

“I’m 13. I’m Online. U Believe Me?”: Implications for Undercover Internet Stings

Michelle Drouin

Indiana University–Purdue University Fort Wayne

Vincent Egan

University of Nottingham

Nicholas Yergens and Elisa Hernandez
Indiana University–Purdue University Fort Wayne

The extent to which 262 adults, recruited from a U.S. university sample and Amazon’s Mechanical Turk, believed the identity of a 13-year-old boy or girl persona in an online chat room was examined. Sixteen undergraduate confederates (aged 19–38) followed 2 basic biographical sketches (“Amber” or “James”) and were instructed to “chat like a teenager” in 1-to-1 private chats. Confederates cycled through conditions in which they provided no age or gender information (control), stated age and gender, stated age and gender and provided a picture (attractive or average), or stated grade only. In all but the control condition, participants estimated the average age of their chat partner to be 13 to 14. Meanwhile, when confederates stated their age and gender, 83% to 88% of participants believed the confederate’s stated age, and 94% to 98% believed stated gender. When asked about cues used to discern age and gender, most participants (98%) used multiple cues, including stated age and gender, content cues, style cues, and picture. Moreover, natural language analyses showed that confederates used significantly fewer analytic and 6-letter words, and displayed less clout than the adults with whom they chatted, and confederates who used more analytic words were perceived as older. These findings contradict Lincoln and Coyle (2013) and suggest that even confederates who are not specially trained law enforcement agents are able to deceive others about their age and gender in online chat rooms. This has important implications for those conducting undercover Internet stings and those who prosecute those cases.

Keywords: online deception, Internet sting, undercover agents, sex offenders, prosecution

Supplemental materials: <http://dx.doi.org/10.1037/law0000149.supp>

There is an ever-growing body of research examining deception in online environments (e.g., Buchanan & Whitty, 2014; Caspi & Gorsky, 2006; Drouin, Miller, Wehle, & Hernandez, 2016; Ellison, Hancock, & Toma, 2012; Guadagno, Okdie, & Kruse, 2012; Hancock & Woodworth, 2013; Toma, Jiang, & Hancock, 2016; Warkentin, Woodworth, Hancock, & Cormier, 2010; Whitty, 2002). The findings across these studies are fairly consistent: People admit to lying online and they assume others are lying, too. Although differences in people’s expectations of lying behavior

across different online forums are partly dependent on how much anonymity and relational distance the forum provides (e.g., Drouin et al., 2016; Toma et al., 2016; Warkentin et al., 2010), an undercurrent of distrust saturates the online world, epitomized by one study participant’s observation, “Everyone lies on the Internet” (Drouin et al., 2016, p. 134).

The default assumption that people are lying online leads to many social and legal issues. In principle, people may be reticent to form online relationships with strangers until unknown others are vetted through proofs of identity like photographs or shared acquaintances. Warkentin et al. (2010) refer to these as *warrants*, and Whitty and Joinson (2009) assert that these types of requests are common when people are trying to establish whether one’s online identity is a truthful representation of their offline self. From a legal perspective, those who become involved in illegal activities online (e.g., sexual solicitation of a minor, planning of other types of criminal activity) can plausibly argue they did not believe the person with whom they were communicating was really who they said they were. In support of this, Lincoln and Coyle (2013) found college participants very accurate at discerning the actual age and gender of their online chat partners, even when those adult chat partners were pretending to be a 13-year-old girl. Given that the online ruse was ineffective, it may be that law enforcement agencies conducting covert sting operations could be easily compromised.

This article was published Online First November 30, 2017.

Michelle Drouin, Department of Psychology, Indiana University–Purdue University Fort Wayne; Vincent Egan, Centre for Forensic and Family Psychology, Department of Psychiatry and Applied Psychology, University of Nottingham; Nicholas Yergens and Elisa Hernandez, Department of Psychology, Indiana University–Purdue University Fort Wayne.

Some preliminary findings from this study were presented at the Midwestern Psychological Association Annual Meeting in Chicago, Illinois (April 2017). Additionally, some of the main findings were presented at the Decepticon Conference at Stanford University (August 2017).

Correspondence concerning this article should be addressed to Michelle Drouin, Department of Psychology, Indiana University–Purdue University Fort Wayne, 2101 E. Coliseum Boulevard, Fort Wayne, IN 46805. E-mail: drouinm@ipfw.edu

In this study, Lincoln and Coyle's (2013) study was replicated using different settings and experimental conditions to more closely mimic the types of interactions that might occur within real police sting operations. The main research aim was to determine the extent to which people believed a false persona encountered in an online chat room.

Identity Deception: A Legal Perspective

The issue of online identity deception is relevant for a variety of legal cases. One type of deception that many are familiar with is the 419 advance-fee online scam, in which someone pretends to be a Nigerian Prince, a stranded traveler, or a government agent in order to get money or personal information from an online target (Isacenkova, Thonnard, Costin, Francillon, & Balzarotti, 2014). These scams use sophisticated social engineering tactics to trick online targets to send money or personal information via e-mail, chat rooms, and even through online dating sites and social networking sites (Muscanell, Guadagno, & Murphy, 2014). Another is "catfishing," in which persons create false online personas for deceptive relationships (Drouin et al., 2016). More relevant to the current inquiry is a variant of online deception not carried out by would-be criminals or persons seeking to boost their mate value in a competitive field, but by law enforcement officials and their vigilante equivalents (Drouin, Boyd, Hancock, & James, 2017; Egan, Hoskinson, & Shewan, 2011): the undercover sting operation. In undercover Internet sting operations, law enforcement or other trained confederates adopt a false online persona in order to identify, gather evidence about, and/or apprehend those engaged in illegal activities. One popular version of this is the undercover sex sting, whereby law enforcement officials pretend to be an adolescent girl or boy in order to identify and apprehend adults who are attempting to solicit children or teenagers online for sexual activity. These proactive investigations seek to identify prospective child abusers before they are able to offend with an actual child. However, the strength of proactivity may become a prosecutorial hurdle, as the legal burden of proof then rests with the prosecution to prove that the defendant believed they were communicating with an actual child and/or that they had intent of committing child sexual abuse (e.g., Wright, 2009). This burden is compounded by empirical work demonstrating that estimation of a young person's age can be difficult even when persons are not seeking to dissimulate (Egan & Cordan, 2009).

Many of those apprehended in these sex sting operations insist that they did not believe they were communicating with an actual child. This has been referred to as the fantasy/role-play defense (Colleluori, 2010; Wright, 2009), and although it is sometimes employed, few such defenses are upheld (Colleluori, 2010). Consider, for example, the case of former U.N. employee Scott Ritter. In 2009, Ritter had a conversation with 15-year-old "Emily," masturbated on camera, and then announced his intent to sign out of the Yahoo chat room they were using. When "Emily" revealed himself as law enforcement, Ritter reportedly replied, "Nah, your not 15. . . . It's all fantasy" (Franklin, 2012). Although Ritter insisted that he always thought he was communicating with an adult (fantasy defense), he had two prior arrests with similar charges, and was sentenced to 18 months to 5.5 years in prison (Franklin, 2012).

However, there have also been cases in which the fantasy/role-play defense has been successfully employed. In a case detailed in Lincoln and Coyle (2013), Darryl Plumridge of Queensland, Australia, was charged in an Internet sex sting in which a middle-aged male police officer posed as a 13-year-old girl (*R v. Plumridge*, 2007). In his defense, Plumridge claimed that he thought he was role playing with an adult male, noting that the content of the messages was not consistent with what a 13-year-old girl would say. For example, it took his chat partner some time to negotiate how to send a picture file, and his chat partner used antiquated phrases like "See ya later alligator." Most notably, the undercover agent inadvertently stated that "she" was at the office rather than stating that she was at school. Defense counsel for the case cited s 218A(8) of the Criminal Code of Queensland, proposing that these and other linguistic cues constituted "evidence to the contrary" of the notion that the person was under the age of 16 (as cited in Lincoln & Coyle, 2013). The judge allowed expert testimony from a psychologist and a linguist, the jury was convinced, and Plumridge was acquitted of the crime.

