

Butler Lies from Both Sides: Actions and Perceptions of Unavailability Management in Texting

Lindsay Reynolds¹, Madeline E. Smith³, Jeremy Birnholtz³, Jeff Hancock^{1,2}

¹Department of Communication

²Department of Information Science

Cornell University

Ithaca, NY 14853

llr448@cornell.edu, jth34@cornell.edu

³Department of Communication Studies

Northwestern University

Evanston, IL 60208

jeremyb@northwestern.edu,

madsesmith@u.northwestern.edu

ABSTRACT

In an always-connected world, managing one's unavailability for interaction with others can be as important and difficult as coordinating mutual availability. Prior studies have identified the butler lie, a linguistic strategy commonly used to manage unavailability, and examined message-level data to examine how message senders' use of butler lies varies across media and situations. This study is the first to examine how butler lies are perceived by those who receive them. Pairs of student participants provided messages sent to each other in real conversations and indicated whether these messages were deceptive or not. These messages were then passed to the partner, who indicated perceived deception and provided an explanation. Results suggest that participants expect butler lies regularly although not as often as they are actually produced, and participants are not very accurate in identifying butler lies. Moreover, detailed analysis of messages and explanations suggests that butler lies play a relational role that is expected by both parties in a dialog.

AUTHOR KEYWORDS

Availability management; butler lies; SMS; texting

ACM Classification Keywords

H.5.3. Group and Organization Interfaces

INTRODUCTION

A significant success of today's communication environment is that social connectivity – the ability to communicate with others – is virtually constant [1,6,21]. Even as the always-on world facilitates easy interaction and novel modes of coordination, however, people increasingly report being overwhelmed or distracted by interaction opportunities [29]. In a recent Pew study, for example, 29% of participants said they had recently turned off their mobile device simply to get a break from using it [26]. Moreover, Leonardi et al. [15] observed telecommuters strategically

using communication technology to avoid rather than enable interaction, thus increasing perceived distance between them and their co-workers.

These problems reflect a fundamental shift in availability management, a phrase we use to refer to activities and social processes related to initiating, concluding or coordinating social interactions. Historically, people were not co-present (either physically or in a mediated environment) most of the time, so availability management was largely a function of coordinating co-presence (either face-to-face or in a particular medium): establishing a time and place for interaction and/or initiating an interaction once co-presence was established [28].

In an always-on world where constant connectivity and virtual co-presence are assumed, however, coordinating co-presence is no longer the key challenge in availability management. Rather, people today frequently seek to avoid communication when it is technically possible, but conversation would be socially awkward, inappropriate or disruptive [21,33]. In these cases, co-presence does not suffice to indicate availability, and additional effort is needed to explain why – despite mediated co-presence and seeming availability – interaction cannot take place now or did not take place in the past [1,30].

Prior work has identified the butler lie, a common linguistic strategy for availability management [12]. Butler lies draw on ambiguity inherent both in communication media (e.g., about one's location or current activity) and in the social conventions for their use (e.g., appropriate response times, locations for conversation). These studies, however, have primarily analyzed messages from the sender's perspective. This leaves open questions about how often these messages occur relative to other types of lies, how they are perceived by others, and their emotional and relational impact.

In the paper that follows, we present an examination of butler lies in SMS text messaging in a student population at a US university. The same messages are analyzed from the perspective of both senders and receivers. We examine the rate of butler lies relative to other types of lies, assess how accurately receivers detect deception in butler lies, and examine the relational function and impact of messages.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA.

Copyright 2013 ACM 978-1-4503-1331-5/13/02...\$15.00.

BACKGROUND

Deception in Availability Management

While not previously labeled as availability management, this concept has been preliminarily explored in recent literature. Aoki and Woodruff [1] argue that technologies should provide space for users to make stories about when and why they want to interact. If, for example, a person does not want to be available for a call, they could tell their interlocutor that they are in a poor phone reception area.

A key insight from Aoki and Woodruff's paper was that people sometimes use deception as a strategy for dealing with the always-on nature of modern communication (see also [4]). Salovaara et al. [24] further found that mobile device users felt a need to explain unavailability, but did not feel these explanations had to be truthful. Specifically, deception has been defined as the intentional misleading of another person [16]. To be clear, we use "lying" to refer to a subcategory of deception that involves deceptive statements by one person, that intentionally create a false belief in another person.

