

15. Deception in computer-mediated communication

Jeffrey T. Hancock and Amy Gonzales

1. Introduction

Why do people lie? The reasons are as varied as human life itself, but there is almost always a reason, and these reasons can be categorized in a variety of ways. St. Augustine, for example, classified lies into eight types of varying severity. Modern psychology has produced smaller taxonomies. In the influential self-presentational framework of deception (DePaulo, Lindsay, et al. 2003; Vrij 2008a), people lie to enhance or protect their own self-image (self-oriented), their partner's (partner-oriented), or some third party (altruistic). According to this view, most everyday conversational lies are part of an effort to manage interactions with others and achieve self-presentational goals (Goffman 1959). Other lies, like those associated with scandals or crime, may be less about self-presentation goals and more about seeking material rewards or avoiding prosecution, but lies are still a means to accomplishing something in the world.

Since lies are useful for accomplishing goals in the world, why not simply lie all the time? A whole host of factors constrain the use of deception. Some of these constraints are external, such as the social stigma associated with being caught lying. Because humans are social animals, it is important to have the trust of other humans. Getting caught lying can damage that trust (Möllering 2006) and even ultimately destroy the ability to communicate with one another (Grice 1989). Other constraints against lying are internal. Most people view themselves as honest, and this self-concept of an honest person can shackle how much one is prepared to lie. In a series of experiments, Ariely and colleagues (Mazar and Ariely 2006) demonstrated that even when there was no possibility of being caught, participants in their experiments only lied by a fraction of the full amount they could have lied, suggesting that the self-concept of honesty prevents people from lying too much.

In this view, the act of lying is a sort of calculus. Lying can be a tool for accomplishing things in the world, but it is restricted by internal and external constraints. The question addressed here is whether communication technologies, like telephones, email, and text messaging, enter into this calculus. And, once the decision to lie has been made, what are the implications for language?

The present chapter has two primary interests. The first is a concern with how communication and information technology affect the pragmatic calculus of deception, and what features of these new communication spaces may enter into this calculus. The focus is on everyday social interactions, which in the present era are

supported primarily by the telephone, email, and various forms of messaging (e.g., instant messaging, text messaging via mobile phones), although the chapter also briefly reviews research examining deception in more anonymous Internet spaces, including chat rooms and newsgroups. The central argument is that these communication spaces shape the pragmatics of deception by varying communication features and motivations for deception (Carlson et al. 2004).

The second interest is in the language people use once the decision to lie with technology has been made. Traditional deception research has examined the verbal cues of deception (Vrij 2008b), with mixed and sometimes contradictory results. However, because digital technologies leave traces of language that previously were evanescent (e.g., the contents of an email or a text message are automatically recorded), deception research has seen a resurgence of interest in verbal cues. In combination with the advances in text analysis tools from computer science that make analyzing and processing textual data extremely fast and efficient, these conversational records are providing material for researchers to make new discoveries about how deception is manifested in language. The second part of this chapter reviews this research.

In this chapter, digital deception is defined as the intentional control of information in a technologically-mediated message to create a false belief in the receiver of the message (Hancock 2007). While this definition is but one of many for defining deception, it contains two key attributes that are necessary to highlight. The first is the idea of a false belief, which means that deception leads a conversational partner to believe something that the speaker does not believe to be true. The second is the notion of intentionality. False statements must be intentional to be deceptive; otherwise they are simply mistakes.

There are, of course, many ways to deceive, and almost any kind of speech act can be used deceptively. Indeed, a number of theories and taxonomies lay out how different speech acts can lead to different types of deception. For example, Information Manipulation Theory (McCornack 1992; McCornack, Levine, Torres, and Campbell 1992) notes that violations of Grice's four maxims lead to different types of deception. The literature on deception and CMC unfortunately has not advanced to this point, with few exceptions (e.g., Twitchell, Nunamaker, and Burgoon 2004). Consequently, the chapter limits the analysis and review to the choice of using deception in mediated interactions and the linguistic implications of this decision.

2. To lie or not to lie

2.1. Everyday technology and everyday lies

The question of how people choose to use different media to accomplish different goals has been of interest to researchers since the earliest days of computer-mediated communication (CMC) research. One of the first theories in the area, Media

Richness Theory (MRT; Daft and Lengel 1984, 1986), was concerned with how managers might use different media to accomplish different tasks. The theory holds that people match the medium to the equivocality of the task they wish to accomplish. Low equivocality tasks are straightforward and simple, such as arranging a place and time to meet for lunch. High equivocality tasks are ambiguous and complex, such as doing a performance review or hiring or firing someone. According to the theory, managers should match low equivocal tasks to lean media and equivocal tasks to rich media.

Lying is relatively ambiguous and complex given that it is an attempt to mislead, making it an equivocal task (but see Bavelas, Black, Chovil, and Mullett's 1990 work on equivocality for a more nuanced description of equivocation in communication). Thus, people who have decided to engage in deception should choose a rich medium over a lean medium, preferring face-to-face to the telephone, and the telephone to email, for example. This makes sense when one considers that having a richer medium would allow the liar to track the partner's responses more carefully (Buller and Burgoon 1996).

The question of whether people lie more in one medium or another was first directly tested by psychologist Bella DePaulo and colleagues (DePaulo, Kashy, et al. 1996) in a large diary-based study that had people record all of their lies and social interactions for seven days. With this method, the researchers counted the number of times participants reported lying in a conversation and divided that by the total number of social interactions participants reported. Although the diary method has a number of limitations, especially the potential for self-report biases, this study provided the first data that allowed for comparisons of lying rates by the same individuals across media. Because the study was conducted in the early 1990s, before email and the web had reached critical mass, the media of interest were limited to face-to-face, the telephone, and letters (unfortunately there were too few letters reported to allow for cross-media comparisons).

Of these media, MRT would predict the most lying face-to-face, followed by the telephone, and lastly letters. However, DePaulo and colleagues had a different prediction than MRT. They noted that people report feeling temporary discomfort when telling a lie, presumably because lying is a negative social act and a breach of trust. They argued that communication technologies that allow for social distance between the liar and the partner should provide some relief for the discomfort they experienced when lying.

