Butler Lies: Awareness, Deception, and Design

Jeff Hancock[1,2], Jeremy Birnholtz[1,2], Natalya Bazarova[1], Jamie Guillory[1], Josh Perlin[1], Barrett Amos[1]

[1]Department of Communication
Cornell University
Ithaca, NY 14853

[2]Information Science Program
Cornell University
Ithaca, NY 14853

{jeff.hancock, jpb277, nnb8, jeg258, jcp57, bea4} @ cornell.edu

ABSTRACT

Instant messaging (IM) is a common and popular way for co-workers, friends, and family to stay in touch, but its "always-on" properties can sometimes lead people to feel overexposed or too readily available to others for conversation. This, in turn, may lead people to deceive others about their actual status or availability. In this paper, we introduce the notion of the "butler lie" to describe lies that allow for polite initiation and termination of conversations. We present results from a field study of 50 IM users, in which participants rated each of their messages at the time of sending to indicate whether or not it was deceptive. About one tenth of all IM messages were rated as lies and, of these, about one fifth were butler lies. These results suggest that butler lies are an important social practice in IM, and that existing approaches to interpersonal awareness, which focus on accurate assessment of availability, may need to take deception and other social practices into account.

ACM Categories and Subject Descriptors

H.5.3. Group and Organization Interfaces

Author Keywords

Computer-mediated communication, deception, interpersonal awareness, instant messaging

General Terms

Design, human factors, theory

INTRODUCTION

As geographically distributed work and social networks have become more common [23], the use of communication and collaboration technologies to stay connected has also increased [27, 34]. The primary purpose of such technologies is to support connections and social interactions between people. We can communicate across distances using a wide variety of technologies, from email and instant messaging to telephone calls and Facebook or Twitter updates. While we derive much value from being able to connect with others so easily, there is also evidence to suggest that many people find themselves overwhelmed and too readily accessible to others. This problem manifests itself in many ways, from a sense of email overload [43] to being unable to avoid interruptions [29, 31], and it has garnered substantial academic and popular attention.

There have been two general approaches in the CHI and CSCW literatures to dealing with this problem of managing our social interactions. The first has been a technological approach, in which designers have attempted to develop technologies that support increased interpersonal awareness about others' activities, and the degree to which they are interruptible or not (e.g., [40]). The second approach has been to emphasize design principles that consider the social practices that emerge around communication technologies, such as telling plausible, but sometimes false, stories to others about one's availability (e.g., [1, 5]). The plausibility of such stories often rests in technical and social ambiguities, such as using poor cell phone reception as an excuse for missing or abruptly ending a call.

In the present study, we examine a particular kind of social practice, deception, that recent research suggests is an important tactic for dealing with our enhanced communication availability. Users of instant messaging, for instance, report using deception in their status indicators to avoid unwanted interruptions, indicating that they are away from their computer when in fact they are not. Users also

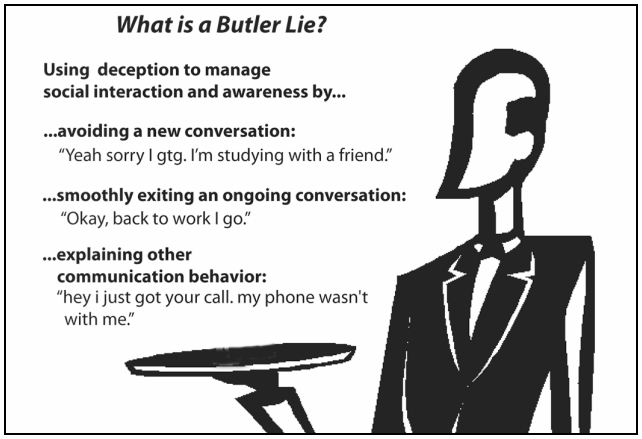

Figure 1. Definition and types of butler lies

CHI 2009, April 4–9, 2009, Boston, MA, USA.

report frequently telling their communication partners that they need to go when in fact they do not [42].

By examining deception in instant messaging, which we define as creating an instant message that deliberately seeks to create a false belief in the recipient of the message (see [20]), we seek to discover insights into how people use deception to manage their interactions, and whether these insights may prove useful for the design of future technologies. In particular, can instances of deception be used as an indicator of a breakdown in a communication system that can provide valuable lessons for design?

The present study has three objectives. First, to examine the general practice of deception in instant messaging, including the frequency of deception and the magnitude of the lies told. Second, we sought to examine how deception in instant messaging is used to manage social interactions, including avoiding interaction, taking leave of interaction, and arranging other interactions. We introduce the concept of *butler lies* to describe these kinds of deceptions that help manage our interactions. Finally, by examining the social practice of deception in managing communication, we note several observations that have key design implications for interpersonal communication technologies.

BACKGROUND

Understanding the role of deception in managing interpersonal awareness and interactions is a critical part of the larger problem of providing awareness information within geographically distributed work and social groups. Indeed, there is substantial evidence to suggest that unplanned, informal interactions are important to effective work and social relationships[28, 38], but that these can be more difficult to support in distributed environments[3, 34]. One key factor undermining these vital interactions in distributed groups is the difficulty in determining who is available to interact at any given moment, or maintaining a sense of interpersonal awareness[3, 15]. Efforts to date have not produced widely adopted technologies that overcome these problems [37].

The notion of interpersonal awareness has been the focus of significant research in the CHI and CSCW communities for two decades. As noted earlier, there have been both technical and social approaches to this issue.

