate reality, a vacation or work experience (the selection of work or vacation alternated with each participant) that never happened but could have happened. The investigator wrote down everything the participant said, while at the same time, coaching the participant to include details that the investigator needed for the prepared questions. We used the same procedure as Ganis et al. (2003, p. 832):

> For example, if they actually took their vacation in Florida, we instructed them to pretend that the vacation took place in another location in the United States (e.g. California); if they traveled there by car, we instructed them to pretend they used another means of transportation (e.g. by plane), and so on. We helped the participants in this process by ensuring that the scenarios they generated were coherent and internally consistent.

However, in this study, the investigator explicitly asked the participant to think of a *plausible* alternative, such as another vacation that he/she was actually considering. Most of the participants could easily think of the place where they could have gone on vacation or another plausible event at work. The investigator probed for details about that alternative, but never actually posed the questions that were used in the study. We created a transcript of this "alternative reality."
Approximately 30 min after this conversation, the investigator emailed the participant the transcript of the alternative reality to memorize *before* they went to sleep that night. We asked them to visualize themselves going on this vacation or working at this place, which would help them to memorize the details of the scenario. We also told the participants that they had to memorize the scenario in detail because they would be asked questions about it when they came into the lab.

We created two sets of 11 questions for each participant, one set for the work scenario and one set for the vacation scenario. Each of these sets was presented to the participant twice, with one presentation requiring true responses and one requiring deceptive responses. The "alternative reality" scenario was used to produce the deceptive responses for the relevant event for a given participant. All questions were recorded on the computer prior to the participant's arriving at the lab, using SoundEdit 16 software. Unlike Ganis et al. (2003), we also avoided asking questions about the time period of the event, because we wanted the alternative reality to be plausible even if it was not what actually occurred at that point in time.

Procedure

The participant came to the lab and completed the informed consent form. After this informed consent had been completed, we administered the nine MRAB tasks, presented on a Palm Tungsten C personal digital assistant (PDA). The investigator began by reading a prepared script, which explained to each participant how to use the PDA, how to position the body in the chair, how to position the thumbs over the PDA buttons, and how to proceed through the nine tasks. The script ended by instructing the participants to complete each of the nine tasks as quickly and as accurately as possible. The investigator remained in the room, behind and out of sight of the participant, for the entire study and answered any questions the participants may have had.

Once the participant completed the MRAB, the investigator sat across from the participant and asked whether he or she had memorized the ''alternative reality'' scenario before going to sleep the previous night. Each participant said that he or she had memorized the alternate scenario. Next, the investigator asked the participant to recite this scenario. If the participant missed a detail, the investigator helped the participant to think about that detail, without asking the participant a question. It was rare that the participant did not mention all the details; when that did occur, the participant quickly remembered upon being told that he or she needed to talk more about a certain aspect of the scenario. Finally, the investigator asked the participant to review and then discuss the written transcripts of both of the true scenarios and the single false alternate scenario.

The investigator then instructed the participant to move close to the microphone (which was placed on a stand on the desk). A video camera was attached to the top of the computer and the keyboard was off to the side of the computer. The participant read the instructions on the computer screen, pressing the space bar to progress to the next page on the screen. The instructions indicated that a second investigator would be watching the participant's face from another room, and would be trying to determine on which of the four blocks the participant had lied. Participants were told that every time the observer detected lies, the participant would lose \$4 from the hourly wage; every time the participant succeeded in deceiving the observer, he or she would gain \$4. When the participant came to the page that described the video camera and the fact that he/she was being watched, the investigator turned the camera on, which produced a clearly audible ''beep'' so that the participant knew that it was operating. The investigator would say, ''You are now being watched.'' The investigator would then leave the room briefly and go across the hall, pretending that she was checking with the ''second lab member'' who was viewing the participant. She would come back in the room and say, ''You are positioned correctly. He can see your face,'' so that the participants would really believe that they were being observed. Although participants were told that their earnings depended in part on successfully deceiving the observer, in fact there was no lab member watching the experiment and all participants were paid the full amount as if they had successfully deceived the observer.

The investigator next described the individual trials. The cue at the beginning of each block indicated which event would be probed, and how the participant should respond. For example, one cue read ''Please answer the following questions about your VACATION experience with FALSE answers.'' Participants were asked to respond to each question with a one word response. They were asked not to say ''um,'' and to remain silent between questions.

The last screen of the instructions asked the participant to paraphrase what he/she would be doing during the study. The investigator would not proceed to the practice trials until this paraphrase was completely correct. For example, the instructions explained that there would be four blocks randomly presented. One would be a block of questions for which the participants had to give true answers about their vacation, one would require false answers about their vacation, one would require true answers about their work, and one would require false answers about their work. If the participants had a memorized, alternate scenario for their work memory, the investigator would again explain that the false answers for the work memory had to come from the alternative reality, and nothing else. For the other event, which did not have a previously formulated alternative reality, the investigator went on to explain that the spontaneous false answers had to make sense. For example, if the participant said he went to Hawaii, he could not say that he saw a polar bear there—and, the participant was reminded, such a response included two words, but an appropriate response should include only a single word.

Once the investigator was satisfied that the participants understood the procedure, participants were allowed to proceed to the practice trials. The final instruction on the screen urged the participants to ''Please respond as quickly and as accurately as possible.''

Participants were presented with four practice questions, one for each block of instructions.