Following Aoki and Woodruff's approach for examining how lying can be used with technology to facilitate coordination, several studies have investigated a particular linguistic strategy for availability management, called the butler lie, which people use to avoid social interaction or to account for a failure to communicate. Butler lies were first documented in a study that asked participants to identify the lies they told in instant messaging (IM). Many butler lies involved coordinating the entry and exit of conversations, with an eye toward managing interpersonal impressions when face-threatening actions occurred, such as ending a conversation prematurely or avoiding interaction [12].

Since that initial study, several other papers have reported on how butler lies are used to manage social interactions in other text-based media, including instant messaging, SMS and BBM text messaging [12,3,23]. The resulting body of work has contributed observations about how and why butler lies are produced. First, butler lies are relatively common. Over the course of three studies, people reported lying in about ten percent of their messages overall. Of these lies, about one fifth of them are butler lies intended to manage availability, suggesting that butler lies are an important strategy for dealing with the pressures and constraints of always-on communication.

Second, butler lies draw on ambiguities that emerge from online communication, such as where someone is, what they're doing, or when they received a message. These ambiguities can be exploited in accounting for past behavior (e.g., "sorry, just got your message") or explain current or future behavior (e.g., "going to bed") in ways that are often not possible in co-present interactions where much more contextual information is shared.

Third, patterns of butler lies are affected by the norms associated with different media. IM butler lies, for example,

are mostly about exiting a conversation (e.g., "Sorry, I have to go eat lunch"), while butler lies in SMS texting are most often concerned with avoiding other social interactions (e.g., "Can't meet up later, busy"), presumably because an important use of SMS among young people and students is managing their social life and interactions [11, 14]. Lastly, an important function of butler lies appears to be the management of interpersonal relationships and impressions. Explanations that people provide for their butler lies frequently involve concerns about avoiding hurting their communication partner's feelings or trying to make themselves look better. The relational importance of butler lies is also highlighted by the observation that when using a closed texting network, such as BBM, participants told more butler lies to those they coordinated with most, suggesting that butler lies are important for managing social interactions in key relationships [23].

Multiple Perspectives on Butler Lies

Several critical questions remain, however, about butler lies and their role in availability management. Perhaps the most important gap in our understanding is how receivers of butler lies perceive them. Are butler lies viewed as deceptive, and if so, are they frequently detected? Do people have expectations around butler lies, such as how often they take place? To date, behavioral research on butler lies has focused exclusively on the liar and their perceptions of their own messages.

This state of affairs, in which we know one person's perspective of their own lying behavior, parallels that of the general deception literature, in which research has focused either on how lies are produced (e.g., [8,13]) or detected (see [31] for a review). We are aware of no research that has examined how both the liar and their interaction partner perceive the same message.

A second question around butler lies that remains understudied is the relationship between butler lies and other types of lies. We know the rate at which butler lies occur from previous research. We do not, however, know whether butler lies occur more or less frequently than other types of lies. Relatedly, we do not know if receivers judge butler messages to be more or less deceptive than other lies.

THE PRESENT STUDY

The present study aims to address these questions by considering multiple perspectives on butler and other types of lies. We accomplish this and build on the prior results described above by examining message-level data from both the senders' and receivers' perspectives, drawing on both quantitative and qualitative analyses.

Terminology

Before presenting our research questions, we must define key terms that will be used throughout the paper. We refer to the producer of a text message as the *sender*, and to the person that received the text message as the *receiver*. We use these terms for their straightforwardness and because

they align with the structure of the technology (in which each message has a clear sender and receiver). At the same time, we acknowledge two drawbacks to these terms. First, and most importantly, language use is a dynamic and collaborative process, and autonomous models of language that assume a sender transmitting a message to a receiver have been shown to be overly simplistic [7]. Second, communication partners are both senders and receivers of messages. Despite these drawbacks, we use these terms to be clear about participant roles for any given text message.