Technical Approaches to Awareness

Early research in this area focused on the development of media spaces that used video to provide awareness of others' presence [15, 17]. Groups of people put video cameras in their offices, and the outputs from these cameras were displayed to all members of the group, such that others could see whether or not a particular person was: a) present, and b) seemingly available for interaction (e.g., not on the phone, asleep, etc.).

While these systems were useful in providing presence awareness information, they presented two significant issues. First, they were generally not seen by users to provide for the subtlety and nuance of initiating conversations and attracting the attention of others [4, 40]. Second, the placement of cameras in people's offices raised significant privacy concerns. Indeed, this concern reflects a fundamental tradeoff between individuals' need for privacy on the one hand [6, 10] and awareness on the other [24].

More recent research has addressed the issues raised in the media spaces work in a number of ways. The advent and rapid spread of instant messaging (IM) technologies has provided the "buddy list" as a very common means of conveying presence awareness information [39]. Buddy lists provide users with a list of their contacts (sometimes sorted into user-specified groups), and an indication of who is online and, presumably, available to chat. Most instant messaging systems also allow users to explicitly specify when they are "busy" or "away," and use recent keyboard activity (or lack thereof) to determine if a user is "idle."

Several research systems developed in the CHI and CSCW communities have provided similar awareness functionality by displaying lists or clusters of contacts, and facilitating interaction. These include systems such as Community Bar [30], Notification Collage [19], and Hubbub [26].

Moreover, these systems often assume that a single status designation applies equally to all one's contacts. As Davis and Gutwin [11] suggest, however, people are often differently available for interruption by different people. An interruption from a co-author as a deadline looms, for example, may be more welcome than one from a friend.

One theme in all of this work, as noted earlier, is the inherent tension between privacy and awareness. People have an individual need for privacy, and yet they also need to be aware of what others are doing, in order to gauge their presence and availability for interaction.

A second theme is accuracy of awareness information. There is an implicit assumption in much awareness research that people desire accurate information about others, and to convey accurate information about themselves. We believe this assumption to be problematic, as do the proponents of more social approaches to awareness described below.

Social Approaches to Awareness

We believe that one of the problems limiting progress in this area is the implicit desire in many awareness and interruption-minimization technologies to reduce ambiguity and provide accurate information about others. In other words, the underlying assumption is that people want accurate and unambiguous information about others with which they can make reliable assessments of availability and that, inversely, people are willing to provide this information about themselves [5].

There is increasing evidence, however, that such assumptions may not be valid. In communication, people frequently draw on the ambiguous properties of certain

media to maintain plausible deniability about having received messages at all [32], avoid confrontation about possible threats to self-presented identify [33], as well as to encourage or discourage rapid response to a query [1]. In their analysis of ambiguity in mediated communication, Aoki and Woodruff [1] argue that systems should provide space for users to make stories about when and why they want to interact. For instance, if they do not want to be available for a call, they could tell their interlocutor that they are in a poor reception area for cell phones and they will call back. That is, they argue that the ambiguity that emerges from communication technology should be used as a resource for design in communication systems rather than a target for elimination (see also [5]).

With regard to awareness technologies, such practices can sometimes be considered deceptive in that false information about one's availability may be provided to colleagues or friends. In the present study, we focus on instances of deception in the management of IM interactions in an effort to understand how and why users resort to deception, and whether these deceptions can provide design insights into supporting interpersonal awareness.

Butler Lies: Deception in Managing Social Interactions

Recent research suggests that deception is frequently used to create a virtual barrier between users and unwanted conversations [42]. In one focus group study, participants were asked about the norms and practices of deception and honesty in online chat and instant messaging. The responses revealed that while adolescents (ages 11-13) reported honest behavior as the norm, older participants (e.g., ages 15–23) reported frequently using deception to avoid interactions (e.g., by choosing status that indicates 'away' when in fact they were not) and by extricating themselves from unwanted or uninteresting conversations (e.g., saying that they need to go to dinner when in fact they do not).

We refer to these types of lies, in which deception is used to manage the entry and exit of social interactions, as *butler lies* (see Fig 1). Butler lies include the strategies of using deception to "avoid interaction" and to "take leave of interaction" described by Camden [9]. We use the term "butler lies" to allude to the social buffering function that butlers provided for their employers. Consider, for example, the buffering function described in the following quote from Roberts' [35] (p. 79) *The House Servant's Directory:*

> In the next place, you should never admit any person or persons into the parlour or drawing room, without first announcing their names to your mistress or master. This you can readily find out by saying, "What name shall I say, ma'am?" or "sir?" Therefore by this way, you will find out whether your employers wish to see them or not. If not, tell them your mistress, master or whoever they wish to see, are engaged, &c. in a polite and civil manner.

As this example illustrates, butlers functioned as a barrier between their masters and persons desiring an interaction. The excuses or justifications they provided if their employer did not wish to meet the persons in the parlour (i.e., that they are "engaged, &c.") provided the social barrier for the master or mistress.

The example also highlights the crucial role of civility and politeness in managing social interactions. Indeed, the influential self-presentation framework of deception [14] emphasizes the role of politeness concerns as a motivation for many of our everyday deceptions (see also [16]). People seek to maintain a sense of "face," a positive image of ourselves and sense of our own autonomy. We also seek to maintain our partner's "face", their own positive self-image, and to not infringe on their autonomy. As Brown and Levinson [7] point out, people use different language strategies to avoid threatening one's own or another's face. Deception is one language strategy (an "off-the-record" strategy in the Brown and Levinson terminology) that we use when committing a face-threatening act.