Before each practice question, participants were given a cue that indicated how they were to answer the following question, and were asked to push the space bar to continue. Following the presentation of the cue, a question was presented auditorially. Once the participant had responded, a fixation point appeared on the screen to signal that the next cue would appear shortly. The fixation point remained on the screen for 1 s. After the practice trials the investigator talked with the participant about any errors in the responses. For example, if the participant said "um," it was caught during the practice trials and the instructions were reiterated by the investigator.

The participants were then presented with four blocks of questions, with each block containing 11 questions, for a total of 44 questions. To stay as close to the Ganis et al. (2003) procedure as possible, the truth blocks came after the lie blocks, but the order of the vacation or work blocks was counterbalanced. All the questions in the blocks were presented in a random order. We presented both lie conditions first in order to avoid potential short-term interference from the truth conditions (such as actively having to inhibit primed responses; see Ganis et al., 2003).

All questions were presented auditorially by the PsyScope program (Cohen, MacWhinney, Flatt, & Provost, 1993), on a Macintosh computer. Upon presentation of the question, the IRT of the participant's spoken reply was registered by the PsyScope program, via a microphone attached to the computer.

Once the actual experimental trials began, the investigator sat behind the participant so that she could see the computer screen and the dot that appeared instantly when the participant made a response. She recorded the verbal responses at the same time they were made and recorded any responses that were invalid (two words instead of one, "um," or deviation from the alternative reality scenario). Error rates later were calculated from this information.

After the four blocks of test trials, the investigator shut down the camera. If the investigator had any questions about the responses, she asked the participant at this time. Finally, we administered the Bors and Stokes (1998) short form of the Raven's APM, along with a debriefing form that asked not only whether participants had any problems during the study, but also questions pertaining to MRAB and biorhythms, such as "What time did you go to sleep last night?" and "Did you have breakfast this morning?". These questions were asked in order to analyze the MRAB data in other, larger databases.

Results

For each of the four conditions (Work/Truth; Work/Lie; Vacation/Truth; Vacation/Lie, where one of the lie conditions was spontaneous and one memorized), we calculated difference scores for each participant by subtracting the mean IRT for the valid truthful responses from the mean IRT for the valid deceptive responses (we did not include IRTs for trials on which participants made errors). We grouped the deceptive responses according to whether the lies were spontaneous or based on the memorized alternative reality.

We began with a two-way analysis of variance, comparing the difference scores for Spontaneous versus Memorized lies for men versus women. There was no effect of gender, $F < 1$, nor an interaction between gender and condition (Wilks' $\lambda = .96$, $F = 1.39$, $p = .24$). Thus, in all subsequent analyses we pooled over gender.

In the Memorized condition, the participants tended to take less time to lie than to tell the truth (with means of 1345 and 1523 ms for the lie and truth conditions, respectively), $t(33) = -1.78$, $p = .08$ (two-tailed). This result is consistent with Greene and colleagues' (1985) finding that participants responded faster when they provided a memorized lie than when they provided the truth. We speculate that participants may have had easier access to the lie than to the truth (which was not as well rehearsed) following the recent rehearsal and intense emphasis on memorization of the alternate scenario; that is, reciting the well-memorized lie may have required less cognitive effort than telling the truth, resulting in lower IRTs.

In contrast, in the Spontaneous condition we did not find even a tendency toward a difference between the two conditions (with means of 1667 and 1599 ms for the lie and truth conditions, respectively), $t < 1$, $p > .1$ (two-tailed). This result also agrees with the findings of Greene et al. (1985), who reported no difference between the response times for participants providing a spontaneous lie and those providing the truth. For the Spontaneous condition, the difference between the cognitive processes required to produce a spontaneous lie and those required to provide a truthful response may not be captured by IRT differences alone. In addition, despite the fact that the alternate scenario was not used in the Spontaneous condition, perhaps having a well memorized story available interfered with participants' performance in the Spontaneous condition. Alternatively, perhaps the participants were fatigued by the time the truth blocks were presented, although there was no evidence of such an effect in Ganis et al. (2003).

Next, for each of the nine MRAB tasks, we calculated each participant's average response time (RT) and error rate (ER). All RT calculations were based only on the trials that the participant answered correctly. We hypothesized not only that the MRAB cognitive measures can predict performance on deception tasks, but also that different subsets of these measures will predict performance for different types of lies. We also included each participant's score on the shortened Raven's form as a means of controlling for general intelligence.

Next, we used a forward stepwise multiple regression analysis to examine the independent contributions of the different MRAB scores, if any, for predicting the IRT difference scores.

The mean Raven's score for the participants is significantly higher ($p = .04$) than the norms reported in Bors and Stokes (1998). In order to examine the impact of the participants' relatively high Raven's scores, we analyzed the data both with and without using the Raven's score as a predictor variable. For the first two stepwise regression analyses performed on data from each condition, we forced the short-form Raven's score in as a predictor at the outset; that is, we selected the Raven's score as a predictor of the IRT difference scores, and allowed the stepwise regression procedure to determine which MRAB scores predicted IRT difference scores after controlling for the Raven's scores. In the second pair of regression analyses, we omitted the Raven's score altogether as an independent variable, and only included the MRAB scores. The stepwise regression procedure then determined which MRAB scores predicted the IRT difference scores without controlling for general intelligence. The results of the stepwise regressions are summarized in Tables 1 and 2.