In fact, this is exactly what DePaulo et al. found when they examined their participants' diaries: People reported lying in a greater proportion of telephone calls than in face-to-face interactions. While this supported the idea that people lie more frequently in more socially-distant media, the results were the opposite of the MRT prediction.

Note that a common assumption of both of these arguments is that media vary along a single dimension, richness or social distance, which is assumed to affect

decisions about deception uniformly. A number of scholars with backgrounds in linguistics (Herring 2001, 2007) and cognitive psychology (Clark and Brennan 1991), however, have noted that communication technologies, like email, have many features that affect language processes; they do not just differ along a single dimension. Clark and Brennan's (1991) influential chapter on grounding in CMC provides one example of the complexity of the relationship between deception and choice of technology. This approach, in which specific features of a technology are expected to have specific effects on language use, has influenced thinking about how features of media shape the pragmatics of deception. In the next section, a model of deception describes three such features and makes predictions about how the degree of recordability, synchronicity, and co-presence influence the likelihood of deception in technology.

2.2. Feature-based approaches to deception and technology

The first and most obvious feature that might affect deception is the recordability of some forms of communication. Relative to face-to-face communication, in which words disappear as they are spoken, most textually-based communications leave some record. Email leaves perhaps the most durable record, with copies not only on servers but on the target's computer. Even instant messaging and text messages, which are impermanent, leave a trail. Indeed, a recent *New York Times* article noted that over half of divorce cases in New York state involved email or text messaging as the primary form of evidence. People's perceptions of the recordability of a medium should factor into the pragmatics of deception and media. The more recordable, the less likely one should use that medium to lie.

A less obvious feature is synchronicity. Research from psychology (e.g., DePaulo et al. 2003) suggests that many everyday lies are the small "white" lies told in conversation when an honest answer would be awkward, such as telling a friend that, in fact, his new hair cut is awful. Lying in conversation can be an important pragmatic choice when trying to protect one's own image (self-centered lies) or to avoid harming the conversation partner (other-centered lies) or some third party (altruistic lies) (Vrij 2008a). Media that allow for real-time synchronous communication should increase the times when these kinds of awkward situations may arise. In contrast, asynchronous media, like email, allow time to craft a response that need not rely on deception to avoid violating one's own or another's self-image.

The third feature initially considered in the Feature-Based Model was whether the communication was co-present. Being in the same physical space as a conversation partner is important for evidentiary reasons. When people are in the same physical space, options for deception are reduced, because lies about where one is, who one is with, and what one is doing are all impossible. In contrast, as soon as communication becomes distal, all of these deceptions are possible.

Taken together, the components of the Feature-Based Model predict that the telephone should involve more lies than face-to-face, which in turn should have more lies than email. Phone conversations are typically not recorded, they are synchronous, and they involve conversation at a distance. Face-to-face conversations are synchronous and not recorded, but they involve, by definition, communication in the same physical space. Email is recordable and asynchronous, although not (necessarily) co-present.

To examine these predictions, we (Hancock, Thom-Santelli, and Ritchie 2004a) replicated the DePaulo et al. (1996) diary study by asking students to record all of the social interactions and lies over seven days. In almost all respects the procedure in our study followed that of the DePaulo et al. study, with one important exception: We had our participants record not only face-to-face interactions and phone calls, but also emails, instant messaging, chat rooms, multi-user domains (MUDs) and newsgroups (too few participants recorded any chat room, MUD, or newsgroup conversations for analysis, so we exclude them here).

Following DePaulo et al., we calculated the lying rate for each medium as the proportion of lies per number of conversations in that medium for each participant. Our data replicated the DePaulo et al. finding that phone conversations involved significantly more lies than face-to-face conversation. The more interesting result, however, was the lying rate in email. Emails involved fewer lies than any other medium. This finding contrasted with the idea that social distance should increase deception. However, the pattern of data fit with the features expected to shape the pragmatics of deception: recordability, synchronicity, and co-presence.

Obviously, not all lies are the same. In addition to the three orientations of lies described above (self-oriented, partner-oriented, and altruistic), people can also lie about different things, and to different people. For instance, in DePaulo et al.'s (1996) influential deception taxonomy, lies can be about facts, feelings, opinions and explanations. Also, relationships based on genuine trust should involve fewer lies than more superficial relationships. Kashy and DePaulo (1998) found this to be the case: People reported lying less often in conversations with close friends and romantic partners than with strangers and acquaintances (see also Ennis, Vrij, and Chance 2008; Whitty and Carville 2008).

In a follow-up study that used the same diary procedure but asked more information about what and to whom people lie, we (Hancock et al. 2004b) reached three conclusions. First, the pattern of deception across media, with the most lies reported on the telephone and the fewest in email, was repeated. Second, technology affected the content of the lie. Phone conversations involved more action lies than any other medium, presumably because the phone (at the time of the study) was the most mobile technology and would allow lies about what people were doing. Face-to-face conversations had the most lies about one's feelings, presumably because synchronous face-to-face conversations are more likely to involve questions related to one's opinions (e.g., "What do you think of my new

shirt?"). Finally, email had the most explanation lies. The asynchronous nature of email allows more time for constructing a sound explanation than synchronous formats, and this extra time presumably allows the writer to construct a lie that can withstand the potential of later scrutiny associated with a highly recordable medium.

Third, the way technology interacted with deception and relationships was relatively straightforward. Lies to family and friends tended to take place most frequently via a mediated channel, while lies to strangers and acquaintances occurred primarily face-to-face. The data are consistent with previous research conducted only in face-to-face settings, suggesting that lies are more likely to be told to strangers and acquaintances (Ennis, Vrij, and Chance 2008; Kashy and DePaulo 1998). These data add the observation that when people lie to people close to them they tend to use technology to communicate the deception. Finally, one last interesting effect was that although students did not tell email lies often, the most common email lie was an explanation lie (e.g., "I couldn't finish my paper because my printer died") to their professors.