Next, we consider two types of deceptive messages. Following previous research, we separate messages into two broad categories: 1) *butler messages*, are those that pertain to the entry, exit or arranging of social interactions [12], 2) *other messages*, which refer to messages that served different purposes. Thus, deceptive butler messages are *butler lies*, and other deceptive messages are *other lies*. In using these terms, we recognize that the other lies category will involve many kinds of lies, many of which have been more extensively studied by others (e.g., [8]).

Research Questions

Our first question concerns the rates at which lying occurs and is expected to occur. That is, how often do senders report lying in butler messages relative to other messages? Given the pressures we note above on managing availability, we wondered whether butler messages would be deceptive more often than other messages. We also wondered how often receivers expect senders to lie to them. In particular, do receivers expect butler messages to be deceptive more or less frequently than other messages?

RQ1: What are the actual (sender perspective) and perceived (receiver perspective) rates of lying in text messaging?

RQ1 examines independently the rates at which senders tell and receivers expect deception in their SMS conversations. By combining these two measures we can assess the accuracy with which receivers detect lies, another area that previous studies have not explored.

Previous work more generally on the detection of deception has shown that people are not particularly accurate at detecting deceptive messages. A recent meta-analysis of laboratory studies, in which participants were presented with equal numbers of deceptive and truthful messages, showed an average accuracy rate of 54%, only slightly above chance [5]. Even under circumstances that would seem most favorable to those assessing lies, such as when participants receive detection training prior to making their assessments, accuracy rates still do not exceed 75% [9].

This 54% accuracy rate may be an artifact of study designs, however. Accuracy rates are impacted by the *truth bias* [19], which is the reliable tendency for people to assess messages as truthful. As such, when a set of messages contains more lies, the truth bias should lead people to make more errors be less accurate. Levine et al. [16]

demonstrated this by varying the proportion of deceptive messages presented to participants, and found that this proportion had a linear relationship with detection accuracy rates. When half of the messages were deceptive, participants' accuracy rates were about 50%, or similar to the chance rates noted above. On the other hand, when no messages were deceptive accuracy improved to 65% [17].

This effect of truth bias on accuracy has important implications for everyday conversation. While deception is considered to be a common element of communication, previous research suggests that only a small fraction of all messages sent involve deception [8,13,3], but this can vary considerably between individuals [25].

Given this background, we wondered if there would be any differences in how accurately receivers perceive butler and other messages as deceptive. If, for example, RQ1 reveals a difference in the proportion of deceptive messages produced by senders between butler and other content, this may affect receiver's accuracy [16]. Additionally, there may be truth bias differences between butler lies and other lies. If RQ1 reveals that receivers are more or less suspicious of butler messages relative to other messages, this may also impact accuracy. We asked:

RQ2: How accurate are receivers when predicting which messages are deceptive?

The previous two research questions focused on the act of producing and perceiving deceptive messages. We are further interested in the emotional impact of these deceptive acts. Past psychological research has identified that people experience discomfort when telling lies in everyday life [32]. Additionally, some have argued that deception can undermine personal relationships [20] while research on butler lies suggests that they can actually be face-saving and supportive of relationships [12].

Our final research question focuses on the potential emotional impact of telling and perceiving butler lies, another area not previously investigated. Past studies on butler lies have yielded data about the senders' emotions when sending a butler lie, but have not investigated the receivers' side. The design of the present study allows us to collect multiple perspectives on the same messages for the first time. Thus we can compare the emotional impact of these types of messages on both senders and receivers.

RQ3: What is the emotional impact of telling and being told butler lies?

Method

We use a combination of qualitative and quantitative methods. Quantitative analyses are used to address the three research questions, and qualitative analyses provide more depth in understanding participant's perspectives.

Participants

Participants included 82 dyads (164 individuals), all students at a large U.S. university. We have demographic

information from 126 participants. Their ages ranged from 18-34 ($M=20.85$), and 72% were female. This is in contrast to the overall undergraduate population (50% female), which indicates possible selection bias. On average, participants had used some form of text-based messaging for 5.86 years. Due to a database failure, demographic information for the other 38 participants was not recorded.

Students were recruited via an on-campus web-based recruitment system, and signed up for the study in pairs, with instructions that their partner should be somebody with whom they exchange texts with on a regular basis and have a friendly – but specifically non-romantic – relationship. Virtually all participants (93.3%) brought friends as partners, so there was little variance in relationship type. All received either course credit or \$10.