First, we note that the results change depending on whether the Raven's score is forced into or excluded from the regression analysis. However, in either case, the crucial result stands: Specifically, we found that different MRAB results predict the IRT difference scores in the Spontaneous versus Memorized conditions.

In order to document further that individual differences in different processes predict specific types of lies, we asked whether either of the sets of predictors identified by the stepwise regression analysis could predict both types of lies successfully. Each IRT difference score was fit to both sets of MRAB results, with the Raven's score excluded from the analyses. To account for the fact that the two types of lies have differing numbers of MRAB predictors, for each fit we calculated the Akaike Information Criterion (AIC; Akaike, 1974). A better fit can always be obtained by adding more predictors; the AIC evaluates the fit of a linear model by using maximum likelihood, but also adds a penalty based on the number of predictors used in the model. When comparing linear models with the same response variable, the model with the lower AIC value fits the data better. The results are presented in Table 3. The set of predictors identified by the stepwise regression in the Memorized condition (AIC = 524) explains the IRT difference scores for memorized lies better

TABLE 1
Stepwise regression analysis for Experiment 1, with IRT difference scores as criteria and MRAB scores as predictors: Raven's scores forced in at the outset

Criteria	*Predictors*	*β*	*Partial* R^2	*Model* R^2	*F*
Memorized	Raven's score	0.31	.17	.17	6.24**
	Mental rotation RT	−0.51	.13	.30	5.50**
	Vigilance RT	0.36	.11	.41	5.17**
Spontaneous	Raven's score	−0.05	.01	.01	0.22
	Spatial working memory RT	−0.77	.09	.10	2.93*
	Verbal working memory RT	0.80	.12	.22	4.36**
	Vigilance ER	0.50	.10	.32	3.96*
	Filtering RT	−0.36	.11	.43	5.12**

Notes: *$p < .10$; **$p < .05$; ***$p < .01$; ****$p < .005$.

TABLE 2
Stepwise regression analysis for Experiment 1, with IRT difference scores as criteria and MRAB scores as Predictors: Raven's scores not included

Criteria	*Predictors*	*β*	*Partial R^2*	*Model R^2*	*F*
Memorized	Mental Rotation RT	−0.50	.22	.22	8.93***
	Vigilance RT	0.37	.09	.31	4.11*
	Cognitive set switching ER	−0.31	.07	.38	3.14*
	Spatial working memory RT	−0.25	.05	.43	2.48
Spontaneous	Spatial working memory RT	−0.78	.09	.09	2.97*
	Verbal working memory RT	0.80	.11	.20	3.97*
	Vigilance ER	0.52	.12	.32	4.92**
	Filtering RT	−0.36	.11	.43	5.21**

Notes: *$p < .10$. **$p < .05$. ***$p < .01$. ****$p < .005$.

than those identified in the Spontaneous condition (AIC = 538). The reverse is true for the Spontaneous IRT difference scores: The set of predictors identified in the Spontaneous condition (AIC = 526) explains the IRT difference scores for spontaneous lies better than the set of predictors identified in the Memorized condition (AIC = 539). Thus, we see that no single set of MRAB scores better predicts the IRT scores for both types of lies.

Discussion

The results revealed that different individual differences predict the IRT measure (the difference in time to initiate a lie versus to tell the truth) for Memorized versus Spontaneous deceptive responses. In particular, we see that Filtering RT (the MRAB Stroop task) was identified as a significant predictor for the Spontaneous, but not the Memorized, IRT difference scores. This finding is evident regardless of whether or not the Raven's scores are included in the equations.

The Memorized lie model estimated with the Raven's score excluded from the analyses explains a higher percentage of the variance than when the Raven's score is forced into the regression (and for the Spontaneous lie models, the Raven's simply doesn't contribute significantly to the variance), so we discuss the models that do not include the Raven's in further detail.

TABLE 3
AIC for IRT difference scores fitted against identified sets of MRAB results for Experiment 1

	IRT Difference scores	
MRAB results	*Memorized*	*Spontaneous*
Mental rotation RT, vigilance RT, cognitive set switching ER, spatial working memory RT	524	539
Spatial working memory RT, verbal working memory RT, vigilance ER, filtering RT	538	526

For the Memorized IRT difference scores, RT on the MRAB Mental Rotation task accounts for the highest portion of the variance explained by the model. Mental Rotation RT was not identified as a significant predictor of the Spontaneous IRT difference scores. However, contrary to what we might have expected, Mental Rotation RT is negatively associated with the Memorized IRT difference scores. That is, participants with worse mental imagery—as reflected in the Mental Rotation scores—had a larger negative difference between the amount of time required to produce lies as compared to the truth. Our speculation (and it is only that) is that the participants may have used mental imagery to provide truthful responses to the questions. Perhaps this mental imagery was not needed to produce the lies in the Memorized condition as those lies were well rehearsed. Along these lines, this mental imagery would have been necessary to produce both the truth and lies in the Spontaneous condition, which could explain why Mental Rotation RT did not significantly predict the difference between lies and truth in that condition. We stress that this is only a post-hoc speculation and whatever the reason for this result, the crucial finding is that a different set of MRAB tasks was identified for the Memorized IRT difference scores than for the Spontaneous IRT difference scores.