2.3. Scenario-based approaches to deception and technology

More recent studies have used scenario methods to examine how people might make decisions about using deception across different communication technologies. For example, in a recent study that directly examined the question of how technology enters the calculus of deception, Whitty and Carville (2008) provided participants with scenarios that presented a situation, in a variety of media, in which a lie could be a potential pragmatic choice (e.g., "You receive an email from a person you do not know well. Within the email they ask you if you think they look attractive. You do not think that they are attractive but you do not want to hurt their feelings so you email them back and tell them that they are attractive"). Participants indicated how likely it was that they would use that medium to tell that lie to either someone close (e.g., a friend) or distant to them (e.g., a stranger).

There were a number of interesting findings, with mixed results for the Feature-Based Model described above. First, consistent with the diary-based results, the phone tended to have more lies than face-to-face. This provides converging evidence across methods that the phone is the most lied-in medium. Also consistent with the results above, when lying to close friends and family people tended to use mediated channels (email, telephone) more than face-to-face. Whitty and Carville highlight an interesting connection with the flaming sometimes observed in CMC (Sproull and Kiesler 1986; Whitty and Carr 2006), in which interactants use excessively insulting and aggressive forms of talk, with this apparent willingness to tell the truth to strangers. It may be that whatever compels flaming with strangers also leads to an increased willingness to be frank and honest online.

There was, however, an important difference between the Feature-Based Model and the results reported above: Participants reported that they would lie in email more than face-to-face. Differences in methodology may underlie this effect. Diary-based studies require participants to record their actual behaviors, rather than what they believe they will do in a hypothetical situation, as required by a scenario-based method. Indeed, when we (Hancock et al. 2004a, b) ask our participants in diary studies where they think they lied the most, many report saying email, despite the fact that their diary shows they lied the least there. These results suggest that people do not have accurate insight into how they use different media for lying.

A second scenario-based study to address the question of media and technology examined the choices that managers would make when needing to communicate a lie (George and Carlson 2005). In this study, managers reacted to a scenario that involved either a severe or less severe deception to either a familiar or unfamiliar recipient by choosing which medium they would use to communicate the lie. Here the Feature-Based Model was directly compared with Media Richness Theory, with the same predictions described above. The data revealed that, regardless of severity of the lie or familiarity of the target, face-to-face was the manager's first choice in lying, followed by the telephone and email. These data are supportive of Media Richness Theory's prediction that participants will choose richer media for deception.

Although email was the least chosen medium for lying, consistent with our diary-based studies, this study was the only one in which the telephone was *not* the most lied-in medium. There are at least three possible explanations that stem from differences in methods and analytic approach. Perhaps the most important is that this procedure asked participants to choose explicitly *a priori* their medium for communicating the lie, which differs from both the diary method (DePaulo et al. 1996; Hancock et al. 2004a, b) and from the scenario procedure employed by Whitty and Carville. In Whitty and Carville's (2006) procedure, participants rated how likely they would be to lie in the medium in that scenario. In this case, the lie was specified for a specific medium in the scenario, and no media choices were offered. In diary studies, it is impossible to know whether lies in a given medium occur because the diarists explicitly and *a priori* chose that medium to conduct the lie, or whether the lie simply emerged *post hoc* while interacting in that medium. Thus, of all the methods presented, the method used by George and Carlson is the only one that assessed the *a priori* choice of media for lying. Under these circumstances, people report that they would hypothetically lie most often face-to-face.

A second difference is that unlike the previous studies, these participants were managers and not students. As such, the differences observed here could relate to how college students and professionals choose to lie differently. Indeed, DePaulo et al. (1996) found that students and non-students tended to lie somewhat differently: Students lied almost twice as often per day as non-students.

The third difference is a statistical one. In the diary studies, lying frequency is reported as a proportion of the total social interactions, rather than the absolute frequency of lies that take place in a medium, to control for differences in how often people use different media. To account for these differences across media, as well as across speakers (some people have many interactions, while others have very few), the rate of deception is calculated as a function of total social interactions. In the George and Carlson procedure, only the absolute frequency is available.

2.4. Anonymous communication spaces

Most everyday interactions involve communications between already acquainted people. With the exception of spam (mass emails by senders attempting to sell products, services, or a scam), emails usually are signed and come from people whose identities are known or easily discerned. Instant messaging primarily takes place between people who are in each other's contact lists, and even the phone now typically provides some identity information in the form of caller ID. In these interactions, forms of *identity-based deception* are much less likely (Hancock 2007).

There is a tremendous amount of concern that deception is rife in anonymous online spaces. One study (Caspi and Gorsky 2006) queried newsgroup users about how much deception they believed was in their online space. Of all the respondents, 73 % believed that deception was very widespread, although only a minority (29 %) reported that they themselves had lied. In an earlier study that examined people's self-reported practices in chat rooms, Whitty (2002) found that lying about oneself was quite common, especially for men. Over half of the respondents reported lying about age and occupation, with over a quarter lying about education and income.

These data, however, do not tell us whether people were choosing to lie more often in anonymous online spaces than in less-anonymous spaces such as face-to-face communication. One study by Cornwell and Lundgren (2001) did exactly that. In this study, the researchers compared how people misrepresented themselves in romantic relationships that formed on the Internet or face-to-face. Their findings revealed that in a few categories, such as age and physical appearance, participants reported lying marginally more online than in their offline relationships. Surprisingly, the overall lying rates were relatively consistent across the two types of communication spaces.

Cornwell and Lundgren took the analysis a step further by examining whether interpersonal factors, such as involvement in the relationship, played a role in deception rates across online and physical relationships. In fact, when involvement was included in their regression analysis, the communication space (online vs. physical) was not significant. Instead, the degree to which participants felt involved in the relationship significantly predicted lying. The more involved they felt in a relationship, the less participants reported lying in that relationship. The

results suggested that relationships in online spaces tended to be less involved than those in physical space.

