It's a Jungle Out There:

A Real-World Analysis of Lying in Instant Messaging

Honors Thesis

Presented to the College of Agriculture and Life Sciences, Communication Department

of Cornell University

in Partial Fulfillment of the Requirements

for the Research Honors Program

By

Joshua Perlin

May 2008

Professor Jeffrey Hancock

Acknowledgements

First and foremost, thank you Jeff Hancock. Your guidance and support have been invaluable in the development of this research and paper. Thank you Barrett Amos, whose long-term vision and work on the Apate system made this research possible, and for helping shape this research over the last few years. Thank you to Natalie Bazarova for her assistance in data analysis, Jeremy Birnholtz for his assistance with methods, and to all the professors in the department who were willing to offer this study in their classes for credit. This work would not have been possible had it not been for the assistance of all of these people, and I am incredibly appreciative of your aid. Last but certainly not least, thank you to the Fredric N. Gabler ’93 Memorial Research Honors Endowment for supporting this research.

Abstract

Analysis of instant messaging deception has never been conducted using real-world data. This study tracked and analyzed deceptions in participants' real-life instant messaging conversations, analyzing them on a message-by-message basis instead of merely a conversation-by-conversation basis. The results show that people frequently deceive in instant messaging, and that the magnitude of these deceptions is positively correlated with instant messaging use. Deceptions are told in clusters, suggesting strategic usage of deception in instant messaging, and although people underestimate how frequently they deceive, they have a very good sense of how much they are straying from their normal deception rates.

Is this person really telling me the truth? Deception, communication and everyday life are all intertwined. Communication is a fundamental part of life, and one of its most basic foundations is the assumption that what's being communicated is truthful. For this reason, it is important to know whether or not the information you are receiving during a communicative act is deceptive.

Albert Vrij (2000, pp. 5-7) put forth one of the most frequently cited definitions of deception to date, stating that deception is a deliberate attempt to create a false belief in a communication partner. Whether the message is successfully deceptive or not, there cannot be any intended forewarning by the communicator, in order for a deception to have occurred. Jokes, sarcastic remarks, and ironic musings are therefore not forms of deception, as they do not have any intent to deceive. Furthermore, messages that are unintentionally misleading are also not described as deceptive, but instead are simply "mistakes" or "errors." (Nyberg, 1993, p. 75) It is important to note that the definition considers lies a subset of deception, since deception can take on many forms that don't necessarily include telling a lie. However, for the purposes of this research, the terms lying and deception will be considered to mean the same thing.

Although lying is an important part of everyday life, how big a part of our interactions is it? DePaulo et al. (1996) found that participants reported deceiving their communication partners almost one-third of the time (31 percent) in their daily social interactions. Complicating this picture is the rise of computer-mediated communication, allowing people to communicate more than ever before and in mediums other than face-to-face. Since deception is so frequently a part of interaction, it has become critical

important to be able to seek out and identify deception in everyday interactions, and in communication mediums other than face-to-face.

Deception and Instant Messaging

Previous studies analyzing deception prevalence offline have done so through either laboratory settings or self-reporting in participants' real-world conversations. Those studies that used the latter often relied on a journal entry method in which participants rated their own lies after conversations, based on their memory of the conversation. (DePaulo, 1996)

This has the advantage of not analyzing data from a potentially contrived experimental environment, but has the disadvantage of relying on participants' memories of conversations. Furthermore, this method can be used to check for the frequency of lies within in a given conversation, but not for the placement of lies within an overall conversation, or to check the fractional amount lies make up in a given conversation. The only study to date to have tracked deception in instant messaging, conducted in 2004 by a Communication student (Barrett, *In Progress*), found six percent of instant messaging are deceptive.

Depending on the medium, deception can be found in many different forms and with different characteristics. Hancock et al. (2004) found that although "digital deception" is very much like offline deception, online and offline deception may have very different properties. Analyzing these properties in the computer-mediated environment of instant messaging is critical to better understanding deception and communication, as instant messaging grows increasingly popular.

Although text-based, sent in discrete unites (messages) and lacking non-verbal cues, instant messaging does have some similarities to face-to-face communication. It is a near-synchronous medium of communication (occurs in near real-time), ideal for studying deception because long but normally paced conversations can occur. Nardi (2000) notes that its conversations tend to occur in a manner similar to that of traditional face-to-face interactions. What's more, Hancock et al. (2004) found that people lie in instant messaging and face-to-face at similar rates.

A 2004 Pew Internet & American Life survey revealed that more than about 53 million American adults used instant messaging programs, and that 11 million of them used instant messaging at work for productivity purposes. At the time, the growth rate of instant messaging was estimated to be 29 percent. Given the small amount of research that has been conducted on deception through instant messaging, and given the medium's importance in everyday interaction as well as business settings, this study sets out to conduct a broad and accurate analysis of deception in instant messaging. In addition, results from instant messaging could also potentially clarify or even point out trends in face-to-face interaction, given the similarities between the mediums.

Research Goals

This research attempted to look at lies, online, uninhibited, for the very first time. Specifically, this research intends to gather a large corpus of data using real-world instant messages that have been rated for their level of deception. By looking at deception in the real world, there can be confirmation of research conducted in the laboratory setting about deception online and offline. A great deal of data could be collected because of the

simplicity and minimally invasive behavior of the program, allowing for a large amount of data over a long period of time to be collected as well.

What's more, this research can determine which parts of a conversation that deceptions are most likely to occur. By analyzing the location of deceptive message(s) within a given conversation as well as the properties of that deceptive message, questions about the motivations and objectives of deceivers can be answered. Certain factors that make people more likely to deceive — such as experience with medium — can also be determined through the questions presented in this research. These findings have broad implications, impacting everyday conversation between friends, business negotiations and collaborations, or even counter-terrorist activities.

Research Questions and Hypotheses

Although a number of studies have looked at the prevalence of deception in face-to-face interaction (DePaulo, 1996) (Hancock et al., 2004), little research has been conducted to define the characteristics of deception in instant messaging. For one, how prevalent is deception in instant messaging? Preliminary research in this area (Barrett & Hancock, *In Progress*) suggests that approximately six percent of messages within a given computer-mediated communication conversation are deceptive. In this study, participants were asked to rate their own lies in a laboratory setting. Therefore, it was hypothesized that deception would occur in instant messaging at about this rate.

Second, are these "big" or "small" deceptions? Not all deceptions are of the same magnitude: some are big and some are small. How deceptive a message is a subjective measure based on what people believe is the idea being used to deceive. Is it the complexity of the deception? How far from the relative "truth" the message seems to be?

To date, no study appears to have analyzed the magnitude of deception within conversations, only focusing on the frequencies or content being discussed. DePaulo at al.'s content analysis study of lies (1996) found that "participants did indeed describe their lies in matter-of-fact ways. They said that their lies were generally not serious ones." (pp. 991) Therefore, given the similarities between instant messaging and face-to-face, and given the widespread usage for communication and social nature of instant messaging, it was hypothesized that people will tell smaller lies most frequently, and bigger lies infrequently.

Third, do different kinds of people deceive about different things? In other words, does gender or experience in the medium affect how frequently or how strongly people deceive one another? Again, this has not been analyzed in the instant messaging medium of communication, and only characteristics by frequency have been assessed at the face-to-face level.

DePaulo et al. (1996) found that women and men lie at about the same frequency, so it was hypothesized that this would also be the case in instant messaging interactions. DePaulo did a content analysis of lies and found that men lied about different kinds of things than women, but did not go so far as to see if the magnitude of those lies differed. Furthermore, it would seem intuitive that given more experience in a medium of communication, a person would feel more comfortable in it and using it. Hancock et al. (2004) found that experience was correlated with lying in email. Therefore, it was hypothesized that the more experience someone has using instant messaging, the more that person will feel comfortable to deceive, and might be inclined to deceive more frequently or tell bigger deceptions.

Fourth, do people have a good sense of how much they actually deceive? Previous studies had relied on people accurately reporting their deceptions even with time delay, allowing for less accurate tracking of deceptions. It is likely that people forget some of their smaller deceptions because deception in communication is so commonplace (DePaulo,1996). Following that line of thought, is it possible that people are not aware of how frequently they deceive in conversations? Given previous studies' (DePaulo et al., 1996, 1998) findings that lies of lesser magnitude are commonplace and therefore probably overlooked, it was hypothesized that people will believe they deceive less than their actual rate of deception.

Sequential Patterns of Deception. Instant messaging produces individual text-based messages. Although it is possible for multiple propositions to be contained within a single message, this is not often the case. This isolated, sequential nature of instant messaging allows for entire conversations to be analyzed by each message, rather than at the overall conversation level.

McCornack (1992) notes that most studies of deception do not look at deception on a message-by-message basis, but at the overall conversation level. These studies focus almost exclusively on testing for the occurrence of deception in a conversation, but do not seek to know where in the conversation it occurred. Furthermore, some studies sought to qualify the content of deceptions or lies (e.g., DePaulo et al, 1996), but few, if any, have analyzed deception at the level of individual messages and try to quantify their magnitude.

Instead of looking at whether people lie in an overall conversation, the present study analyzes how individuals deceive within a conversation, and if it can be used to

detect other lies. That is, given the location of one lie in a stream of messages, can it be used as the "anchor point" to detect other lies? There are three sequential possibilities for how lies fit into a message stream. First, that lie messages occur randomly in the message sequence, and therefore, there is no clustering of deceptive messages in conversations. This is unlikely, however, as people will probably not insert deceptive messages randomly in their conversation given that deception appears to be a strategic communicative act (Burgoon, 1994).