Identity Deception: An Empirical Perspective

There is an empirical basis for the claims made using the fantasy/role-play defense. For more than two decades, researchers have been studying cases of online identity deception. In one of the first studied incidents, Van Gelder (1991) described a case of a prominent male psychiatrist who created an online persona, "Joan," and then proceeded to develop online relationships with many women (some of whom shared computer-mediated sexual experiences with Joan) via an online discussion board. Since that time, various other examples of this type of deception have been highlighted as case studies in the research literature (e.g., Feldman, 2000; Joinson & Dietz-Uhler, 2002). In less extreme cases, 25% to 60% of chat room users, online daters, and discussion board participants admit to lying about some aspect of their identity (i.e., their age or gender) at some point within online forums (Caspi & Gorsky, 2006; Whitty, 2008; Whitty & Joinson, 2009). In line with these statistics, a recent study by Drouin et al. (2016) showed that 77% to 88% of adults believed that others were at least sometimes lying about their age, and 39% to 76% were at least sometimes lying about their gender on social media, online dating, chat rooms, and on sexual chat sites.

Although many people believe that others lie online, Caspi and Gorsky (2006) noted that few people believe that others had deceived *them* online. Moreover, only 20% of those who admitted to online deception thought that the person they were communicating with *did not* believe them. Thus, a paradox seems to exist, whereby online deception appears to be rampant; however, few feel that they, personally, are being lied to online. This bias is in keeping with the common cognitive distortions associated with subjective overconfidence (Keren, 1997). As only 20% of those who were lying online thought that their deception had been detected, the following empirical question arises: How sensitive are people to others' online identity deception?

In one of the only known empirical studies on the topic, Lincoln and Coyle (2013) addressed this question experimentally and found that, generally, people are quite sensitive to others' online identity deception. Lincoln and Coyle analyzed the data from 20 pairs of undergraduate and postgraduate students (aged 18–38)

from Australia who participated in an experimental protocol in which they engaged in an online conversation with an unknown chat partner. In their study, participants were recruited in pairs, and were directed separately into different rooms on a college campus where they accessed an online chat portal. Participants were assigned, based on whether they arrived first or second, to be a deceiver (Group A) or a receiver (Group B) in their chat conversation. Those assigned to the deceiver role were “instructed to play the role of a 13-year-old female while participating in the online chat: that is, they deliberately tried to deceive their interlocutor in Group B (‘receivers’) as to their age and sex” (p. 297). After being instructed that the discussion topic was open (provided communication was not defamatory), the receiver and deceiver then engaged in a chat for a maximum of 30 min. Afterward, receivers were asked the age and sex of their chat partner, and how they determined this information. They found that 16 of 20 of the receivers correctly guessed their deceiver’s actual gender, and 15 of 20 correctly guessed their deceiver’s actual age (within a 5-year bandwidth). None of the receivers guessed that their chat partner was under the age of 16, indicating *none* of the deceivers were successful in convincing their chat partner that they were a 13-year-old girl. Although adolescent participants in May-Chahal et al.’s (2014) cited using both content (i.e., what the person wrote about) and style cues (e.g., how messages were written) to discern age and gender among adults and teens, Lincoln and Coyle found that most of their adult participants used content rather than style cues to discern their chat partner’s age and gender. Based on these findings, they suggested certain types of linguistic content (e.g., content that denotes a high IQ) could “give away the game” for undercover agents.

Although the Lincoln and Coyle (2013) study is important to the field of online deception, there are several methodological issues that limit the generalizability of its findings. First, Lincoln and Coyle recruited college students and brought them into rooms around campus, and participants used the university’s intranet to engage in their chat. College samples are frequently employed in social science research; soliciting participants on campus and having them engage in a chat through the university’s intranet may have led receivers to surmise they were communicating with someone who was also college-aged. Second, the deceivers were untrained and not provided with a backstory—they were simply told upon arrival to impersonate a 13-year-old girl. With no solid backstory and no experience in reading chats between stings and adults, these deceivers had to concoct spontaneous lies, which may be less plausible than prepared dissimulation (e.g., O’Hair, Cody, & McLaughlin, 1981). Third, it is unclear in the Lincoln and Coyle study whether the deceivers stated that they were a 13-year-old girl, and it does not appear that the deceivers sent a photo as proof of their identity. However, it is very clear (and mandatory) that law enforcement officials state their age in Internet sex sting operations, and also common for undercover agents to display or send photo(s) of a young teen as proof of their identity (<http://perverted-justice.com/?con=full>). Because of these methodological limitations, it is impossible to determine whether the receivers in the Lincoln and Coyle study were actually sensitive to the deceivers’ cues to their real identities, or whether situational factors (e.g., setting and method) influenced their assessments of their chat partner’s age and gender.

Another important aspect of these conversations that Lincoln and Coyle (2013) ignored was their linguistic properties. Linguistic evidence proved vital in *R v. Plumridge* (2007), as, according to defense counsel, the language used by the undercover agent provided cues to Plumridge that his chat partner was over the age of 16. As this issue remains unexplored, the question remains: What types of linguistic cues, aside from obvious slips from an undercover agent that denote adulthood (e.g., “I’m at work”), might cue a target that they are interacting with someone older? There are marked cognitive and social changes between early adolescence (12–14) and mid-adolescence (15–17; see Meschke, Peter, & Bartholomae, 2012, for a review). A study by Simmonds, Hallquist, Asato, and Luna (2014) showed that frontocortical white matter connections mature at around 15 to 16 in most individuals, which corresponds to the greater relative cognitive and behavioral maturity at this age compared with earlier adolescence. Thus, one might expect to see patterns of intellectual or social maturity in the language of adults compared with young adolescents.

One promising method for studying the content of communication is computerized natural language analysis using the Linguistic Inquiry and Word Count (LIWC) program (Pennebaker, Boyd, Jordan, & Blackburn, 2015). The LIWC program processes natural language data and quantifies it in terms of word use patterns into approximately 80 dictionary-based categories. Currently, it is the gold standard in computerized linguistic analysis, having been employed in a variety of disciplinary settings for the past 20 years (e.g., Bantum & Owen, 2009; Iosub, Laniado, Castillo, Fuster Morell, & Kaltenbrunner, 2014; Tani, Peterson, & Smort, 2016). LIWC has been used in forensic settings for studies of online deception, online grooming, and recidivism among sex offenders (Black, Wollis, Woodworth, & Hancock, 2015; Bond & Lee, 2005; Drouin, Boyd, & Greidanus Romaneli, 2017; Drouin, Boyd, Hancock, et al., 2017).

Corresponding to the changes in cognitive and social development from early adolescence to adulthood, Drouin, Boyd, Hancock, et al. (2017) identified three relevant LIWC categories to aid differentiation: clout (social dominance), analytic words, and six-letter words in the chat transcripts. Clout is a composite variable, comprised of several LIWC categories including personal pronouns (Pennebaker et al., 2015), which vary in usage dependent on social standing (Kacewicz, Pennebaker, Davis, Jeon, & Graesser, 2014). The analytic category contains word dimensions that depict formal, logical thinking (Pennebaker et al., 2015), and the six-letter word category comprises the percentage of the target’s words containing six or more letters. As confederates were pretending to be 13 years old, they were instructed to use unsophisticated language (i.e., more abbreviations and slang), in line with linguistic analyses that show that these features are regarded by teens as language markers of adolescence (May-Chahal et al., 2014). Consistent with the 600-plus chats of online sex stings archived on Perverted Justice website (<http://perverted-justice.com/?con=full>), confederates were instructed to display less social dominance and reasoning than their adult chat partners. Supporting this method, a recent linguistic analysis of chat transcripts showed that undercover agents pretending to be teenagers (aged 11–15) had significantly lower clout scores than the adult offenders with whom they chatted (Drouin, Boyd, Hancock, et al., 2017), and college students with lower cognitive performance (as measured by GPA) used fewer words from the analytic words category than did higher

performers in their college admissions essays (Pennebaker, Chung, Frazee, Lavergne, & Beaver, 2014).

Current Study

In this study, Lincoln and Coyle's (2013) study was replicated to examine the extent to which people were sensitive to others' online deception, expanding upon it so that the methods more closely resemble those employed in actual sex sting operations. For example, both a college sample and a sample recruited from Amazon's Mechanical Turk (MTurk) were employed, and all data collection was conducted online. Additionally, instead of using random college students as neutral "deceivers," a short backstory (see Appendix in the online supplemental materials) was created, which the confederates used as a basis for their online persona. Finally, five different experimental conditions ranging from "no statement of age or gender" to "stating age and gender and sending a picture" were employed to examine how stating one's age and gender might influence perceptions of a confederate's age and gender. As an exploratory variable, an attractive-picture condition and an average-picture condition were included, as it is often mentioned that those who are catfishing (i.e., lying about their online identity) use attractive pictures of others in their online subterfuge (e.g., Moss, 2015).