Butler lies are one type of strategy for dealing with the face-threatening acts associated with entering or exiting social interactions. Butler lies about avoiding an interaction, or lies related to leaving a conversation that the partner wants to continue, are designed to maintain our own face (not coming across as mean or haughty) as well as our partner's (that we respect and like them).

Given that technologies like instant messaging increase the potential for initiating and engaging in conversations with others while at the same time reduce interpersonal awareness, we should see butler lies used frequently as a social practice for managing interactions in IM.

The Present Study

In the present study we examine the practice of lying in instant messaging, paying particular attention to a specific type of deception, butler lies. As noted, previous research based on focus-group results suggests that participants believe that deception is relatively common in instant messaging [42]. In a diary-based study, in which participants were asked to indicate whether a conversation contained a lie or not, IM contained fewer lies than the telephone but more lies than email. Lying rates in IM were approximately the same as those observed for face-to-face conversations [21]. Approximately one-fifth of instant messaging conversations involved a lie.

We developed software that allowed us to examine the practice of deception at the message level, allowing for a more fine-grained analysis. Participants used our software to rate each of their instant messages at the time of sending for the presence of deception. In particular, they were asked to rate whether or not the message was deceptive, and if it was, to rate the magnitude of the deception on a scale from 1 (a little deceptive) to 5 (very deceptive).

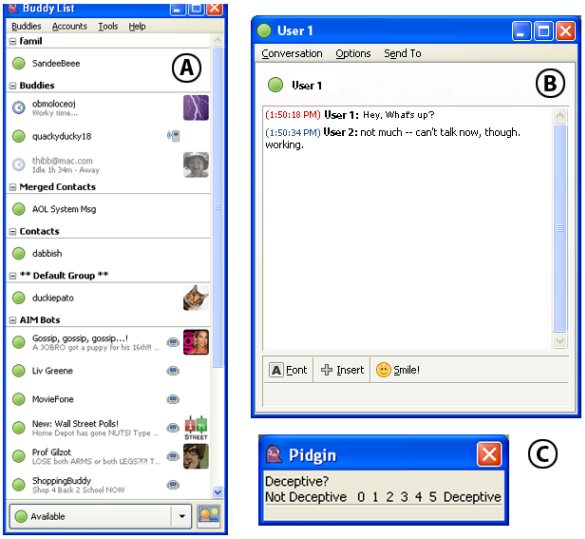

Figure 2. The Apate IM Client: (A) illustrates the buddy list, (B) illustrates the IM messaging window, and (C) illustrates the window that popped up as soon as the participant hit enter to send their message. Once the participant rated the message by typing the appropriate number the pop up window would disappear and the message would be sent.

Our first research question was:

RQ 1: How frequently does deception take place in IM and what is the magnitude of the deceptions?

We were most interested in a particular form of instant messaging deception, butler lies. We coded and analyzed all the lies identified in the instant messages to determine whether they were a butler lie, and if so, whether the butler lie was designed to manage the entry or exit of the current conversation, or to manage another social interaction (e.g., a meeting in the future). By connecting these data to participant characteristics, we were also able to examine whether age, gender, or IM experience play a role in butler lie production. Our second research question was

RQ 2: How are butler lies used in instant messaging? How frequently are they observed, and how are they used to manage the entry and exit of social interactions, or to arrange other social interactions? What kind of people use butler lies the most?

Finally, we were interested in determining whether butler lies could be used as a source of insight for improving the design of interpersonal awareness systems. Deception may be a valuable indicator of breakdown in awareness/ communication systems that can provide valuable lessons for design. Understanding how and why people use deception in managing their communication can help us to understand flaws in existing awareness technologies and improve them. Moreover, this may also help us to identify the conditions under which deception is most likely to occur, and this knowledge can also be used to improve awareness tools. Our final research question asks:

RQ 3: Can butler lies be used as a resource for improving the design of interpersonal communication systems?

METHODS

Participants. A total of 50 participants took part in the study. Two were excluded for failing to follow instructions, and 5 were excluded because they recorded fewer than 20 messages, our minimum criterion for inclusion. The final sample comprised 43 participants consisting of 18 male and 24 female undergraduates (one person failed to identify their gender) that received course credit at a large northeastern university. Participants ranged from age 18 to 29 (average = 20.5) and had 1 to 15 years of experience with instant messaging (average = 8.6 years).

The Apate Research System. The Apate Research System consists of two parts, an online, web-based *Apate Experiment Manager* (see *apate.net)* and the *Apate Instant Messaging (IM) Client*, a modified Pidgin client. Combined, the two allow researchers to examine the social dynamics of instant messaging by requiring users to make a rating every time they send a message (for instance, whether or not the message was deceptive, how much they liked their partner at that moment, etc), a level of granularity not reached in prior online deception research.

The *Apate Experiment Manager* is a PHP and MySQL-based system that allows experimenters to set up the parameters of an experiment, from handling participant recruiting, scheduling, and reminders, to setting up surveys related to the study, and finally, controlling how the instant messaging client will function (e.g., what rating system will be used, the anchors of the scale, etc.).

This tool automates all tasks related to running an experiment using IM, including the experimental procedures, instructions, surveys, emails, participant information, data collection, and basic data analysis tools. Pre-experiment and post-experiment surveys were administered through the online web interface run by the Apate Experiment Management System, and it also sent reminder emails and instructions to participants.