Second, lies may spread out across a conversation. That is, lies will not occur together in the message stream but instead will be spaced out, a strategy akin to hiding lies amongst truths. This strategy makes it difficult to track where a lie is occurring, thereby potentially making it less likely that a person could catch the liar in the act. This strategic use of deception seems more likely than the first, given the logic in masking deception and the advantages of using it instead of lying without assessing the usage of lies. However, from a cognitive point of view, it may be difficult to hold in memory 1) when the last lie was told and 2) when to insert the next lie in the message stream.

A third possibility is that lies cluster together in the message stream. Activation and priming research (Bargh, 1996) suggest that once one lie is told, the liar may be primed to tell more lies within the conversation. What's more, this may be especially true if the lie is of a large magnitude, since the person may be inclined to back up the lie with many smaller, supporting lies (DePaulo, 2004). Thus, if a lie of large magnitude can be detected, the messages immediately following and within the same topic should be also be suspect. That is, lies will cluster in the stream of messages.

Methods

Participants. A total of 35 people took part in the study. Three did not collect data because of technical error, and were excluded. Four others collected fewer than 20 rated instant messages, so their results were discounted as outliers for a number of analyses (N=28). The final sample consisted of 15 male and 17 female undergraduates taking a number of Communication classes at Cornell University. Particpants were members of the College of Agriculture and Life Sciences, Industrial and Labor Relations, College of Engineering, and College of Arts and Sciences, ranged from age 18 to 22 (average = 20.6) and had 5 to 15 years of experience with instant messaging (average = 8.5 years). Participants received credit for their participation.

The Apate Client. For the current study, a group of Cornell students in a computer science class modified the open source instant messaging client Pidgin. The program was modified to have a window pop up after the participant sends a message, asking the participant to rate how deceptive the message is on a scale from 0 (not deceptive) to 5 (highly deceptive). Participants were instructed to rate their truths as 0 and their lies on a scale from 1 to 5.

Once the message was rated, it was logged for later analysis by a remote server. The entire conversations, on a message-by-message basis, were logged on the participants' computer, with all message contents, ratings, screenames, timestamps and keystroke data recorded for each message. At a specific time every night, the modified Pidgin program would allow the remote server to retrieve this log file for data collection and later analysis, thereby simplifying the process of data collection without significant intrusion on the participant. This system allowed participants to rate their messages in

near-synchronous time, giving the data unprecedented accuracy in its rating of deceptive messages without the inaccuracies of recalling deceptions through memory. It also allowed participants to rate deception on the message level, again at unprecedented accuracy, so that the conversation could be analyzed from both the overall conversation level and the message-by-message level, as well as be measured for magnitude in both a conversation and on a message-by-message basis.

Apate: Experiment Management System. The data collection and procedures were handled through a custom-built automated system called Apate (Amos & Perlin, 2008). This system automated all tasks related to the experiment, including the experimental procedure, instructions, surveys, emails, participant information (for credit), collected data, and a link to installing the modified Pidgin client (with instructions, including how to uninstall the client).

Pre-experiment (see Appendix B) and post-experiment (see Appendix I) surveys were administered through the online web interface run by Apate. Apate would send reminder emails for participants, and instructed them how to install the modified version of Pidgin, how to uninstall upon completion, and provided links and instructions for each step of the experimental procedure (see Appendices C -K)

Finally, in order to analyze online deceptions outside the laboratory, participants were required to download this modified instant-messaging client to their hard drive, install it and set it up using instructions provided. After the experiment was completed the program could be removed and all records deleted from the participant's computer.

Recruitment. Participant recruitment was conducted in two ways: through a signup sheet outside the Communication Department office, and through announcements

in a number of Communication classes. The latter required participants to contact the experimenter directly via email, while the former required participants to provide their email address. Participants received extra credit for their participation in the experiment.

Procedure

Participants who expressed interest through recruitment were added to a participant list on Apate, at which time the system would send an email with a hyperlink to a consent form. From this point, participants could read the consent form and decide whether or not to participate. Upon giving consent electronically, participants were stepped through (via the Apate web client) a pre-experiment questionnaire, information and questions on how to rate deceptive messages consistently and accurately, and then given instructions and a link to install and use the modified Pidgin client.

After installing Pidgin, participants would then use the client as their regular instant messaging interface, rating messages as they were sent, 0 (not deceptive) to 5 (very deceptive). Automatic reminder emails were sent nightly, telling participants to continue their participation. Sent/received messages, along with ratings and other message information, were collected nightly. Data was “annonymized” automatically by stripping out all names (using the 5,000 most common first and last names in the U.S.) and changing screenames in analyzed data. Raw and “annonymized” data were both kept on the server.

Participants were obligated to use the client for four days, at which time the server sent an email with a link to completing the post-experiment survey (through the Apate web client) and finishing the experiment. Upon completion, the server stopped collecting

information from the participant, and the participant received a thank you email, including w instructions on how to remove Pidgin from their system.

Coding. Since a number of the pre-experiment and post-experiment data was entered in text format, a range of entries could be used for different answers, such as "how many years have you been using IM" and "approximate number of people on your buddy list." In this case, if multiple years were entered (ie. "4 - 5"), the largest number was used for data analysis.

Results

In all, 10,651 messages were collected from 32 participants. 4,821 of those messages were sent and rated by participants (45.3 percent), while the rest were messages written by partners. Broken down by gender, 46.5 percent of messages were rated by 15 males and 53.5 percent were rated by 17 females.

All but one participant had more than five years of experience with instant messaging, and more than half of the messages rated (56.8 percent) were done so by participants with nine or more years of experience. Overall, 63.99 percent of rated messages came from participants who used instant messaging daily or more than once per day, and 91.9 percent of messages were rated by participants who used instant messaging more than once per week.

Participants rated how similar to normal instant messaging use their instant messaging habits were during the study (1 to 5, "not at all similar" to "similar," respectively). Participants rated their use, on average, to be 3.6, with a median of 4 and a

mode of 4, suggesting that the captured instant messaging behavior was typical of participants' usual instant messaging behaviors.

Frequency. Of the 4,821 rated messages, 552 (11. 5 percent) were self-reported as deceptive (ranked 1 through 5) by participants. Messages written by males involved deception 13.7 percent of the time, while messages written by females involved deception 9.5 percent of the time.

To examine how participants' characteristics such as gender and experience in instant messaging affect lying behavior, the rate of deception was calculated for each participant. Calculated at the participant level, the average rate of deception was 15.3 percent, somewhat higher than the overall message analysis. This difference is due to averaging each participant rather than all the messages as a whole.

A comparison across gender showed that the rate of deception was not significantly different between males (M = .17, SD = .13) and females (M = .14, SD = .11), $t(29) = 0.46$, *ns*. The next analysis examined whether experience had an impact on lying frequencies. Three measures of experience in instant messaging were asked of participants: years of instant messaging, frequency of instant messaging use, and approximate size of the participant's "buddy list" (a manually administered list of screenames to whom the person speaks). None of the experience measures had a significant effect on deception frequency.

Magnitude. Ratings of all deceptive messages (ratings 1 to 5) were examined. More than half of messages were rated 1 or 2, the lowest magnitude ratings. Looking at all of the deceptions, the greater the magnitude of the lie, the fewer number of messages

received that rating. The one exception to this was the highest rating, 5, which had a slight increase from 4.

Males and females rated their lies almost identically. Out of the rated messages marked as deceptive, males (M = 2.53, SD = 1.42) and females (M = 2.51, SD = 1.43) had almost the exact same average magnitude of deception and variation. To examine the relationship of experience on lying behavior, three indicators of instant messaging experience were analyzed again: years using instant messaging, frequency of instant messaging, and the self-reported size of the participant's buddy list.

The data shows that the number of years using instant messaging (r = .138, *ns*) and number of instant messaging buddies (r = .164, *ns*) were not significantly correlated with lie magnitude. In contrast, the reported frequency of instant messaging use was positively correlated with lie magnitude (r = .38, p = .05), suggesting that the more often people use instant messaging, the greater the magnitude of lies told.

Experience. Buddy list size was positively correlated with increased frequency of instant messaging use. In other words, the more frequently participants used instant messaging, the more buddies participants had on their buddy list.

Sequential Analysis. To evaluate whether lies and truths occur randomly in the message stream, a nonparametric runs test for binary data (Gibbons, 1985) was employed. The runs test is used to test the independence of two elements in a sequence, with a run defined as a segment of the sequence with adjacent, similar responses (lies or truths). The runs test compares the actual number of runs with the expected number of runs, using a z-score distribution. The negative z-score (too few runs to be in random)

indicates clumping, the positive z-score (too many runs to be random) indicates more alternating, and a value close to zero indicates randomness.

To obtain a two-valued data sequence, lies of different magnitude were re-coded into a single value "lie." Individual runs tests were done on 29 participants, yielding individual *p*-values. Based on Goutis, Cassella, & Wells's (1996) methods, the omnibus p-value was derived to assess a multivariate hypothesis of lie/truth distribution independence by combining *p* values from independent tests. This approach relies on the fact that under the null hypothesis a *p*-value from a continuous test statistics follows a uniform distribution from 0 to one. It is often used in the meta-analysis to evaluate evidence from individual studies. The results revealed a pooled value based on the sum of –2 log p-values derived from individual tests gave $\chi(58) = 78.275$, $p = .039$, suggesting that lies tended to cluster together in the message stream more than would be expected by chance.

Beliefs. Participants were also asked to rate how frequently they believed they had lied. On average, they believed they had lied in 8.0 percent of their messages. This overall estimate of lie frequency was significantly lower than the frequency calculated by their message-by-message identification of lies of 15.4 percent, $t(15)=3.047$, $p=0.008$, demonstrating an substantial underestimation bias. Nonetheless, despite this difference, participant estimates were strongly correlated ($r = .901$, $p < .0001$) with their actual rates. The results suggest that participants have a sense of how much they are lying, but that they underestimate how much they are lying.

