

Deception with Technology

Jeffrey T. Hancock and Jamie Guillory

Department of Communication, Cornell University, Ithaca, NY, USA

Butlers, Sock Puppets, and the Chinese Water Army

When it comes to deception, there is never a shortage of examples. And with the widespread adoption of technology in human communication, it can sometimes seem that deception is more rampant than ever. We have seen famous football players tricked into believing they are dating a woman that does not exist. Recent work on deception and technology has revealed *butler lies*, in which people use technology to create a social buffer and avoid or put off communication, such as the ubiquitous “on my way” when in fact the person has not yet left their house. There are also *sock puppets*, in which glowingly positive book reviews are in fact written by their high-profile authors; or, more generally, fake reviews for hotels, restaurant, and any other business online (Ott, Cardie, & Hancock, 2012). And more worrisome yet are large-scale deception phenomena, such as the so-called *Chinese Water Army*, in which individuals are paid by governments or corporate organizations to spread propaganda in an organic, grassroots way. Here in North America this is often referred to as *astroturfing* (for review, see Hancock, 2012).

At the same time, there is a steady drumbeat of scandals in which deceptions are revealed by digital technology. When Mitt Romney told a group of donors at a closed-door meeting about 47% of the nation being dependent on the government, it was a digital video recorder that captured the moment and allowed it to be shared worldwide. In fact, almost any major political scandal of the last decade has involved some digital record, from email to a text message. The ways in which lying takes place, and how it is detected, are being transformed by technology along with all the other aspects of human communication. But, as we shall see, our assumptions about how technology and deception affect one another are not always as simple as they seem, and we propose here that there are several key factors that need to be considered to

understand deception in the networked age, including biases and heuristics unique to deception, the transformation of the interaction space, and our own views of ourselves as good, honest people.

The Cues Heuristic and Implications for Deception Online

One of the most basic questions we can ask about deception in the networked age is whether communication technology changes how often people lie. The answer for most people is obvious: technology leads to more deception (Keyes, 2004). The basic reasoning for this answer, which we call the cues heuristic, goes like this: there are fewer cues online so people lie more because they can get away it.

Although this seems like a fairly straightforward argument, there are three important assumptions in this heuristic. The first assumption is that there are fewer cues online than face to face (FtF). Email has fewer cues than the phone, for example, because there are no vocal cues, while the phone has fewer cues than FtF because there are no nonverbal cues, and so on. The second is that there are reliable cues for detecting deception. The belief that liars give off cues that can indicate their deception is near universal. It follows, therefore, that because technologically mediated communication has fewer cues than FtF interaction it should be more difficult to detect deception. So, email lies should be more difficult to detect than phone lies, which in turn should be harder than Skype-based lies. The last assumption is that when people can get away with it, they will lie more. In the case of lies online, because there are fewer cues that make it harder to detect deception, people will take advantage of this fact and lie more.

The belief in this cues heuristic is wide spread, as we discovered in a survey of the population on their beliefs about deception behavior (Toma, Jiang, & Hancock, 2013). To examine beliefs about deception online, we first asked people how much they lie in different media, including FtF, the phone, messaging, and email. We later asked how much they thought that others lie in these media. While people believe that others always lie more than they themselves do, a typical self-other asymmetry, the difference between self and other was larger the fewer cues there were. That is, the asymmetry was bigger in online media like email and messaging. In a second study, we replicated this pattern: the asymmetry between self and other for deception increased for online media (e.g., messaging and email) compared to richer media such as the phone or FtF. The reason for the intensified asymmetry was the different motivations people ascribed to their own lies versus others' lies. Consistent with self-other asymmetries in all types of social psychological judgments (e.g., Pronin, 2006), people ascribe altruistic motivations for their own deception using online media (e.g., "I lie to protect myself online"), but selfish motivations for others' deceptions online (e.g., "Others lie to take advantage of the lack of cues").

Unfortunately, each of the assumptions underlying the cues heuristic is problematic. The first assumption, that online communication has fewer cues, appears accurate on its face. After all, text-based communication, which makes up the majority of online communication, has fewer vocal and nonverbal cues than FtF. But this

assumption ignores two of the most important features of online communication for deception: recordability and social networks. As we will discuss later in the chapter, online communication leaves digital traces that can be easily copied, backed up, shared, searched, and analyzed. This property of online communication has crucial implications for deception, both how it is practiced and how it is detected. In general, leaving a record of a lie is a bad idea. Networks also have important implications for deception. Networks constrain what a person can say because others that have ground truth information may encounter the message (Warkentin, Woodworth, Hancock, & Cormier, 2010). Consider, for example, the person that says they have a computer science degree on linkedin.com and who has their college advisor in their network. If in fact the claim is not true, the claim may be called into question by the advisor. This possibility should constrain deceptions in networked environments.

The second assumption, that cues are useful for detecting deception, is based on several myths surrounding deception and does not take into account recent meta-analyses of several decades of deception detection research. Bond and DePaulo (2006) found that over many studies involving deception detection, in which participants make judgments about the veracity of messages (both verbal and nonverbal), accuracy rates hover around 54%, statistically higher than chance but practically no better than flipping a coin. Why are accuracy rates so low? The primary reason is that there are, in fact, very few reliable cues to deception that can be used by people to detect deception. This view is supported by another DePaulo meta-analysis (DePaulo et al., 2003) that examined cues to deception across hundreds of studies. The data revealed that very few nonverbal cues are reliable signals, and those that are, such as vocal pitch and pupil dilation, are difficult for humans to detect given their subtlety.

These empirical findings stand in stark contrast to most people's beliefs about the cues to deception. For instance, a study that examined beliefs about cues to deception in countries around the world found that one of the most common cues cited as important for detecting deception was eye gaze. Participants in almost every country reported that eyes are the window to the truth, and that liars have shifty eyes. People also list many other nonverbal cues as important across the world, including fidgeting and other body movements. In fact, eye gaze is not one of the reliable cues identified in the meta-analysis, nor is fidgeting (DePaulo et al., 2003). Other beliefs about deception that do not match empirical findings include microexpressions, brief facial expressions that reveal suppressed emotions (Ekman, 1985). While most people believe that microexpressions indicate deception, recent work has shown that there are no peer-reviewed publications presenting evidence that this is the case (Porter & ten Brinke, 2008).