Based on the findings in Lincoln and Coyle (2013), we expected the following:

Hypothesis 1: Adults should be sensitive to online deception and should be able to discern the actual age and gender of online chat partners within a 5-year window.

Hypothesis 2: Individuals should use content cues more than style cues to make these age and gender determinations.

Additionally, based on previous computerized linguistic analyses related to cognitive and social maturity (e.g., Drouin, Boyd, Hancock, et al., 2017; Pennebaker et al., 2014), we expected the following:

Hypothesis 3: Confederates should display less clout, use fewer analytic words, and use fewer six-letter words than the adults with whom they communicated, and there should be a positive relationship between the confederate's clout, analytic word, and six-letter word scores and the estimated age of the confederate.

Method

Participants

Participants ($N = 262$; women = 152, men = 110) were recruited from both a psychology department subject pool at a midsized Midwestern university ($n = 136$) and Amazon's MTurk worker pool, including only residents in the United States ($n = 126$). In both settings, participants were recruited to an institutional review board (IRB) approved "academic study on online interactions," for which they would be "communicating with another person." Although a large number of participants completed the initial online consent form and a few demographic questions (University, $N = 283$; MTurk, $N = 339$), many participants (55%

University and 63% MTurk) did not complete the 30-min online chat.¹ A dropout analysis for the university and MTurk samples showed that those who did not complete the survey were not significantly different than those who did complete the survey in terms of age, ethnicity, sexual orientation, or frequency of computer device usage (all $ps > .44$). Only participants who completed the chat were included in this sample. The average age of the participants was 25.69 years ($SD = 8.52$). In terms of ethnicity, most of the participants identified themselves as Caucasian (76%), 9% were Hispanic, 6% Asian, 6% African American, and 4% were Biracial, Native American, or Other. Additionally, most of the sample characterized their sexual orientation as heterosexual (90%), 5% identified as gay/lesbian, 4% identified as bisexual, and 2% indicated their sexual orientation was something other than those listed (e.g., asexual or pansexual). On average, participants reported using any computer device (i.e., cell phone, computer, tablet) an average of 27.63 hr per week ($SD = 29.38$).

In terms of sample comparability, previous studies have found that MTurk samples produce reliable and valid data (Buhrmester, Kwang, & Gosling, 2011; Rouse, 2015) and perform similarly to other samples on a variety of tasks (e.g., Buhrmester et al., 2011; Casler, Bickel, & Hackett, 2013; Feitosa, Joseph, & Newman, 2015; Hauser & Schwarz, 2016); however, there are often sample characteristic differences between college and MTurk samples (e.g., Hauser & Schwarz, 2016). In this sample, the MTurk participants were significantly older than the college participants ($M = 31.41$, $SD = 8.64$ vs. $M = 20.38$, $SD = 3.51$, respectively), $t(1,260) = 13.35$, $p < .001$, and the MTurk sample was significantly more balanced than the college sample in terms of gender (52% male vs. 33% male, respectively), $\chi^2(1, N = 262) = 9.19$, $p = .002$, and race (67% Caucasian vs. 84% Caucasian, respectively), $\chi^2(1, N = 262) = 9.59$, $p = .002$. There were no differences between the samples in sexual orientation or in number of hours using computer devices each week ($ps > .05$).

Procedure

All participants were recruited online for an anonymous, IRB-approved study entitled "Online Interactions." After completing an online consent form, participants completed some demographic questions and were then directed to the chat platform "Chatzy," where they were greeted in a private chatroom by an administrator and a confederate posing as a 13-year-old boy (James) or a 13-year-old girl (Amber). Although the vast majority of online sex offenders are men who seek to groom males or females, both men and women participants were included in this study because online identity deception is broader than sexual offending and encompasses many different fraudulent online activities (e.g., Nigerian Prince and money transfer schemes). The administrator and confederate role for each chat was conducted by a single, trained research assistant. The 16 research assistant confederates in this study were all undergraduates (five men and 11 women) aged 19

¹ For the most part, people exited the survey after the instruction on the survey that directed them to the online chat forum, likely because they realized they would need to input a password after a 30-min chat was completed. Additionally, many did not complete the first part of the survey at a time when a chat was available (they did not follow instructions) and did not have the option of completing a chat.

to 38 ($M = 22.71$, $SD = 3.83$). As part of a different project, these assistants were required to read and double-check accuracy of at least 10 chat transcripts between adults and undercover agents involved in sex sting operations that had been taken from the Perverted Justice website via a web scraping process. Thus, they were familiar with the language used by undercover agents pretending to be 11 to 14 year olds. As a starting point for this study, the research assistants each developed a backstory for 13-year-old boy or girl containing common features of the various personas portrayed during actual stings (e.g., single parent, homeschooled). Then, they voted as a group on the one boy (James) and one girl (Amber) persona that contained features which they all agreed were consistent with the Perverted Justice chats and a teenage identity (see Appendix in the online supplemental materials). All of the research assistants then posed as either James or Amber within chats and also acted as the administrator.

Once the participant entered the chat room, the administrator gave both participants the following instructions, adapted from the instructions provided in Lincoln and Coyle (2013):

You are going to have an instant messaging conversation with another person. This person's age could range anywhere from a child to an elderly person. The topic of your discussion is open—you may chat about anything you want—but you must not engage in communication that could be considered defamatory, racist, or sexist, or in any other way discriminatory. Also, please do not ask your conversation partner about his or her age or gender.² I will leave the conversation window for 30 minutes, and once the 30 minutes is complete, I will come back in and give you the link to the survey again.

After 30 min, the administrator reentered and stated, “Your 30-minute chat is now complete. Thank you both for participating.” The administrator then provided a link to the remainder of the survey along with a password and a unique ID number to both the participant and the confederate that they needed to complete the rest of the survey. This same research assistant was also the confederate posing as either James or Amber (based on a random assignment for each case) during the Chatzy conversation and cycling through each of the following experimental conditions.

Control. In the control condition, the confederate provided no age information or any explicit cues to their age (like what grade they were in or school they attended). Confederates in this condition were also instructed that if their chat partner asked their age, they were supposed to reply “I don’t think we are supposed to do that.”

Stated age and gender. In this condition, the confederates stated age and gender within the first few sentences of the chat, stating either “I’m 13/f” or “I’m 13/m,” depending on whether they were portraying a girl or boy, respectively.

Stated Age and Gender + Attractive picture. Confederates stated their age and gender but also, within the first few minutes of the chat, sent a link to a picture (headshot) of an attractive teenage boy or girl, stating, “Here’s a picture of me.” These photos were purchased for commercial use from an online photo marketplace and were judged to be attractive (i.e., greater than 5 on a 0–10 scale) by a group of 10 research assistants (see Figure 1).

Stated Age and Gender + Average picture. Confederates stated their age and gender but also, within the first few minutes of the chat, sent a link to a picture (headshot) of an average-looking teenage boy or girl stating, “Here’s a picture of me.” These photos



Figure 1. Pictures provided to participants: (A) “Average girl”; (B) “Attractive girl”; (C) “Average boy”; (D) “Attractive boy.” See the online article for the color version of this figure.

were purchased for commercial use from an online photo marketplace and were judged by a group of 10 research assistants to be average or below average in attractiveness (i.e., less than 5 on a 0–10 scale) by a group of 10 research assistants (see Figure 1).

Grade information only. Participants were instructed to not mention their age or gender but to instead state their grade in school (i.e., eighth grade) or that they were going into high school next year. Again, confederates in this condition were instructed that if their chat partner asked their age, they were to reply “I don’t think we are supposed to do that.”

After participants completed the chat, they were directed back to the anonymous survey website by the administrator to complete the remainder of the survey (see the Measures section for a description of survey items). Once they completed the survey, all of the participants were debriefed and told the true nature of the study and informed that they had actually been communicating with an adult member of the research team. They were also asked to not reveal the true nature of the study to anyone else so that the integrity of the research could be maintained. After completing the survey, participants in the Psychology Department subject pool were credited with one research credit, and participants in the MTurk pool received \$1.50 in their Amazon.com account.