The *Apate IM Client* is an open source instant messaging client based on Pidgin. The program was modified so that a window pops up after the participant sends a message, requiring the participant to rate the message on a numeric scale. Each message is then sent to its intended recipient and a copy of the conversation is sent to the *Apate.net* server, including message content, message ratings, screen names, timestamps, and keystroke data. This entire process is completely transparent to the conversation partner.

To address privacy concerns, all data pertaining to partners were anonymized, including hashing all screennames and automatically removing instances of the 5000 most common proper first and last names.

Participants needed to download and install the Apate IM Client on their machine. After the experiment was

completed, the program could be removed and all records deleted from the participant's computer.

Procedure

Participants that volunteered for the project were sent an initial email with a link to a web-based consent form. After giving consent, participants received instructions about the study. Second, they completed a questionnaire asking for age, gender, major, how many years and how often they use IM, the size of their buddy list, and who they interact with most often in IM. Third, they read instructions for rating instant messages as deceptive or not deceptive. Instructions detailed the definition of deception (i.e., any statement that creates a false belief in the receiver that the sender knows to be untrue) and distinguished deception from humor (e.g., "His head is the size of the moon" is not literally true, but is not deceptive as it does not create a false belief). Participants then completed a 10-item quiz testing their ability to distinguish between truthful messages, deceptive messages, and humorous messages. The system provided feedback on items rated incorrectly.

Once the quiz was completed they could download the *Apate IM Client*. Participants were asked to use this client as their regular instant messaging interface. As shown in Figure 2, the Apate IM client required participants to rate each message on a scale from 0 (not deceptive) to 1 (a little deceptive) to 5 (very deceptive). The message would not be sent until it had been rated.

Participants were asked to use the client for four days, after which the server sent an email with a link to completing the post-experiment survey (through the Apate Experiment Manager) and finishing the experiment. The post-experiment survey included the following items: whether their IM use was fairly similar to their average IM use; whether they used other IM clients besides the one provided to chat during this study; how frustrated they were with the instant messenger provided; to identify who they talked to the most over the course of their conversations – local friends, friends at a distance, family, colleagues, a spouse/boyfriend/girlfriend.

Message Rating	# of Msgs	% of Msgs	# of Lies	% of Lies
Not Deceptive	6234	89.1%	n/a	n/a
1	213	3.1%	201	29.3%
2	185	2.6%	165	24.1%
3	146	2.1%	130	19%
4	77	1.1%	69	10.1%
5	141	2.0%	120	17.5%

Table 1. Breakdown of messages by participant rating and lies coded as jocular.

Upon completion, the server stopped collecting information from the participant, and the participant received a thank you and debriefing email, including instructions on how to remove the Apate IM Client from their system.

Message Coding

The messages rated as lies by the participants were coded in three phases: 1) for jocularity, 2) for butler lies, 3) whether the butler involved the entry or exit of the current interaction, or arranging or discussion of other communicative behavior. Each phase of coding followed a hierarchically arranged coding scheme using if->then decision trees, which is available from the authors.

The first phase involved coding for jocularity because self-rated lies sometimes involved jokes, sarcasm or irony that were clearly not deceptive. A lie was coded as jocular if the message was clearly not intended to create a false belief in the partner, a characteristic required by our definition of deception. Some examples of messages self-rated as a lie, which were subsequently coded as jocular or not deceptive, include: "because time travel hasn't been invented yet," "if you want to give him your first born, then yeah, do that", and "you're a dirty bum". Inter-rater reliability for the jocularity coding was high (Kappa = .93). These messages were excluded from further analysis.

In the second phase, the remaining lies were coded as butler lies if the deception was about managing the communication. Specifically, butler lies related to entering or exiting a conversation, or to arranging or discussing another communication event, such as planning a future engagement or explaining why a past event occurred (e.g., failing to answer a phone call). Inter-rater reliability was acceptable (Kappa = .79).

Finally, butler lies were categorized into three types: 1) as entering the conversation (e.g., "hey i just got your missed call"), 2) exiting the conversation (e.g., "okay last question and then I have to get to class"), or 3) arranging or discussing other social interactions (e.g., "nice! well let me know and we can meet up hopefully"). Inter-rater reliability was high (Kappa = .98).

RESULTS

General Deception Patterns in IM

The participants produced a total of 6996 messages. Of these messages, participants rated 762 as lies and the remainder as not involving deception. Seventy-seven of the self-rated lies were coded as jocular and were excluded from further analyses of lies. Thus, the total number of lies was 685, which represents 9.79% of all IM messages.

The breakdown of the ratings for all messages is described in Table 1. In terms of lie magnitude, 53.4% of all lies were rated as "small deception" (1 and 2 on the scale of 1 to 5), while 27.6% were rated as "large deception" (4 and 5 on the deception magnitude scale). The average magnitude of lies told in IM was close to the mid-point (2.62 on the 1 to 5

scale). Taken together with the range of magnitude for lies described in Table 1, the data suggest that the lies we observed were frequent and relatively small in magnitude.

Butler Lies – Frequency, Type and Magnitude

Of the 685 lies 132 (19.3%) were identified as butler lies, suggesting that about one-fifth of the lies told in IM are butler lies. Stated in terms of the total number of messages produced by the participants, butler lies accounted for almost 2 out of every 100 IM messages, or 1.89%.

The majority of butler lies involved exiting interactions (40.9%) and arranging other social interactions (37.9%). A chi-square test, $\chi^2(2, N = 132) = 8.91$, $p = 0.01$, revealed that these exit and arranging butler lie types were produced significantly more often than butler lies that managed the entry of social interactions (21.2%).