Procedure

Participants arrived in pairs and the procedure was explained to them, clearly indicating that they would be entering messages in an online survey and that their partner would then see these messages. Partners were then separated and seated at individual computers for the duration of the study. After completing an online consent form, they read a short tutorial, which provided definitions and examples for deceptive content and butler content. The deception tutorial provided a clear definition of deception (based on the definition cited above) along with several examples and counter-examples. The butler tutorial was developed for this study, but used a similar approach and was tested with members of our laboratory.

Next, participants completed a short questionnaire including demographic questions and questions about their text messaging behavior (e.g., how long they have used text messaging, people with whom they most often exchange text messages). They then completed Phases 1 and 2.

Phase 1: Using a web-based questionnaire, each participant entered the last 15 text messages that she had sent to her study partner (or fewer if she had not sent 15). The number of messages entered per participant ranged from 3 to 15 (Mode = 15), with only 16 participants providing fewer than 15 messages. For each message participants indicated whether the message was deceptive (measured using a 6-point scale, anchored by 0, “not deceptive,” and 5, “extremely deceptive”). If a message was marked as deceptive, participants explained why the message was deceptive, and indicated how bad they would feel if the receiver found out the message was deceptive (measured using a 5-point scale, anchored by 1, “would not feel bad at all,” and 5, “would feel extremely bad”).

Each participant also answered questions about their relationship with their study partner, both in terms of their categorical relationship (e.g., “family member,” “acquaintance,” “close friend”) as well as the closeness of their relationship (5-point scale, anchored by 1, “not close at all,” and 5, “very close”).

Phase 2: In this phase, participants rated the messages just submitted by their partners. For each message read, participants indicated whether they believed the message was deceptive (using the same scale noted above). If they believed a message was deceptive, they were asked to explain why they thought so and to indicate how bad they would feel if they found out the message actually was deceptive (using the same scale described above).

Message Coding

Of the 2,341 messages we collected, 239 were marked deceptive by senders and 215 were judged deceptive by receivers. We coded these messages for jocularity to exclude messages that were clearly not intended to create a false belief in the recipient (e.g., the message “lol” was not technically true because the participant was not actually laughing out loud, but was not likely intended to mislead the recipient; see [12]). All messages marked as deceptive by senders or receivers were coded for jocularity by two independent coders. The coders together rated a subsample of messages until they reached 80% agreement, and then independently rated the remaining messages. Percent agreement between coders was 92.5% for sender ratings and 89.0% for receiver ratings. Twenty messages initially rated as deceptive by senders and 25 by receivers were coded as jocular and treated as non-deceptive.

Senders indicated that 450 of the messages contained butler content. To verify that butler lies had been correctly identified, we then coded the 219 deceptive messages for butler content, as defined above. For example, “I like your dress.” was other content while “I’m on my way.” was considered butler content. The coders first rated a sample of these messages and after they reached 80% agreement they continued independently. Final percent agreement between coders was 77.3%. We then tested the rate of agreement between coders and participant ratings for the butler lies and the percent agreement was 71.7%. As this was an adequate rate of agreement, participants’ butler content ratings were therefore used for the non-deceptive messages. In total, 474 messages were coded as butler content and as 1,867 other content.

RESULTS

We divide our presentation of results according to the research questions presented above, followed by a qualitative discussion of messages and their explanations.

RQ1: What are the actual and perceived rates of lying in text messaging?

Rate of lying

Our first research question focused on the actual (sender perspective) and perceived (receiver perspective) rates of lying in text messaging. First we looked at the percentage of deceptive and non-deceptive messages containing butler and other content. In our corpus of 2,341 messages, 219 (9.4%) were coded as deceptive based on the senders’ ratings and explanations. One hundred sixteen (6.2%) of the

1,867 other messages and 103 (21.7%) of the 474 butler messages were deceptive.