Taken together, these data suggest that anonymous spaces can reduce involvement in romantic relationships, and that this can increase lying. A good deal of research, however, has indicated that interpersonal relationships can be as intimate and involved as face-to-face relationships (see Parks and Floyd 1996; Whitty and Carr 2006). If this is the case, then at least in romantic relationships, the nature of the relationship may be more important than the anonymous nature of the communication space. This would be similar to face-to-face research that suggests people are less likely to lie to closer intimates (Kashy and DePaulo 1998).

However, there are many other types of deception in anonymous spaces than misrepresenting oneself to a romantic partner. In one of the first treatments of deception online, Donath (1999) described several different types of deception, including identity concealment, category deception, trolls, and impersonation. Here we focus on gender switching, when people present a gender online that is different from their offline gender (Herring 2003).

Gender switching is of particular importance for this chapter because of its linguistic implications. For instance, several studies (reviewed in the next section) have examined whether people retain characteristics of their gendered language when acting as the other gender (Herring and Martinson 2004). Here, the question is whether (and to what extent) men and women take advantage of the lack of visual and vocal cues in text-based CMC to perform gender deception (Danet 1998). In self-reports, this type of deception appears to be quite common. For example, Whitty (2002) found that about one-quarter of her male respondents reported pretending to be a female at some point in public chat rooms. However, as with the everyday lies described in the previous section, the self-report data and observational data do not necessarily line up. In her review, Herring (2003) reports on her examination of Internet Relay Chat (IRC) channels and her observation that gender deception was not common in this anonymous space. Others have also observed that gender switching online is relatively rare, perhaps because pretending to be someone one is not is difficult to maintain over time (Roberts and Parks 1999).

2.5. Implications for deception in anonymous space

Clearly, because technology allows people to interact anonymously, it also enables deception. However, comparisons to date of how often deception is used as a pragmatic choice in anonymous versus acquainted conversation reveal that simply because deception is easier does not mean that deception is more likely to occur. In fact, in a recent survey on misbehaving online, Selwyn (2008) found that participants reported approximately the same level of lying in online and offline interactions. The best predictor of whether a participant lied in online communication was whether they reported lying in face-to-face communication.

As with everyday, identifiable interactions, deception in anonymous spaces appears to be strategic and motivated, although some of these motivations may be different from those in traditional face-to-face deception. In particular, misrepresenting one's attractiveness and concealing aspects of one's identity appear to be common in anonymous spaces, and although gender switching is more easily performed in anonymous CMC spaces than face-to-face, it occurs infrequently. Deception appears to be an important pragmatic resource for crafting a representation of the self or for experimenting with aspects of the self, such as gender. Although this is not unique to anonymous online spaces, they certainly appear to provide more opportunities to use deception to accomplish self-presentational goals.

3. Linguistic implications of deception in CMC

Once the pragmatic decision is made to lie with communication technology, in particular in text-based CMC (e.g., email, instant messaging, chat rooms, newsgroups, text messaging, etc.), communicators must make linguistic choices about how to express their deception. The following section focuses on this issue, laying out recent trends in the research on verbal cues of deception.

3.1. Theoretical and empirical shift towards language in deception

Most previous deception research has been grounded in theories that focus primarily on the non-verbal cues associated with deception (Ekman 1985). For example, non-verbal "leakage" cues are assumed to reveal hidden emotions that are manifest in unconscious and uncontrolled movements of the face or body. This approach emphasized non-verbal cues because non-verbal behavior was assumed to be less controllable than speech (Vrij 2008b).

Since the early 1990s, however, theories of deception have begun to consider the linguistic aspects of deception. For example, Information Manipulation Theory (McCornack 1992) draws on Grice's Cooperative Principle and assumes that when people lie they violate one of the Cooperative Principle's four maxims of quality (veridicality of an utterance), quantity (amount of information in an utterance), relevance (relatedness to prior utterances), and manner (clarity of an utterance). These violations are assumed to have detectable linguistic manifestations.

Other theories have begun to emphasize the cognitive, emotional, and motivational consequences of deception on language use, such as Criteria-Based Content Analysis (CBCA; Köhnken 1996) and Reality Monitoring theory (Johnson and Raye 1981). CBCA provides 18 verbal cues that are associated with the cognitive and motivational aspects of truthful accounts, which are assumed to include more detail, be more logically coherent, contain more spontaneous corrections, and include more quoted speech. Similarly, Reality Monitoring theory assumes

that descriptions of real memories of an event differ from imagined or fabricated memories, such that descriptions of real memories will contain more perceptual and contextual information than false memories (Vrij 2008b).

Taken together, these theories suggest that deception should be reflected in language, and recent empirical evidence provides some support for this assumption (see Vrij 2008b for reviews of CBCA research and Masip, Sporer, Garrido, and Herrero 2005 for reviews of Reality Monitoring research). A recent meta-review of studies examining verbal features of deception revealed that the majority of studies examined found support for verbal differences between deceptive and truthful language, prompting a leading deception scholar to conclude that attending to language features can lead to more reliable deception detection than non-verbal cues (Vrij 2008b).

In addition to this theoretical turn towards verbal features, a second trend bolstering verbal research on deception is the recent advance in the computerized analysis of language. Sophisticated computer programs can now parse language into syntactic components and assign words to various semantic and psychological categories rapidly and automatically. Previous research on deception has relied on human coders in assessing language patterns. While human judgment is critical for assessment of features such as plausibility, humans tend to be unreliable judges of more fine-grained aspects of language, such as identifying and counting prepositions, pronouns, articles, etc. Furthermore, given the vast amounts of text now available through CMC, computerized text analysis allows for much larger corpora of language to be examined and compared.

3.2. Linguistic aspects of deception in CMC

Building on the theoretical and empirical advances on language and deception described above, a number of studies have used computerized text analysis programs to identify linguistic features of deception in CMC (Hancock, Curry, Goorha, and Woodworth 2008; Keila and Skillicorn 2005; Zhou, Twitchell, Burgoon, and Numamaker 2003). These studies use programs that rely on word counting approaches, in which each word is assigned to linguistic (e.g., pronouns) or psychological categories (e.g., positive or negative affect). For a review of these kinds of programs, see Tauszick and Pennebaker (2010).