One important question about the nature of butler lies was how participants perceived the lies in terms of their magnitude. Did participants view them as little "white lies" as suggested by [42]? Or did participants view them as more serious lies? To examine this question we compared the perceived magnitude of butler lies with other types of lies. If they are small white lies, then they should be rated as less important than other lies. If they are more important lies, they should be rated as higher in magnitude.

A hierarchical regression, with lies nested within participants, showed that there was a significant difference in perceived lie magnitude between butler lies and other types of lies, F (1, 657)=8.21, p = .004. Butler lies on average were ranked higher in magnitude (M=2.82, SE = .16) than other types of lies (M=2.47, SE=.13). This suggests that participants considered butler lies more serious than other types of lies. Interestingly, this effect held for each of the three types of butler lies. The comparison of lie magnitude by type of butler lie (entry of social interaction, exit of social interaction, and to arrange other interactions) did not reveal a significant difference in lie magnitude between them, F (2, 122)=1.19, p=.31.

Butler Lies – Gender, Age and IM Experience.

Another important question is whether characteristics of the user, such as gender, IM usage patterns, and experience, affected how butler lies were used. To examine this question we analyzed the occurrence of butler lies in relation to the participant characteristics derived from the pre-experimental and post-experimental questionnaires, including age, gender, number of years using IM, size of buddy list, and IM usage. The frequency of butler lies was modeled with a negative binomial distribution, which accounts for non-normality and over-dispersion in the count data. Regression coefficients were adjusted for differences in the total number of messages each participant produced.

The results demonstrated that butler lies were used ubiquitously, with no significant difference in butler lie occurrences across age, F (1,39)=.05, p=.83; gender, F (1,40)=3.47, p=.07; number of years using IM, F (1,38)=.54, p=.47; a size of the buddy list, F (1,39)=2.37, p=.13; how often participants use IM, F (6,35)=.82, p=.56, who they generally talk most often to in IM, F (3,38)=.83, p=.49; and who they talked to most often during the experimental procedure, F (3, 37)=.61, p=.61.

We did detect one potential difference across gender. If participant IM experience (i.e., the number of years using IM and buddy list size) was held constant, a marginal gender effect emerged, F (1, 35) = 4.0, p =.05, suggesting a trend for males to use more butler lies than females.

Butler Lies in Detail

To better understand the nature of butler lies in everyday usage, we looked more closely at the IM transcripts.

Smooth exits

First, as noted above, most butler lies were in the "Exiting a Conversation" category. In other words, people frequently lied in order to bring the conversation to a close.

Many participants rated statements of 'BRB' ("be right back") as lies. Presumably this meant that they were saying to their interaction partner that they would, in fact, be right back, but in reality had no intent to return to the conversation. Indeed of the 22 total messages in the entire data set containing "brb," 4 are butler lies. The self-reported magnitude of these lies were 2, 3, 4 and 5, respectively.

Similarly, some participants used more specific statements implying that they were about to or had to do something else, even though they did not. This builds on the informal and transient nature of IM conversations in that the mention of another activity (that is likely to be perceived as legitimate by the interaction partner) can serve as reason to bring a conversation to a close. These lies often related to school or other work:

> "Anyway, I have to go so I can write a paper. I'll ttly" (5)
>
> "okay, back to work I go." (4)
>
> "I have a prelim tomo tho so I gotta study for ittt and then head to a meeting" (3)
>
> "alright work time here too" (2)

Note that in these lies and those presented below, the magnitude score is listed in parentheses.

Lies about specific activities were sometimes temporal in nature. In other words, they were butler lies that were plausible because of the time of day when they occurred, referring to activities such as sleep or meals:

> "alright im gonna go get some dinner..i'll speak to you later" (3)
>
> "im gonna go grab some lunch now i guess" (4)
>
> "and exhuasted...going to sleep in a bit" (2)
>
> "sleep! Time"(2)

Interestingly, one participant used a butler lie describing ongoing activity (as opposed to planned future activity):

"yeah sorry i gtg, im studying with a friend" (5)

Some participants also used generic expressions indicating that they had to leave, but did not name a specific activity, though these were far less common:

"not much..gotta go later" (3)

"aight man gotta run enjoy your last few weeks" (1)

One interesting aspect of comparing the examples discussed thus far is that the butler lies with the highest magnitude ratings (i.e., all of those rated 4 or 5) tended to describe specific activities. To be sure, some of the lies involving specific activities were also rated lower, but it is worth noting that none of the non-specific butler lies were rated higher than 3. This seems to indicate that people felt lying about specific events was worse or more salient than the generic lies, even though specific lies were more common.

Lying About Promised Communication

Another way that butler lies were used to exit conversations was to lie about implied future interaction between the speaker and partner. For example, one participant rated "talk to you laterrr" a '4' in magnitude, suggesting that he or she did not actually intend to talk to their partner later.

While the extent to which ritualized farewells such as "ttyl" actually imply future interaction is debatable (see [25]), we also observed more specific butler lies about implied future communication, such as the following:

"thanksss okay see u i a couple weeks then a:)" (4)

"and i will email u as well" (5)

These lies imply that participants sometimes ended conversations with the implication that they would talk to their partner again, even if they did not intend or wish to do so. While scholars have argued that these are ostensible speech acts [25], in which neither participant believes the statements to be sincere, our participants judged these statements to be deceptive.