Many of these myths about deception detection are portrayed regularly in the popular media. In an insightful study illuminating the problems with these deception myths, participants were asked how much they watched the popular deception detection TV show *Lie To Me* and then completed a deception detection task (Levine, Serota, & Shulman, 2010). The results indicated that those who watched the show performed significantly worse at detecting deception than people who did not watch the show.

Unfortunately, these misplaced beliefs in the cues to deception, namely that there are reliable nonverbal cues, suggest to people that it is more difficult to detect deception online. In fact, there is very little evidence that this is the case. When we

experimentally manipulated deception detection across Internet chat and FtF, we saw no difference in lie detection accuracy unless the liar was highly motivated (Hancock, Woodworth, & Goorha, 2010).

The third assumption inherent in the cues heuristic is that people will lie more if they can get away with it. This assumption is also deeply flawed, it turns out. Recent work suggests that people usually only lie when they have a reason (Levine, Kim, & Hamel, 2010). Although there is a small percentage of the population that lie regularly for no apparent reason (e.g., pathological liars, psychopaths), most people require some motivation to lie. Other work from behavioral economics confirms this view (Ariely, 2012), with many studies showing that people may cheat a little when given the opportunity, but that large deceptions are rare even when it is impossible to be caught.

Why do people only lie for a reason, even when they are unlikely to get caught? One reason is that lying is difficult. As Sir Walter Scott noted, “Oh what a tangled web we weave when at first we seek to deceive.” Maintaining a fabricated story is extremely difficult, and this remains true online, as the hoaxer in the Manti Te'o case learned. Another reason is that people have a strong aversion to being considered a liar (Anderson, 1968). When people lie or cheat too much, it can trigger a shift in self-identity from good, honest person to liar or cheat. Ariely (2012) and colleagues' work suggest that self-identity as an honest person is a strong constraint that prevents most people from lying too much, even when they can get away with it. So, even if someone could get away with a lie more easily online, it is not necessarily the case that they would lie more simply because they could.

The cues heuristic, though a powerful and intuitive one, is based on faulty assumptions and clouds our judgments about how deception interacts with technology. In *The Transformation of the Interaction Space* we examine a more fruitful approach to deception and technology that takes a more holistic look at deception and communication technology.

The Transformation of the Interaction Space

A more useful way to approach deception and technology than focusing simply on cues is to consider the broader transformation that technology has wrought on human interaction. Consider that humans have been using spoken language to talk with one another for at least 50,000 years, although probably much longer (Dediu & Levinson, 2013). During this period there have been at least 100 billion humans, a species evolving over thousands of generations by using language to communicate FtF. Writing emerged much more recently, approximately five to six thousand years ago. For all of the humans that lived before the invention of writing, the only way they could communicate was if the other person was in the same physical location speaking at the same time the other person was listening. That is, before the invention of writing interactants needed to be co-present and co-temporal. One consequence of this is that for the 80 billion or more humans who lived before the advent of writing, not one of their words left a record. The evanescence of spoken language means that each word that they spoke disappeared as they said it. In fact, this is the case for most people throughout human history. Even after the invention of the printing press,

most of the world's population was illiterate. Indeed, it was not until after World War II that the world literacy rate exceeded 50%.

When we step back to consider this history of interpersonal human communication, it puts into perspective the radical changes that communication technology have brought to the human interaction space. People can now talk to others across distance and time, and they regularly do so. People write a tremendous number of words in email, texts, social media and other forms of online interaction. Eric Schmidt, former CEO of Google, estimates that more information is created every 2 days as existed up until 2003. On a personal level, this means that many adults in the United States produce more written words in a day than most humans throughout history produced in their entire lifetime.

These changes to the interaction space are having a profound effect on how deception and deception detection operate, but as noted before, this effect is not simply a matter of cues. Here we describe three ways that the transformation of the interaction space has modified deception in human communication: (1) the introduction of the display as a mode of interaction, (2) the effect of certain communication features introduced by these changes, such as the behavioral traces we now leave behind in digital contexts, and (3) the presence of warrants or connections in virtual space to a person's real-world identity.

The display

When we interact with others online, we do so through a display of some kind, such as a computer monitor, a mobile display, or, in the near future, our glasses. The introduction of the display to human communication is perhaps one of the most studied aspects of computer-mediated communication. In the context of deception, the introduction of the display means that interlocutors have much more control over their interaction space. As many others have now noted, people can control their self-presentation to a greater degree because they can choose and edit what they present, a term called selective self-presentation (Walther, 1996). People who seek to deceive can take advantage of this increased control to develop a number of different techniques for deceiving others.

The most well-known of these is spam, which can be deceptive in a variety of ways. The 419 spam, for example, involves a message from Nigeria or other African country that supposedly is from a wealthy individual who needs help getting their money out of the country, help that will lead to great riches for the spam recipient. Another common form of deception that relies on the control afforded by the display is phishing, which involves a website that looks like a bank or other trustworthy institution but is in fact controlled by a fraudster. The goal is to lure the user into providing personal information that can be used to access their bank accounts or steal their identity (Dhamija, Cassidy, Hallam-Baker, & Jacobsson, 2006).

Both spam and phishing are made possible by the control that the fraudster has over the user's display. In the spam case, for example, the fraudster can present as a wealthy individual from a far off land. In a nonmediated world, this would be much more difficult (though certainly not impossible) given both the physical and social indicators

of cultures and wealth. In the phishing case, the fraudster can modify the user's display to show trustworthy images, logos and text that mimic, for instance, a bank. Both spam and phishing attacks use their control over the user's display to leverage a variety of psychological principles common to most scams, such as distraction, social compliance, greed, and scarcity (Stajano & Wilson, 2011). Phishing attacks that manipulate the user's display and rely on these scam principles can be highly effective, with up to 90% of users falling for well executed attacks (Dhamija, Tygar, & Hearst, 2006).

While these examples highlight how deception can be facilitated with the control of the display, the control of the display can also be used to modify the likelihood of dishonesty. Research in social psychology and behavioral economics suggests that there are subtle cues that can prime behaviors related to deception (Ariely, 2012). Several studies, for example, have tested the effect of surveillance primes, such as a pair of human eyes, on changing peoples' pro-social behavior (for review, see Nettle et al., 2013). In one study, people gave more money to an honesty jar for coffee cream when a photo of human eyes was affixed to the jar than a photo of a flower (Bateson, Nettle, & Roberts, 2006). More recently, surveillance primes have been used to reduce bicycle theft by nearly two-thirds.

