Lies in Conversation:
An Examination of Deception Using Automated Linguistic Analysis

Jeffrey T. Hancock (jeff.hancock@cornell.edu)
Department of Communication, Cornell University
320 Kennedy Hall, Ithaca, NY 14850 USA

Lauren E. Curry (lec26@cornell.edu)
Department of Communication, Cornell University
320 Kennedy Hall, Ithaca, NY 14850 USA

Saurabh Goorha (sg278@cornell.edu)
Department of Communication, Cornell University
320 Kennedy Hall, Ithaca, NY 14850 USA

Michael T. Woodworth (mwoodwor@dal.ca)
Department of Psychology, Dalhousie University
Life Sciences Center, Halifax, NS, B3K 1L4, Canada

Abstract

The present study investigated changes in both the sender's and the receiver's linguistic style across truthful and deceptive dyadic communication. A computer-based analysis of 242 transcripts revealed that senders used more words overall, increased references to others, and used more sense-based descriptions (e.g., seeing, touching) when lying as compared to telling the truth. Receivers naïve to the deception manipulation produced more words and sense terms, and asked more questions with shorter sentences when they were being lied to than when they were being told the truth. These findings are discussed in terms of their implications for linguistic style matching.

Introduction

Maxims such as "honesty is the best policy" and "let the truth be told" reinforce the notion that telling the truth is the best way to communicate. When telling everyday lies, then, deceivers must be careful to assume a position of sincerity in order to make their partners believe them and avoid being viewed in a negative light. In fact, this feat might not be very difficult to accomplish. Previous research suggests that it is quite difficult to catch a liar as deception detection rates in many experiments are not much better than chance (Vrij, 2000).

In general, there are three methods for trying to detect deceit. The first method focuses on vocalic and physical nonverbal behaviors (e.g., movements, smiles, voice pitch, speech rate, stuttering, and eye gaze) (Vrij, 2000). The second method involves measuring physiological responses with various technologies, such as polygraph machines (Vrij, Edward, Roberts, & Bull, 2000).

The third method is concerned with the content of what is said (e.g., verbal behavior, as well as a study of linguistic properties of liars' texts). For example, previous research suggests that liars tend to make less sense and tell less plausible stories (e.g., making discrepant and ambivalent statements), among other verbal characteristics (for review, see DePaulo, Lindsay, Malone, Mulenbruck, Charlton, & Cooper, 2003).

The present study employs automated linguistic analysis, in which a computer program is used to analyze the linguistic properties of texts, to examine the verbal content of deceptive and truthful conversations. As Pennebaker, Mehl, and Niederhoffer (2003) note, words used in daily interactions reveal both psychological and social aspects of peoples' worlds. Certain words and parts of speech can be markers of emotional, psychological, and cognitive states. Given that deceiving others likely involves changes in emotional or psychological states, linguistic cues detected using automated techniques may indicate lying in conversation.

Linguistic Indicators of Deception

A review of the relatively small literature concerned with automated linguistic analyses of deception indicates that, to date, at least four main types of linguistic cues have been associated with deception: 1) word counts 2) pronoun usage, 3) words pertaining to feelings and the senses, and 4) exclusive terms (Burgoon, Buller, Floyd, & Grandpre, 1996; Burgoon, Bliar, Qin, & Nunamaker, 2003; Newman, Pennebaker, Berry, & Richards, 2003; Pennebaker et al., 2003).

Consider first differences in word counts across deceitful and truthful messages. Previous studies have found that senders offer fewer details when lying than when telling the truth (Burgoon et al., 2003; DePaulo et al., 2003; Vrij, 2000). Senders may offer fewer details because they are less familiar with what they are discussing, or because they are trying to avoid providing details that may be inconsistent

with their fabrication. As such, senders may be expected that deceptive interactions would be characterized by fewer words on the part of the sender.

With regard to pronoun usage, Newman et al. (2003) observed that individuals consistently used first person singular pronouns less frequently when lying than when telling the truth. Using first person pronoun words such as "I," "me,"or "my" involves taking ownership of a statement, and deceivers may refrain from using these first person pronouns due to either a lack of personal experience or a desire to dissociate themselves from the lie being told. The findings regarding the use of second and third person pronouns are less consistent. Some studies have found that liars are less likely to use second and third person pronouns (Newman et al., 2003) while other studies have found that liars in fact use more second and third person pronouns (Ickes, Reidhead, & Patterson, 1986). According to Ickes et al. (1986), senders who are careful about constructing deceptive messages will exhibit an increased other-focus and therefore a higher use of second and third person pronouns. Finally, DePaulo et al. (2003) also found that liars are more likely to use third person pronouns in their deceptive interactions.