The first study to examine the linguistic features of CMC was a careful and sophisticated analysis of asynchronous communication conducted by Zhou et al. (2003). In this study, participants logged into a system on which they could leave messages for, and receive messages from, a partner. Their task was to complete the desert survival task, in which they needed to choose a set of items to take with them when hypothetically crashed in the desert. One of the partners was instructed to lie about their preferences for the items. Using a collection of textual analysis tools, Zhou and colleagues found that messages from deceptive partners were character-

ized by more sentences with more words (including verbs and noun phrases). This finding was contrary to face-to-face studies that show that liars use fewer words when lying. One reason for this may be the asynchronous nature of the interaction: Online liars can take their time to craft the lies. Liars also produced less lexical and content diversity (e.g., a lower type-token ratio), suggesting that their messages were less semantically complex and less detailed than those of truth tellers. These findings were consistent with theories that suggest lies tend to be impoverished relative to truthful accounts. The liars also used more emotional language than truth tellers, including more positive and negative affect words (e.g., sad, delighted). Finally, with respect to pronoun usage, liars used more third-person references than truth tellers, but no differences were found for self-reference.

More recently, Hancock et al. (2008) examined the linguistic patterns of deception in synchronous CMC. Participants interacted over instant messaging and discussed four topics. On two of the topics, one of the participants was asked to lie to his or her partner but to tell the truth on the other topics. Their partners were asked to simply engage in the conversation and were unaware of the deception manipulation.

The hypotheses in this study were derived from those laid out by Zhou et al. (2003) described above and from the Newman-Pennebaker (NP) model (Newman, Pennebaker, Berry, and Richards 2003), an empirically-derived model of deception that predicts several language features that change relative to truthful language. According to the NP model, lies should involve fewer first person singular pronouns (“I”), as liars try to distance themselves psychologically from their lie. Second, lies should involve fewer exclusive words (“except”, “but”), as lies tend to be less complex than truthful statements. Third, lies should include more negative emotion terms (e.g., “sad”, “anxious”, “afraid”), which reflect the guilt and anxiety related to being deceptive. Last, lies should contain more motion verbs (e.g., “go”, “fly”), which help move the story along and distract the listener.

In general, the data were supportive of the NP model’s main predictions. The results revealed that liars, like those in Zhou et al.’s study, produced more words when lying than when telling the truth, suggesting that CMC liars may use more words than face-to-face liars to accomplish their deception. Unlike the Zhou et al. study, liars used fewer self-referencing pronouns and more other-oriented pronouns, consistent with the idea of psychological distancing. Finally, liars tended to use more causal terms (e.g., “because”, “therefore”, etc.) when lying compared to telling the truth, suggesting that they were potentially constructing more coherent and plausible stories.

Unlike most deception studies, Hancock and colleagues also examined the language of the partner. Even though the partner was blind to the deception manipulation, the partner’s language differed across deceptive and truthful topics. The differences often followed the senders. For instance, both the liar and the partner increased their use of first person pronouns during deception. This kind of lin-

guistic style matching occurred more frequently during deceptive parts of the conversation than during the truthful parts. A number of theories describe why and how interlocutors in face-to-face communication tend to synchronize their language, from syntactic usage (Pickering and Garrod 2004) to coordinating referents (Brennan and Clark 1995) (for a review, see Burgoon, Stern, and Dillman 1995). The partners, however, were unable to use this information to improve their deception detection accuracy: Partners performed at chance when trying to detect the lies. These data suggest that, while these changes in language patterns did not help the partner in detecting a lie, the partner's language patterns may provide some clues as to whether they are being lied to or not.

While the above-mentioned studies took place in the laboratory, Keila and Skillicorn (2005) conducted an analysis of deception in the real world, namely the Enron email corpus provided by the Enron trial (the corpus contains approximately 500,000 emails relating to a scandal at a large corporation), in which executives were sued for deliberately misleading investors. The Enron corpus consists of approximately 500,000 emails, and the researchers applied a subset of the NP model to the data in an attempt to classify deceptive emails. The dimensions employed were first person singular, negations, and exclusive words. Using these three dimensions, along with statistical techniques for classifying messages along the dimensions, Keila and Skillicorn successfully identified the most problematic (i.e., deceptive) emails with respect to the court case.

Taken together, a few trends emerge across these three studies. The first is that when people lie in CMC, they tend to use more words than when they are telling the truth, a finding that is the opposite of research on face-to-face deception. As noted, this may be due to the written nature of CMC, which provides more time and allows for editing. A second trend is changes in pronoun usage. Relative to telling the truth, lies include fewer self-references but more other-references. Emotional terms also seem to be different across lies and truths, with more emotional terms, especially negative ones, observed during deception. Exclusive words (e.g., "except", "however"), which may indicate the complexity of the grammatical structure, tend to decrease during CMC deception. Last, the linguistic style-matching data from Hancock et al. (2008) suggest that the partner's language should also be considered.

3.2.1. *Gender deception and language in CMC*

A number of studies have examined the language of gender in the context of CMC (Hall 1996; Herring 1993, 2003; Thomson, Murachver, and Green 2001). In general, these studies suggest that men and women use language in different ways along some dimensions. Females tend to use more hedges, justifications, personal pronouns, and expressions of emotion, while males tend to use assertions, rhetorical questions, sarcasm, profanity, and challenges. An important question, however,

is whether these differences persist when speakers are being deceptive about their gender. Since gender cues are presumably given off unconsciously (Goffman 1959), they should be difficult for speakers to control.

Several studies have looked at the detection of gendered language (Savicki, Kelley, and Oesterreich 1999; Thomson, Murachver, and Green 2001). When Thomson and Murachver presented students with emails and asked them to identify the gender of the authors, the students were relatively successful in distinguishing female emails from male ones. The authors concluded that a combination of linguistic features led to the decisions, rather than any one feature alone.