An important question is whether these kinds of primes can also be used to shape deception-related behavior online. In one study that suggests this is possible, Haley and Fessler (2005) examined the effect of surveillance primes on pro-social sharing behavior in an economic game. In the game, one of the participants had the ability to decide on how to share a pot of money with a partner. This person was in one of two relevant conditions: either a surveillance prime, implemented as a set of cartoon eyes, was present on the display or it was not. People in the condition with a surveillance prime on their display shared more of the pot with their partner than people without a surveillance prime on their screen, suggesting that surveillance primes on a computer screen can alter behaviors related to cheating.

We recently demonstrated that this effect can also take place on Facebook and with interpersonal lies (Spottswood & Hancock, 2013). We inserted images of eyes in the advertising stream on Facebook—for example, ads for eye drops that prominently displayed eyes—or images that did not involve eyes. We then gave participants the opportunity to lie to a friend either on a wall post (in public) or in a message (in private). We found that the presence of the eyes had an important effect, but it was not simply to reduce deception. Instead, we found that the eyes increased the likelihood that the participant would “do the right thing,” namely follow the norm for public or private behavior. When in public the eye spots led participants to lie more often to protect their friend, but in private the eyes led participants to tell the truth more often, presumably to give their friends more honest feedback. The takeaway for us is that the interface can have important, and sometimes subtle, effects on how we lie or are lied to.

Communication features and deception

With the transformation of our interaction space comes changes to the way that we use deception in computer-mediated contexts. Research to date suggests that differences in the way that people lie across communication media are affected by specific

features of these different media. Hancock, Thom-Santelli and Ritchie's (2004a) feature-based model of deception and technology suggests that there are three important features of communication contexts that should influence deception. First, they argue that the recordability of a communication medium, or the degree to which communication in a given space has a durable record, should influence deceptive behavior. As communication in a given medium becomes more recordable, they argue that the likelihood of deception decreases. If we consider the recordability of different types of communication media, email provides the most permanent record, with copies of messages available not only on the individual's personal computer, but also on servers. In comparison, instant message conversations and SMS text messages appear less permanent, although this is mostly an illusion, as many political scandals involving text messaging have revealed. Almost any form of online interaction, from email to text messages to Facebook posts, leave a digital trail that can be retained, searched, and copied.

A second feature of communication media that is important to deception is the synchronicity of communication. Synchronicity refers to the degree to which a conversation is occurring in real time. Telephone conversations, for example, are highly synchronous as conversation partners are exchanging messages in real time. Email correspondence is often far less synchronous as delays between message exchange can be anywhere from seconds to days or even weeks, with a general expectation that responses will not be read immediately following message receipt. Hancock et al. (2004a) suggest that deception should occur more frequently in highly synchronous media. This is because highly synchronous media provides less time to craft honest responses. Many of the lies that we tell in conversation are small lies told to maintain relationships and protect our own image (DePaulo et al., 2003) (e.g., "No, you don't look fat in that dress, dear!"). In real-time communication, we often need to rely more on these types of lies to avoid situations where the truth might create an awkward or unwanted situation. When communication is less synchronous, we have more time to craft honest responses that need not rely on deception to be plausible.

A third feature of communication media that may dictate deceptive behavior is whether communicators are co-present. If a person is physically present, they are unable to lie about certain things, such as where they are, who they are with or what they are doing. Thus, computer-mediated communication contexts provide a greater number of opportunities to be deceptive about one's location, activities, and the company one is keeping.

The three parts of the feature-based model suggest that people should lie more in phone conversations than FtF conversations. In addition, FtF conversations should have more lies than email conversations (phone > face-to-face > email). Phone calls are not recorded, are not co-present and are highly synchronous. FtF conversations are not recorded, but do involve co-presence and are highly synchronous. Email, on the other hand, is recorded, but does not involve co-presence or high synchronicity.

Several studies have tested these predictions of the feature-based model. In an initial study, Hancock et al. (2004a) asked students to use a diary method to record all social interactions and lies over a 7-day period. Participants in this study made

records of FtF interactions, as well as phone calls, emails, and instant messages. The authors used participants' data to calculate lying rates by media, as well as proportions depicting the number of lies per number of conversations by media for each individual participant. Similar to previous research (DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996), they found that people lied more in phone conversations than FtF conversations. They also found that emails involved the fewest number of deceptions across media. These findings are in line with predictions from the feature-based model.

In a follow-up study, the authors (Hancock, Thom-Santelli, & Ritchie, 2004b) used the same procedure, but sought more detailed information about the types of lies that were told in daily interactions. In this study, the authors requested more detailed information about the specific lies that people told and to whom these lies were told. In this study, the authors observed the same patterns of deception as the initial study. In addition, they found that the features of the media used influenced the content of lies. Phone conversations had more lies about what people were doing, which makes sense given that this was the most mobile form of media assessed at the time of the study. On the other hand, FtF conversations included the greatest number of lies related to feelings, which the authors explain was likely because synchronous conversations often involve discussions of opinions. Email included the greatest number of lies related to explanations, presumably because the asynchronous nature of email allows more time to construct of these types of lies. In addition to differences in the types of lies people told, the recipients of these lies differed across media. People lied more often to friends and family via mediated communication channels (i.e., phone and email). Alternatively, people told more lies to strangers or acquaintances in FtF conversations, which is consistent with findings from previous research (DePaulo & Kashy, 1998; Ennis, Vrij, & Chance, 2008).

A more recent study conducted by Whitty and colleagues (2012) examined deceptive behaviors across a broader range of communication media. In addition to exploring deception across phone, FtF and email, they also asked participants to report on interactions and deceptions in IM, social networking websites, and SMS text messages. Using the feature-based model they predicted that people would lie more on the phone, followed by FtF interactions and IM (due to similarities in synchronicity and recordability), and then email, SMS texting and social networking websites (due to similarities in co-presence, recordability, and asynchronicity). They were also interested in the spontaneity of deception used in these contexts. While many of the lies that people tell are spontaneous, research suggests that a proportion of lies are planned (e.g., DePaulo et al., 1996, 2003). For example, if a person knows that they are going to arrive late to meet a friend without a legitimate excuse, they are likely come to the destination with a prepared lie that provides a reason for their lateness rather than coming up with a lie on the spot.