Research examining verbal cues associated with feelings and sense terms (e.g., see, touch, listen, etc.) suggests that deceivers tended to use more expressiveness, which includes both negative and positive forms of emotion, compared to truth-tellers (Burgoon et al., 2003). In addition, senders may be more likely to use sense words in an effort to create a detailed story to avoid eliciting skepticism from the deceiver (Burgoon et al., 2000).

Finally, previous research also suggests that liars use fewer exclusive words than truth-tellers (Newman et al., 2003). Exclusive words include prepositions and conjunctions such as "but," "except," "without," and "exclude." These words require a deceiver to discuss what is in a category and what is not. As such, deceivers may find it a more complex task to invent what was done versus what was not done (Newman et al., 2003).

Deception, Conversation and the Receiver

Although the literature on automated approaches to linguistic analysis of deception suggests that word counts, pronouns, feeling words, and exclusion words may predict deception, previous research is limited in two important ways. First, previous research has been limited primarily to analyses of deception in the context of monologues rather than in conversational contexts. For example, Newman et al. (2003) conducted five studies in which participants discussed a given topic by writing about it, talking about it to a video camera, or by typing their views on it. In these cases, only the liar's behavior was analyzed because there was no target of those lies present during the studies. Given that most lies tend to occur during conversations with others (DePaulo, Kashy, Kirkendol, & Epstein, 1996), and given the fact that language use in conversation differs in important ways from language use in monologues (Clark,

1996; Schober & Clark, 1989), the focus of previous research on monologue-based deception may limit its applicability to everyday conversation.

A second, and related, weakness is that previous research on linguistic predictors of deception has focused almost exclusively on the sender (i.e., the teller of the deception or the truth). For example, Newman et al. (2003) examined only a sender's handwriting, videotapes, and typed transcripts. In no case were the reactions of receivers (i.e., the targets of deceptive messages) studied. However, in conversations there is a reciprocal exchange between senders and receivers that can have important effects on deceptive behavior (Burgoon et al., 1996; Burgoon, Buller, & Floyd, 2001). As such, it may be important to look at both parties when examining interactions. If senders alter their behavior in systematic ways when lying versus when they are telling the truth, as previous research suggests, then an important question that remains to be addressed is whether receivers will also behave differently when lied to than when they are told the truth.

One possible outcome is that receivers will engage in linguistic style matching, which refers to the degree to which two people in conversation adjust their own speaking behavior, or style, to match their partners' behavior (Niederhoffer & Pennebaker, 2002). The observation that people vary their words on a turn-by-turn level when in conversations with others is assumed to reflect the coordination processes inherent in natural conversations (Grice, 1989; Niederhoffer & Pennebaker, 2002). Indeed, participants in conversations have been known to exhibit similar types of concurrent behaviors (both vocal and nonvocal), kinesics, proxemics, facial expressions, and word usage, regardless of topic content (Niederhoffer & Pennebaker, 2002).

If, as linguistic style matching suggests, people in conversation adjust their linguistic behavior to that of their partners, then any differences in linguistic behavior by senders across deceptive and truthful communication should also be observed in the receiver's behavior. That is, receiver's behavior should mirror the behavior of the sender in terms of word usage and linguistic variables across deceptive and truthful communication. If receivers engage in linguistic style matching during deceptive interactions, then receivers, like senders, should produce fewer words, fewer first person pronouns, more second and third pronouns, more exclusive words and negations, and more words pertaining to the senses.

Finally, there may also be receiver activities that do not simply match those of the sender. For example, if a receiver becomes suspicious of the sender's truthfulness, the receiver may probe their partner more frequently, perhaps by asking additional questions.

The present study examined the linguistic styles of senders and received engaged in truthful and deceitful conversations. The conversations were conducted in a text-based, computer-mediated setting, in which participants exchanged synchronous messages. A computer-mediated

communication setting was used in the present study for several reasons. The first is that the transcripts were created automatically as the participants interacted. The second is that, because the interaction was entirely text-based, all of the information exchanged by the participants during their interaction was captured in the transcripts (Hancock, in press).