For the Spontaneous IRT difference scores, ER on the MRAB Vigilance task accounts for the highest portion of the variance explained by the model. Vigilance ER was not identified as a significant predictor of the Memorized IRT difference scores. Higher ER on this task is associated with larger differences between truth and lie IRTs, and thus Vigilance RT is related to the participants' responding faster with the truth or more slowly with lies. It is possible that vigilance reflects the participants' being poised for the cue, and ready to respond to the truth; if so, they would be faster if they were more vigilant to the cue and slower if they had to overcome such priming. But again, this is merely a post-hoc account.

We stress that the key finding here is that individual differences in different processes assessed by MRAB predicted the time to produce the different sorts of lies. This finding is *prima facie* evidence that different processes are used in the two conditions. Moreover, the fact that our version of the Stroop task (''Filtering'') predicted times only in the Spontaneous condition, as expected on the basis of the fMRI results, can be treated as a form of validation of the methodology.

EXPERIMENT 2

Experiment 1 served to show that a novel behavioral method could be used to document that different sorts of lies rely on different underlying processes. In this experiment we use this method to chart new territory, to ask whether lies could differ along another dimension. There is ample evidence that the ''self'' is represented differently in the brain than ''others.'' For example, words about the self are processed by different brain regions than words about others (Craik et al., 1999; Gusnard, Akbudak, Shulman, & Raichle, 2001; Kelley et al., 2002). Moreover, it is well known that the monkey brain has distinct regions for representing space around the self versus other regions of space (Berti & Frassinetti, 2000; Gross & Graziano, 1995).

Thus, in this experiment we used the same individual-differences approach to investigate whether there is a difference between lying about oneself versus lying about another person. If in fact the results document this distinction, we then can compare these findings to those from the previous experiment. It is possible that there are at least four systems that underlie deception. All of the responses in this study were spontaneous, but it is possible that spontaneous lies about facts regarding the self or others (which may be essentially semantic memories) both differ from spontaneous lies about episodic memories of events.

Method

Participants

Forty-two participants (22 females and 20 males) volunteered to participate for pay. Two participants (one female and one male) were removed from the sample for failing to meet the study prerequisites. Seven participants (five females and two males) were removed from the sample for having fewer than seven valid responses for at least one of the four conditions. A response given by a participant was considered valid if, for that question, the participant both adhered to the cue and knew the truthful response to the question. The final sample consisted of 16 females (mean age 21.7 years; range 18–33 years) and 17 males (mean age 23.6 years; range 18–35 years). Eleven females and 12 males were Harvard students, and the remaining participants were from the community or other local colleges. None of these people had participated in the previous experiment. The mean score on the Bors and Stokes (1998) short form of the Raven's APM for the females was 8.63 (range 1–12) and for the males was 8.06 (4–12). The mean number of years of education for the females was 14.25 years (range 12–21 years) and for the males was 14.06 years (range 12–18 years); 12 females and 16 males were right-handed. The inclusion criteria noted in Experiment 1 were also applied in Experiment 2.

Materials

We prepared two sets of 24 questions; one set required a single-word response about oneself (used in the ''Self'' condition), and the other set required a single-word response about President George W. Bush (used in the ''Other'' condition). President Bush was chosen to be the subject of the ''Other'' questions because we assumed that most participants would know a wide range of information about him. The questions asked the participants about factual conditions. For example, two questions were ''What is your home state?'' and ''How many children does Bush have?'' The questions were paired in that, for each fact, participants were asked both about themselves and about President Bush. For example, participants were asked both ''What is your job title?'' and ''What is Bush's job title?''.

However, once the study had been completed but before performing the data analysis, we decided that five of the ''Self'' questions did not strictly ask about the participants themselves. We eliminated these questions because each of these items asked the participants to provide information not strictly about themselves but about a person close to them (mother, father, significant other, co-workers, siblings). Upon further reflection, we realized that these questions could have been construed by the participants as querying an ''other'' (e.g., a participant could have interpreted the question ''What is your father's first name?'' not as a question about themselves but as a question about another person, namely their father). However, this was ambiguous, and thus we could not reclassify these five questions as ''Other'' questions because some participants may have actually viewed them as ''Self'' questions. Because this ambiguity was not present for the ''Other'' questions, we removed the responses to these five questions for each participant, but not for the five corresponding ''Other'' questions.

Procedure

The procedure was the same as in Experiment 1, except as noted below. After the participant completed informed consent, we administered the MRAB, as in Experiment 1. As before, the investigator remained in the room, behind and out of sight of the participant, for the entire study and answered any questions the participants may have had.

Once the participants completed the MRAB, we presented the deception task with the PsyScope software on a Macintosh computer, beginning with instructions to the participants that they would now answer questions about either themselves or another person with either false or truthful responses. For each trial, the computer screen displayed a cue that instructed the participant how to answer the question that would follow the cue. The cue was either the letter ''T,'' indicating that the participant was to answer the following question with a truthful response, or the letter ''F,'' indicating that the participant was to answer the question with a false response. The four conditions, Self/Truth, Self/Lie, Other/Truth, and Other/Lie, were intermingled. The participants were presented with eight practice questions, two for each of the four conditions. Following the practice trials, the investigator asked the participants whether they had any questions. Once the investigator was satisfied that the participants understood the procedure, the participants were presented with 48 questions, 12 for each of the four conditions. Questions were presented in a random order, except that no more than three in a row could refer to the same person (self vs. other) and no more than five in a row could have the same type of response cue. Participants were asked to respond to each question with a single word.

After the cue had remained on the screen for 4 s, a question was auditorially presented by the computer. Upon presentation of the question, the IRT of the participant's spoken reply was registered by a microphone attached to the computer. Once the participant had responded, a fixation point appeared on the screen to signal that the next cue would appear shortly. The fixation point remained on the screen for 1 s.