In a 1-week diary study, participants reported their daily interactions and deceptions, and provided information about the content of their lies, the reason for the lie, the type of lie, the referent of the lie (i.e., who did you tell this lie to?), and the referent's closeness to the individual (i.e., how close are you with the person you lied to?) and gender. In line with findings from earlier studies (Ennis et al., 2008; Hancock et al., 2004a, 2004b; Kashy & DePaulo, 1998), Whitty et al. (2012) observed a greater

number of deceptions over the phone compared to FtF interactions. People were also more likely to lie over IM than SMS texting. No other differences in deception were observed across the other types of media, though the authors noted the low frequency of deceptions in IM, email, and social networking websites made interpretation of this null finding difficult.

In addition to replicating findings from earlier studies, the authors observed some interesting differences in the way that people told planned versus spontaneous lies. As might be expected, participants in the study were more likely to tell planned lies in certain forms of asynchronous media (e.g., SMS texting) than synchronous media (e.g., FtF interactions). They also found that participants rated planned lies as more serious lies than spontaneous lies.

The aforementioned studies provide evidence for the feature-based model of deception and technology. These findings suggest that people take features of the communication media (e.g., recordability, synchronicity, co-presence) into consideration when making decisions about whether or not to lie. We see that these features are also important when people consider the recipients and content of their lies.

The Warranted Self-Presentation Perspective

Deception is almost always used to accomplish goals, such as appearing attractive or interesting for a romantic partner or appearing competent to an employer. Self-enhancing deceptions are common, and typically motivated by a desire to present oneself positively to others. In the self-presentational framework of deception, DePaulo et al. (2003) suggest that lies are used to manage our self-presentations. This framework suggests that self-presentation lies are motivated by the psychological rewards that come with conveying a more positive self-presentation to others. Lies are a self-presentational strategy that allow the individual to appear more positively than their actual traits suggest, but also allow them to avoid the consequences that they might experience from revealing their true characteristics (e.g., disapproval, ridicule, etc.).

DePaulo argues that self-presentational lies can be self or other-oriented. Self-oriented lies are those told to protect or enhance one's self-presentation. For example, in a self-oriented lie, a job seeker might lie about their how long they worked at an organization to appear more experienced in hopes of being hired. Other-oriented lies are those told to protect others' feelings while maintaining a positive impression of the self. For example, in an other-oriented lie, a person might tell a friend that they are on their way to a meeting place when they have not yet left the house, which both conveys to the friend that they have not forgotten their plans while simultaneously maintaining the person's identity as a considerate friend.

DePaulo's research provides a detailed picture of lies told to accomplish self-presentation goals, but what happens when our self-presentations move to online contexts? As we have described, the online environment has evolved dramatically over our traditional FtF form of interaction. While we discuss the opportunities for deception in various media earlier, changes to communication capabilities provide abundant opportunities for self-presentation. For example, we craft detailed profiles on social

networking websites, often sharing details about our personal and professional lives in great detail. In addition, many of our online activities are tracked and shared with members of our social network. Many websites and mobile-based applications allow users to share their online activities with members of their social network, providing unique information about our identity.

These traces of our online activity provide members of our social network with a highly detailed picture of our characteristics and behaviors that would be difficult, if not impossible, to match with information gleaned in offline contexts (e.g., Backstrom & Kleinberg, 2014). These changes in the amount of and accessibility to self-presentational context raise the question of whether deception for self-presentation operates differently online. Does the online environment lead to radical changes in the way that deception occurs, or do the same principles for deception in offline contexts apply online?

While some of the unique features of online communication that enable deception are described earlier, these new opportunities for self-presentation online, such as social networking websites, present important new constraints to deception. Perhaps the most important of these constraints is described in what we call the warranted self-presentation perspective. In explaining this perspective, it is important to first outline the features of online communication that constrain deception. Social network profiles make self-presentations publicly available and link individuals to the profile who can verify whether information is deceptive or not. Researchers both on and offline have long demonstrated the importance of our social relationships in fostering honesty between individuals (Resnick, Kuwabara, Zeckhauser, & Friedman, 2000; Resnick & Varian, 1997; Zimmerman & Kurapati, 2002). For example, recommender systems on websites like eBay help to ensure that transactions in these environments remain honest by providing users, who have no previous history of a seller, with valuable information about the seller's trustworthiness. Affordances that allow for social links between the virtual representation of self that one encounters online and the physical, offline self can improve the likelihood of honesty in the online environment.

Walther and Parks (2002) call these links between the virtual and physical self *warrants*, a concept first coined by Stone (1995). Stone (1995) originally defined a warrant as the construction and maintenance of a connection between the physical self and a discursive representation of the self. Stone's (1994) original conception of warrants comes from virtual systems theory and views "the production and maintenance of a discursive space and a physical space as 'warranting'" (p. 181). Here Stone (1994) sees individuals as composed of both physical and discursive components (i.e., intangible elements of the individual). Discursive components are grounded in the physical body via warranting. Walther and Parks introduced warranting to the communication field by stating that we can "conceive of a warrant [as] connecting the self with an on-line presentation" (p. 551).

The warranted self-presentation perspective extends the idea of a warrant by suggesting that senders or creators of content use warrants when crafting self-presentations. While previous work has viewed warranting as a receiver-oriented process, in which receivers make judgments about an individual, here we examine warranting from the sender or content producer perspective. In the present chapter, we argue that the

presence of a warrant should also affect the production of deception. More specifically, the central argument is that senders should lie differently when warrants are present than when they are not as warrants represent a connection between the physical self and online self-presentations. Here, we suggest that warrants should constrain how individuals use deception (e.g., frequency of lies, types of lies) to accomplish self-presentational goals.

Next we examine the warranted self-presentation perspective in a number of online contexts and show that many of the same deception principles apply in online contexts, but there are some unexpected changes in how lies are implemented to accomplish self-presentation goals in online settings.