² Inclusion of this instruction served a way to make the topic of age and gender salient but also ensure treatment fidelity (so that those participants assigned to the “no age or gender” condition would not ask for this information).

Measures

Before the chat, participants were asked a small number of demographic questions (e.g., age, sex, and ethnicity). After the chat, participants answered the following questions, in the order discussed.

Age of chat partner. Participants were asked, “What do you think was the age of your chat partner?” and were asked to enter their responses in whole numbers. Responses ranged from 11 to 65.

Gender of chat partner. Participants were asked, “What do you think was the gender of your chat partner?” They were given two options (0 = female, and 1 = male).

Cues used to estimate age and gender. After participants entered the age and selected the gender of their chat partner, participants were asked four questions on how they determined age, and three questions about how they determined gender. For age, participants were asked,

To what extent did you use the following to make this [age] estimation? (1) “The CONTENT of their chat messages (i.e., what the person talked about),” (2) The STYLE of their writing in their chat messages (how the person said what they said), (3) “Their STATED AGE,” and (4) “Their PICTURE.”

For each of the four questions, participants responded on a 5-point Likert scale (1 = *not much*, 5 = *a great deal*). After selecting a gender, participants were asked the same questions on the same scale, except for the stated age question. For analyses, these answers were combined into three composite variables (reliability for use of content cues, $\alpha = .67$; use of style cues, $\alpha = .65$; use of picture, $\alpha = .92$). Use of stated age remained a single variable.

Explanation for assigning specific age. Participants were asked, “Please explain in a few sentences some examples of what you used to make the age and gender estimations of your chat partner. How did you know their age and gender?” Using directed content analysis, two trained research assistants coded each of the qualitative responses for how the participants estimated the age and gender of their chat partner. These codes corresponded to the conceptual framework (and experimental design) that content, style, stated age and gender, and a picture could contribute to age and/or gender determination in online chat contexts. For each variable, binary coding was employed so that if the person mentioned using that method, it was coded as a 1, and if they did not use that method, it was coded as a 0. Additionally, as the MTurk and Psychology pools typically contain people 18 years old and over, we also coded for whether participants cited the MTurk or Psychology pool as a reason for assigning age or gender (0 = *not mentioned*, 1 = *mentioned*). Finally, after reviewing all of the qualitative comments, we noticed that many people were using the screen name (“James” or “Amber”) to make age or gender determinations. Therefore, we also coded for use of screenname (0 = *not mentioned*, 1 = *mentioned*). Table 2 provides for examples of comments that fit into each of the categories as well as the initial interrater agreement between the raters for category. Cohen’s κ values for these categories ranged from .70 to .97, which, according to Landis and Koch (1977), fall in the “substantial” to “almost perfect” range. In the final analysis, differences between coders were resolved by discussion.

Attractiveness and believability of picture received. Participants who received a picture were asked to rate the attractiveness of the photo they saw on a 11-point Likert scale (0 = *not at all attractive*, 10 = *extremely attractive*). They were also asked to rate the believability of the pictures on a 5-point Likert scale (1 = *not at all*, 5 = *extremely*).

LIWC language analysis. As a starting point for the language analysis, all chats were separated by speaker into separate text files, so that the confederates and adult participants’ language could be compared. All separated transcripts were of sufficient length for inclusion (average word count for confederates = 281.95 [$SD = 93.70$]; average word count for participants = 327.58 [$SD = 139.87$]). Scores for each of the three LIWC categories (clout, analytic, and six-letter words) were calculated by processing uploaded chat transcripts through the LIWC program. The LIWC output consists of category scores. For example, the six-letter word score represents the percentage of total words that fit within that category. Thus, a score of .05 would mean that 5% of the words in that transcript are in the six-letter word category. Clout and analytic scores reflect empirically based constructs (combining multiple LIWC categories), population-normed, and scaled from 0 to 100. Higher scores on these measures indicate that more words from these categories were present in the transcript.

Analysis Plan

Using SPSS Version 24, we first conducted preliminary analyses to examine whether there were any significant differences in age estimations between the MTurk and Psychology pool participants, between those who chatted with a boy or girl persona, and between confederates. We also conducted analyses to determine whether participant demographics (i.e., age and gender) were related to age estimations. Next, we calculated prevalence and central tendency statistics for participants’ estimates of their chat partner’s age and gender, using ANOVAs with Bonferroni post hoc comparisons to examine whether this was affected by participant condition. As part of these analyses, we also examined whether the use of an attractive or average picture affected these age estimations, and the extent to which participants were able to discern the *actual* age and gender of their chat partner. Next, we analyzed the cues participants used to estimate age and gender, again using ANOVAs to determine whether cue use differed by condition. These quantitative analyses were supplemented by qualitative analyses, in which participants’ comments on how they determined age and gender were coded and quantified. Finally, we conducted the LIWC analyses, examining whether there were significant differences in the language used by confederates and participants, and whether age estimations were correlated with measures of clout, analytic words, and six-letter words.

Results

Preliminary Analyses

Across conditions, there were no significant differences between the two samples (i.e., college students and MTurk participants) in their estimations of age of their chat partner in each of the conditions (all $ps > .19$) or between those who chatted with a boy or a girl persona (all $ps > .05$). Additionally, there were no significant

differences in participants' age estimations by confederate, $F(15, 141) = .860, p = .61$; the mean age estimated across all confederates in these conditions was 13.83 ($SD = 2.87$; range = 13.00–15.75). Finally, across all conditions, the correlation between participant age and chat partner estimated age was nonsignificant ($r = .09, p = .16$), and the difference between men and women's age estimations was also nonsignificant, $t(260) = 1.75, p = .07$. Therefore, none of these variables were considered in subsequent analyses.

Discernment of Chat Partner's Age and Gender

The primary research question was how sensitive people were to others' online deception when they had adopted a fictitious online persona. As a control condition, we measured participants' estimates of chat partner's age and gender in absence of any stated age, gender, or grade information. In this condition, participants estimated the mean age of their chat partner as 19.35 ($SD = 9.08$, range = 11–65). However, aside from two participants who estimated the age of their chat partners as 55 and 65 years old, most control participants (96%) estimated that their chat partner was 28 years old or younger. Additionally, more than half (54%) of the participants estimated that their chat partner was under the age of 18, and 33% estimated that their chat partner was under the age of 16. Estimates of gender in the control condition show 87% believed the gender portrayed by the confederate (see Table 1). In other words, when a name ("James" or "Amber") was provided in conjunction with a gender-stereotypical online persona (e.g., "James likes to play football and catch snakes with his brother"), the vast majority of participants believed the confederate.

Next, we examined whether age or gender estimates varied by condition, finding this upheld (see Table 1). Bonferroni post hoc comparisons showed that this difference was attributable to differences between the control group and the experimental conditions. Notably, participants in all four experimental conditions (even those in the condition in which the confederate did not state their age but instead provided only grade information) estimated similar ages for their chat partner, with means between 13 and 14 years (range = 13.30–14.07 years). Table 1 shows that when the confederate stated their age, most participants (83%–88%) believed the confederate and estimated their age to be 13 years old or less. When only grade information was provided, 49% estimated the age of their chat partner as 13 years old or less, and a further 33%

estimated the age of their chat partner as 14 years old. With regard to belief of the gender portrayed, there were no significant differences between groups: Most participants believed the gender portrayed, though this was highest when they stated their age and gender and sent an attractive picture, and lowest in the control condition.

The next aim was to examine whether the attractiveness of the picture affected the believability of the picture or the participants' age or gender estimations. Although participants who saw an attractive picture rated it as more attractive than those who saw the average picture ($M = 5.37, SD = 2.17$ vs. $M = 2.90, SD = 2.34$, respectively), $t(1,66) = 3.56, p = .001$, both groups rated the picture they saw as equally believable ($M = 2.42, SD = 1.27$ vs. $M = 2.40, SD = 3.03$, respectively), $t(1,66) = 3.56, p = .001$. Moreover, as shown in Table 1, there were no significant differences between these two conditions in their age estimations or belief of age or gender.