Methods

Participants (*n* = 66) were upper-level students at a northeastern American university, and they participated for credit in various courses. Participants were randomly paired to form 33 same-sex, unacquainted dyads (15 male and 18 female).

Participants were recruited for a "study of how unacquainted individuals communicate about various conversation topics." Upon reporting to the laboratory, participants were led separately to remote rooms where they completed an initial set of forms, including informed consent.

The general procedure was adapted from Burgoon et al. (2001). All participants were told that they would be having a conversation with an unknown partner. They were instructed that they would discuss 5 topics, which were then provided to the participants on a sheet of paper. The first topic was always "When I am in a large group, I…" This initial topic was designed to allow the participants to become comfortable interacting with their partner, and was not included in any analyses. After this topic, participants began a discussion of the four experimental topics which included: "Discuss the most significant person in your life", "Talk about a mistake you made recently", "Describe the most unpleasant job you have ever had to do" and "Talk about responsibility." There was no time limit and participants were asked to discuss each topic until they had exhausted it and understood each other's responses.

One of the two participants was randomly assigned to the role of sender, and the other to the role of receiver. Senders were asked to sometimes deceive their partners. In particular, they were instructed "to NOT tell 'the truth, the whole truth, and nothing but the truth'" (Burgoon et al., 2001) on two topics, and to be truthful on the other two topics. The two topics in which the whole truth was not to be told were marked with an asterisk on the sheet of paper given to the sender.

Examples of lies were given to the senders, and it was emphasized that the senders should try to produce lies that were fairly substantial (e.g., saying that they went on a vacation when in fact they did not) rather than small lies (e.g., saying that they went on a vacation from August 4th to the 10th when they actually went from August 5th to the 11th). Senders had approximately five minutes to plan their stories. Receivers were blind to the deception manipulation and were told that they were going to have a conversation with another person and that their role was to keep the conversation going. The same list of topics in the same

order was given to the receivers but without any asterisks marking topics.

The sequence in which the topics were discussed, and the order in which the sender lied, was counterbalanced across 16 orders. After the initial ice-breaking topics, senders were instructed to lie on either the next two topics or on the last two topics. Half of the senders followed a truth-first, deception-second order. The remainder followed a reverse order. Because topics followed a diagram-balanced Latin square order within truth and within deception, all topics appeared within a given time period.

Participants discussed the topics in a text-based, computer-mediated setting and performed the task at isolated computer terminals. Participants used one of two desktop computer stations while the experimenter monitored and recorded the interaction from a third station. Once participants were seated at their terminals, the experimenter briefly demonstrated the use of the computer interface, Netmeeting, in which participants typed their message in a private composition window and hit enter to send their message to a shared chat window (see Figure 1). Message transmission was virtually instantaneous.

Once participants finished the discussion task, they were asked to complete a series of questionnaires based on their conversation. The data from these questionnaires are not reported here. After completing the post-interaction questionnaires, each member of the dyad was brought to a common room, and introduced to his or her partner and they were fully debriefed.

Figure 1. Screenshot of the Netmeeting Interface
(A= Sender, B= Receiver)
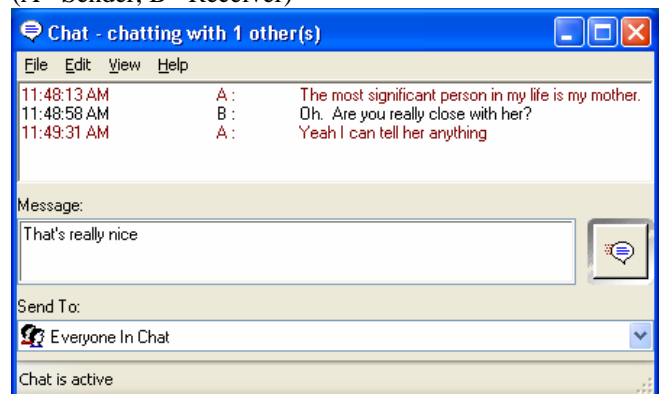
Automated Linguistic Analyses

Both sender and receiver transcripts were converted into separate text files separated by topic. Each dyad produced eight different transcript files: two deception discussions and two truthful discussions for each sender, and two deception discussions and two truthful discussions for each receiver, which produced a total of 264 transcripts.