Discussion

The present study analyzed deception at the message-by-message level during participants' actual everyday conversations using instant messaging, which to the best of our knowledge has not been done before. DePaulo et al. (1996) reported participant deception but not at the message level, only asking for recollection of how many lies were told overall. Thus, many lies could have been told in one conversation, but might have been counted as one deception instead of many, as was the case in this study. Given this, and the recording as part of day-to-day activity not in the laboratory, we consider this study to be a more accurate indicator of instant messaging deception. In looking at messages, we considered deception at the message level in terms of lie frequency, magnitude, sender characteristics, and sequential patterns.

Frequency. The observed overall frequency of deception in instant messaging, 11.5 percent, as well as the average rate of deception of 15.3 percent, is consistent with prior research that showed prevalence of lying in everyday conversation. This supported the hypothesis that lying would be pervasive in instant messaging. The deception frequency is more than double the observed rate in previous preliminary research (6 percent), which could be a consequence of rating not in the laboratory setting with goals in mind, as well as a better rating scale to self-rate lies.

It is difficult to compare these averages with any studies of the past because no study has ever looked at lies at the message-by-message basis. Therefore, it is difficult to know without further analysis if the deception rate presented does compare favorably to previous studies on a per conversation basis. Regardless, the average rate of instant messaging deception is rather high, and near estimates of face-to-face deception.

Hancock et al.'s feature-based model of communication outlines that, "the more synchronous and distributed, but less recordable, a medium is, the more frequently lying should occur." (pp. 130) In the case of instant messaging, the feature-based model predicts that lying will occur in instant messaging frequently because partners are physically distributed and, typically not recorded, and is relatively synchronous. The data were largely consistent with this prediction as deception was frequently reported in the medium.

Magnitude. Overall, people told small magnitude lies most often, and lies of larger magnitude less frequently. Overall, 31.9 percent of lies were rated the lowest possible magnitude (1) and more than half the lies told were rated 1 or 2. This is consistent with DePaulo et al.'s findings that small lies are told the most frequently as part of normal conversation.

Interestingly however, nearly half (47.3 percent) of the lies told were rated 1 or 5. This shows that participants tended to self-report their lies as being small or big. This could be because they did not associate their leis as being moderate deceptions. Although it is possible that participants only told very big or very small lies, it seems more likely that unfamiliarity with the rating system and an unwillingness to spend too much time focusing on the magnitudes of the lie caused them to rate high or low.

Gender. Males lied in an average of 17.0 percent of their messages, while females deceived in 13.9 percent of their messages on average. However, gender did not play a significant role in deception frequency or magnitude. As hypothesized, these findings are consistent with DePaulo et al.'s findings in frequency of face-to-face deception. DePaulo (1996) did note differences in content of lies between males and females, which could

have lead to indicators of lie magnitude difference, depending on the topics. However, no significant differences in the lying rates or magnitudes between genders were found.

Experience and Frequency. Experience was measured in three ways: how many years of experience participants reported having with instant messaging, how frequently participants used instant messaging normally, and how many people were on participants' buddy list. The logic behind this last measure is that given more people to talk to currently (since it cannot be known how many people were spoken to in the past), participants should have more experience in the medium, as years of experience and frequency of instant messaging use may not have as strong an impact on experience as number of people with whom there was (probable) interaction.

No measure of experience influenced the frequency of deception. When compared with Hancock and colleagues' findings, this can be interpreted in one of two ways. First, Hancock et al found that experience appeared to be correlated positively with rate of deception in email, "suggesting that more experienced email users lied more frequently in email than less experienced users." (Hancock et al., p. 132) This can be interpreted to mean that experience should have a positive correlation in other, text-based communication mediums like instant messaging.

Looking at the results from this perspective, one reason for this difference in results may have had to do with the sample of people selected: participants in the Hancock et al. study could be experienced with any number of a multiple mediums of communication and not necessarily instant messaging, whereas participants in this study all had at least five years of experience in instant messaging. This is likely because an instant messaging study appealed to participants who used instant messaging frequently.

Therefore, the sample of people used in the Hancock study may have been a more accurate reflection of the general instant messaging population.

However, Hancock et al. did not find a correlation between experience in instant messaging despite finding it in email, suggesting that perhaps experience doesn't have the same effects in instant messaging as it does in other modes of text-based communication. Therefore, it seems that these findings support the conclusions of Hancock et al., that no difference in frequency can be found from experience in instant messaging.

Experience and Magnitude. In terms of magnitude, neither buddy list nor years of instant messaging experience correlated to the magnitude of lies told. Self-reported frequency of day-to-day and week-to-week instant messaging habits, however, did have a significant positive correlation on the magnitude of lies, suggesting that the more participants used instant messaging, the greater the size of their lies.

Since the magnitude results correlate positively to a measure of experience in the medium, these data are consistent with channel expansion theory (Carlson & Zmud, 1999). Channel expansion theory hypothesizes that people become more skilled and comfortable with their medium the more they use it, and this would include being more adept and willing to deceive.

Experience. An important set of questions to ask, however, is do these three measures really demonstrate "instant messaging experience," and if not, what are they actually measuring? What is experience in the case of instant messaging? Based on email and other text-based forms of communication, it should be amount of interaction in the medium. Emails have number of messages and size/language usage, as well and sender

and receiver data. This data should be sufficient to determine experience in instant messaging, since the mediums have similar characteristics, with the critical difference being the size of messages sent on a message-by-message basis. (The speed of instant messaging is also different, though this is probably trivial for the purposes of experience measures.) However, unlike in instant messaging, (all of this information) emails are saved on the served by default, while instant messages are not unless logging is enabled. Thus, it is difficult to find a true measure of experience in instant messaging without meticulous, detailed logging, especially over time with people obtaining new computers and chatting from many different computers as well.

The limits of instant messaging as a medium make it difficult to find one true measure of experience for it. If two measures are good at quantifying experience in the medium, they will probably be correlated with each another. The only set of experience variables that were correlated with one another were buddy list size and frequency of instant messaging. This shows that those who use instant messaging more frequently tend to have more buddies or “friends” in the medium. However, years of instant messaging experience did not have a correlation to either of the other experience measures, suggesting that it is not enough to have been using instant messaging for a significant period of time to have “experience” in the medium. Intuitively, years of instant messaging usage doesn’t necessarily measure the quantity of usage, only the amount of time since usage began, which makes it a poor measure of experience in instant messaging.

Furthermore, buddy list size is probably not a particularly accurate indicator of experience on its own. Although probably a reliable indicator that at least conversation

has occurred in the past with that screename, the buddy list does not necessarily quantify anything other than the awareness of a person's presence in instant messaging. Finally, frequency of instant messaging usage doesn't guarantee talking in instant messaging. Therefore, although no one measure outside of logging is perfect for measuring history, putting multiple measures together probably can help paint an accurate picture of participant instant messaging habits though time. Given the correlations, it seems that buddy list size and instant messaging usage are a good combination of indicators. Ideally, this would be used in tandem with message logging, since the major drawback of the other two measures is that neither guarantees interaction between buddies.

Lie Clustering. Sequential analysis of messages found that participants tended to cluster their lies together in conversations, not spread them out strategically or randomly disperse them. The reasons for this could be numerous and related to one another. For one, lies are likely told for a strategic reason, to logically serve a purpose in conversation (Burgoon, 1994). Therefore, it would make sense that if multiple lies are told, they should support one another.

A second reason might be related to the content of the lie itself, with some deceptions needing others that support them or are related. Deception is not necessarily limited to just a single message, it can be about an idea over multiple messages. Therefore, it would make sense for some deceptive messages to be grouped together to support a deceptive idea, or for some deceptions to be related to one another's content.

Furthermore, although no analysis was conducted to determine the magnitude of lies in a cluster, given that there are a number of smaller lies and fewer bigger lies, there are three sets of possibilities about how lies of different magnitudes are clustered. One is

that there is no discernable pattern in the dispersion of people's lie magnitude in a conversation. Another possibility is that smaller lies are clustered around, before, or after bigger lies. A final possibility is that smaller lies are clustered together, and bigger lies are clustered. Further analysis is necessary to determine which of these three possible clustering patterns was observed in this study.

Another reason may be that the cognitive load of lying is very high, making it easier for people to lie all at once. The reasons for cognitive load (Zuckerman & DePaulo) could include not wanting to worry about self-presentation, not wanting to deal strategically with lying over the course of a conversation, or just opportunistically jumping at the "opening" in a conversation for the person to lie.

Beliefs. Looking retrospectively through a post-experiment survey, participants were relatively accurate in estimating whether they had lied more or less. They consistently underestimated how often they lied, however. This means that participants could accurately guess if they had lied more or less than normal, but underestimated how often they lied in total. This could be explained by the size of the lies, since with so many little lies being told, it may be difficult to keep track of deceptions retrospectively if they are not seen important, and large in magnitude or memorable.

This finding does not bode well for previous diary studies of lying in face-to-face that relied on participant awareness and memories of conversations and deceptions. For one, it is possible that the act of rating in real-time made participants more aware of their own deceptions than they would normally been, regardless of the medium. After all, there is a big difference in rating messages after the fact, and rating messages in real-time and then being asked to recall your own ratings. Even if this is not the case, however,

previous diary studies may have further underreported lying rates because they do not take into account lies at the message-by-message level. Lies in multiple messages could have been reported as one “conversation” lie, and combined with underreported deception totals from memory, may have resulted in a significant underreporting in diary studies of deception.

Conclusions

The results of this study suggest that people deceive frequently in use instant messaging, and this finding remains consistent with theories and studies that demonstrate lying as a pervasive activity in social interaction. Males and females appear to deceive at a similar rate. Although it is difficult to truly measure instant messaging experience, amount of instant messaging use has a direct relation to the size of lies told, suggesting that experience plays a role in instant messaging deception magnitude.