Once all of the questions had been presented, participants were asked to rate President Bush on a 10 point scale, with ''1'' representing the viewpoint ''I strongly dislike Bush'' and ''10'' representing the viewpoint ''I strongly like Bush.'' We asked the participants to rate President Bush in order to ensure that our sample did not overwhelmingly differ from the general population in their opinion of President Bush and that emotionality did not play a strong role in their response times. The mean rating for President Bush was 3.95 points (range 1–10 points).

In order to induce the participants actually to lie (that is, to attempt knowingly to deceive someone else), we used the same cover story as in Experiment 1: We told them that they would be paid according to how convincing their lies were, as assessed by an observer in another room (who was supposedly monitoring their face, via the video camera mounted on the computer). In addition, participants were told that the observer would not know which question they were answering. This instruction was given to ensure that participants would not have to be concerned that the observer's knowledge of President Bush would interfere with the success of their deceptive responses.

As before, participants were told that they would win or lose money depending on how convincingly they lied; however, in this study participants were told that every time the observer detected a lie, the participant would lose 25 cents from the hourly wage, and every time the participant succeeded in deceiving the observer, he or she would gain 25 cents. In fact, all participants were paid the full amount as if they had successfully deceived the observer.

Following the computer-administered task, we asked the participants to give us truthful responses to all of the questions they had just been asked. We did this in order to determine whether the participants had correctly followed the cues, which allowed us to compute error rates for each participant in each condition. In addition, we asked the participants to indicate questions for which they did not know the correct response. If a participant did not respond correctly to the cue for a question, or did not know the truthful response, we removed his or her IRT for that question from the sample.

Finally, the investigator then administered the short-form APM for 15 min, debriefed the participant (as in Experiment 1) and paid the maximal possible amount.

Results

For each of the four conditions (Self/Truth, Self/Lie, Other/Truth, Other/Lie), we calculated difference scores for each participant by subtracting the mean IRT for valid truthful responses from the mean IRT for valid deceptive responses.

We began with a two-way analysis of variance, comparing the IRT difference scores for Self versus Other condition for men versus women, and found no effect of gender, $F < 1$; we also did not find any interaction between gender and condition (Wilks' $\lambda = .93$, $F = 2.16$, $p = .15$). Thus, we pooled over gender in all subsequent analyses. Participants took more time to respond to questions in the Other condition, mean: 1358 ms, than in the Self condition, mean: 1167 ms, $t(32) = 2.04$, $p < .005$. Moreover, we did not find a relationship between participants' ratings of Bush and either their mean IRT for the Other condition ($r = -.14$, $p > .1$) or their Other IRT difference scores ($r = .09$, $p > .1$). However, we now found that the participants did in fact take longer to respond with a lie than the truth, and found this in both the Self, lie mean: 1244; truth mean: 1099, $t(32) = 3.05$, $p < .005$, and Other, lie mean: 1441; truth mean: 1281, $t(32) = 2.43$, $p = .02$), conditions.

Next, for each of the nine MRAB tasks, we calculated each participant's RT and ER. We hypothesize not only that these cognitive measures would predict performance on deception tasks, but also that different subsets of these measures would predict performance for different types of lies. And as with Experiment 1, we also included each participant's score on the short-form APM as a means of controlling for general intelligence.

We again used forward stepwise regression analyses to examine the independent contributions of the different MRAB scores, if any, for predicting the IRT difference scores. As in Experiment 1, the mean Raven's score for the participants is again significantly higher ($p < .005$) than the norms reported in Bors and Stokes (1998). Therefore we again performed all analyses both with and without using the Raven's score as a predictor variable. For the first of the two stepwise regressions performed on each condition, we forced the Raven's score in as a predictor as a means of controlling for general intelligence; the second pair of regressions omitted the Raven's score altogether as a independent variable, and only the MRAB scores listed above were used. The results of the stepwise regressions are summarized in Tables 4 and 5.

First, we note that the results change somewhat depending on whether the Raven's score is forced

TABLE 4
Stepwise regression analysis for Experiment 2, with IRT difference scores as criteria and MRAB scores as predictors: Raven's scores forced in at the outset

Criteria	*Predictors*	*β*	*Partial R^2*	*Model R^2*	*F*
Self	Raven's score	−0.37	.17	.17	6.42**
	Spatial working memory RT	0.48	.23	.40	11.45****
Other	Raven's score	−0.23	.00	.00	.02
	Spatial working memory RT	0.37	.24	.24	9.32****
	Mental rotation ER	−0.50	.14	.38	6.79**
	Vigilance RT	0.47	.10	.48	5.59**
	Spatial working memory ER	−0.38	.11	.59	7.23**

Notes: *p <.10; **p <.05; ***p <.01; ****p <.005.

TABLE 5
Stepwise regression analysis for Experiment 2, with IRT difference scores as criteria and MRAB scores as predictors: Raven's scores not included

Criteria	*Predictors*	*β*	*Partial* R^2	*Model* R^2	*F*
Self	Spatial working memory RT	0.47	.26	.26	11.06****
	Vigilance RT	0.51	.08	.34	3.67*
	Perceptual reaction time RT	−0.38	.09	.43	4.67**
Other	Spatial working memory RT	0.36	.24	.24	9.64****
	Mental rotation ER	−0.41	.11	.35	5.09**
	Vigilance RT	0.52	.13	.48	7.18**
	Spatial working memory ER	−0.32	.08	.56	5.34**

Notes: *$p < .10$; **$p < .05$; ***$p < .01$; ****$p < .005$.

into or excluded from the analysis. However, in either case our fundamental point is made: again, we see that different MRAB results predict the different IRT difference scores.