Deception in online resumes

Deception for self-presentation purposes offers the reward of appearing more positive to others, but deceptive activities bring with them a host of consequences. One particular context where these consequences can be severe is in the professional world. High-profile examples of these consequences abound. Take Janet Cooke, for example, who lost the Pulitzer Prize in 1981 after it was discovered that she lied about her education on her resume (Kidwell, 2004). While these consequences should deter people from using extreme deceptions in resumes, the motivation to appear competent still drives deception in resumes. Research from a resume consulting service estimates that 43% of resumes contain significant inaccuracies (Cullen, 2006). In experimental research, 90% of people lied on a resume-like scholarship application, despite the fact that they were offered no explicit incentives to be deceptive (George, Marett, & Tilly, 2004). How might the online environment influence the way that people produce self-presentation deceptions in online resumes?

Most of our professional activities can now be conducted online, with people searching and applying for jobs and creating personal profiles in professional networking websites like LinkedIn. This shift in professional activities, such as sharing resumes and job applications online, has important implications for how deception should occur. While individuals are motivated to provide positive self-presentations, the warranted self-presentation perspective suggests that the publicness of resumes (i.e., accessibility of profiles to colleagues and friends) in social networking websites makes people more accountable for information shared online. When a person creates a LinkedIn profile, the site provides default settings making the profile public, creating a potential audience to which the communication partner must explain deceptions (profiles can be made private upon request). Public settings should increase the possibility that an employer might discover deceptions. Traditional resumes, on the other hand, are confidential and are not widely shared outside organizations (Rousseau, 2001). Though it's common for employers to contact references to review truthfulness of resumes, traditional resumes are limited in their ease of accessibility to others, with far fewer potential viewers to verify veracity.

The likelihood of being caught in a lie about previous employment should be higher for publicly available LinkedIn profiles than for traditional resumes. Though profile publicness does not guarantee relevant audiences (e.g., supervisors) will view

profiles, profile creators should alter deceptive behavior to be consistent with information known by potential audiences. Socially connected displays of information on these sites should constrain deception, as being detected has serious consequences (e.g., exposure of deceptions by network members).

In a recent study, we found that publicness does indeed influence deceptive behaviors in LinkedIn profiles (Guillory & Hancock, 2012). Our participants created resumes in one of three resume settings: a traditional (offline) resume, a private LinkedIn profile, or a public LinkedIn profile. Private profiles made the content available only to the participant and the experimenter and were used in the study to ensure that differences in deception were not fueled by the profile being online. Public profiles in the study made the content available to anyone online via a web search. Undergraduate student participants in the study were asked to create a profile for an entry-level job with qualifications that would be very difficult for this demographic to meet. Participants were not specifically encouraged to be deceptive in the study, but were told that the best fitting resume for the job would receive a \$100 gift card. At the conclusion of the study, participants were asked to use a retrospective identification technique to identify all deceptions in their resume.

On the surface, it would seem that making LinkedIn profiles public would reduce the frequency of deceptions that people used in profiles. Findings from this study suggest that it is not so simple. Instead, deception occurred at similar frequencies across the three resume conditions—participants lied on average three times regardless of condition. Specific types of information provided on a resume, however, were affected differentially by the presentation format. When contrasted with traditional resumes, LinkedIn resumes were less deceptive about claims that could be verified by people in a person's social network, such as an applicant's prior work experience and job responsibilities. On the other hand, creators of LinkedIn profiles were more deceptive about difficult to verify information, such as interests and hobbies, in which only the self-presenter knows the truth (e.g., favorite book, favorite hobby).

These findings make sense from the warranted self-presentation perspective: deception should be constrained by the people in a person's social network who have access to profile information. Deceptions about verifiable information, such as work experience or educational background, pose major risks to a person's reputation and future career if made public and are more likely to be classified as lies. On the other hand, when the truthfulness of resume claims is difficult to assess objectively, as is the case with hobbies or interests, there less risk of being caught in a lie and these claims are less likely to be identified as lies (Fielder & Walka, 1993; Walther & Parks, 2002). The falsifiability heuristic claims that when a person shares information about the self that is more objectively verifiable, such as observable behaviors, these claims are perceived as more credible and are less likely to be identified as lies (Fielder & Walka, 1993). Therefore, it makes sense that participants in this study practiced deception in a strategic manner by lying about different kinds of information based on whether profile claims were publicly available.

These strategic considerations surrounding deception are understandable, as research shows that the discovery of deceptions about verifiable resume claims entails consequences such as loss of jobs and awards and damage to one's professional

reputation (McCornack & Levine, 1990). To avoid these consequences, participants in our study who created public resumes lied about information that was not widely known to members of their social network and therefore less job-relevant. In one example, participants noted that the job they were applying for included a great deal of travel and lied in their resumes about interests in travel and learning new languages to appear to be an apt candidate for this position, without lying about information that would be verifiable by network members. Though unverifiable information may be less relevant to obtaining a job, our research suggests that people can still use it strategically to accomplish self-presentational goals.

These findings counter the assumption that the online environment encourages excessive deception, and alternatively suggests that features of social networking websites that make information publicly available actually help to encourage more honesty in the context of online resumes compared to traditional paper resumes. In this professional context, the consequences for deceptive behavior are often grave (e.g., Kidwell, 2004) and could involve damage to one's professional reputation or even job loss. Will deceptive behaviors be influenced in the same way in a broader range of contexts with less serious consequences surrounding deception?

Warranted self-presentation deceptions across communication contexts

How does the presence of warrants to a greater or lesser degree influence deception in different online communication contexts? A recent study (Warkentin et al., 2010) set out to explore the way that differing levels of anonymity in online communication contexts affected the frequency and types of deception that people use. This research used the warranted self-presentation perspective to predict that deception would occur less frequently in communication contexts with fewer warrants. Three types of warrants were considered: (a) name warrants, or references to a person's real name; (b) photo warrants, which included recognizable photographs of a person that could be visually linked to their real-world identity; and (c) acquaintance warrants, which were defined by the degree to which a person's communication partners knew about his or her identity offline.