People seem to have a solid overall grasp on their deviations from normal lying habits, but do not seem to recall how often they are lying. This shows that diary studies, which have made up a majority of deception studies in the past, may not be accurate in their conversation frequency reports. This may be further compounded by the fact that these studies have looked at lies from a conversational level, but not on a message level, which could cause lying rates to be further underreported in face-to-face.

Clustering of deceptive messages was observed, and this supports the idea of people lying strategically. Furthermore, the results show that if one deceptive message can be detected, it’s likely that others have surrounded it in the conversation. This has numerous implications for tracking and catching those deceiving using instant messaging.

Works Cited

Amos, B. & Perlin, J. (2008). Apate Experiment Management System [Computer Software]. Ithaca, NY: Cornell University. Retrieved April, 2008. Available from http://apate.net/

Bargh, J., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: direct effects of trait construct and stereotype activation on action. *The Journal of Personality and Social Psychology, 71* (2), 230-244.

Barrett, T. *Work in Progress*.

Buller, D. B., & Burgoon, J. K. (1994). Deception: Strategic and non-strategic communication. In J. A. Daly & J. M. Wiemann (Eds.), Strategic interpersonal communication (pp. 191-223). Hillsdale, NJ: Lawrence Erlbaum.

Carlson, J. R., & Zmud, R. W. (1999). Channel expansion theory and the experiential nature of media richness perceptions. Academy of Management Journal, 42, 153-170.

Daft, R. L., & Lengel, R. H. Information richness: A new approach to managerial behavior and organization design. In B. M. Staw & L. L. Cummings (Eds.), *Research in organizational behavior* (Vol. 6, pp. 191-233). Greenwich, CT: JAI Press, 1984.

DePaulo, B. M., Ansfield, M. E., Kirkendol, S. E., & Boden, J. M. (2004). Serious lies. *Basic and Applied Social Psychology, 26* (2&3). 147-167.

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,* 979-995.

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

Gibbons, J. D. (1985). *Nonparametric Methods for Quantitative Analysis.* Columbus, Ohio: American Sciences Press, Inc. Goutis, C., Casella, G., & Wells, M. T. (1996). Assessing evidence in multiple hypotheses. *Journal of the American Statistical Association,* 91, *1268-1277.*

Hancock, J.T., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: The impact of communication technologies on lying behavior. *Proceedings, Conference on Computer Human Interaction,* 6, 130-136. New York, ACM.

McCornack, S. (1992). Information manipulation theory. *Communication Monographs*, (59). 1-16.

Nyberg, D. (1993). *The Varnished Truth.* Chicago: The University of Chicago Press.

Nardi, B., Whittaker, S., Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, Philadelphia, Pennsylvania, United States, pp. 79–88.

Pew Internet & American Life. (2004, Sep.1). How Americans Use Instant Messaging. Retrieved September 10, 2007, from the World Wide Web: http://www.pewinternet.org/pdfs/PIP_Instantmessage_Report.pdf

Vrij, A. (2000). *Detecting Lies and Deceit: The Psychology of Lying and its Implications for Professional Practice.* New York: Wiey.

Table 1

Ratings as Percent of Overall Messages

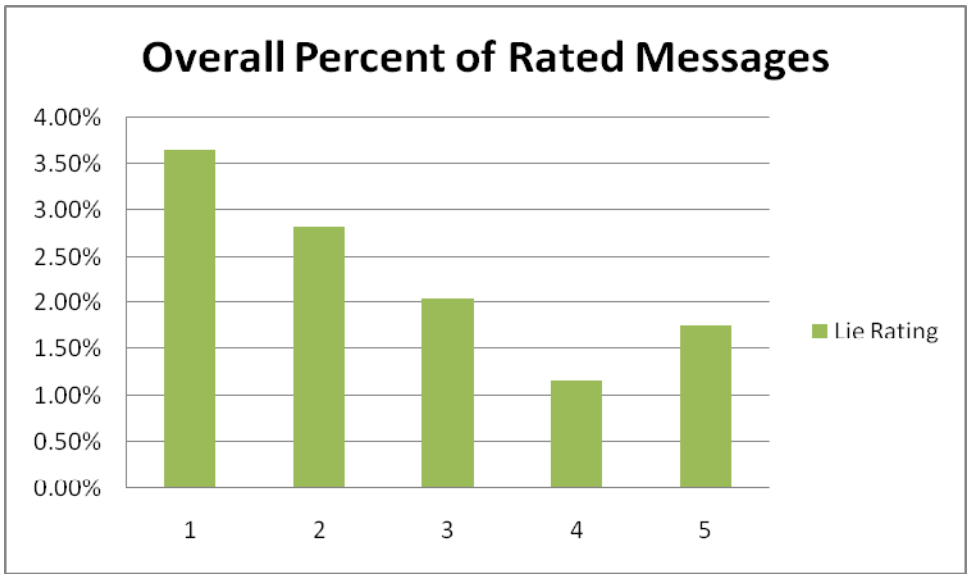

Table 2

Percent of Rated Messages, by IM Usage Frequency

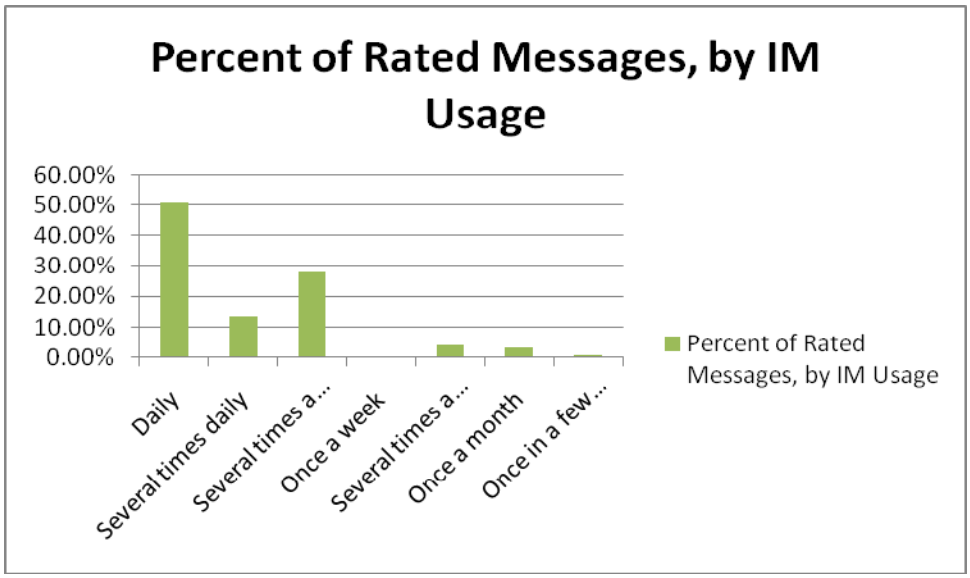

Table 3

Percent of Messages, by Rating

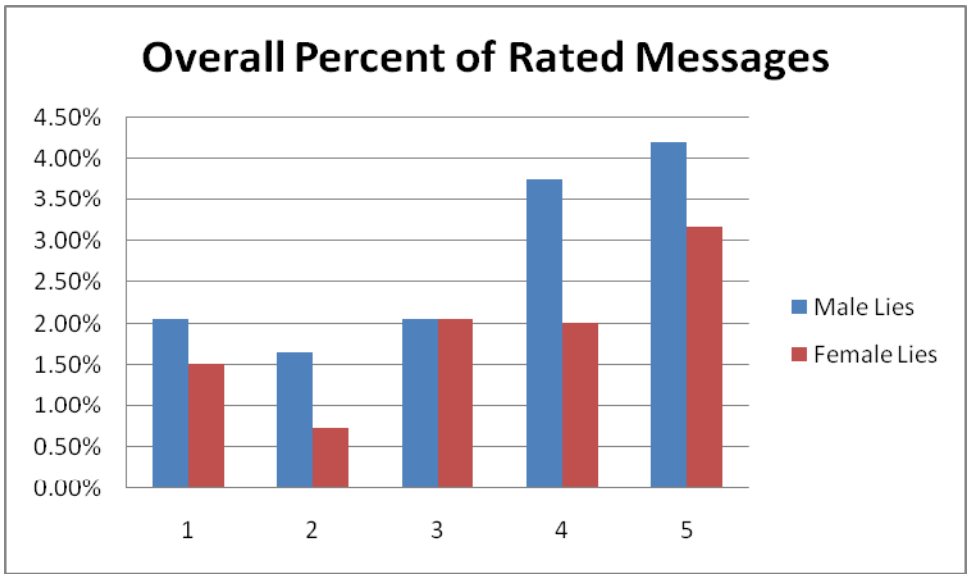

Table 4

Percent Rated Messages, by Instant Messaging Experience

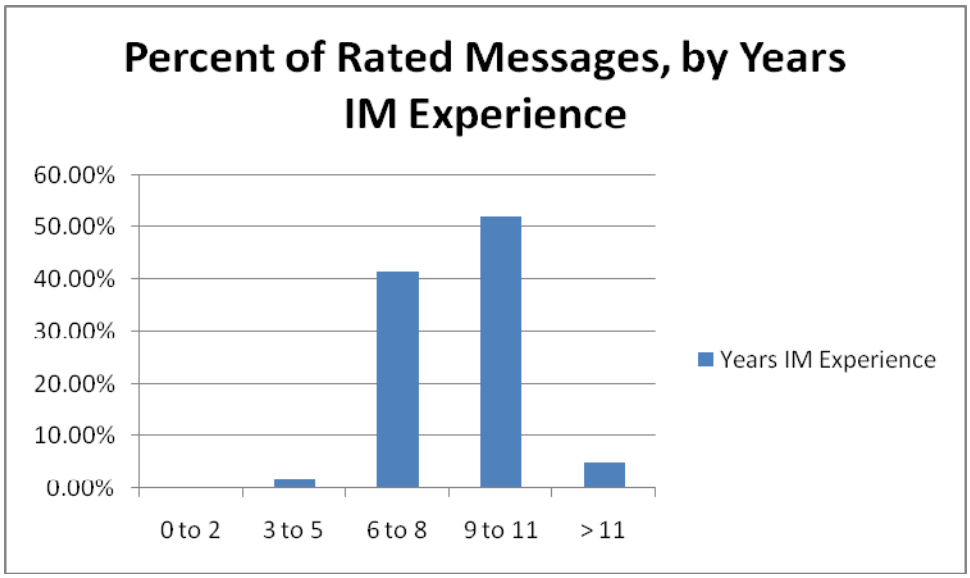

Table 5

Rated Messages Breakdown

Message Types	**Number**
No lies (0)	4269
Lies (1-5)	552
Total Rated (0-5)	4821
Lies %	11.45%