In order to document further that different sets of MRAB results predict different types of lies, we asked whether either of the sets of predictors identified by the stepwise regression analysis could predict both types of lies successfully. Each IRT difference score was fit to both sets of MRAB results with the Raven's score excluded from the analyses. To account for the fact that the two types of lies have differing numbers of MRAB predictors, for each fit we calculated the AIC. The results are presented in Table 6. The set of predictors identified by the stepwise regression in the Self condition (AIC = 454) explains the IRT difference scores for lies about the self better than the set of predictors identified in the Other condition (AIC = 460). Furthermore, the set of predictors identified in the Other condition (AIC = 471) explains the IRT difference scores for lies about President Bush better than the set of predictors identified in the Self condition (AIC = 480). As in Experiment 1, we find that no single set of MRAB scores better predicts both types of lies.

TABLE 6
AIC for IRT difference scores fitted against identified sets of MRAB results for Experiment 2

	IRT difference scores	
MRAB scores	*Self*	*Other*
Spatial working memory RT, Vigilance RT, Perceptual reaction time RT	454	480
Spatial working memory RT, Mental rotation ER, Vigilance RT, Spatial working memory ER	460	471

Discussion

The stepwise multiple regressions where the Raven's scores were forced into the analysis at the outset (thereby removing the contribution of general intelligence) accounted for a larger proportion of overall variance, and thus we will focus on these analyses. We found that three MRAB scores predicted the Other, but not Self, IRT difference scores: ER for the Mental Rotation task, RT for the Vigilance attention task, and ER for the Spatial Working Memory task. For these regression analyses, no MRAB score uniquely predicted the Self IRT difference scores; Spatial Working Memory RT was the only MRAB score that predicted Self lies, but this measure also predicted Other lies. (In contrast, when the Raven's scores were excluded from the regressions, ER on the Perceptual/Motor Control task was a significant predictor for the Self, but not Other, IRT difference scores.) But again, we stress that the important result is that different sets of variables predicted performance for the two types of lies. Moreover, the sets of predictors were not only distinct for the two conditions in this study, but also were distinct from those in Experiment 1.

One puzzle here, however, is that the Filtering task did not predict performance, even though all lies were spontaneous. One possible reason is that knowledge about the self, and even President Bush, is much better learned than knowledge about the work or vacation incident. If so, then participants may not have had as much conflict between the reality and lie, because the two were so easy to distinguish. Furthermore, the spontaneous lies told in this experiment may not have been directly comparable to the spontaneous lies told in Experiment 1. The participants in Experiment 1 were told that their spontaneous lies had to fit into a coherent story. No such instruction was given for Experiment 2. When lying in Experiment 1, participants may have had to inhibit false responses that were not consistent with each other, and such inhibition would not have been necessary in Experiment 2.

It would be useful to conduct the present study using fMRI. If the present interpretation is correct, we would not expect to find anterior cingulate activation in this task.

GENERAL DISCUSSION

An individual's specific basic cognitive strengths and weaknesses should predict his or her performance on various complex tasks that draw on these abilities. For example, performance on a simple mental rotation task should predict performance on a more complicated procedure that requires the use of mental imagery transformations. Using the MRAB, we collected information about individual differences in a wide range of cognitive abilities. We found that these measures did indeed predict performance in four types of deception tasks: producing a memorized set of lies about a life episode; spontaneously creating a coherent set of lies about a life episode; telling a set of isolated lies involving self-knowledge; and telling a set of isolated lies involving knowledge about another person. The fact that variations in performance of different MRAB tasks predicted the time required to produce different types of deception provides evidence that different types of lies arise from distinct cognitive processes.

The present study provided convergent evidence for the conclusions reached by Ganis et al. (2003) in their fMRI study of isolated spontaneous versus memorized coherent lies. In addition, we extended the logic of that study to examine lies about the self versus others—and again found evidence that distinct processes underlie the two types of lies. We also note that lies may differ on more dimensions than the ones considered here. Differences in the subject of the lie may cause different areas of the brain to be accessed; for example, telling a lie about the distant past may require different cognitive processes than telling a lie about more recent events. In addition, lies can differ not only in subject but also in degree. Telling a lie that is a slight exaggeration of the truth may be a cognitively different task than creating a lie that is a complete fabrication.

The lies told in the studies reported here involved unemotional topics; participants were asked to lie in response to questions such as "What is your name?" and "What is your home state?". It is plausible that lying about an emotional topic requires more cognitive effort than lying about a neutral topic. When attempting to deceive another regarding an emotional topic, often the deceiver has to conceal not only the truth, but also the emotions associated with the truth. For example, a man who is trying to deceive his spouse into believing that he is not having an extramarital affair will have to conceal not only the truth about his whereabouts but also any guilt he may have over the affair.

Furthermore, the context surrounding deception may be relevant for the cognitive processes that are evoked. Lies may be told to avoid serious consequences as opposed to being told in a situation where there are only minor consequences to being discovered. Finally, the deceiver's attitude towards the lie may be important. In some situations the deceiver may view telling a lie as being more moral than telling the truth, for example when a lie is told to avoid hurting someone's feelings. This situation may be cognitively different from one in which the deceiver feels that the lie is morally wrong. And so on. These other types of deception may also depend on at least some distinct underlying processes.