A more recent study examined how people performed gender online when being deceptive about their actual (offline) gender (Herring and Martinson 2004). In this study, the authors examined the transcripts from a gender-switching game, in which some participants were deceptive about their gender and others were not. The game took place in a synchronous chat space. The researchers found that participants produced stereotypical content when performing the opposite gender, but they also gave off stylistic cues to their actual gender. That is, when females played males they talked about overtly stereotypically male topics or content (e.g., cars), while males playing females talked about stereotypically female content (e.g., shopping). Nonetheless, there were some stylistic cues associated with their gender that participants were unable to control, such as gendered differences in pronoun usage, hedges, and rhetorical questions. Judges in the game only guessed participants' gender correctly approximately half the time, in part, perhaps, because they based their judgments primarily on the stereotypical content rather than on language use.

3.3. Style versus content in the language of CMC deception

A comparison of the language of deception in CMC interaction and the language of gender deception highlights an important distinction about the kind of language that is most likely to reflect deception. As noted in the gender research, participants were unable to hide the stylistic features of their language, such as pronoun usage. A review of the deception in CMC results suggests that stylistic features such as pronoun usage are also what liars find most difficult to control. This is no coincidence, as stylistic words tend to be less consciously accessible. Words can be divided roughly into two broad categories: content words and function words (Chung and Pennebaker 2007). Content words refer to terms that have semantic value and include word categories such as nouns, adjectives, and lexical verbs. In contrast, function words convey primarily grammatical information and include pronouns, prepositions, articles, conjunctions, and auxiliary verbs.

Whereas the average native English speaker has a vocabulary of well over 100,000 words, fewer than 400 are common function words (Baayen, Piepenbrock,

and Gulikers 1995). This seemingly trivial percentage (less than 0.04%) of vocabulary accounts for over half of the words in daily speech. Despite the frequency of their use, function words are the hardest to master when learning a new language. Indeed, they are virtually invisible in daily reading and speech (Chung and Pennebaker 2007).

Recent research shows that the analysis of function words, such as articles, auxiliary verbs, prepositions, and pronouns, provides important clues about the social relationship between a speaker and the audience, the psychological state of the speaker, and information about the speaker's social role and status in the community (Chung and Pennebaker 2007; Herring and Paolillo 2006). Function words may also be important in the language of deception in CMC, whether it be lies about gender or some other topic.

3.4. Future directions

The growing body of research described above suggests that linguistic and discourse patterns can be extracted from deceptive language in CMC (Hancock et al. 2008; Herring and Martinson 2004; Zhou et al. 2003). This program of research, however, is at its nascence, and there are at least two important directions in which this research needs to move. First, the vast majority of research on deception in CMC has focused only on English from Western cultures. The scant research that has examined deception patterns across cultures suggests that there are some culture-specific differences in how deception is perceived, but that there are also some principles of deception that may be universal across cultures. For example, Zhou and Lutterbie (2005) find that the collectivism-individualism dynamic should affect whether pro-social lying designed to maintain harmonious relationships is perceived as acceptable or not, with people from collectivist cultures viewing these pro-social lies as more acceptable than people from individualistic cultures.

An important question is whether the linguistic markers identified in English hold or change across cultures. There are important linguistic differences, such as differences in obligatory evidentiality between English and Arabic or differences in pronoun use between English and Chinese, that may affect how deception is signaled in language. Zhou and Lutterbie (2005) advocate a top-down and bottom-up approach, in which bottom-up language patterns are identified statistically without reference to psychological expectations, while top-down approaches guide specific analyses (e.g., speakers psychologically distance themselves from their lies, which theoretically should be the case across cultures).

The second direction is concerned with improved ecological validity. Most deception research takes place within fairly static and highly controlled laboratory studies, which may or may not generalize to the real world. This is especially the case for CMC research. Most studies involve participants telling true or false stories to each other in a chat room (Hancock et al. 2008; Zhou et al. 2003). Criti-

cally, the lies generated in most lab studies are mocked up, and liars have little motivation to succeed in their lies.

Some non-CMC research has begun to examine deception in naturalistic settings. For example, in a project examining deception in political speech, Markowitz, Hancock, and Bazarova (2011), compared false statements (e.g., claims that Iraq had weapons of mass destruction or direct links to al Qaeda) and non-false statements (e.g., that Saddam Hussein had used gas on his own people) produced by officials in the Bush administration in the run-up to the 2003 Iraq war. Consistent with the NP model of deception's predictions, false statements contained substantially reduced rates of first-person singular ("I") and exclusive terms ("except", "but") but more negative emotion terms and action verbs. Using this simple model, we were able to classify approximately 76% of the Bush administration's statements correctly as either false or not false.

CMC deception seems to be a good fit for naturalistic studies. Like the Enron analysis by Keila and Skillicorn (2005), future research needs to identify archival cases in which the ground truth of deception has been established, and the language records have been preserved. More of this kind of research should lead to major advances in understanding the pragmatic and linguistic implications of deception in CMC.

4. Conclusion

The first of this chapter's central questions was concerned with whether, and if so, how, technology enters into the calculus of deception. The studies reviewed suggest that the communication spaces created by various technologies can shape speakers' decisions to use deception as a pragmatic tool in accomplishing their objectives. This seems to be the case in both everyday interpersonal interactions, with technologies such as email, instant messaging, and the telephone, and in more anonymous online spaces, such as chat rooms and news groups.

The relationship between deception and technology is clearly not simple, but in our view, it is systematic. Equally importantly, many more questions remain to be answered as technologies evolve. For example, new technologies, such as the avatars in graphically-based virtual environments, may make deceptions more sophisticated. As Donath (2007: 53) asks, "as behavioral software becomes more sophisticated, are we creating avatars that will be increasingly attractive and seemingly friendly but are in fact the ideal mask behind which a dishonest or manipulative person can operate?" Indeed, a recent study suggests that people feel more comfortable lying when using an avatar than when lying in text-based communication (Galaxhi and Fui-Hoon Nah 2007).