Table 6

Lie Magnitude by Rated/Unrated Messages

Rating (0/1 - 5)	**STDDEV**	**MEAN**	**MEDIAN**	**MODE**
Avg Rating (0-5)	0.94	0.29	0	0
Avg Lie (1-5)	1.42	2.53	2	1

Table 7

Message Ratings by Years Instant Messaging Experience

Years of IM Experience	% of Rated	Average Rating	Average Lie
0 to 2	0.00%	—	—
3 to 5	1.89%	0.19	1.70
6 to 8	41.28%	0.32	2.43
9 to 11	51.73%	0.27	2.62
> 11	5.10%	0.28	3.00
Total	4821		

Table 8

Message Ratings by Instant Messaging Usage

	% of Rated	Avg Rating (0-5)	Avg Lie (1-5)
Daily	50.43%	0.18	2.46
Several times daily	13.57%	0.35	2.64
Several times a week but less than daily	27.92%	0.46	2.78
Once a week	0.00%	No Data	No Data
Several times a month but less than weekly	4.07%	0.44	1.87
Once a month	3.40%	0.16	1.37
Once in a few months	0.62%	0.30	1.80
TOTAL	4821		

Table 9

Message Breakdown by Gender

	All Msgs	**Rated**	**% Rated**
Male	4824	2242	46.50%
Female	5827	2579	53.50%
Total	10651	4821	100.00%

Table 10

Percent Message Ratings by Gender

Male			*Female*		
Rating	**Total**	**% of rated**	**Rating**	**Total**	**% of rated**
5	46	2.05%	**5**	39	1.51%
4	37	1.65%	**4**	19	0.74%
3	46	2.05%	**3**	53	2.06%
2	84	3.75%	**2**	52	2.02%
1	94	4.19%	**1**	82	3.18%
0	1935	86.31%	**0**	2334	90.50%
Total	2242	100.00%	**Total**	2579	100.00%

Table 11

Message Ratings by Gender

Males	**STDDEV**	**MEAN**	**MEDIAN**	**MODE**
All Rated Msgs (0-5)	1.02	0.35	0	0
Deceptive Only (1-5)	1.42	2.53	2	1
Females	**STDDEV**	**MEAN**	**MEDIAN**	**MODE**
All Rated Msgs (0-5)	0.86	0.24	0	0
Deceptive Only (1-5)	1.43	2.51	2	1

Table 12

Participant Variable Analysis for Lie Frequency

Descriptive Statistics			
	Mean	Std. Deviation	N
Total Lies by total senders' messages	.1536426	.11750561	30
Age	20.62	1.879	32
Years of using IM	8.56	1.983	32
Number of buddies on the buddy list	171.13	94.063	31
IM frequency	5.06	1.605	32
Estimate percent lies	.08	.072	17

Correlations							
		Total Lies by total senders' messages	Age	Years of using IM	Number of buddies on the buddy list	IM frequency	Estimate percent lies
Total Lies by total senders' messages	Pearson Correlation	1.000	.015	.060	-.303	-.155	.901**
	Sig. (2-tailed)		.937	.751	.110	.414	.000
	N	30.000	30	30	29	30	16
Age	Pearson Correlation	.015	1.000	.206	-.276	-.088	.079
	Sig. (2-tailed)	.937		.259	.134	.631	.764
	N	30	32.000	32	31	32	17
Years of using IM	Pearson Correlation	.060	.206	1.000	-.052	-.092	-.241
	Sig. (2-tailed)	.751	.259		.781	.615	.352
	N	30	32	32.000	31	32	17
Number of buddies on the buddy list	Pearson Correlation	-.303	-.276	-.052	1.000	.463**	-.321
	Sig. (2-tailed)	.110	.134	.781		.009	.209
	N	29	31	31	31.000	31	17
IM frequency	Pearson Correlation	-.155	-.088	-.092	.463**	1.000	-.043
	Sig. (2-tailed)	.414	.631	.615	.009		.869
	N	30	32	32	31	32.000	17
Estimate percent lies	Pearson Correlation	.901**	.079	-.241	-.321	-.043	1.000
	Sig. (2-tailed)	.000	.764	.352	.209	.869	
	N	16	17	17	17	17	17.000
**. Correlation is significant at the 0.01 level (2-tailed).							

Table 13

Participant Variable Analysis for Lie Magnitude

Descriptive Statistics			
	Mean	Std. Deviation	N
Mean of lie magnitude	2.2937	.82019	28
Age	20.62	1.879	32
Years of using IM	8.56	1.983	32
Number of buddies on the buddy list	171.13	94.063	31
IM frequency	5.06	1.605	32
Estimate percent lies	.08	.072	17

Correlations							
		Mean of lie magnitude	Age	Years of using IM	Number of buddies on the buddy list	IM frequency	Estimate percent lies
Mean of lie magnitude	Pearson Correlation	1.000	-.072	.138	.164	.379*	.117
	Sig. (2-tailed)		.715	.483	.413	.047	.678
	N	28.000	28	28	27	28	15
Age	Pearson Correlation	-.072	1.000	.206	-.276	-.088	.079
	Sig. (2-tailed)	.715		.259	.134	.631	.764
	N	28	32.000	32	31	32	17
Years of using IM	Pearson Correlation	.138	.206	1.000	-.052	-.092	-.241
	Sig. (2-tailed)	.483	.259		.781	.615	.352
	N	28	32	32.000	31	32	17
Number of buddies on the buddy list	Pearson Correlation	.164	-.276	-.052	1.000	.463**	-.321
	Sig. (2-tailed)	.413	.134	.781		.009	.209
	N	27	31	31	31.000	31	17
IM frequency	Pearson Correlation	.379*	-.088	-.092	.463**	1.000	-.043
	Sig. (2-tailed)	.047	.631	.615	.009		.869
	N	28	32	32	31	32.000	17
Estimate percent lies	Pearson Correlation	.117	.079	-.241	-.321	-.043	1.000
	Sig. (2-tailed)	.678	.764	.352	.209	.869	
	N	15	17	17	17	17	17.000
*. Correlation is significant at the 0.05 level (2-tailed).							
**. Correlation is significant at the 0.01 level (2-tailed).							

Appendix A

Consent Form

You are invited to participate in a research study examining deceptive online messages. We ask that you read this form and ask any questions you may have before agreeing to be in the study.

Background Information: The purpose of this study is to better understand deception in instant messaging.

Procedures: If you agree to be in this study, we will ask you to do the following: Using a modified instant messenger program called Apate, we would like you to rate each message that you send as either not deceptive or deceptive (on a 0 - 5 scale, with 0 meaning no deception, 1 meaning slightly deceptive and 5 meaning very deceptive). You will rate messages during your regular instant messaging. The modified interface has been designed so that it would be comparable to your normal instant messenger. You will rate your messages for 4 days.

Risks and Benefits of Being in the Study: As is the case with all online communication, there is the risk that your messages will be read, intercepted, or modified by a third party. Like regular instant messaging, your messages will be transmitted in "clear text," with no attempt to encrypt, validate, or otherwise protect the data. This includes the deceptiveness rating which you send with each message.

There is also a risk, although very unlikely, that the Apate instant messaging software you will install may cause harm to your computer, physically, or virtually. This is the same risk you take when you install any program onto your computer. Apate is a modified version of "Pidgin." Pidgin is a well established program, and we have experienced no difficulties during our initial testing. You can find more information about the software used in this study at http://www.pidgin.im/. As an open source program you will have the opportunity to download the source code to this program. There are no direct benefits to participating in the study, although you will receive course credit if you are in a course that awards credit.

Compensation: You may be offered extra credit in participating courses with the consent of your instructor.

Voluntary Nature of Participation: Your decision whether or not to participate will not affect your current or future relations with Cornell University or with other cooperating entities. If you decide to participate, you are free to withdraw at any time without affecting those relationships with the University or with other cooperating entities. The extra credit will be yours either way.

Confidentiality: The records of this study will be kept private. In any sort of report we might publish, we will not include any information that will make it possible to identify you. Research records will be keptin a secure database; only the researchers will have access to the records. Proper names in the messages will be removed automatically from the text, although it is possible that some unusual names may remain.

Contacts and Questions: The researcher(s) conducting this study are Josh Perlin and Jeff Hancock. Please ask any questions you have as soon as possible. If you have questions later, you may contact them at jcp57@cornell.edu; or at jth34@cornell.edu. If you have any questions or concerns regarding your rights as a subject in this study, you may contact the University Committee on Human Subjects (UCHS) at 607-255-5138, or access their website at http://www.osp.cornell.edu/Compliance/UCHS/homepageUCHS.htm.

This consent form was approved by the UCHS on 08/02/2006.