To explore whether participants were sensitive to the actual age and gender of their chat partners, we also examined congruency between the actual age and gender of the confederates, and the participants' estimates of age and gender of their chat partner. Of the 262 participants, only four (2%) guessed the exact age of their chat partner; three of these participants were in the control condition, and one was in the Stated Age + Attractive Picture Girl condition. To more directly compare the current results with those from Lincoln and Coyle (2013), we also examined how many participants guessed their confederate's actual age within a 5-year bandwidth. To allow the most leeway in this comparison, this 5-year bandwidth was extended both above and below the confederate's actual age, providing a 9-year bandwidth. For example, when the confederate's actual age was 19, any estimated age from 15 to 23 would count as a correct response. With these more liberal parameters, 39 of 262 participants (15%) were able to discern their chat partner's age within a 9-year bandwidth. Most of these (67%) were in the control condition; however, even in this condition, fewer than half of the Participants 26 of 54 (48%) were able to correctly discern their chat partner's actual age within a 9-year bandwidth. We then examined whether participants could detect when the online persona portrayed was inconsistent with the confederate's actual gender. Of the 118 chats in which the confederate was posing as a person whose gender was inconsistent with his or her own (in 61 chats female confederates posed as a

Table 1

Participants' Assessments of Age and Gender of Chat Partner and Cues They Used to Determine Age and Gender of Chat Partner by Condition

Outcome Measure	Control <i>M (SD)/N (%)</i>	Stated age & gender <i>M (SD)/N (%)</i>	Stated age & gender + Attractive picture <i>M (SD)/N (%)</i>	Stated age & gender + Average picture <i>M (SD)/N (%)</i>	Grade information only <i>M (SD)/N (%)</i>	<i>F</i> or χ^2
Age of chat partner	19.35 _a (9.08)	14.07 _b (3.53)	13.72 _b (2.01)	13.30 _b (1.29)	13.94 _b (1.64)	14.58*
Estimated age ≤13	7 (13%)	72 (84%)	33 (83%)	29 (88%)	24 (49%)	$\chi^2(4, N = 262) = 93.24^*$
Believed gender	47 (87%)	80 (93%)	39 (98%)	31 (94%)	46 (94%)	$\chi^2(4, N = 262) = 4.13$

Note. Means with different subscripts indicate a significant difference between groups at the $p < .05$ level. Control, $n = 54$; Stated age & gender, $n = 86$; Stated age & gender + Attractive picture, $n = 40$; Stated age & gender + Average picture, $n = 33$; Grade information only, $n = 49$. Believed gender = Participant indicated that their chat partner was the gender consistent with the online persona portrayed by the confederate during the chat.

* $p < .001$.

Table 2
Qualitative Analysis Categories for Reasons People Gave for Assigning Age and Gender, Percentage of Participants Who Mentioned Something From That Category, Representative Comments, and Interrater Reliability for Each Category

Category	n/% mentioning ^a	Representative comment(s)	Cohen's κ
Content cues	163 (62%)	"I think she was a girl because she likes listening to Taylor Swift and plays volleyball. I know guys play volleyball but they don't normally listen to Taylor Swift." "He talked about liking snakes and chilling with little bro oh and likes Harry Potter and Hunger Games."	.77
Style cues	99 (38%)	"Well they did tell me her age, but their grammar wasn't good. She said 'wats up' instead of 'hello' or 'HI.'" "He typed in acronyms (ie. idk and hbu). Many teens I've spoken to speak the way they text."	.70
Stated age/gender	125 (79%)	"When we first started talking, the chat partner told me their name and age."	.97
Picture	24 (33%)	"Well she sent a picture of herself so that helped with the gender. The picture also helped me estimate her age."	.91
Grade information	37 (76%)	"My partner mentioned that she was in eighth grade, so that was my main information source re: their age"	.87
Screenname	88 (34%)	"The username they selected in the chat was James so that made the gender quite easy to guess."	.92
MTurk/Psychology study	4 (2%)	"Thought it was strange that someone so young would be doing a psychology experiment like this."	.80
Didn't believe	15 (6%)	"He said his age was 13 and that he was a male. But having said that i do not believe anything this person was telling me."	.72

Note. MTurk = Amazon Mechanical Turk.

^a The percentage of those who received that cue who used that cue to discern age or gender.

boy, in 57 chats male confederates posed as a girl), 111 (94%) participants stated their chat partner was the gender that they portrayed in the chat. In other words, in the cases in which the confederate was posing as a person whose gender was inconsistent with their actual gender, only seven participants (6%) correctly guessed the true gender of the confederate.

Cues Used to Estimate Age and Gender

With regard to the cues participants used to determine age and gender of their chat partners, cues used to determine age and gender varied by condition (see Figure 2). A series of ANOVAs showed that groups differed significantly in their use of content

cues, $F(4,255) = 3.73, p = .006$, and the picture, $F(4,256) = 47.13, p < .001$, to determine age and gender, and their use of stated age, $F(4,257) = 54.99, p < .001$, to determine age. However, there were no significant differences between groups in their use of style, $F(4,256) = 2.22, p = .07$. Bonferroni post hoc comparisons found that participants who were in the Stated Age and Gender + Attractive Picture condition were less likely to use content cues to estimate age than those in the control condition, Stated Age and Gender condition, or Grade-Information-Only condition ($ps < .05$). Additionally, as expected, in conditions in which participants received a picture, they were more likely to use the picture to determine age and gender than those who did not




Figure 2. Average cue use to determine age and gender by condition.

receive a picture ($ps < .05$). Finally, to estimate age, in conditions in which the confederate stated their age as opposed to conditions in which only grade information was provided, or no age or grade information was provided, participants were significantly more likely to use stated age to make their age determinations ($ps < .05$).

To determine which cues were most used to estimate age within each condition, we conducted a series of repeated measures ANOVAs, including only those cues that were provided by the confederate during the chat. In the control group, participants were just as likely to use content cues as they were to use style cues to estimate the age of their chat partner, $F(1, 53) = 0.96, p = .33$. However, participants in the grade-information-only group were more likely to use content cues than style cues, $F(1, 48) = 8.98, p = .004$. When the participants were provided with a stated age (but no picture), there was a significant difference in the cues they used to estimate age, $F(2, 85) = 9.89, p < .01$. Pairwise comparisons showed that participants who were provided with a stated age used it to estimate the age of their chat partner significantly more than they used the content or style of their messages ($ps < .01$). Participants in these conditions also used the content more than the style of their chat partner's messages to estimate age ($p = .001$). Finally, when the participants were provided a stated age and a picture (attractive and average picture groups combined, as there were no significant differences between these groups in their use of the picture to make age and gender determinations; $p > .05$), there were again significant differences in the cues they used, $F(3, 70) = 12.02, p < .001$. Pairwise comparisons showed that participants who were provided with a stated age and a picture used stated age to determine the age of their chat partner significantly more than the content or style of their messages or the picture provided (all $ps < .01$).

Finally, to supplement these quantitative analyses, we also examined all of the qualitative comments participants provided about how they determined age and gender. As shown in Table 2, the qualitative analyses supported the quantitative findings: People used stated age and grade the most, and then content cues, and, finally, picture and style cues to make age and gender determinations. Notably, people used a combination of cues to determine age and/or gender of their chat partner; of the 259 participants who gave any response to the qualitative question, 256 (98%) listed multiple cues that helped them determine age and gender. Additionally, as shown in Table 2, only 2% of the participants mentioned that it was odd that a child would be in a college psychology experiment or part of the MTurk worker pool, and only 6% mentioned anything about not believing what their chat partner said.

Natural Language Analyses

Finally, we analyzed the output files from a LIWC analysis to determine whether (a) there were any consistent natural language patterns that distinguished the confederates from the participants, and (b) whether age estimations were related to the clout, analytic, or six-letter word scores of the confederate. As Table 3 shows, the confederates had significantly lower scores than the participants for each of the categories measured (i.e., clout, analytic words, and six-letter words), and the effect sizes for these differences were moderate to large. However, only one of these language dimensions was significantly related to the estimation of age: Confeder-

Table 3

Descriptive Statistics and Significant Differences for the LIWC Clout, Analytic Words, and Six-Letter Words in the Chat Between the Confederate and the Participant

LIWC Category	Confederate <i>M (SD)</i>	Participant <i>M (SD)</i>	<i>t</i>	<i>d</i>
Clout	42.77 (15.83)	52.98 (20.44)	-5.97*	-.56
Analytic	17.01 (11.47)	24.48 (13.82)	-7.07*	-.59
Six-letter words	8.00 (2.33)	11.45 (2.63)	-18.60*	-1.39

Note. $N = 246$ chats valid for analysis.