Current methods of detecting deception rely on the implicit assumption that deception is a homogenous process with consistent peripheral markers. These methods aim to identify one or more of these peripheral "hallmarks" of deception. One way to circumvent difficulties in recording from the periphery has been to turn to the brain. However, the results of these studies suggest that lying is not a single activity. The fact that successful production of different types of lies depends on different cognitive processes suggests that different forms of deception may be marked not only by different peripheral signals, but also by different patterns of neural activity (Ganis et al., 2003). If so, then an effective, reliable method of deception detection will need to take account of the many dimensions along which lies can vary and the differing cognitive processes that create those lies.

REFERENCES

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723.

Berti, A., & Frassinetti, F. (2000). When far becomes near: Remapping of space by tool use. *Journal of Cognitive Neuroscience*, *12*, 415–420.

Bors, D. A., & Stokes, T. L. (1998). Raven's advanced progressive matrices: Norms for first-year university students and the development of a short form. *Educational and Psychological Measurement*, *58*, 382–398.

Burgoon J. K., & Buller, D. B. (1994). Interpersonal deception: III. Effects of deceit on perceived communication and nonverbal behavior dynamics. *Journal of Nonverbal Behavior*, *18*, 155–184.

Bush, G., Whalen, P. J., Rosen, B. R., Jenike, M. A., McInerney, S. C., & Rauch, S. L. (1998). The counting Stroop: An interference task specialized for functional neuroimaging—validation study with functional MRI. *Human Brain Mapping*, *6*, 270–282.

Cody, M. J., Marston, P. J., & Foster, M. (1984). Deception: Paralinguistic and verbal leakage. In R. N. Bostrom (Ed.), *Communication Yearbook 8* (pp. 464–490). Beverly Hills, CA: Sage.

Cohen, J., MacWhinney, B., Flatt, M. & Provost, J. (1993). PsyScope: An interactive graphical system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, and Computers*, *25*, 257–271.

Craik, F. I. M., Moroz, T. M., Moscovitch, M., Stuss, D. T., Winocur, G., Tulving, E., et al. (1999). In search of the self: A positron emission tomography study. *Psychological Science*, *10*, 26–34.

deTurck, M. A. & Miller, G. R. (1985). Deception and arousal: Isolating the behavioral correlates of deception. *Human Communication Research*, *12*, 181–201.

Ekman, P. (1992). *Telling lies.* New York: WW Norton.

Ekman, P. (2001). *Telling lies: clues to deceit in the marketplace, politics, and marriage*. New York: WW Norton.

Farrow, T. F. D., Reilly, R., Towhida, A. R., Herford, A. E., Woodruff, P. W. R., & Spence, S. A. (2003). Sex and personality traits influence the difference between time taken to tell the truth or lie. *Perceptual and Motor Skills*, *97*, 451–460.

Ganis, G., Kosslyn, S. M., Stose, S., Thompson, W. L., & Yurgelun-Todd, D. A. (2003). Neural correlates of different types of deception: An fMRI investigation. *Cerebral Cortex*, *13*, 830–836.

Greene, J. O., O'Hair, H. D., Cody, M. J., & Yen, C. (1985). Planning and control of behavior during deception. *Human Communication Research*, *11*, 335–364.

Gross, C. G., & Graziano, M. S. A. (1995). Multiple representations of space in the brain. *The Neuroscientist*, *1*, 43–50.

Gusnard, D. A., Akbudak, E., Shulman, G. L., & Raichle, M. E. (2001). Medial prefrontal cortex and self-referential mental activity: Relation to a default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, *98*, 4259–4264.

Horvath, F., Jayne, B., & Buckley, J. (1994). Differentiation of truthful and deceptive criminal suspects in behavior analysis interviews. *Journal of Forensic Sciences*, *39*, 793–807.

Iacono, W. G. (2000). The detection of deception. In J. T. Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds.), *Handboook of psychophysiology* (2nd ed., pp. 772–793). Cambridge, UK: Cambridge University Press.

Iacono, W. G., & Lykken, D. T. (1997). The validity of the lie detector: Two surveys of scientific opinion. *Journal of Applied Psychology*, *82*, 426–433.

Johnson, S. C., Baxter, L. C., Wilder, L. S., Pipe, J. G., Heiserman, J. E., Prigatano, G. P. (2002). Neural correlates of self-reflection. *Brain*, *125*, 1808–1814.

Kelley, W. M., Macrae, C. N., Wyland, C. L., Caglar, S., Inati, S. & Heatherton, T. F. (2002). Finding the self? An event-related fMRI study. *Journal of Cognitive Neuroscience*, *14*, 785–794.

Kozel, F. A., Revell, L. J., Lorberbaum, J. P., Shastri, A., Elhai, J. D., Horner, M. D., et al. (2004). A pilot study of functional magnetic resonance imaging brain correlates of deception in healthy young men. *Journal of Neuropsychiatry and Clinical Neuroscience*, *16*, 295–305.

Langleben, D. D., Schroeder, L., Maldjian, J. A., Gur, R. C., McDonald, S., Ragland, J. D., et al. (2002). Brain activity during simulated deception: An event-related functional magnetic resonance study. *NeuroImage*, *15*, 727–732.