Statement of Consent: By typing your name in this box you are agreeing that you are over the age of 18, have read the above information, and consent to take part in this study.

[Digital Signature] [Date]

Statement of Consent: In addition to my consent to participate, I consent to the inclusion of my messages in a publication. I understand that my messages will be edited to protect the identities of myself and others prior to publication. This is not required to receive credit for participation in this study.

☑ Yes, you may use my messages

Appendix B

Pre –Experiment Questionnaire

Please take a moment to answer this questionnaire before proceeding with the experiment.

Q1)
Age

Q2)
Gender

- Male
- Female

Q3)
Major(s)

Q4)
How many years have you been using IM

Q5)
Approximate number of people on buddy list

Q6)
How often do you use IM

- Daily
- Several times daily
- Several times a week but less than daily
- Once a week
- Several times a month but less than weekly
- Once a month
- Once in a few months

Q7)

Who do you talk to the most in your instant messages

- Co-Workers
- Family
- Boyfriend/Girlfriend/Spouse
- Friends in Ithaca
- Friends not in Ithaca
- Other

Q8)

Who do you talk to the second-most in your instant messages

- Co-Workers
- Family
- Boyfriend/Girlfriend/Spouse
- Friends in Ithaca
- Friends not in Ithaca
- Other

Appendix C

Experimenter Participants Tracking Page

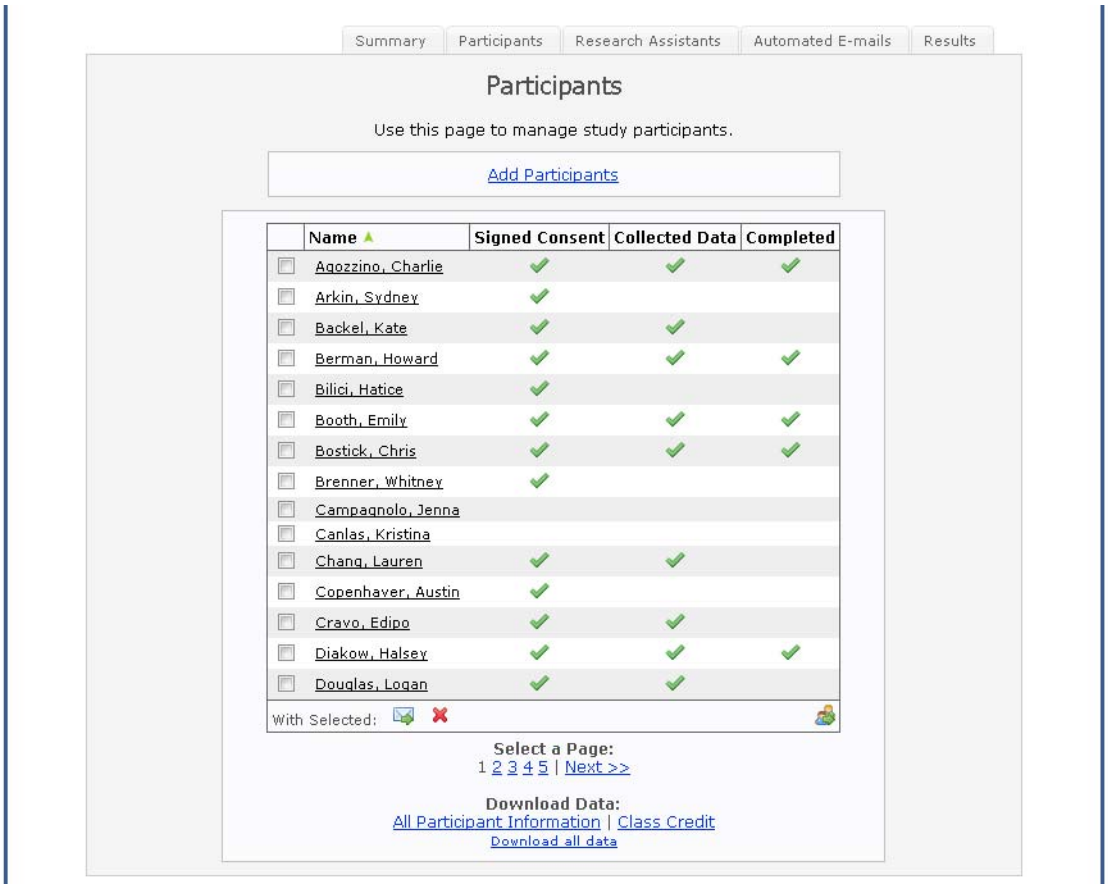

Appendix D

Experimenter Results Page

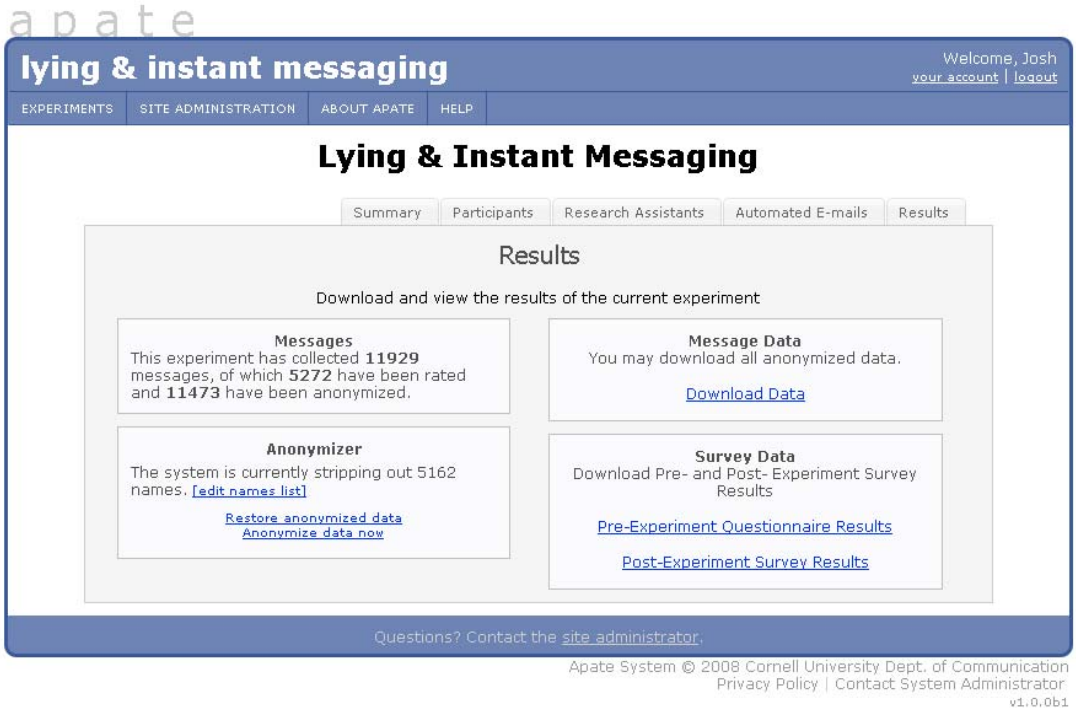

Appendix E

Experimenter Summary Page

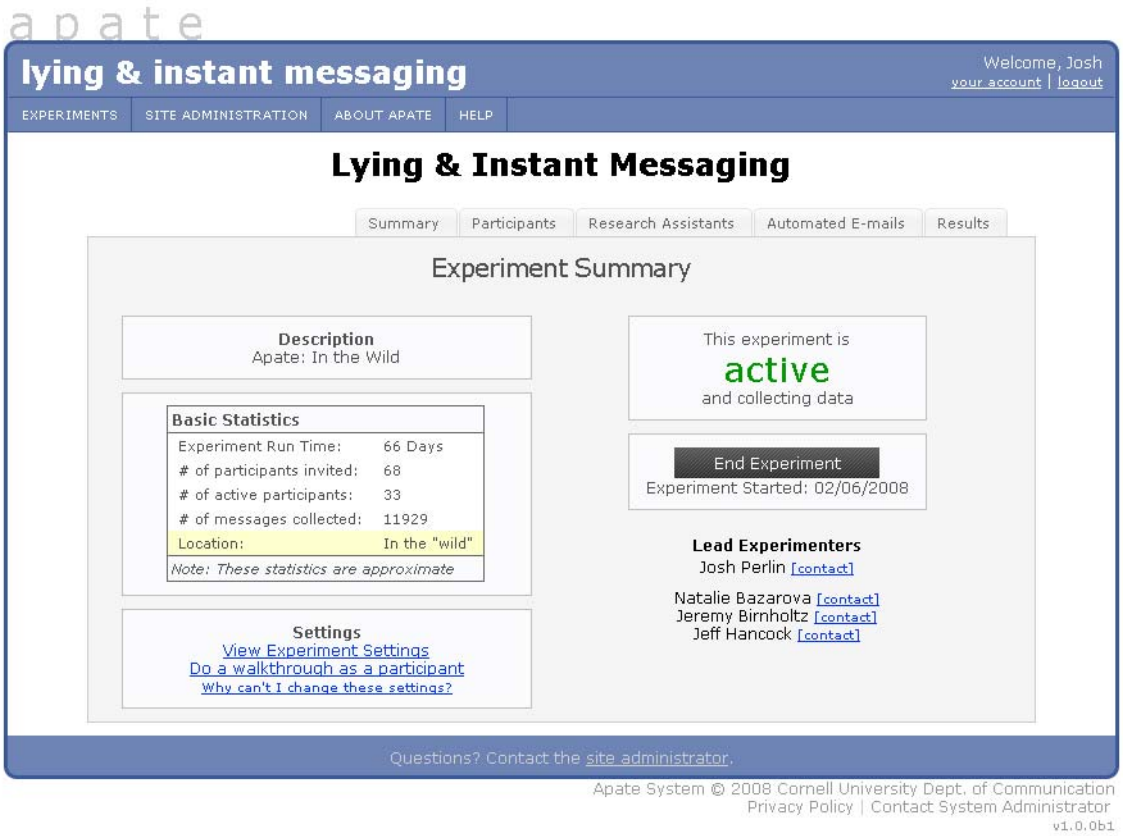

Appendix F

Consent Form Screenshots

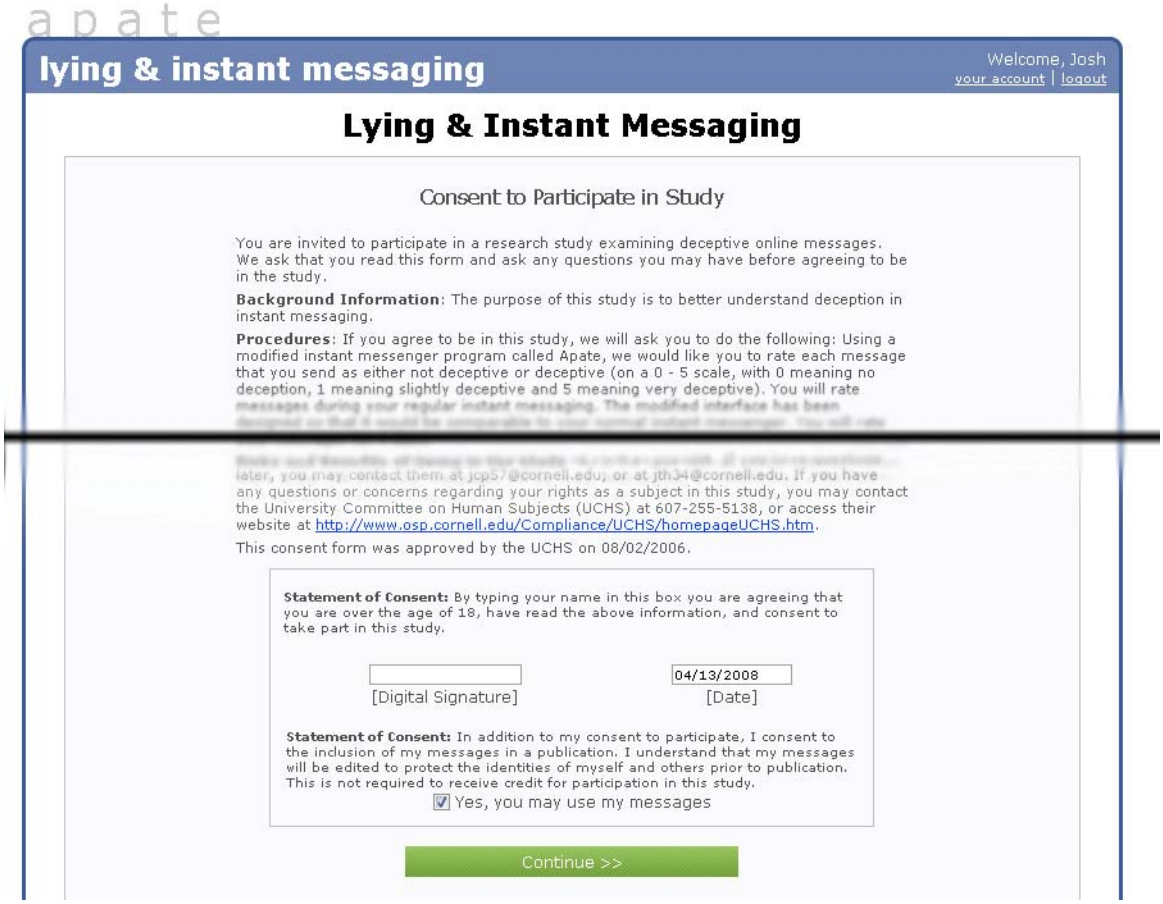

Appendix G

Experiment Instructions

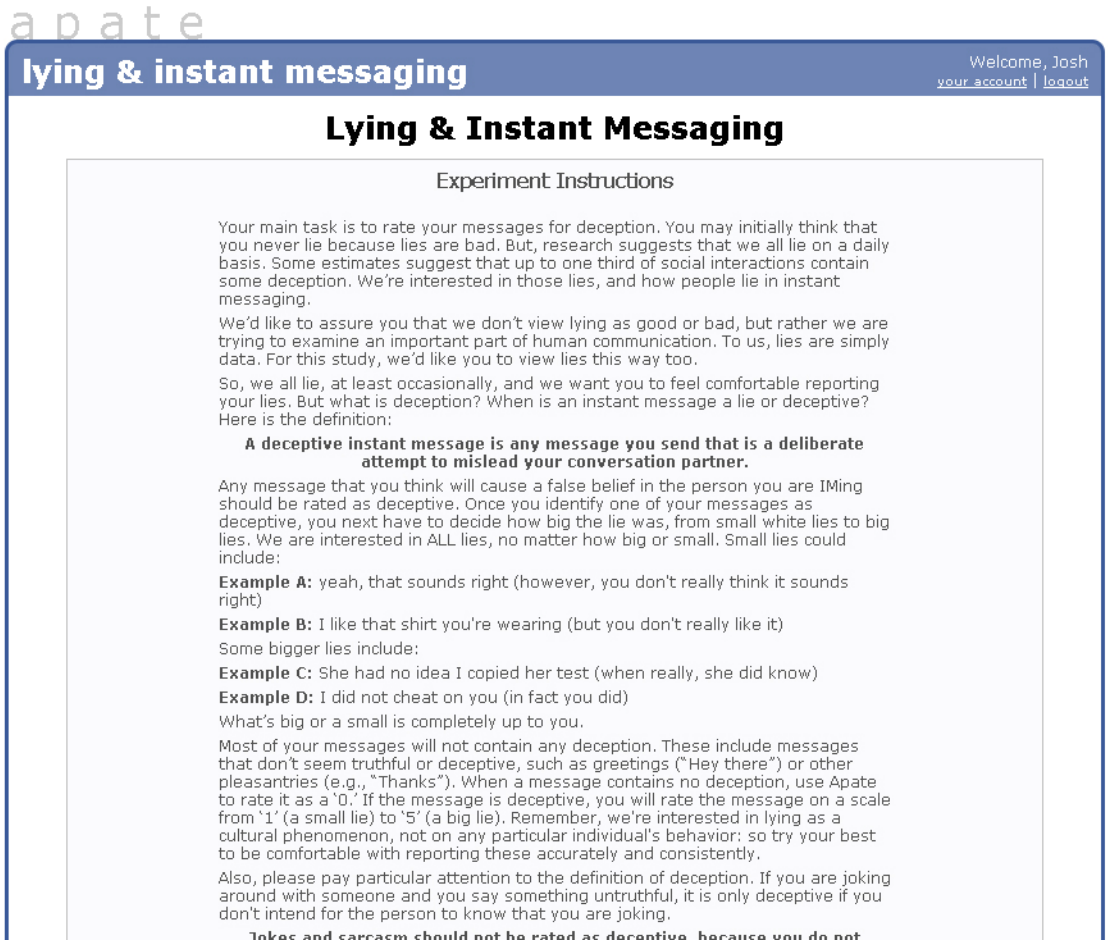

Appendix H

Pre-Experiment Screenshot

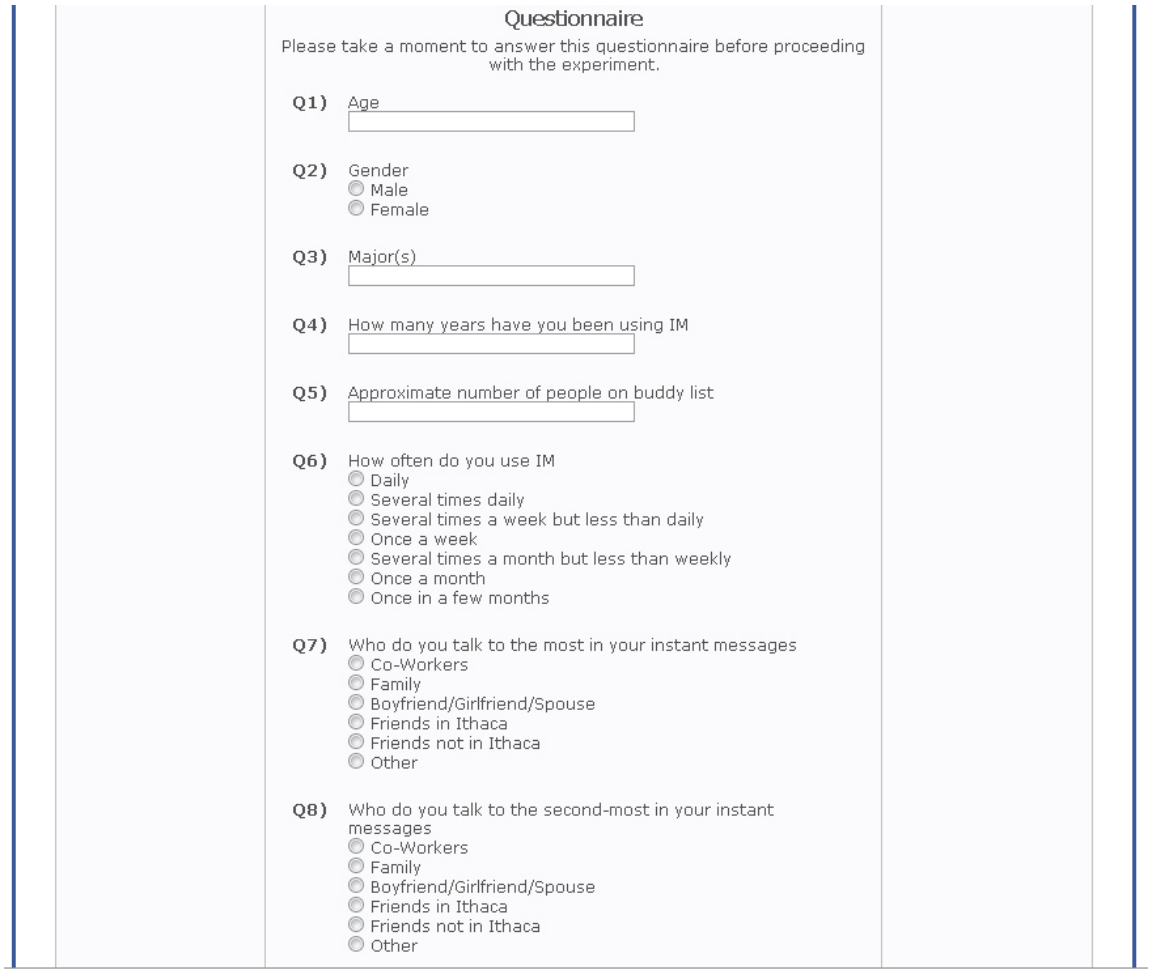

Questionnaire

Please take a moment to answer this questionnaire before proceeding with the experiment.