Our second question was concerned with the linguistic traces of deception in CMC. Although this line of research has just begun, it suggests some potential in

being able to identify patterns of deception in language, with liars online using more words, less self-referencing and more other-referencing language, and producing less grammatically complex messages. Whether these patterns stand up to more ecologically valid and diverse tests remains to be answered. Nonetheless, the fact that stylistic features (such as function words) tend to be critical in both interpersonal deception and gender-based deception suggests that this class of words may be especially important.

This final point has relevance for the question of whether to lie online or not. Currently, there are no reliable methods for identifying deception in CMC, so this kind of issue does not enter into the question of whether one should lie online. However, given the potential to identify deceptive patterns, and the fact that much of what people say now on the Internet will be around for many years, in the future it may be necessary to consider even more carefully the question of whether or not to lie.

References

- Baayen, R. Harald, Richard Piepenbrock, and Léon Gulickers
1995 The CELEX Lexical Database [CD ROM]. Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.
- Bavelas, Janet B., Alex Black, Nicole Chovil, and Jennifer Mullett
1990 Truths, lies, and equivocations: The effects of conflicting goals on discourse. *Journal of Language and Social Psychology* 9(1–2): 129–155.
- Brennan, Susan E. and Herbert H. Clark
1995 Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(6): 1482–1493.
- Buller, David B. and Judee K. Burgoon
1996 Interpersonal deception theory. *Communication Theory* 6(3): 203–242.
- Burgoon, Judee K., Lesa A. Stern, and Leesa Dillman
1995 *Interpersonal Adaptation: Dyadic Interaction Patterns*. New York: Cambridge University Press.
- Carlson, John R., Joey F. George, Judee K. Burgoon, Mark Adkins, and Cindy White
2004 Deception in computer-mediated communication. *Group Decision and Negotiation* 13(1): 5–28.
- Caspi, Avner and Paul Gorsky
2006 Online deception: Prevalence, motivation, and emotion. *Cyberpsychology and Behavior* 9(1): 54–59.
- Chung, Cindy K. and James W. Pennebaker
2007 The psychological functions of function words. In: Klaus Fiedler (ed.), *Social Communication*, 343–359. New York: Psychology Press.
- Clark, Herbert H. and Susan E. Brennan
1991 Grounding in communication. In: Lauren B. Resnick, John Levine, and Stephanie D. Teasley (eds.), *Perspectives on Socially Shared Cognition*, 127–149. Washington, DC: APA.

- Cornwell, Benjamin and David C. Lundgren
 2001 Love on the Internet: Involvement and misrepresentation in romantic relationships in cyberspace vs. realspace. *Computers in Human Behavior* 17(2): 197–211.
- Daft, Richard L. and Robert H. Lengel
 1984 Information richness: A new approach to managerial behavior and organization design. In: Barry M. Staw and Larry L. Cummings (eds.), *Research in Organizational Behavior* 6, 191–233. Greenwich, CT: JAI Press.
- Daft, Richard L. and Robert H. Lengel
 1986 Organizational information requirements, media richness and structural design. *Management Science* 32(5): 554–571.
- Danet, Brenda
 1998 Text as mask: Gender, play and performance on the Internet. In Steven G. Jones (ed.), *Cybersociety 2.0: Revisiting CMC and Community*, 129–158. Thousand Oaks, CA: Sage.
- DePaulo, Bella M., Deborah A. Kashy, Susan E. Kirkendol, Melissa M. Wyer, and Jennifer A. Epstein
 1996 Lying in everyday life. *Journal of Personality and Social Psychology* 70(5): 979–995.
- DePaulo, Bella M., James J. Lindsay, Brian E. Malone, Laura Muhlenbruck, Kelly Charlton, and Harris Cooper
 2003 Cues to deception. *Psychological Bulletin* 129(1): 74–118.
- Donath, Judith S.
 1999 Identity and deception in the virtual community. In: Marc A. Smith and Peter Kollock (eds.), *Communities in Cyberspace*, 25–59: London: Routledge.
- Donath, Judith S.
 2007 Virtually trustworthy. *Science* 317: 53–54.
- Ekman, Paul
 1985 *Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage*. New York: W. W. Norton.
- Ennis, Edel, Aldert Vrij, and Claire Chance
 2008 Individual differences and lying in everyday life. *Journal of Social and Personal Relationships* 25(1): 105–118.
- Galanxhi, Holtjona and Fiona Fui-Hoon Nah
 2007 Deception in cyberspace: A comparison of text-only vs. avatar-supported medium. *International Journal of Human-Computer Studies* 65(9): 770–783.
- George, Joey F. and John R. Carlson
 2005 Media selection for deceptive communication. In: *Proceedings of the Thirty-Eighth Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Press.
- Goffman, Erving
 1959 *The Presentation of Self in Everyday Life*. New York: Doubleday.
- Grice, H. Paul
 1989 *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hall, Kira
 1996 Cyberfeminism. In Susan C. Herring (ed.), *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*, 147–170. Amsterdam/Philadelphia: John Benjamins.