McCornack, S. A. (1997). The generation of deceptive messages: Laying the groundwork for a viable theory of interpersonal deception. In J. O. Greene (Ed.), *Message production: Advances in communication theory* (pp. 91–126). Mahwah, NJ: Lawrence Erlbaum Associates.

Mehrabian, A. (1971). Nonverbal betrayal of feeling. *Journal of Experimental Research in Personality*, *5*, 64–73.

National Research Council. (2003). *The Polygraph and Lie Detection*. Committee to Review the Scientific Evidence on the Polygraph. Division of Behavioral and Social Sciences and Education. Washington, DC: The National Academies Press.

Seymour, T. L., Seifert, C. M., Shafto, M. G., & Mosmann, A. L. (2000). Using response time measures to assess "guilty knowledge". *Journal of Applied Psychology*, *85*, 30–37.

Shephard, J. M., & Kosslyn, S. M. (2005). The MiniCog Rapid Assessment Battery: Developing a "blood pressure cuff for the mind". *Aviation, Space and Environmental Medicine*, *76*(Suppl. 6), B192–B197.

Sporer, S. L., and Schwandt, B. (2006). Paraverbal indicators of deception: A meta-analytic synthesis. *Applied Cognitive Psychology*, *20*, 421–446.

Stiff, J., Corman, S., Krizek, B., & Snider, E. (1994). Individual differences and changes in nonverbal behavior: Unmasking the changing faces of deception. *Communication Research*, *21*, 555–581.

Vrij, A. (1994). The impact of information and setting on detection of deception by police detectives. *Journal of Nonverbal Behavior*, *18*, 117–136.

Vrij, A. (2008). *Detecting lies and deceit: Pitfalls and opportunities* (2nd ed.). Chichester, UK: Wiley.
Zuckerman, M., DeFrank, R. S., Hall, J. A., Larrance, D. T., & Rosenthal, R. (1979). Facial and vocal cues of deception and honesty. *Journal of Experimental Social Psychology*, *18*, 378–396.

APPENDIX

Here we outline the tasks used in the MiniCog Rapid Assessment Battery. Each of the nine tasks in the battery relies on response times and error rates to assess a particular type of cognitive ability. Together these tasks provide a profile of the participant's attention, working memory, reasoning and perceptual and motor control skills.

Attention

The MRAB contains three tasks that are designed to assess differing types of attention.

Vigilance

The Vigilance task is aimed at assessing an individual's ability to maintain concentration while waiting for a specific event to occur. During this task, the user is given a target stimulus and instructed to press a button whenever that stimulus appears amid other stimuli.

Filtering

The Filtering task is a version of the classic Stroop task. The user is presented with a digit repeated on the Palm Pilot screen multiple times. Users are asked to respond according to how many times the digit appears on the screen (not to the meaning of the digit itself). To perform well on this task, the user must ignore, or filter out, the meaning of the digit. The task contains both "congruent" (i.e., the number of times a digit appears is the same as the meaning of the digit) and "incongruent" (i.e., the number of repetitions is not the same as the meaning of the repeated digit) trials.

Divided attention

The Divided Attention task assesses an individual's ability to focus on two unrelated features of stimuli at once. In this task, the user must attend to both the colors and shapes of stimuli. For example, the user might be asked to press one button if presented with a stimulus that is either black or a triangle (or both) and another button if the stimulus is either white or a circle (or both).

Working memory

The MRAB is designed to assess two types of working memory: verbal working memory and spatial working memory.

Verbal working memory

The Verbal Working Memory task assesses an individual's ability to remember and manipulate verbal information. In this task, the user is presented with a sequence of digits, one at a time. For each digit, the user must determine whether the current digit is the same as one presented two digits earlier in the sequence.

Spatial working memory

The Spatial Working Memory task is designed to assess an individual's ability to remember and manipulate spatial information. In this task, the user is presented with a sequence of digits appearing one at a time in various locations on the screen (the digit itself is otherwise unchanging). For each location, the user must determine whether the current location is the same as the one presented two stimuli earlier.

Reasoning

The MRAB contains three tasks designed to assess reasoning abilities.

Cognitive set switching

The Cognitive Set Switching task assesses an individual's ability to switch between different cognitive sets. On each trial of this task, the participant is presented with four letters. One of the letters differs from the other three in either its case or identity, and the user must identify which letter does not belong. The same criterion is used for a sequence of trials and then changes after a random number of trials, at which point the user must adapt to using the new criterion to determine which letter is out of place.

Three term reasoning

The Three Term Reasoning task assesses an individual's verbal reasoning skills. The user is presented with the first two terms (e.g., ''Mary is not shorter than Katherine; Mary is not as tall as Anne'') and asked to answer true or false to the third term, which is a question about the implications of the first two terms (e.g., ''True or False: Anne is the tallest'').

Mental rotation

The Mental Rotation task assesses an individual's spatial reasoning skills. In this task, the user is given a pair of stimuli, with one stimulus rotated relative to the other. The user is asked to determine whether the stimuli are identical or mirror images of each other. To perform this task, the user must be able to mentally rotate the images in order to compare them.

Perceptual and motor control

Finally, the MRAB contains a task to assess an individual's reaction time.

Perceptual/motor reaction time

In the trials of the Perceptual/Motor Reaction Time task, a stimulus appears randomly above one of the four Palm Pilot buttons. The user is asked to press that button as quickly as possible.