To examine rates of deception, we used a linear mixed model with rate of lying as the dependent variable. In this model, messages were nested within participants, participants were a random variable nested within pairs, and pairs were also set as a random variable. Relationship closeness and content type (butler vs. other) were included as fixed variables. The results of this model suggest that, when controlling for relationship closeness, senders produced a significantly higher proportion of lies in butler content ($M=27.1\%$, $SE=1.5\%$) than in other content ($M=6.4\%$, $SE=2.0\%$), $F[1,168.08]=47.06$, $p<0.001$ ¹. Thus the rate of lying in messages about social interactions (entering, exiting, or arranging), or butler messages, was higher than the rate of lying in other messages. However, the difference in the severity of deception between butler and other messages was not significant.

Perceived rate of lying

Next we looked at the percentage of messages perceived to be deceptive by receivers. Again, we looked at the difference between messages containing butler content and other content. Of all messages, receivers perceived 192 (8.2%) as deceptive, including 23 (9.1%) of the 474 butler messages and 149 (8.0%) of the other messages.

We used the same linear mixed model described above to examine the perceived rate of deception as a function of content type (butler vs. other) while controlling for relationship closeness. The difference between the perceived rate of lying for messages containing butler and other content was marginally significant, $F[1,230.41]=2.75$, $p<0.1$. On average, receivers perceived more butler messages ($M=10.4\%$, $SE=1.5\%$) to be deceptive than other types of messages ($M=8.0\%$, $SE=1.4\%$). These data suggest that receivers perceive butler messages as marginally more likely to be deceptive than other messages.

Figure 1 shows the average rates of sent and perceived lying for butler and other messages. Several important patterns, supported by the statistical analysis above, are visible. First, the receivers' estimate for the rate of lying in other messages was remarkably similar to what the senders reported actually producing (8.0% vs. 6.4%). In contrast, the receivers substantially under-estimated the rate of lying for butler messages, estimating less than half the actual rate that the sender reported producing (10.4% vs. 27.1%), suggesting that senders were lying in their butler messages more than was anticipated by receivers.




Figure 1: Comparison of actual and perceived rates of lying

While these data tell us about the rates of actual and perceived lies in the partner's text messages to one another, they do not tell us about how *accurately* the receiver detected lies. This is addressed in the next section.

RQ2: How accurately do receivers predict deception?

Our second research question focused on the accuracy of receivers when perceiving deceptiveness in text messages. We examined the sender and receiver indicators of actual and perceived deceptiveness for each message.

Overall, receivers correctly judged 2024 of the messages for an overall accuracy rate of 86.5%, including 373 (78.9%) of the 474 butler messages and 1,650 (88.4%) of the 1,867 other messages.

We used the same linear mixed model as above to test for accuracy as a function of content type while controlling for relationship closeness. The difference in accuracy rates between other ($M=87.1\%$, $SE=1.6\%$) and butler messages ($M=71.1\%$, $SE=1.8\%$) was statistically significant, $F[1,222.47]=18.82$, $p<0.001$, suggesting that receivers were less accurate in detecting butler lies than other lies.

When these data are considered along with the results from RQ1, we see that receivers substantially underestimated the rate at which butler lies are produced, and consequently that they are significantly less accurate at detecting them.

RQ3: What is the emotional impact of lies?

Our third research question focused on senders' emotions when sending deceptive messages and receivers' emotions when they believed they were lied to. We asked senders how bad they would feel if their partner found out they were lying, and receivers how bad they would feel if they found out their partner was lying. We refer to this here as the emotional impact rating.

We again used a linear mixed model with average emotional impact rating as the dependent variable. In this model, messages were nested within participants,

¹ Note that the denominator degrees of freedom in linear mixed models are estimated using a Satterthwaite's approximation, which can yield non-integer degrees of freedom [27].

participants were a random variable nested within pairs, and pairs were also set as a random variable. Relationship closeness, content type (butler vs. other content), and role (sender or receiver) were included as fixed variables. The results of this model indicate that senders feel significantly worse about deceptive messages ($M=1.75$, $SE=.08$) than receivers ($M=1.53$, $SE=.07$), $F[1,303.09]=8.3$, $p<.005$. There was no effect of content type and no interaction between content type and role.

We then separated these deceptive messages into those that were detected as deceptive by receivers, and those that were not, as we were interested in whether that affected the emotional impact of lies. For non-detected deceptive messages, there were no significant differences between content type and role, and no significant interaction.