Q1) Age

Q2) Gender
- Male
- Female

Q3) Major(s)

Q4) How many years have you been using IM

Q5) Approximate number of people on buddy list

Q6) How often do you use IM
- Daily
- Several times daily
- Several times a week but less than daily
- Once a week
- Several times a month but less than weekly
- Once a month
- Once in a few months

Q7) Who do you talk to the most in your instant messages
- Co-Workers
- Family
- Boyfriend/Girlfriend/Spouse
- Friends in Ithaca
- Friends not in Ithaca
- Other

Q8) Who do you talk to the second-most in your instant messages
- Co-Workers
- Family
- Boyfriend/Girlfriend/Spouse
- Friends in Ithaca
- Friends not in Ithaca
- Other

Appendix I

Post-Experiment Questionnaire

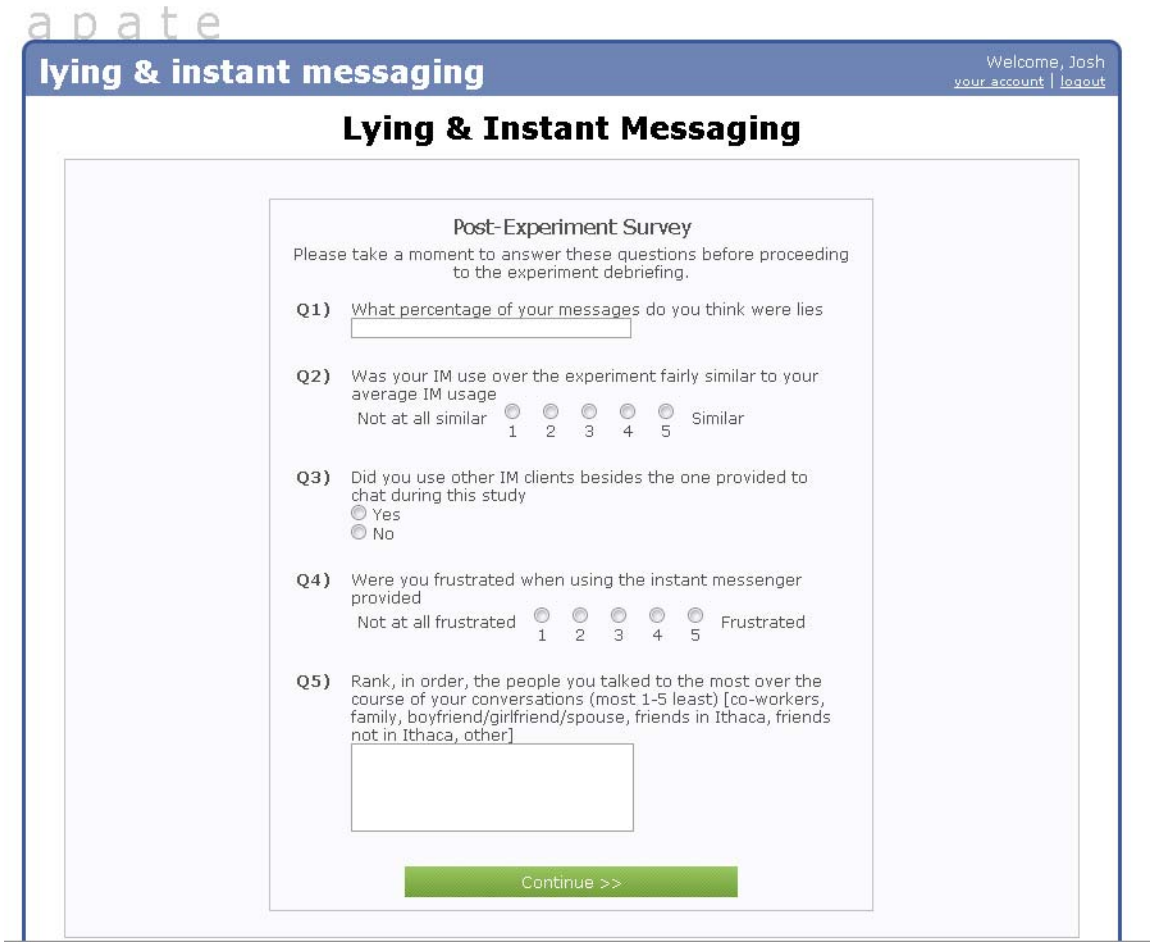

Appendix J

Lie Self-Rating Reliability Test

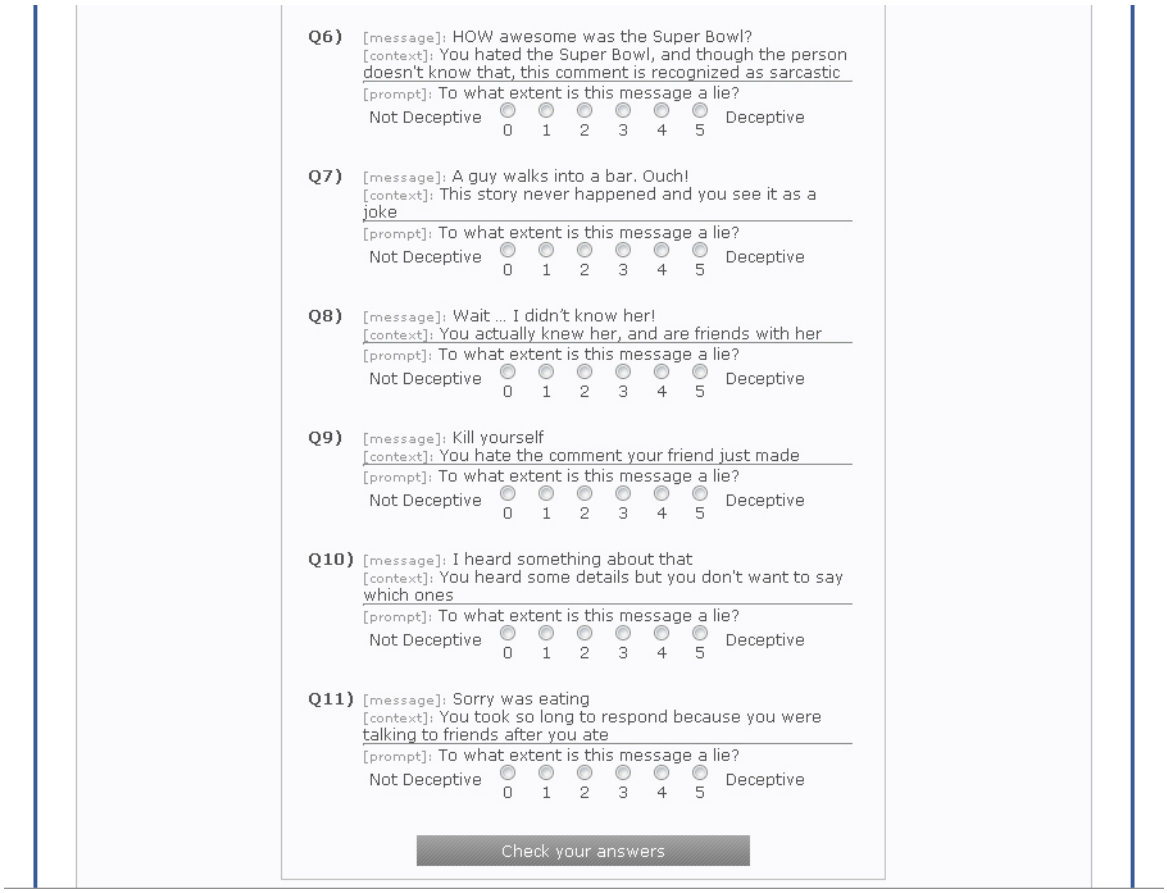

Q6) [message]: HOW awesome was the Super Bowl?
[context]: You hated the Super Bowl, and though the person doesn't know that, this comment is recognized as sarcastic
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Q7) [message]: A guy walks into a bar. Ouch!
[context]: This story never happened and you see it as a joke
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Q8) [message]: Wait ... I didn't know her!
[context]: You actually knew her, and are friends with her
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Q9) [message]: Kill yourself
[context]: You hate the comment your friend just made
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Q10) [message]: I heard something about that
[context]: You heard some details but you don't want to say which ones
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Q11) [message]: Sorry was eating
[context]: You took so long to respond because you were talking to friends after you ate
[prompt]: To what extent is this message a lie?
Not Deceptive 0 1 2 3 4 5 Deceptive

Check your answers

Appendix K

Apate Install instructions