In order to explore the role that these different types of warrants play in deceptive behaviors, they surveyed participants about their deceptive behaviors in email, IM, forums, chat rooms, and social networking websites. Consistent with their predictions, they found that more warrants led to fewer lies. They also found that people told less serious lies in the presence of a more warrants. When they considered the influence of the specific type of warrant on deception, they found that acquaintance warrants were the main factor driving the reduction in both the frequency and magnitude of deception.

These findings suggest that warrants in online communication act as a major factor influencing deception across a wide variety of online communication contexts, and provide some support for a more sender-oriented perspective of warrants. That is, in addition to affecting how we judge the credibility of information online, the presence of warrants also constrain the production of deceptive information.

Love and Lies: Deception in Online Dating

The self-presentational goal of appearing attractive is particularly salient in the context of online dating. In wanting to appear attractive in these online self-presentations, daters may consider using deception in their dating profiles to improve other daters' perceptions of their attractiveness. While self-presentation deception is perceived as occurring frequently in these spaces, it should be buffered by daters' desire to appear honest given that they hope to 1 day meet a romantic partner in person. This prospect of anticipated future interaction should constrain people from using extreme deceptions that would be revealed as lies if they were to meet the romantic partner in person (e.g., claiming to have a "toned and athletic figure" when the individual is overweight).

Toma, Hancock, and Ellison (2008; Ellison, Hancock, & Toma, 2012; Hancock & Toma, 2009; Toma & Hancock, 2010) conducted a series of studies exploring deceptive processes in online dating. In the first of these studies, Toma et al. (2008) compared content from online dating profiles, such as profile holders' stated height, to their observed characteristics, namely how tall they were as measured in the lab, to determine how much the daters' self-presentation in their dating profile deviated from their actual characteristics. Because many online daters hope to eventually meet romantic partners in person, Toma et al. (2008) also predicted that lies in profiles would be small in magnitude. Deception in online dating profiles should also be driven in part by gender differences. In addition, based on characteristics that men and women seek in mates, they predicted that men would lie more about social status, occupation, education, and height (as being tall is associated with higher social status in men). In contrast, women should lie more about age and weight, which is an indicator of physical attractiveness and valued by men.

In line with predictions, Toma et al. observed that the magnitude of deceptions that people shared was rather small. The average deviation in height was less than 1 in. (2.09% of participants' actual height), the average deviation in weight was less than 10 pounds (5.5% of participants' actual weight), and the average deviation in age was 0.55 years (1.4% of participants' actual age). While men and women subtle in their deceptions about age, weight and height, key gender differences emerged. For example, consistent with expectations from evolutionary psychology (e.g., Buss & Schmitt, 1993), men were more likely to over-report their height than women, while women were more likely to under-report their weight than men.

The findings from this initial study illustrate the tension between the self-presentation goal of appearing attractive to other daters and the desire to appear honest in anticipation of future interactions. In a follow-up study, Hancock and Toma (2009) explored deception in the context of photographs from online dating profiles. While the process of selective self-presentation has been explored extensively in computer-mediated communication contexts, few studies have explored how these processes occur in nontextual aspects of computer-mediated communication (e.g., photographs). Photographs may be altered in many ways. People can choose flattering camera angles and scenery. In more extreme cases, people can use tools, such as Adobe Photoshop, to edit the content of photos (e.g., removing blemishes or other unattractive physical features from photos).

Similar to the initial study, Hancock and Toma (2009) predicted that photographic deception would differ by gender based on characteristics that men and women value in potential mates (Lance, 1998; Woll & Cozby, 1987). Because men should value physical attractiveness more than women, they predicted that women's photographs would contain more deceptions. Because men should value youthfulness in female mates, they also predicted that women would provide younger photographs than men.

In order to study photographic deception, which is more difficult to objectively verify than something like height or weight, participants were first asked to rate their photographs for inaccuracies. Next participants posed for a photograph to obtain a record of their everyday appearance. They then asked independent judges to compare the everyday photograph of the individual to the profile photograph. Consistent with predictions, women's photos contained a greater number of discrepancies between the everyday photograph and the profile photograph than men's photographs. The data also suggest that women chose profile photographs that were significantly older (based on self-ratings) than men's photographs, which is consistent with the prediction that women would attempt to present a younger version of the self to other daters.

While these studies identify deceptive patterns in the context of online dating, they provide little information about how possessing certain qualities will lead men and women to engage in different deceptive strategies in this context. In another study, the authors (Toma & Hancock, 2010) explored the role that physical attractiveness plays in deception in online dating. The authors predicted that in cases where individuals were less physically attractive, they would compensate by using more deception, both in photographs and self-descriptions (e.g., height, weight, age). They also predicted that this tendency would be stronger for women, consistent with findings from their previous two studies (Hancock & Toma, 2009; Toma et al., 2008) and theory that suggests men value physical attractiveness more than women in mates (Lance, 1998; Woll & Cozby, 1987). In addition, they predicted that more attractive daters would want to take advantage of this positive quality by posting more photographs to their dating profile. In line with their predictions, Toma and Hancock (2010) found that less attractive participants were more prone to engaging in deception, both in photographs and self-descriptions (e.g., height, weight, age), while attractive daters were indeed likely to take advantage of their attractiveness by posting more photographs than unattractive daters.

Findings from this series of studies provide strong evidence that self-presentation deception is an important tactic in online dating profiles and that these deceptions are driven by a desire to appeal to specific characteristics that are desirable to men and women, respectively. In addition, this research suggests that having (or not having) certain qualities, such as higher levels of attractiveness, leads to different patterns of deception in online dating profiles.

On My Way: Butler Lies

The cases discussed earlier focus on the use of deception to accomplish self-presentational goals that make the individual appear more attractive, be it for a job or potential mate. Recall that DePaulo would refer to these lies as self-oriented lies, or those that are

told to protect or enhance one's self-presentation. There are also cases where people use lies that have an other-oriented component. Other-oriented lies are those told to protect others' feelings in order to maintain positive self-presentations. These types of lies come up quite often when people are attempting to manage their availability (i.e., initiating, concluding, or coordinating social interactions). Text-based communication conveniently provides opportunities to access our social contacts regardless of the time or location. With this convenience comes the burden of constantly being available to people in our social network, even at times when communication is unwanted or difficult. As a way to deal with this constant availability, we have observed that people have developed deceptive strategies to avoid social interaction or account for failure to communicate. We refer to these strategies as *butler lies*. Butler lies are told to manage social interaction and awareness either by avoiding a new conversation, smoothly exiting an ongoing conversation, or explaining other communication behavior.