For detected deceptive messages, however, on average senders felt significantly worse about deceptive messages ($M=2.09$, $SE=.16$) than receivers ($M=1.40$, $SE=.16$), $F[1, 34.82]=12.18$, $p<.001$ (see Figure 2). There was no significant effect of content type and no significant interaction between content type and role.

Overall, this pattern of results indicates that senders feel worse about lying than receivers feel about being lied to, but only when the lie is detected. In general, these data are consistent with previous research indicating that people experience negative emotions when lying [32]. The present data indicate that, in fact, people also feel worse when they are the target of the lie. Although senders felt worse than receivers when telling detected lies, their emotional responses were indeed correlated for both butler ($r=.25$, $p<.05$) and other lies ($r=.27$, $p<.05$), suggesting that the emotional impact of telling lies in text messaging was felt similarly for both senders and receivers.

Given these important emotional effects associated with butler lies, we next sought to examine the relational effects and functions of butler lies by examining the explanations that our participants provided about their lies and their perceptions of their partner's lies.

Analysis of Messages and Explanations

To better understand how people perceived messages as butler lies, we examined their explanations for why they perceived messages to be deceptive. We found that, generally, there were two themes in their explanations.

Relational Coordination

Many messages were rated deceptive based on a combination of the receiver's understanding of their partner and expectations around the situation.

In some cases, butler lies were used simply to expedite the coordination of interaction, such as ensuring co-presence for a later conversation. One participant, for example, texted "*I'm walking in the upstairs door right now.*" Their partner correctly identified this message as deceptive, noting "I don't think he was actually walking in the door. I




Figure 2: Emotional impact ratings for detected lies

think he was a little further away." This exactly reflects the sender's explanation, "I was not yet near the door. i could see it... but I was still about 50 yards from being there. Written to assure the message receiver that I would be there shortly plus by the time she saw it, I might actually be at that place." Neither of them indicated they would feel bad if this was known to be deceptive, as both rated the emotional impact as '1,' the low point on the scale, reflecting the implicit purpose of expediting co-presence at the door and the low stakes of the deception.

As participants tended to come to the lab with people they knew relatively well, their explanations often referenced prior experience with the partner. This experience also affected how butler lies were used and interpreted. In one case, for example, the sender texted "*I'm going to bed.*" The receiver rated the message as deceptive, explaining, "knowing this person i highly doubt this is a true statement." In fact, the message was deceptive, with the sender explaining "...it was already near midnight and he hadn't shown up. i was pissed and wanted him to think i didn't care if he was coming or not." Here the explicit message was to convey unavailability for interaction, but the implicit message – which draws on a shared understanding of expected bedtimes – was to convey frustration and possible trouble in the relationship. The sender's explanation makes clear that they were acting strategically and in anticipation of a particular response from the receiver, and the sender's emotional impact rating ('1') reflects this intent to convey frustration. The receiver did not provide an emotional impact rating.

There were also cases where implicit relational messages were more positive or routine. One participant, for example, texted "One last game sweetie (sic)," with emotional impact rating of '3' His partner rated this as deceptive, noting that "usually he says one last video game and ends up playing two or three more games," but rated the emotional impact as '1'. The sender did indicate that the message was deceptive, explaining that "I texted saying that I would be playing one more game to express the fact that I would be leaving soon. In reality I don't believe I was going to come

back soon. Most likely it would still take some time for me to come back.” Both sender and receiver explanations here reflect an understanding that this message is not literally true, and some shared sense that the actual purpose of the message was to convey that the sender would continue playing games, but be available later. That the sender feels worse about the deception than the receiver, moreover, is consistent with the above analysis, and suggests that the recipient is willing to accept the deceptive message.

As noted above, many true messages were also perceived as deceptive. This was often true when the message involved a common butler lie. One participant, for example, texted “*Hii sorry I didn't have ym hpone (sic)*,” which was not marked as deceptive. The receiver, however, marked it as deceptive, rating emotional impact ‘1,’ and explained, “I bet she had her phone just mentally texted back (read it and forgot to respond).” This combination of the plausible explanation and ambivalence about the deception (i.e., the receiver would not feel bad) suggest a willingness to accept the message whether it is true or not.