All transcripts were analyzed using the Linguistic Inquiry and Word Count (LIWC) program (Pennebaker, Francis, & Booth, 2001). This text analysis program was used to create empirically derived statistical profiles of deceptive and truthful communications (Pennebaker et al., 2003), and it

Table 1. Means and (Standard Errors) of the linguistic output variables by role and truth condition.

	Sender		Receiver	
	Lie M (SE)	Truth M (SE)	Lie M (SE)	Truth M (SE)
Word Count	156.11 (17.07)	125.08 (11.20)	157.36 (16.56)	119.61 (10.96)
Words / sentence	10.20 (.97)	9.03 (.53)	8.21 (.42)	9.07 (.59)
1st Person Pronouns	8.01 (.35)	8.52 (.34)	8.08 (.45)	8.92 (.41)
2nd Person Pronouns	2.41 (.31)	2.82 (.32)	2.64 (.32)	2.25 (.22)
3rd Person Pronouns	3.30 (.33)	2.46 (.18)	2.57 (.31)	2.43 (.27)
Negations	2.19 (.21)	1.77 (.16)	2.27 (.19)	2.20 (.21)
Senses	2.47 (.16)	2.09 (.19)	2.49 (.18)	2.18 (.22)
Exclusive Words	4.01 (.27)	4.18 (.32)	3.63 (.22)	3.86 (.31)
Questions	15.88 (2.27)	16.39 (2.32)	15.33 (1.53)	10.84 (1.34)

Note: all statistics represent the percentage of total words in the transcript, with the exception of Word Count, Word per Sentence and Questions variables, which represent absolute totals.

has been used in studies to predict outcome measures like social judgments, personality, psychological adjustment, and health. LIWC analyzes transcripts on a word-by-word basis, including punctuation, and compares words against a file of words divided into 74 linguistic dimensions. For the purposes of this study, only variables relevant to the hypotheses or of potential interest to deception were included, which left 8 variables within the four categories mentioned above: word counts; pronouns; emotion words and words pertaining to the senses; and exclusive words and negations. In addition, question frequency was also analyzed.

LIWC produces the percentage of each variable type by dividing the frequency of the observed variable by the total number of words in the sample. Word counts were not reported as percentages, but as frequency totals.

Results

A 2 (discussion type: truthful vs. deceptive) x 2 (role: sender vs. receiver) repeated measure type General Linear Model (GLM) procedure was conducted on each dependent variable. Table 1 contains the descriptive statistics for each variable.

Overall, more words were produced during deceptive discussions than during truthful discussions $F (1,32) = 7.11$, $p < .05$. The increase in word count for deception was equivalent for both senders and receivers, $F (1,32) < 1$, ns, and no interaction was observed, $F (1,32) < 1$, ns, suggesting that both senders and receivers used more words when the sender was lying than when the sender was telling the truth.

An analysis of the number of words used per sentence revealed no main effect of discussion type or role. However, a significant interaction between discussion

type and role was observed, $F (1,32) = 4.07$, $p < .05$. Simple effects analyses conducted at each level of discussion type revealed that during truthful discussion senders and receivers produced the same number of words per sentence, $F (1,32) < 1$, ns. In contrast, during deceitful discussion, receivers used marginally fewer words per sentence than senders, $F (1,32) = 3.81$, $p = .06$. Considered together, these data suggest that receivers used shorter utterances when being lied to than when they were being told the truth, while senders used the same number of words per sentence regardless of discussion type.

The next set of analyses examined pronoun usage. No significant effects were observed for first person pronouns (e.g., "I," "we," "self"). Similarly, no effects were observed for the usage of second person pronouns (e.g., "you"). An analysis of third person pronouns referring to others (e.g., "he," "she," "they"), however, revealed a main effect of role, $F (1,32) = 4.68$, $p < .05$. Senders used third person pronouns more frequently than receivers. In addition, senders were significantly more likely to discuss others when lying as compared to when they were telling the truth, $F (1,32) = 4.57$, $p < .05$.

With regard to the production of exclusive words and negations, no reliable effects were observed. Regardless of discussion type, senders and receivers produced the same number of exclusive words and negations.