In an initial study, we explored butler lies in the context of instant messaging (IM) (Hancock et al., 2009). Participants in this study were asked to rate the deceptiveness of all outgoing instant messages from their personal computer during a 4-day period. Many of the lies that participants told were related to either to avoid starting a new conversation, to exit an ongoing conversation, or to arrange or explain other interactions or failure to interact with a person. Our analysis of the content of these butler lies suggests that these lies were other-oriented self-presentational deceptions that were told to manage interpersonal impressions when the individual wanted to engage in an activity that could potentially hurt the feelings of their communication partner, such as ending a conversation early or avoiding a conversation altogether.

While understanding how people use butler lies in IM conversations provides a small window into people's availability management strategies, communication via mobile phones, primarily SMS text messaging, helps us to understand how people manage their availability in wider range of contexts. In a second study, we explored butler lies in the context of SMS text messaging and found that while these messages occurred in similar frequency to those that we observed in IM conversations, the types of butler lies used differed from IM. While people primarily used butler lies as a strategy for exiting conversations in IM, but lies in the SMS context were used primarily to avoid social interactions (e.g., "Can't meet up later, gotta work"), which can be attributed to the fact that SMS is often used to coordinate social interactions (Grinter & Eldridge, 2003). In a second text messaging study, we compared Blackberry Messenger (BBM) with SMS (Reynolds et al., 2011). This follow-up study replicated findings from the previous two studies and also demonstrated the significance of people's relationships in driving the use of butler lies. Participants tended to tell more butler lies to the people they made plans with the most, which implies that these types of lies are an essential component of relationship management.

Findings from these studies demonstrate several important points about the availability management strategies. First, we see that people have difficulty managing their availability given the constraints of communication technology and have developed their own strategies to manage availability in these situations. Second, these studies show us that a major part of our availability management involves maintaining our

relationships, highlighting once again that deception is often a reasonable response to technology that takes into consideration relational factors, rather than technology simply increasing deception.

Conclusion

Whether its butler lies, sock puppets or the Chinese Water Army, there is no question that technology and deception are influencing how people communicate with one another, from the everyday interactions on mobile texting to the politically significant scandals revealed by the digital traces left behind by communication technology. It is also clear that the relationship between deception and technology is not a straightforward one. Although powerful heuristics lead us to believe that technology facilitates deception, it is much more complicated than that.

In the present chapter, we have argued that understanding our biases around deception, and around how technology affects us, are important to understand and diagnose in order to understand the impact of technology on deception. We also offer an important alternative approach that considers transformations to the interaction space wrought by new technologies. These transformations include the changes to the visual presentation of information, the features of novel communication spaces (e.g., recordability), the ways in which social networks now link and constrain us online, and the importance of relational concerns in understanding how and why we lie with technology. By approaching deception in this way we believe that we can advance our understanding of this very old human phenomena in these very novel times.

References

- Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, *9*, 272–279.
- Ariely, D. (2012). *The honest truth about dishonesty*. New York, NY: HarperCollins.
- Backstrom, L., & Kleinberg, J. (2014). Romantic partnerships and the dispersion of social ties: A network analysis of relationship status on Facebook. In S. Fussell, W. Lutters, M. Ringel Morris, & M. Reddy (Eds.), *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW2014)*, Baltimore, MD. New York, NY: ACM.
- Bateson, M., Nettle, D., & Roberts, G. (2006). Cues of being watched enhance cooperation in a real world setting. *Biological Letters*, *2*, 412–414.
- Bond, C. F., Jr., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review*, *10*, 214–234.
- Buss, D. M., & Schmitt, D. P. (1993). Sexual strategies theory: An evolutionary perspective on human mating. *Psychological Review*, *100*, 204–232.
- Cullen, L. T. (2006). Getting wise to lies. *Time*, *167*, 59–60.
- Dediu, D., & Levinson, C. (2013). On the antiquity of language: The reinterpretation of Neanderthal linguistic capacities and its consequences. *Frontiers in Psychology*, *4*, 397.

- DePaulo, B. M., & Kashy, D. A. (1998). Everyday lies in close and casual relationships. *Journal of Personality and Social Psychology*, 74, 63–79.
- DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology*, 70, 979.
- DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, 129, 74–118.
- Dhamija, R., Cassidy, P., Hallam-Baker, P., & Jacobsson, M. (2006). Phishing phishing. In A. L. Car et al. (Eds.), *Proceedings of the ACM 15th international conference on World Wide Web (WWW2006)* (pp. 749–749). Edinburgh, UK: ACM.
- Dhamija, R., Tygar, J. D., & Hearst, M. (2006, April). Why phishing works. In P. M. Aoki et al. (Eds.), *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 581–590). New York, NY: ACM.
- Ekman, P. (1985). *Telling lies: Clues to deceit in the marketplace, politics and marriage*. New York, NY: Norton.
- Ellison, N., Hancock, J. T., & Toma, C. (2012). Profile as promise: A framework for conceptualizing the veracity of online dating profiles. *New Media & Society*, 14, 42–59.
- Ennis, E., Vrij, A., & Chance, C. (2008). Individual differences and lying in everyday life. *Journal of Social and Personal Relationships*, 25, 105–118.
- Fielder, K., & Walka, I. (1993). Training lie detectors to use nonverbal cues instead of global heuristics. *Human Communication Research*, 20, 199–223.
- George, J., Marett, K., & Tilly, P. (2004). Deception detection under varying electronic media and warning conditions. In J. F. George, K. Marett, & P. Tilley (Eds.), *Proceedings of 37th Hawaii International Conference on System Sciences*. New York, NY: IEEE Computer Society.
- Grinter, R., & Eldridge, M. (2003, April). Wan2tlk? Everyday text messaging. In P. Korhonen (Eds.), *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 441–448). New York, NY: ACM.
- Guillory, J., & Hancock, J. T. (2012). The effect of LinkedIn on deception in resumes. *Cyberpsychology, Behavior, and Social Networking*, 15, 135–140.
- Haley, K. J., & Fessler, D. M. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior*, 26(3), 245–256.
- Hancock, J. T., Birnholtz, J. P., Bazarova, N. N., Gullory, J., Perlin, J., & Amos, B. (2009). Butler lies: Awareness, deception, and design. In D. R. Olsen Jr. et al., (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 517–526). New York, NY: ACM.
- Hancock, J. T. (2012). *The future of lying*. Technology, Education & Design (TED), Winnipeg, MN. Retrieved from http://www.ted.com/talks/jeff_hancock_3_types_of_digital_lies
- Hancock, J. T., Thom-Santelli, J., & Ritchie, T. (2004a). Deception and design: The impact of communication technology on lying behavior. In M. Tscheligi (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 129–134). New York, NY: ACM.
- Hancock, J. T., Thom-Santelli, J., & Ritchie, T. (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.
- Hancock, J. T., & Toma, C. L. (2009). Putting your best face forward: The accuracy of online dating photographs. *Journal of Communication*, 59, 367–386.
- Hancock, J. T., Woodworth, M. T., & Goorha, S. (2010). See no evil: The effect of communication medium and motivation on deception detection. *Group Decision and Negotiation*, 19, 327–343.