* $p < .001$.

ates who used more analytic language in their conversations with participants were rated as significantly older than those who used less analytic language ($r = .13, p = .048$); neither clout nor six-letter words used by the confederate were related to the age estimations ($ps > .05$).

Discussion

Proactive undercover Internet police investigations may require law enforcement or other trained officials to pose as a person who is very different than their offline self. This situation creates a very important empirical and legal question: Is it possible to see through this online façade and detect the real characteristics of an online chat partner? In Lincoln and Coyle's (2013) experiment, the answer was apparently definitive: Identity deception is difficult online, and most people can discern the real age and gender of online chat partners. However, as a number of methodological issues limit the generalizability of Lincoln and Coyle's findings, we sought to replicate their study using settings and methods that more closely resemble the conditions of actual undercover sting operations.

Overall, the results starkly contradict Lincoln and Coyle's (2013) findings. Contrary to Hypothesis 1, for the most part, the adults in this experiment (undergraduates and MTurk workers) *did* apparently believe their online chat partner's online persona. When the participants stated their age, 84% to 88% of participants believed their chat partners, and even when only grade information was provided (e.g., "I'm in 8th grade, going into high school next year"), 82% of the participants estimated that their chat partner was 13 or 14 years old. Moreover, even in the absence of any age or grade information, 13% of participants estimated that their chat partners were 13 or younger, and 33% estimated their age as under 16. The findings for gender were even more compelling. Across all of the conditions, 87% to 98% of the participants believed the gender portrayed in the chat. In cases when the confederate's gender was different than the online persona, only 6% of participants were able to discern the confederate's actual gender. Notably, none of the participants in Lincoln and Coyle's study stated that their chat partner's age was under 16, and 80% were able to discern their chat partner's actual gender, findings that they used to claim "individuals can discern the age and gender of their interlocutor within a very brief period. Thus, for some individuals, the defense employed by the accused in the matter of *R v. Plumridge* (2007) has scientific justification, with caveats" (p. 300). These caveats are that errors in content may be more likely to expose the ruse than stylistic or syntactic errors.

With regard to cues, the analyses partially supported Hypothesis 2, that people would use content cues more than style cues to estimate age and gender. In line with Lincoln and Coyle (2013), the quantitative analyses showed that when grade information was provided alone, or when the confederate only stated their age and gender, participants used content cues more than style cues to determine age and gender. However, in all of the other conditions, content and style cues were used equally. More importantly, when a stated age was provided, participants used this cue significantly more than either content or style. Notably, stated age was not a cue examined in either Lincoln and Coyle's or May-Chahal et al.'s (2014) studies; these results suggest that this cue is very important for age estimation by adult interlocutors. In line with this, law enforcement officials conducting undercover stings should make age information salient in the beginning of an online chat, and perhaps repeat their purported age at other times throughout the chat, especially when their content cues (e.g., sexual experience) may suggest a higher level of maturity. Meanwhile, pictures—which appear to be mainstays of undercover sex stings—were the least-used cue to age and gender. The qualitative analyses generally supported these trends.

More importantly, the qualitative analyses showed that people do not rely upon a single cue for age and gender. Instead, online chatters use multiple cues to discern the identity of their chat partner, which makes it especially important for law enforcement agents involved in undercover stings to maintain a consistent persona across cue types. Any slip-ups across cues (e.g., talking about activities that a young teen would not engage in, using antiquated sayings or too little slang, or sending a picture that is not believable) may breach the undercover agent's ruse. This is salient, as offenders are increasingly sensitized to potential *agent provocateurs* in legislations in which such law enforcement strategies are legal, and information and optimal strategies to commit an offense and avoid detection are commonly shared on Internet forums (Lorenzo-Dus & Izura, 2017). Understanding offender expertise in apprehension avoidance is a poorly researched topic but likely to prove useful for forensic practitioners, as law enforcement can use this information to refine detection strategies (Grant & Macleod, 2016; Moeller, Copes, & Hochstetler, 2016).

Finally, in examining the linguistic properties of the chats, Hypothesis 3 was partially supported. Confederates who portrayed 13-year-olds displayed significantly less clout and used significantly fewer analytic and six-letter words than did adult participants. Further, those confederates who used more analytic language were perceived as older than those who used less analytic language. These findings are important for two reasons. First, the significant difference between confederates and the participants lends support to the current methodology—confederates aimed to use less sophisticated language than adults would, and they succeeded. Second, this analysis adds to a growing body of forensic linguistic analysis studies (e.g., Black et al., 2015; Drouin, Boyd, Hancock, et al., 2017) and points to the potential utility of these types of analyses in the courtroom. As an example, in cases in which the undercover agent displays high levels of clout, analytic words, and six-letter words, it could serve as supplemental evidentiary support for claims that the accused offender thought that their chat partner was older.

As detailed in the introduction, there are various methodological issues that may explain the incongruence between the current

findings and those of Lincoln and Coyle (2013). The utilization of a diverse sample (comprised of both college students and MTurkers), an exclusively online setting, trained research assistant confederates, and experimental conditions in which the confederates shared their age, gender, and/or a picture may have contributed to the disparate findings. However, as this experimental setting more closely resembles those employed in actual Internet sting operations, the current findings may have more naturalistic utility. In contrast to Lincoln and Coyle, the current results suggest that despite the widespread distrust of others that plagues the Internet (e.g., Caspi & Gorsky, 2006; Drouin et al., 2016; Hancock & Woodworth, 2013; Henderson & Gilding, 2004), many people may believe a fake online persona encountered in a chat room, at least within the context of an experiment. This was true even though there were reasons for participants to doubt that they were communicating with a 13-year-old, like the fact that the college students signed up for this experiment through an online experiment portal for introductory psychology students, or the fact that MTurk participants signed up through the MTurk worker portal. Within both systems, users would likely assume that they were communicating with other adults, and it may be that deception sensitivity differs depending on the nature of the online forum persons are using. The MTurk Participation Agreement: <https://www.mturk.com/mturk/conditionsofuse> states specifically that you must be at least 18 years old to participate in their worker pool (Amazon Mechanical Turk, n.d. Participation Agreement). However, rather than reducing the generalizability of these results, these conditions may increase their generalizability, as some offenders use the defense that they met the undercover agents on websites that are supposed to be adult only, and thus assumed they were adults (Colleluori, 2010). Results suggest that even these assumptions can be overridden.

Importantly, the analysis on use of cues offers one point on which the current findings and Lincoln and Coyle's (2013) findings converge: Content cues appear to be very important in discerning the identity of a chat partner. Thus, our findings support Lincoln and Coyle's suggestion that covert operatives need to be very sensitive to the content of their chats. They should, according to Lincoln and Coyle, be familiar with people in the group they are trying to impersonate. Although stated age was the most utilized cue to age in the 30-min chat scenario, it is possible that those involved in illegal activities online may be warier of one's stated identity because of the inherent risks involved in illegal activities. Thus, especially in prolonged conversations spanning hours, days, or even months, suspicious online offenders may be especially attuned to content, attending to any cue that might make them doubt the authenticity of online agents. The qualitative analysis that shows that most participants use a variety of cues to make their age and gender determinations provides further support for this assertion. In the real world, those engaging in illegal activity online would likely consider all of the cues available (i.e., content, style, stated age and gender, and pictures) to ascertain their chat partner's true identity. The fact that not all of the participants believed their chat partner suggests that there may have been cues to deception in the chat to which some people were more attuned or sensitive.