The next analyses examined the use of words that pertained to the senses (e.g., "see," "touch," "listen"). Participants used significantly more sense words during deceptive conversations than during truthful ones, $F(1,32) = 5.34$, $p < .05$. No effect of role was observed, $F(1,32) < 1$, $n.s.$, nor did role interact with discussion type, $F(1,32) < 1$, $n.s.$

The last analysis was concerned with the number of questions asked during the interactions. A main effect of discussion type was observed, $F(1,32) = 4.02$, $p < .05$. More questions were observed during deceptive discussions than during truthful discussions. This main effect, however, was qualified by a marginally reliable interaction between discussion type and role, $F(1,32) = 3.24$, $p = .08$. Simple effects analyses conducted at each level of role revealed that while senders asked the same number of questions across deceptive and truthful discussion types, $F(1,32) < 1$, *ns.* , receivers asked more questions during deceptive discussions than truthful ones, $F(1,15) = 9.58$, $p < .01$. Considered together, these data suggest that receivers were more likely to ask questions when they were being lied to than when they were being told the truth

Discussion

The primary objective of the present study was to examine the linguistic behaviors of both senders and receivers during dyadic communication that involved both deceptive and truthful discussions. The first question of interest was determining whether the senders' linguistic behavior changed when the sender was being deceptive relative to when the sender was being truthful. The data suggest that, overall, when senders were lying to their partners, they 1) produced more words, 2) used more "other" pronouns (e.g., "he," "she," "they"), and 3) used more terms that described the senses (e.g., "see," "hear," "feel") than when they were telling the truth.

In general, this linguistic profile is consistent with previous research suggesting that senders attempt to construct a more cohesive and detailed story in order to seem believable (Burgoon et al., 1996). For example, the increased number of words observed in the deceptive discussions may reflect the senders' attempts to convey a more complete story when attempting to deceive. Similarly, senders may have increased their use of sense words to enhance the believability of the deception (e.g., "He *saw* her do it."). Finally, the use of other-focused pronouns during deceptive discussions reveals the senders' attempts to shift the focus away from themselves (DePaulo et al., 2003; Ickes, 1986).

The present data, however, differs with previous research in several important respects. For example, previous research suggests that liars tend to use fewer words than truth tellers (Burgoon et al., 2003; DePaulo et al., 2003; Vrij, 2000). Why, then, did senders in the present study produce more words during deceptive discussions than during truthful discussions? One possibility is that the senders in the present study were engaged in conversation with a partner, whereas previous research has focused primarily on deception in monologue formats (e.g., Newman et al., 2003). It is possible, for instance, that senders engaged in conversation used more words in an effort to convince suspicious or skeptical receivers (e.g., Burgoon et al.,

2001). Indeed, receivers in the present study asked more questions when they were being lied to than when they were telling the truth (see below), which may have required senders to use more words to address the additional questions.

Similarly, previous linguistic analyses of deception suggest that senders use more negative emotion terms (e.g., Newman et al., 2003; Vrij, 2003). This difference is perhaps not surprising given the differences in discussion topics between the present study and the Newman et al. (2003) study. As noted above, Newman et al. asked participants to lie or tell the truth about highly emotional topics, such as abortion, which may have been more likely to elicit strong emotional verbal content than the more mundane topics employed in the present study (e.g., "Talk about a mistake you made recently.").

The second question of interest was whether the linguistic style of the receivers changed systematically according to whether or not their partners were lying. The data suggest that, in fact, receivers' linguistic profile changed across deceptive and truthful discussion topics. In particular, when being lied to, receivers 1) used more words, in shorter sentences, 2) used more sense words, and 3) asked more questions than when they were being lied to. These observations are particularly striking given the fact that receivers were blind to the deception manipulation.

These data provide relatively robust support for the linguistic style matching model (Niederhoffer & Pennebaker, 2002). First, receivers matched changes in the sender's total word production and use of sense words. Second, like the senders, receivers' use of emotion words and exclusion words did not change across deceptive and truthful conditions. Considered together, the present data suggest that receivers engaged in linguistic style matching.