Receivers were not always willing to accept deceptive messages, however, and their explanations sometimes revealed frustration or the desire to confront their partners. One participant, for example, texted “*ok sry 4 my horrible txting abilities i am done that awful wrk sht! (sic)*” This seemed to frustrate her partner, who rated the emotional impact ‘2’ and said that “The fact that she didn't text me back for a long period of time because she was busy with ‘wrk sht’ was deceptive, because she had time to eat two different leisurely lunches with friends as she later told me but not 10 seconds to text me back?”. The deception here was subtle, moreover. The sender marked the message as deceptive with an emotional impact of ‘2’ though she was, in fact, working on a worksheet, but also admitted to having time for replying to texts from others (but does not mention lunch): “I wasn't able to respond to several of her texts that morning because I had a worksheet to do. However, i did send a few texts which needed to be immediately received from some other people, while I was doing my work.”

This is interesting in that there is a mutual recognition that the message sender is conveying the message that she is too busy to interact with the recipient, but not too busy to interact with others. The recipient is thus unhappy not that the sender could not talk to the recipient, but that the sender had time to talk to others and not the recipient. This highlights an important tension in the management of social inattention: that people juggle many different priorities, and thus different levels of availability for different others. Successful balance of these tensions seems to require either that the process of managing these interactions be shielded from view (e.g., via ambiguity about what is actually taking place), or reflect a mutual agreement about priorities. Otherwise the result is the type of frustration revealed here.

Factually False Messages

In some cases, participants knew messages to be factually false, so could classify them as deceptive outright. Often this was because they had other sources of information available when they received the message. One participant, for example, texted “Yo, come up to 205. [Fred] just got here, he just started drinking. [Bob]'s here too.” His partner rated this as deceptive, noting that “I knew for a fact that our friend [Fred] had been drinking for awhile.” In fact, the sender’s explanation confirms this, saying that “I said that [Fred] had just started drinking, but he had actually started drinking 3 hours earlier and was hammered....”

In another case, the recipient had additional information at the time of the study that they did not have when they received the message. One participant deceptively said “i am walking about want me to meet u sumwhere and take u” (explanation: “I was walking about previously, but was actually currently in my room.”) Her partner correctly rated this message as deceptive, explaining “A few minutes later she said she was in [her dorm] so she probably wasn't ‘walking about.’” This is, of course, somewhat problematic in that the rating of deceptiveness is retrospective, rather than “in-the-moment.” Thus, this is more an artifact of our study design than a reflection of participants’ immediate perceptions upon receiving the messages. This occurred only a small number of times in our data set, however, so should not affect the results above.

DISCUSSION

We have presented evidence from quantitative and qualitative analyses of butler lies used in a population of American university students, from both the senders' and receivers' perspectives. The data suggest that, proportionally, butler lies occur more frequently in everyday text messaging than other types of lies, and that receivers have less accurate perceptions about the expected rate of deception – as measured via the proportion of messages that they identified as deceptive – for butler lies than other types of lies. Receivers also tended to be less accurate at detecting which butler messages were deceptive than other types of lies. On average, senders reported feeling worse about their deceptive messages than the receivers. We saw from the explanations participants provided that messages were perceived to be deceptive when considered along with additional information about the situation or the individual, or when they contained verifiably false facts.

Our first major finding is that a significantly higher proportion of butler messages were deceptive when compared with other messages. The implication is that deception is a common practice in availability management via text messaging, and appears to occur more frequently than other types of everyday deception [8]. This suggests evolving social norms around availability management practices, a reaction consistent with recent work suggesting

that people feel pressure to be connected constantly and commonly seek to avoid communication [1,33].

We next looked at people's expectations around deception. Receivers expected marginally more butler lies than other types of lies, suggesting that they were more suspicious, or less truth biased, for butler messages than other messages. Despite this reduced truth bias, receivers significantly underestimated the actual frequency of butler lies, with the expected rate of deception less than half that of the actual rate. Moreover, receivers detected butler lies less accurately than other lies. The implication of these results, particularly when combined with analysis of the explanations participants provided, is that people expect to be lied to about availability management, even if they do not always know when. This points toward the possibility of a more nuanced approach to understanding deception in availability management, in which people's expectations and overall relational goals are considered alongside the literal message text and deceptive practices.