Despite specific attention to methodological issues in the development of this study, there were some procedural limitations. One potential limitation is the use of a paid MTurk sample. Although

MTurk workers and college students perform similarly or better on a variety of tasks (Buhrmester et al., 2011; Casler et al., 2013; Feitosa et al., 2015; Hauser & Schwarz, 2016). Hamby and Taylor (2016) observed that MTurk samples are paid for their participation and are motivated to complete tasks quickly, so are more likely to give satisficing or insufficient responses (i.e., a “good enough” answer) than those in a university sample. This potential problem was perhaps circumvented in this study because both groups participated in a 30-min chat with a confederate in order to move on to the survey, which was comparatively shorter than the chat they had just completed. Further, as attention checks, all of the participants included in this study completed the 30-min chat and correctly inputted a random number and a password that were provided by the administrator during the chat. An additional limitation was that the chats only lasted 30 min, whereas in many sex stings, the chats last days or even months (<http://perverted-justice.com/?con=full>). It is quite possible that over longer conversations, law enforcement agents might make more mistakes (as in *R v. Plumridge*, 2007) or give more cues to their actual offline identities. Thus, in actual cases, it would be important to examine the chats in full to determine whether there are content, stylistic, or language pattern markers that “give up the game.” Also, in most of these chats, participants chatted with confederates about general interests (e.g., sports, food, music, and leisure activities). It is possible that when engaging in potentially inappropriate and/or sexual conversations with a minor, individuals might be more attuned to age and deception cues because of the possibility of being caught. As such, transcripts from police operational “stings” should be analyzed thoroughly for evidence that the target did not believe the undercover officer. Additionally, the participants were in a research study context, which may have made them think that they were less likely to be deceived than in the real world. However, this was also the case in Lincoln and Coyle’s (2013) study, but most of them believed the confederate was lying. Thus, there is no evidence that an experimental setting increases trust in another person’s online identity. Finally, the brief training the confederates undertook is minimal compared with the training law enforcement agents must go through to conduct these operations. Rather than weakening the results, this difference may actually strengthen them, as the current study demonstrates that even with minimal training and an appropriate backstory, people can convince others that they are someone who they are not.

Conclusion

Although it is commonly believed people lie online, when adults encountered a trained confederate posing as a 13-year-old boy or girl in an online chat room, they generally believed the online persona, especially when the confederate stated his or her age. Though the plausibility of a 13-year-old participating in a psychology experiment or MTurk study is low, few participants (2%) noted this in their qualitative responses, and only 6% overall stated that they doubted their chat partner’s honesty. Thus, the findings of Lincoln and Coyle (2013), used to support the fantasy/role-play defense that accused online offenders sometimes use when their cases are brought to trial (e.g., *R v. Plumridge*, 2007), is challenged. Although it is reasonable to expect that if online agent mistakenly or inadvertently gives cues that they are older (e.g., through age-inappropriate content or use of sophisticated language

or social dominance), the fantasy/role-play defense is plausible and defensible, results indicated that false online identities—even implausible ones—are often believed.

References

- Amazon Mechanical Turk. (n.d.). Participation Agreement. Retrieved from <https://www.mturk.com/mturk/conditionsofuse>
- Bantum, E. O., & Owen, J. E. (2009). Evaluating the validity of computerized content analysis programs for identification of emotional expression in cancer narratives. *Psychological Assessment, 21*, 79–88. <http://dx.doi.org/10.1037/a0014643>
- Black, P. J., Wollis, M., Woodworth, M., & Hancock, J. T. (2015). A linguistic analysis of grooming strategies of online child sex offenders: Implications for our understanding of predatory sexual behavior in an increasingly computer-mediated world. *Child Abuse & Neglect, 44*, 140–149. <http://dx.doi.org/10.1016/j.chiabu.2014.12.004>
- Bond, G. D., & Lee, A. Y. (2005). Language of lies in prison: Linguistic classification of prisoners’ truthful and deceptive natural language. *Applied Cognitive Psychology, 19*, 313–329. <http://dx.doi.org/10.1002/acp.1087>
- Buchanan, T., & Whitty, M. T. (2014). The online dating romance scam: Causes and consequences of victimhood. *Psychology, Crime & Law, 20*, 261–283. <http://dx.doi.org/10.1080/1068316X.2013.772180>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6*, 3–5. <http://dx.doi.org/10.1177/1745691610393980>
- Casler, K., Bickel, L., & Hackett, E. (2013). Separate but equal? A comparison of participants and data gathered via Amazon’s MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior, 29*, 2156–2160. <http://dx.doi.org/10.1016/j.chb.2013.05.009>
- Caspi, A., & Gorsky, P. (2006). Online deception: Prevalence, motivation, and emotion. *CyberPsychology & Behavior, 9*, 54–59. <http://dx.doi.org/10.1089/cpb.2006.9.54>
- Colleluori, A. J. (2010). Defending the internet sex sting case. GPSOLO January/February 2010. Retrieved from http://www.americanbar.org/newsletter/publications/gp_solo_magazine_home/gp_solo_magazine_index/colleluori.html
- Drouin, M., Boyd, R. L., & Greidanus Romaneli, M. (2017). Predicting recidivism among internet child sex sting offenders using psychological language analysis. *Cyberpsychology, Behavior, and Social Networking*. Advance online publication. <http://dx.doi.org/10.1089/cyber.2016.0617>
- Drouin, M., Boyd, R. L., Hancock, J. T., & James, A. (2017). Linguistic analysis of chat transcripts from child predator undercover sex stings. *Journal of Forensic Psychiatry & Psychology, 28*, 437–457. <http://dx.doi.org/10.1080/14789949.2017.1291707>
- Drouin, M., Miller, D. A., Wehle, S., & Hernandez, E. (2016). Why do people lie online? “Because everyone lies on the Internet.” *Computers in Human Behavior, 64*, 134–142. <http://dx.doi.org/10.1016/j.chb.2016.06.052>
- Egan, V., & Cordan, G. (2009). Barely legal: Is attraction and estimated age of young female faces disrupted by alcohol use, make up, and the sex of the observer? *British Journal of Psychology, 100*, 415–427. <http://dx.doi.org/10.1348/000712608X357858>
- Egan, V., Hoskinson, J., & Shewan, D. (2011). Perverted justice: Content analysis of the language used by offenders detected attempting to solicit children for sex. In R. M. Clarke (Ed.), *Antisocial behavior: Causes, correlations, and treatments* (pp. 119–134). Hauppauge, NY: Nova Science Publishers.
- Ellison, N. B., Hancock, J. T., & Toma, C. L. (2012). Profile as promise: A framework for conceptualizing veracity in online dating self-presentations. *New Media & Society, 14*, 45–62. <http://dx.doi.org/10.1177/1461444811410395>