- Lance, L. L. (1998). Gender differences in heterosexual dating: A content analysis of personal ads. *The Journal of Men's Studies*, 6(3), 297–305.
- Levine, T. R., Kim, R. K., & Hamel, L. M. (2010). People lie for a reason: Three experiments documenting the principle of veracity. *Communication Research Reports*, 27, 271–285.
- Levine, T. R., Serota, K. B., & Shulman, H. C. (2010). The impact of Lie to Me on viewers' actual ability to detect deception. *Communication Research*, 37, 847–856.
- Keyes, R. (2004). *The post-truth era: Dishonesty and deception in contemporary life*. New York, NY: St. Martin's Press.
- Kidwell, R. E. (2004). "Small" lies, big trouble: The unfortunate consequences of résumé padding, from Janet Cooke to George O'Leary. *Journal of Business Ethics*, 51, 175–184.
- McCornack, S. A., & Levine, T. R. (1990). When lovers become leery: The relationship between suspicion and accuracy in detecting deception. *Communication Monographs*, 57, 219–230.
- Nettle, D., Harper, Z., Kidson, A., Stone, R., Penton-Voak, I. S., & Bateson, M. (2013). The watching eyes effect in the Dictator Game: It's not how much you give, it's being seen to give something. *Evolution and Human Behavior*, 34, 35–40.
- Ott, M., Cardie, C., & Hancock, J. T. (2012). Estimating the prevalence of deception in online review communities. *Proceedings of International World Wide Web Conference. Lyon, France. 2012 (WWW2013)*.
- Porter, S., & ten Brinke, L. (2008). Reading between the lies: Identifying concealed and falsified emotions in universal facial expressions. *Psychological Science*, 19, 508–514.
- Pronin, E. (2006). Perception and misperception of bias in human judgment. *Trends in Cognitive Sciences*, 11, 37–43.
- Resnick, P., Kuwabara, K., Zeckhauser, R., & Friedman, E. (2000). Reputation systems. *Communications of the ACM*, 43, 45–48.
- Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40, 56–58.
- Reynolds, L., Gillette, S., Marder, J., Miles, Z., Vodenski, P., Weintraub, A., ... Hancock, J. T. (2011). Contact stratification and deception: Blackberry messenger versus SMS use among students. In P. Hinds (Eds.), *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work* (pp. 221–224). New York, NY: ACM.
- Rousseau, D. M. (2001). Schema, promise and mutuality: The building blocks of the psychological contract. *Journal of Occupational and Organizational Psychology*, 74, 511–541.
- Spottswood, E., & Hancock, J. T. (2013, June). The eyes have it: How publicness and surveillance primes influence prosocial communication on Facebook. *International Communication Association 63rd Annual Conference*, London, UK.
- Stajano, F., & Wilson, P. (2011). Understanding scam victims: Seven principles for systems security. *Communication of the ACM*, 54, 70–75.
- Stone, A. R. (1994). Split subjects, not atoms; or, how I fell in love with my prosthesis. *Configurations*, 2, 173–190.
- Stone, A. R. (1995). *The war of desire and technology at the close of the mechanical age*. Cambridge, MA: MIT Press.
- Toma, C. L., & Hancock, J. T. (2010). Looks and lies: The role of physical attractiveness in online dating self-presentation and deception. *Communication Research*, 37, 335–351.
- Toma, C. L., Hancock, J. T., & Ellison, N. B. (2008). Separating fact from fiction: An examination of deceptive self-presentation in online dating profiles. *Personality and Social Psychology Bulletin*, 34, 1023–1036.

- Toma, C. L., Jiang, L. C., & Hancock, J. T. (2013, June). *Lies in the eyes of the beholder: Self-other asymmetry in beliefs about deception across interpersonal media*. Paper presented at the 63rd Annual Conference of the International Communication Association, London, UK.
- Walther, J. B. (1996). Computer-mediated communication impersonal, interpersonal, and hyperpersonal interaction. *Communication research*, 23, 3–43.
- Walther, J. B., & Parks, M. R. (2002). Cues filtered out, cues filtered in: Computer-mediated communication and relationships. In I. M. Knapp & J. A. Daly (Eds.), *Handbook of interpersonal communication* (3rd ed., pp. 529–563). Thousand Oaks, CA: Sage.
- Warkentin, D., Woodworth, M., Hancock, J. T., & Cormier, N. (2010). Warrants and deception in computer-mediated communication. In K. Inkpen (Eds.), *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work* (pp. 9–12). New York, NY: ACM.
- Whitty, M. T., Buchanan, T., Joinson, A. N., & Meredith, A. (2012). Not all lies are spontaneous: An examination of deception across different modes of communication. *Journal of the American Society for Information Science and Technology*, 63, 208–216.
- Woll, S. B., & Cozby, P. C. (1987). Videodating and other alternatives to traditional methods of relationship initiation. *Advances in Personal Relationships*, 1, 69–108.
- Zimmerman, J., & Kurapati, K. (2002, April). Exposing profiles to build trust in a recommender. In L. Terveen et al. (Eds.), *CHI'02 Extended Abstracts on Human Factors in Computing Systems* (pp. 608–609). New York, NY: ACM.