There were, however, a number of linguistic variables on which receivers and senders diverged. While senders used more other pronouns when lying than when telling the truth, the receivers' use of other pronouns did not differ across discussion types. This observation may reveal the unique motivation of senders to distance themselves from their deception. Perhaps more importantly is the observation that receivers asked more questions and used fewer words per sentence when they were being lied to than when they were being told the truth. These surprising data suggest that the receivers were skeptical of the senders during deceptive conversations. Because senders did not produce more questions when they were being deceptive, the change in the receivers' question-asking behavior does not simply reflect linguistic style-matching. Instead, these data suggest that although receivers were not explicitly aware that their partner was lying to them (i.e., they were blind to the deception manipulation), they were implicitly aware that they were being lied to.

An important limitation of the present study, however, is that participants interacted in a text-based computer-mediated environment. An important question is how these verbal behaviors we observed in text-based conversations will be altered when nonverbal channels of communication are available. While additional research will be required to address this question, communication via the Internet is becoming increasingly ubiquitous. Indeed, millions of people use text-based forms of communication on a daily basis, and previous research suggests that people do tell lies during computer-mediated interactions, such as Email and Instant Messaging, although not as frequently as they do over the phone or in face-to-face contexts (Hancock, Thom-Santelli, & Ritchie, 2004). As such, the present data provide important insights into interpersonal deception in this new communication domain.

Finally, there may not necessarily be a classification of specific words to predict deception, as previous research by Pennebaker et al. (2003) and Newman et al. (2003) may suggest, but deception may be more reliably predicted by looking at the methods of constructing lies. The present research advances our understanding of how linguistic behavior changes according to the truthfulness of the discussion. Lies that take place during conversation tend to include more words, more other-directed pronouns, and more sense words than truths. Equally important, if a receiver is being lied to by someone who fits this linguistic profile, he or she may be more likely to use more overall words and sense terms, and to ask more questions.

Acknowledgments

References

Burgoon, J.K., Bliar, J.P., Qin, T., Nunamaker, J.F. (2003). Detecting deception through linguistic analysis. *Intelligence and Security Informatics, 2665,* 91-101.

Burgoon, J.K., Bonito, J.A., Bjorn, B., Ramirez, A., Dunbar, N.E., Miczo, N. (2000). Testing the interactivity model: Communication processes, partner assessments, and the quality of collaborative work. *Journal of Management Information Systems, 16 (3)*, 33-56.

Burgoon, J. K., Buller, D. B., & Floyd, K. (2001). Does participation affect deception success? A test of the interactivity principle. *Human Communication Research, 27,* 503-534.

Burgoon, J.K., Buller, D.B., Floyd, K., & Grandpre, J. (1996). Deceptive realities: Sender, receiver, and observer perspectives in deceptive conversations. *Communication Research, 23),* 724-748.

Clark, H. H. (1996). *Using language.* Cambridge, UK: Cambridge University Press.

DePaulo, B.M., Kashy, D.A., Kirkendol, S.E., & Epstein, J.A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology, 70,* 979-995.

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

Friedman, H.S., & Tucker, J.S. (1990). Language and deception. In H. Giles & W.P. Robinson (Eds.), *Handbook of language and social psychology.* New York: Wiley & Sons.

Grice, P. (1989). *Studies in the way of words.* Cambridge, MA: Harvard University Press.

Hancock, J.T. (in press). Verbal irony use in computer-mediated and face-to-face conversations. *Journal of Language and Social Psychology.*

Hancock, J.T., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: Effects of communication technology on lying behavior. In *Proceedings of CHI 2004.*

Ickes, W., Reidhead, S., Patterson, M. (1986). Machiavellianism and self-monitoring: As different as "me" and "you." *Social Cognition, 4,* 58-74.

Niederhoffer, K.G. & Pennebaker, J.W. (2002). Linguistic style matching in social interaction. *Journal of Language and Social Psychology, 21,* 337-360.

Newman, M.L., Pennebaker, J.W., Berry, D.S., & Richards, J.M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin, 29*, 665-675.

Pennebaker, J.W., Mehl, M.R., & Niederhoffer, K. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology, 54,* 547-577.

Schober, M. F., & Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, *21*, 211-232.

Vrij, A. (2000). *Detecting lies and deceit: The psychology of lying and the implications for professional practice.* Chichester, England: John Wiley & Sons.

Vrij, A., Edward, K., Roberts, K.P, & Bull, R. (2000). Detecting deceit via analysis of verbal and nonverbal behavior. *Journal of Nonverbal Behavior, 24,* 239-263.