Lying to Excuse or Justify Communicative Behavior

Often participants used deception as an apparent way to excuse past or future communication behavior without disclosing the true reason for that behavior. In these three examples, participants were lying in making excuses for missing particular activities or events. Note in particular that each of them is rated as 5, indicating that the participants felt these were serious lies.

"i forgot about it completely" (5)

"i have lots of work, so i can't" (5)

"i had a conference all week and got home late last night and passed out" (5)

There were also occasions where the behavior being excused was non-participation in the current conversation:

"sorry i walked off" (2)

What is interesting here is that this is one of the few cases where a butler lie is being used to smooth the starting (or, in this case, restarting) of a conversation. Here, the lie serves the purpose of accounting for a long pause in the conversation, and makes for a natural restarting point.

Apate IM Client Evaluation

Given that participants needed to rate every message before sending it, we wanted to assess a) how similar their IM usage was compared to before the study and b) how frustrating they found using the system. The majority of the participants reported that their IM use during the experiment was relatively similar to their typical IM use, M=3.81 (on a 1 to 5 scale), SD=1.17. Most participants interacted with multiple parties during the experiment, including friends, family, girlfriend/boyfriend, and co-workers. Seventy-two percent of the participants conversed most frequently with their friends, both local and at a distance. Thirteen participants (31% of the total number of participants) reported using other IM clients besides the one provided to chat during the experimental study.

With respect to how frustrating the Apate IM Client was, on average, participants reported a medium level of frustration (M=2.55, SD=1.19, measured on a 1 to 5-point scale).

DISCUSSION

In the present study we examined deception in instant messaging, specifically focusing on how deception can be used for managing social interaction in IM. Our second objective is to determine how this type of deception, which we call *butler lies*, can provide insights for designers of systems to support interpersonal awareness.

We developed a novel research tool, the *Apate System*, that allowed us to track deceptions in IM at the message level by allowing participants to identify whether each sent message contained deception or not. This approach has several advantages over previous methods, including 1) improved granularity compared to diary-based approaches (e.g.,[21]), which only assess whether an entire conversation included a deception or not, and 2) the ability to track deception in participants' day-to-day IM behavior with their typical communication partners rather than performing specific tasks with strangers in laboratory experiments (e.g., [8, 22]).

Indeed, examining the practice of deception in everyday IM allowed us to confirm several aspects of deception observed in previous studies of everyday deception in face-to-face contexts (e.g., [12]) and in other online context, such as online dating profiles [41]. As with face-to-face lying, the general practice of deception in IM was frequent, with approximately one out of every ten messages involving deception. These everyday deceptions varied widely in their nature and magnitude. Consider, for example, the following

conversation, in which the conversation partner accuses the participant of being at a bar with someone else:

A: we need to talk
B: about hwat? (0)
A: i saw you with another guy last night
B: n you didnt! (5)
A: yeah u were with [name]
B: no i wasnt (5)
A: i saw yuo at [name of a local bar]
B: i dint go to [name of bar] last night (5)

As this example illustrates, the lies reported by our participants were not the typical made-up lies observed in laboratory studies, but represented the rich diversity of how deception is really used in everyday life.

We were specifically interested in how deception is used for managing social interactions in IM. As we expected, these butler lies were frequently observed, accounting for almost twenty percent of all lies told in IM. The ubiquity of these lies is highlighted by the finding that there were no differences in butler lie production across age or IM experience, suggesting that all of our participants used butler lies frequently. Although the results suggest differences for using butler lies across gender, with no *a priori* prediction for the gender effect, we leave it to future research to explore gender differences in coordinating social interactions in IM.

The least observed type of butler lie was those related to entering conversations. Why weren't entry butler lies observed more often? One possible reason is that these deceptions are handled outside of the actual conversation. Vandeen Abeele and Roe [42] note that users report frequently using their IM status to block interruptions from unwanted conversations by indicating that they are away from their computer when in fact they are not. This may explain why the entry butler lies were the least observed in the conversational content of IM.

The two most frequently observed types of butler lies were exiting interactions and arranging or discussing other communications. Butler lies concerned with exiting IM conversations confirm our expectation that people use social practices, such as deception, to regulate and manage the types of social interactions made possible by new technologies. Our data are also consistent with observations by Schegloff and Sachs [36] about the practices of closing conversations, in which interlocutors frequently indicate future meetings that most likely will not take place.

The fact that butler lies were frequently used to manage other social interactions, including falsely arranging future interactions or justifying past communication behavior, is consistent with previous research identifying the use of instant messaging for coordinating other communication activities [32]. Consider, for example, the following butler lie, in which a participant makes an excuse for why they did not take a friend's call earlier:

hey i just got your missed call (2)
my phone wasn't with me (2)

These types of butler lies provide further evidence that users develop social practices around the use of communication technologies, supporting social approaches to awareness (e.g., [1,5]). These views argue that users draw on the ambiguities inherent in communication technology to create stories that explain their actions [1]. Our study suggests that deception can be an important tool in crafting these stories.

Implications for Design/Practice

It is clear from our results that butler lies are regularly used, at least by some people, as a strategy in managing interpersonal awareness and social interaction. This is particularly true for exiting conversations, and raises questions about the ways in which we provide awareness information, and adds another dimension to the existing tension between privacy and awareness [24]. As noted above, approaches such as those by Fogarty, et al. [18] and Begole [2] attempted to reduce privacy concerns by providing estimates of availability based on aggregations of many data sources, rather than sharing the details of these individual data points. What such an approach does not account for, however, is the desirability of ambiguity. Our results suggest that even an aggregate estimate of availability could be problematic if it conflicts with a false "brb" or "got to go do work" statement.