- Feitosa, J., Joseph, D. L., & Newman, D. A. (2015). Crowdsourcing and personality measurement equivalence: A warning about countries whose primary language is not English. *Personality & Individual Differences*, 75, 47–52. <http://dx.doi.org/10.1016/j.paid.2014.11.017>
- Feldman, M. D. (2000). Munchausen by Internet: Detecting factitious illness and crisis on the Internet. *Southern Medical Journal*, 93, 669–672. <http://dx.doi.org/10.1097/00007611-200007000-00005>
- Franklin, K. (2012, March 11). Internet stings: Does the fantasy defense hold water? *Psychology Today*. Retrieved from <https://www.psychologytoday.com/blog/witness/201203/internet-stings-does-the-fantasy-defense-hold-water>
- Grant, T., & Macleod, N. (2016). Assuming identities online: Experimental linguistics applied to the policing of online paedophile activity. *Applied Linguistics*, 37, 50–70. <http://dx.doi.org/10.1093/applin/amv079>
- Guadagno, R. E., Okdie, B. M., & Kruse, S. A. (2012). Dating deception: Gender, online dating, and exaggerated self-presentation. *Computers in Human Behavior*, 28, 642–647. <http://dx.doi.org/10.1016/j.chb.2011.11.010>
- Hamby, T., & Taylor, W. (2016). Survey satisficing inflates reliability and validity measures: An experimental comparison of college and Amazon Mechanical Turk samples. *Educational and Psychological Measurement*, 76, 912–932. <http://dx.doi.org/10.1177/0013164415627349>
- Hancock, J., & Woodworth, M. (2013). An “Eye” for an “I”: The challenges and opportunities for spotting credibility in a digital world. In B. S. Cooper, D. Griesel, & M. Ternes (Eds.), *Applied issues in investigative interviewing, eyewitness memory, and credibility assessment* (pp. 325–340). New York, NY: Springer. http://dx.doi.org/10.1007/978-1-4614-5547-9_13
- Hausser, D. J., & Schwarz, N. (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48, 400–407. <http://dx.doi.org/10.3758/s13428-015-0578-z>
- Henderson, S., & Gilding, M. (2004). “I’ve never clicked this much with anyone in my life”: Trust and hyperpersonal communication in online friendships. *New Media & Society*, 6, 487–506. <http://dx.doi.org/10.1177/146144804044331>
- Iosub, D., Laniado, D., Castillo, C., Fuster Morell, M., & Kaltenbrunner, A. (2014). Emotions under discussion: Gender, status and communication in online collaboration. *PLoS ONE*, 9, e104880. <http://dx.doi.org/10.1371/journal.pone.0104880>
- Isacenkova, J., Thonnard, O., Costin, A., Francillon, A., & Balzarotti, D. (2014). Inside the scam jungle: A closer look at 419 scam email operations. *EURASIP Journal on Information Security*, 2014, 4–18. <http://dx.doi.org/10.1186/1687-417X-2014-4>
- Joinson, A. N., & Dietz-Uhler, B. (2002). Explanations for the perpetration of and reactions to deception in a virtual community. *Social Science Computer Review*, 20, 275–289. <http://dx.doi.org/10.1177/08939302020003005>
- Kacewicz, E., Pennebaker, J. W., Davis, M., Jeon, M., & Graesser, A. C. (2014). Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33, 125–143. <http://dx.doi.org/10.1177/0261927X13502654>
- Keren, G. (1997). On the calibration of probability judgments: Some critical comments and alternative perspectives. *Journal of Behavioral Decision Making*, 10, 269–278. [http://dx.doi.org/10.1002/\(SICI\)1099-0771\(199709\)10:3<269::AID-BDM281>3.0.CO;2-L](http://dx.doi.org/10.1002/(SICI)1099-0771(199709)10:3<269::AID-BDM281>3.0.CO;2-L)
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33, 159–174. <http://dx.doi.org/10.2307/2529310>
- Lincoln, R., & Coyle, I. R. (2013). No-one knows you’re a dog on the internet: Implications for proactive police investigation of sexual offenders. *Psychiatry, Psychology and Law*, 20, 294–300. <http://dx.doi.org/10.1080/13218719.2012.672274>
- Lorenzo-Dus, N., & Izura, C. (2017). “cause ur special”: Understanding trust and complimenting behaviour in online grooming discourse. *Journal of Pragmatics*, 112, 68–82. <http://dx.doi.org/10.1016/j.pragma.2017.01.004>
- May-Chahal, C., Mason, C., Rashid, A., Walkerdine, J., Rayson, P., & Greenwood, P. (2014). Safeguarding cyborg childhoods: Incorporating the on/offline behaviour of children into everyday social work practices. *British Journal of Social Work*, 44, 596–614. <http://dx.doi.org/10.1093/bjsw/bcs121>
- Meschke, L., Peter, C., & Bartholomae, S. (2012). Developmentally appropriate practice to promote healthy adolescent development: Integrating research and practice. *Child & Youth Care Forum*, 41, 89–108. <http://dx.doi.org/10.1007/s10566-011-9153-7>
- Moeller, K., Copes, H., & Hochstetler, A. (2016). Advancing restrictive deterrence: A qualitative meta-synthesis. *Journal of Criminal Justice*, 46, 82–93. <http://dx.doi.org/10.1016/j.jcrimjus.2016.03.004>
- Moss, C. (2015, January 21). Strangers have been using this woman’s photos to catfish people online for 10 years. *Tech Insider*. Retrieved from <http://www.businessinsider.com/strangers-have-been-using-this-womans-photos-to-catfish-people-for-10-years-2015-1>
- Muscanel, N. L., Guadagno, R. E., & Murphy, S. (2014). Weapons of influence misused: A social influence analysis of why people fall prey to Internet scams. *Social and Personality Psychology Compass*, 8, 388–396. <http://dx.doi.org/10.1111/spc3.12115>
- O’Hair, H. D., Cody, M. J., & McLaughlin, M. L. (1981). Prepared lies, spontaneous lies, Machiavellianism, and nonverbal communication. *Human Communication Research*, 7, 325–339. <http://dx.doi.org/10.1111/j.1468-2958.1981.tb00579.x>
- Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). *The development and psychometric properties of LIWC2015*. Austin, TX: University of Texas at Austin.
- Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I. (2014). When small words foretell academic success: The case of college admissions essays. *PLoS ONE*, 9, e115844. <http://dx.doi.org/10.1371/journal.pone.0115844>
- Perverted Justice. (n.d.). Convictions. Retrieved from <http://perverted-justice.com/?con=full>
- R v. Plumridge. (2007). 590 AA Application before McGill DCJ, unreported (Queensland District Court, 2007).
- Rouse, S. V. (2015). A reliability analysis of Mechanical Turk data. *Computers in Human Behavior*, 43, 304–307. <http://dx.doi.org/10.1016/j.chb.2014.11.004>
- Simmonds, D. J., Hallquist, M. N., Asato, M., & Luna, B. (2014). Developmental stages and sex differences of white matter and behavioral development through adolescence: A longitudinal diffusion tensor imaging (DTI) study. *NeuroImage*, 92, 356–368. <http://dx.doi.org/10.1016/j.neuroimage.2013.12.044>
- Tani, F., Peterson, C., & Smorti, M. (2016). The words of violence: Autobiographical narratives of abused women. *Journal of Family Violence*, 31, 885–896. <http://dx.doi.org/10.1007/s10896-016-9824-0>
- Toma, C. L., Jiang, L. C., & Hancock, J. T. (2016). Lies in the eye of the beholder: Self-other asymmetry in beliefs about deception across interpersonal media. *Communication Research*. Advance online publication. <http://dx.doi.org/10.1177/0093650216631094>
- Van Gelder, L. (1991). The strange case of the electronic lover. In C. Dunlop & R. Kling (Eds.), *Computerization and controversy: Value conflicts and social choice* (pp. 533–546). Boston, MA: Academic Press.
- 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: Association for Computing Machinery.

- Whitty, M. T. (2002). Liar, liar! An examination of how open, supportive and honest people are in chat rooms. *Computers in Human Behavior, 18*, 343–352. [http://dx.doi.org/10.1016/S0747-5632\(01\)00059-0](http://dx.doi.org/10.1016/S0747-5632(01)00059-0)
- Whitty, M. T. (2008). Revealing the “real” me, searching for the “actual” you: Presentations of self on an internet dating site. *Computers in Human Behavior, 24*, 1707–1723. <http://dx.doi.org/10.1016/j.chb.2007.07.002>
- Whitty, M. T., & Joinson, A. N. (2009). *Truth, lies, and trust on the Internet*. London, UK: Routledge, Psychology Press.
- Wright, R. G. (2009). Internet sex stings. In R. G. Wright (Ed.), *Sex offender laws: Failed policies, new directions* (pp. 117–158). New York, NY: Springer.

Received May 23, 2017

Revision received August 30, 2017

Accepted September 4, 2017 ■

Members of Underrepresented Groups: Reviewers for Journal Manuscripts Wanted

If you are interested in reviewing manuscripts for APA journals, the APA Publications and Communications Board would like to invite your participation. Manuscript reviewers are vital to the publications process. As a reviewer, you would gain valuable experience in publishing. The P&C Board is particularly interested in encouraging members of underrepresented groups to participate more in this process.

If you are interested in reviewing manuscripts, please write APA Journals at Reviewers@apa.org. Please note the following important points:

- To be selected as a reviewer, you must have published articles in peer-reviewed journals. The experience of publishing provides a reviewer with the basis for preparing a thorough, objective review.
- To be selected, it is critical to be a regular reader of the five to six empirical journals that are most central to the area or journal for which you would like to review. Current knowledge of recently published research provides a reviewer with the knowledge base to evaluate a new submission within the context of existing research.
- To select the appropriate reviewers for each manuscript, the editor needs detailed information. Please include with your letter your vita. In the letter, please identify which APA journal(s) you are interested in, and describe your area of expertise. Be as specific as possible. For example, “social psychology” is not sufficient—you would need to specify “social cognition” or “attitude change” as well.
- Reviewing a manuscript takes time (1–4 hours per manuscript reviewed). If you are selected to review a manuscript, be prepared to invest the necessary time to evaluate the manuscript thoroughly.

APA now has an online video course that provides guidance in reviewing manuscripts. To learn more about the course and to access the video, visit <http://www.apa.org/pubs/authors/review-manuscript-ce-video.aspx>.