- Hancock, Jeffrey T.
2007 Digital deception: When, where and how people lie online. In: Katelyn McKenna, Tom Postmes, Ulf-Dietrich Reips, and Adam N. Joinson (eds.), *Oxford Handbook of Internet Psychology*, 287–301. Oxford, UK: Oxford University Press.
- Hancock, Jeffrey T., Lauren E. Curry, Saurabh Goorha, and Michael T. Woodworth
2008 On lying and being lied to: A linguistic analysis of deception. *Discourse Processes* 45(1): 1–23.
- Hancock, Jeffrey T., Jennifer Thom-Santelli, and Thompson Ritchie
2004a Deception and design: The impact of communication technologies on lying behavior. In: *Proceedings of the 2004 ACM Conference on Human Factors in Computing Systems* 6, 130–136. New York: ACM.
- Hancock, Jeffrey T., Jennifer Thom-Santelli, and Thompson Ritchie
2004b What lies beneath: The effect of communication medium on deception production. Paper presented at the Annual Meeting of the Society for Text and Discourse, Chicago, IL.
- Herring, Susan C.
1993 Gender and democracy in computer-mediated communication. *Electronic Journal of Communication* 3(2).
<http://www.cios.org/EJCPUBLIC/003/2/00328.HTML>
- Herring, Susan C.
2001 Computer-mediated discourse. In: Deborah Schiffrin, Deborah Tannen, and Heidi Hamilton (eds.), *The Handbook of Discourse Analysis*, 612–634. Oxford, UK: Blackwell Publishers.
- Herring, Susan C.
2003 Gender and power in online communication. In: Janet Holmes and Miriam Meyerhoff (eds.), *The Handbook of Language and Gender*, 202–228. Oxford, UK: Blackwell.
- Herring, Susan C.
2007 A faceted classification scheme for computer-mediated discourse. *Language@Internet* 4, article 1.
<http://www.languageatinternet.org/articles/2007/761>
- Herring, Susan C. and Anna Martinson
2004 Assessing gender authenticity in computer mediated language use: Evidence from an identity game. *Journal of Language and Social Psychology* 23(4): 424–446.
- Herring, Susan C. and John C. Paolillo
2006 Gender and genre variation in weblogs. *Journal of Sociolinguistics* 10(4): 439–459.
- Johnson, Marcia K. and Carol L. Raye
1981 Reality monitoring. *Psychological Review* 88(1): 67–85.
- Kashy, Deborah A. and Bella M. DePaulo
1998 Who lies? *Journal of Personality and Social Psychology* 70: 1037–1051.
- Keila, Parambir S. and David B. Skillicorn
2005 Detecting unusual and deceptive communication in email. IBM Centers for Advanced Studies Conference, Richmond Hill, Ontario, Canada, October 17–20.

Köhnken, Günter

1996 Social psychology and the law. In: Gün R. Semin and Klaus Fiedler (eds.), *Applied Social Psychology*, 257–282: London: Sage.

Markowitz, David, Hancock, Jeffrey T., and Bazarova, Natalya N.

2011 The language of presidential lies: How words reflect lies about war, personal scandal and state secrets. Paper presented at the 97th Annual Meeting of the National Communication Association, New Orleans.

Masip, Jaume, Siegfried L. Sporer, Eugenio Garrido, and Carmen Herrero

2005 The detection of deception with the Reality Monitoring approach: A review of the empirical evidence. *Psychology, Crime, and Law* 11: 99–122.

Mazar, Nina and Dan Ariely

2006 Dishonesty in everyday life and its policy implications. *Journal of Public Policy and Marketing* 25(1): 1–21.

McCornack, Steven A.

1992 Information Manipulation Theory. *Communication Monographs* 59(1): 1–16.

McCornack, Steven A., Timothy R. Levine, H. I. Torres, and D. M. Campbell

1992 When the alteration of information is viewed as deception: An empirical test of information manipulation theory. *Communication Monographs* 59(1): 17–29.

Möllering, Guido

2006 *Trust: Reason, Routine, Reflexivity*. Oxford, UK: Elsevier.

Newman, Matthew L., James W. Pennebaker, Diane S. Berry, and Jane M. Richards

2003 Lying words: Predicting deception from linguistic style. *Personality and Social Psychology Bulletin* 29(5): 665–675.

Parks, Malcolm R. and Kory Floyd

1996 Making friends in cyberspace. *Journal of Communication* 1(4): 80–96.

Pickering, Martin J. and Simon Garrod

2004 Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27(2): 169–225.

Roberts, Lynne D. and Malcolm R. Parks

1999 The social geography of gender-switching in virtual environments on the Internet. *Information, Communication, and Society* 2(4): 521–540.

Savicki, Victor, Merle Kelley, and E. Oesterreich

1999 Judgments of gender in computer mediated communication. *Computers in Human Behavior* 15(2): 185–194.

Selwyn, Neil

2008 A safe haven for misbehaving? An investigation of online misbehavior among university students. *Social Science Computer Review* 26(4): 446–465.

Sproull, Lee and Sara Kiesler

1986 Reducing social context cues: Electronic mail in organizational communication. *Management Science* 32(11): 1492–1512.

Tausczik, Yla and James W. Pennebaker

2010 The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology* 29(1): 24–54.

Thomson, Rob, Tamar Murachver, and James Green

2001 Where is the gender in gendered language? *Psychological Science* 12(2): 171–175.

- Twitchwell, Douglas P., Jay F. Nunamaker, and Judee K. Burgoon
2004 Using speech act profiling for deception detection. In: Hsinchun Chen (ed.), *Intelligence and Security Informatics*, 403–410. Berlin: Springer Berlin.
- Vrij, Aldert
2008a *Detecting Lies and Deceit: The Psychology of Lying and the Implications for Professional Practice*. Chichester, UK: Wiley.
- Vrij, Aldert
2008b Nonverbal dominance versus verbal accuracy in lie detection: A plea to change police practice. *Criminal Justice and Behavior* 35(10):1323–1336.
- Whitty, Monica T.
2002 Liar, liar! An examination of how open, supportive and honest people are in chat rooms. *Computers in Human Behavior* 18(4): 343–352.
- Whitty, Monica T. and Adrian N. Carr
2006 *Cyberspace Romance: The Psychology of Online Relationships*. Basingstoke, UK: Palgrave Macmillan.
- Whitty, Monica T. and Siobhan E. Carville
2008 Would I lie to you? Self-serving lies and other oriented lies told across different media. *Computers in Human Behavior* 24(3): 1021–1031.
- Zhou, Lina and Simon Lutterbie
2005 Deception across cultures: Bottom-up and top down approaches. In: Paul Kantor, Gheorghe Muresan, Fred Roberts, and Daniel Zeng (eds.), *Intelligence and Security Informatics*, 465–470. Springer-Verlag: Berlin.
- Zhou, Lina, Douglas P. Twitchell, Tiantian Qin, Judee K. Burgoon, and Jay F. Nunamaker
2003 An exploratory study into deception detection in text-based computer-mediated communication. In: *Proceedings of the Thirty-Sixth Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Press.