This, in turn, raises several implications for designers. One is that systems should support graceful exit from IM conversations. One simple way to do this would be to allow for the display of different status information for different individuals. As one wraps up a conversation with an interaction partner, for example, a "busy" status could be displayed only to that partner, but not necessarily to all others. This would support the butler lies people already tell, but also allow for easier availability to other contacts.

A second implication is that systems could allow for conversations to gradually fade out over time or as participants appear to lose interest. This is in contrast to current designs where conversational persistence is the norm. It is also the reverse of the "virtual approach" metaphor called for in prior research [4, 40].

This idea could be implemented in several ways, all of which would require additional research and testing. One simple approach would be to give IM windows a finite lifespan, and force re-opening of windows to continue an "expired" conversation. Such an approach, while possibly annoying for desirable long conversations, could also force people to consider whether or not a conversation is worth continuing, a topic that current interfaces do not consider.

A second approach would be to use indicators of conversational activity and both parties' level of interest in the conversation (i.e., frequency and rate of utterances, how often the IM window is in primary focus, linguistic cues in the chat itself, etc.) to either: a) take a more context-sensitive route in an "expiring windows" approach, or b) visualize interest in the conversation, so that participants can gauge whether or not it is appropriate to continue.

To be sure, such approaches are fraught with potential problems of social nuance, honesty, privacy and getting users to understand them in the first place. At the same time, however, our results make clear that current technologies do not appear to effectively support exiting conversations. Supporting these activities more effectively will require some exploration of this design space.

A second interesting implication from our findings is how few of the butler lies we observed occurred at the start of an interaction. Indeed, the overwhelming majority occurred at the end of conversation, when the participants needed a clean/smooth exit. This raises the question of whether some participants would have preferred to avoid these conversations in the first place, but had no way to avoid the interruption. Current IM clients provide some means for avoiding interruption, but this may also merit additional consideration.

Limitations

The method we used to examine deception in IM involves several important limitations. First, by asking participants to report on their own lies we may have changed their lying behavior. This problem is inherent to most deception research, although many previous studies have gainfully used similar types of self-report methods [9, 13, 21].

A related limitation is that self-reports on lying may be problematic given that the method requires asking people to be truthful about their own lying behavior. Although this is also an inherent limitation to self-reports, some previous research has demonstrated that people are typically capable of reporting honestly about their deceptive behavior [41]. Additionally, the question posed to participants addressing the magnitude of deception should have addressed not only how untrue a statement was, but also how serious it was.

Lastly, requiring participants to make a rating on each of their messages may have been burdensome and frustrating. Indeed, participants reported the Apate IM Client to be relatively frustrating but not excessively so (average of 2.55 on a 5-point scale). Additional evaluation work of the Apate IM Client, however, is needed.

CONCLUSIONS

The results from this study advance our understanding of how IM users employ deception in managing their social interactions. Earlier work on the entry and exit of IM conversations had not examined the intentional lies that participants use to extract themselves from conversation or to explain why previous entries failed. Nardi et al [36], for example, discuss the value of ambiguity in IM because it provides plausible deniability and the ability to ignore incoming messages. When they discuss the messages themselves, however, there is no mention of truthfulness or intent to deceive others. Similarly, Tang's [40] careful analysis of the management of interaction in various communication systems notes the possibility that messages will be ignored and the difficulty of leave-taking in some online interaction, but there was no evidence or consideration of intentional deception reported.

By introducing the concept of butler lies, the present study is the first to highlight the importance of deception as a social practice in entering and exiting current interactions, and for arranging or justifying other communicative behavior. While these results support social approaches to the general issue of interpersonal awareness, they also serve as a resource for developing design principles when creating novel technical solutions to problems associated with interpersonal awareness in mediated communication.

ACKNOWLEDGMENTS

We thank our participants, reviewers, and the National Science Foundation (IIS #0843215) for support.

REFERENCES

1. Aoki, P. and Woodruff, A., Making space for stories: ambiguity in the design of personal communication systems. In *ACM CHI* (2005), 181-190.
2. Begole, J., Matsakis, N. and Tang, J., Lilsys: Inferring Unavailability Using Sensors. In *ACM CSCW* (2004), 511-514.
3. Bellotti, V. and Bly, S., Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team. In *ACM CSCW* (1996), 209-218.
4. Birnholtz, J., Gutwin, C., Ramos, G. and Watson, M., OpenMessenger: Gradual Initiation of Interaction for Distributed Workgroups. In *ACM CHI*(2008), 1661-1664.
5. Boehner, K. and Hancock, J., Advancing Ambiguity. In *ACM CHI*(2006), 103-107.
6. Boyle, M. and Greenberg, S. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM Transactions on Computer-Human Interaction*, *12,* 2 (2005). 328-370.
7. Brown, P. and Levinson, S. *Politeness: Some universals in language usage*. Cambridge U. Press, Cambridge, 1987.
8. Burgoon, J., Stoner, G.M., Bonito, J.A. and Dunbar, N.E., Trust and deception in mediated communication. In *HICSS*(2003), 44.
9. Camden, C., Motley, M.C. and Wilson, A. White lies in interpersonal communication: a taxonomy and preliminary investigation of social motivations. *Western Journal of Speech Communication*, *48,* (1984). 309-325.

10. Clement, A. Considering privacy in the development of multi-media communications. *CSCW*, *2,* (1994). 67-88.
11. Davis, S. and Gutwin, C., Using Relationship to Control Disclosure in Awareness Servers. In *Graphics Interface '05*(2005), 75-84.
12. DePaulo, B.M., Kashy, D.A., Kirkendol, S.E., Wyer, M.M. and Epstein, J.A. Lying in everyday life. *Journal of Personality and Social Psychology*, *70,* (1996). 979-995.
13. DePaulo, B.M., Kashy, D.A., Kirkendol, S.E. and Wyer, M.W. Lying in everyday life. *Journal of Personality and Social Psychology 70,* (1996). 979-995.
14. DePaulo, B.M., Lindsay, J.J., Malone, B.E., Muhlenbruck, Charlton, K. and Cooper, H. Cues to deception. *Psychological Bulletin*, *129,* (2003). 74-118.
15. Dourish, P. and Bly, S., Portholes: Supporting awareness in a distributed work group. In *ACM CHI*(1992), 541-547.
16. Feldman, R.S., J.A., F. and Happ, B.R. Self-presentation and verbal deception: do self-presenters lie more? *Basic and Applied Social Psychology*, *24,* (2002). 163-170.
17. Fish, R., Kraut, R. and Root, R., Evaluating Video as a Technology for Informal Communication. In *ACM CHI*, (1992), 37-48.
18. Fogarty, J., Hudson, S.E., Atkeson, C.G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J.C. and Yang, J. Predicting Human Interruptibility with Sensors. *12,* 1 (2005). 119-146.
19. Greenberg, S. and Rounding, M., The notification collage: posting information to public and personal displays. In *ACM CHI*(2001), 514-521.
20. Hancock, J. Digital deception: when, where and how people lie online. in Harrington, B. ed. *Deception: methods, contexts and consequences (Santa Fe Institute)*, Stanford Press, Palo Alto, CA, 2007.
21. Hancock, J., Thom-Santelli, J. and Ritchie, T., Deception and design: The impact of communication technologies on lying behavior. In *ACM CHI* (2004), 130-136.
22. Hancock, J.T., Curry, L., Goorha, S. and Woodworth, M.T. On lying and being lied to: a linguistic analysis of deception. *Discourse Processes*, *45,* (2008). 1-23.
23. Hinds, P. and Kiesler, S. *Distributed Work*. MIT Press, Cambridge, MA, 2002.
24. Hudson, S.E. and Smith, I., Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *ACM CSCW*, (1996), 248-257.
25. Isaacs, E. and Clark, H.H. Ostensible invitations. *Language in Society*, *19,* (1990). 493-509.
26. Isaacs, E., Walendowski, A. and Ranganathan, D., Hubbub: A Sound-enhanced Mobile Instant Messenger that Supports Awareness and Opportunistic Interactions. In *ACM CHI* (2002), 179-186.
27. Kraut, R., Brynin, M. and Kiesler, S. *Computers, Phones and the Internet: Domesticating Information Technology*. Oxford U. Press, Oxford, UK, 2006.
28. Kraut, R., Egido, C. and Galegher, J., Patterns of Contact and Communication in Scientific Research Collaboration. In *ACM CSCW*(1988), 1-12.
29. Mark, G., Gudith, D. and Klocke, U. The cost of interrupted work: more speed and stress. In *ACM CHI* (2008). 107-110.
30. McEwan, G. and Greenberg, S., Supporting social worlds with the community bar. In *ACM GROUP* (2005), 21-30.
31. McFarlane, D.C. and Latorella, K.A. The scope and importance of human interruption in human-computer interaction design. *Human Computer Interaction*, *17,* 1 (2002). 1-61.
32. Nardi, B., Whittaker, S. and Bradner, E., Interaction and outeraction: Instant Messaging in action. In *ACM CSCW* (2000), 79-88.
33. O'Sullivan, P.B. What you don't know won't hurt me: impression management functions of communication channels in relationships. *Human Communication Research*, *26,* 3 (2000). 403-431.
34. Olson, G.M. and Olson, J.S. Distance matters. *Human-Computer Interaction*, *15,* (2001). 139-179.
35. Roberts, R. *The House Servant's Directory*. Charles S. Francis, New York, 1827.
36. Schegloff, E.A. and Sachs, H. Opening up closings. *Semiotica*, *8,* 4 (1973). 290-327.
37. Schmidt, K. The problem with 'awareness'. *CSCW*, *11,* (2002). 285-286.
38. Shiu, E. and Lenhart, A. How Americans use instant messaging, Pew Internet & American Life Project, Washington DC, 2004.
39. Smale, S. and Greenberg, S., Broadcasting information via display names in instant messaging. In *ACM GROUP*(2005), 89-98.
40. Tang, J. Approaching and leave-Taking: Negotiating Contact in Computer-Mediated Communication. *ACM TOCHI*, *14,* 1 (2007). 1-26.
41. Toma, C., Hancock, J.T. and Ellison, N. Separating fact from fiction: an examination of deceptive self-presentation in online dating profiles. *Personality and Social Psychology Bulletin*, *34,* (2008). 1023-1036.
42. Vanden Abeele, M. and Roe, K. White cyberlies: The use of deceptive instant messaging statuses as a social norm Paper presented at the Conference of the *International Communication Association*, Montreal, Canada, 2008.
43. Whittaker, S. and Sidner, C., Email overload: exploring personal information management. In *ACM CHI*(1996), 276-283.