Considering the detection result in more detail, why did receivers perform so poorly at detecting butler lies relative to other lies? There are several possible explanations for this. One is that butler lies are simply lower stakes than other lies and thus harder to detect, similar to "little white lies." Previous research has pointed out that lies not involving high stakes (e.g., punishment or significant reward) are difficult to detect [10]. This explanation, however, seems unlikely given that the emotional impact between butler lies and other lies was virtually identical, suggesting that the stakes of butler lies are not necessarily lower than other kinds of lies.

Our findings point to a second, more plausible explanation, that accuracy rates are directly related to the increased proportion of deceptive butler messages produced by the sender and the receiver's failure to fully anticipate this higher rate of deception. As Levine et al. [16] suggested, the truth bias plays a major role here in that people expect most statements to be true. Our receivers expected about 92% of butler messages to be truthful, when in fact senders reported telling truthful butler messages only about 73% of the time. This mismatch generated by the truth bias, we argue, underlies our participants low accuracy in detecting butler lies relative to other lies. In terms of the bigger picture, this pattern of results suggests that people have not yet adjusted their expectations around the frequency with which butler lies are used in managing availability.

All of this is especially interesting in light of our finding that senders were more emotionally impacted by lies than receivers. On the one hand, this is consistent with previous research that has shown people to experience negative emotions when lying [32]. What is novel here, however, is that our data provided multiple perspectives on the same messages, which allowed us to discover that 1) liars not only feel bad, they actually feel worse than the target of the lie, and 2) the emotional impact of the lie on the sender and

receiver are correlated, suggesting that both parties are aligned in their perceptions of how emotionally negative the deceptive message may be.

Again, this points toward the possible utility of a more nuanced approach to deception in availability management. Indeed, prior work has shown that deception is an everyday behavior [8] and our findings suggest that people expect to be lied to, although they have trouble identifying the specific deceptive messages. If a receiver expects to be lied to, particularly in situations where they know the liar and understand possible implicit relational messages, it is reasonable that they would not be upset when it happens. The sender cannot be sure that the receiver is expecting deception, however, so it is reasonable for them to be more adversely emotionally impacted. Further research is needed to fully understand the emotional effects of butler lies.

Implications for Design

It is clear from our findings and prior work [12,3,23] that butler lies are a commonly used linguistic strategy for availability management that draw on ambiguities resulting from features of communication technologies and the social norms surrounding their use. Our analyses here strongly suggest that senders and receivers of messages recognize the social value of these messages, even when receivers believe them to be false. Moreover, receivers are not very accurate in detecting butler lies.

A naïve response to these results would be a push to design systems that make it easier for people to detect when they are being deceived, thus improving accuracy and reducing deception (but ignoring social value issues discussed throughout this paper). Interestingly, this is exactly what designers of recent commercial systems have done. Despite the potential value in the types of ambiguity commonly drawn on in butler lies, these appear to be under threat. BBM and Apple's iMessage both provide notification when a text message has been read, for example; and Lotus Notes provides a similar feature ("return receipt") in its email system. Automatic location sharing is also becoming more common, and already accompanies status updates and other posts made via GPS-enabled mobile devices on Facebook. Location-aware messaging is also available on some systems, and location sharing is facilitated by systems like foursquare, Facebook, Latitude, Glympe and Twist.

Our results suggest that designers should not – and users likely will not – take this additional information at face value. Reynolds et al. [23] observed BBM users, for example, who intentionally delayed opening messages to avoid generating "read" notifications; and Birnholtz et al. [2] observed an office in which it was not normatively acceptable to use the return receipt feature in Notes. This builds directly on Aoki and Woodruff's [1] notion of making space for stories, and our results build further on these ideas by showing people expect to be deceived some of the time, and are careful consumers of these messages. They also value the social niceties conveyed and enabled by

these practices, however. In other words, they value having “space for stories.”

Our primary design suggestion, therefore, is to allow for more discretion and latitude in the use of these features. At a minimum, this would mean allowing users to turn features on or off selectively. Allowing for some manipulation or obfuscation of automated information could be useful as well. One could set “read” notifications, for example, to appear only after a certain amount of time has passed; and also set “read” messages to appear as “unread.” In this way, one could look briefly at a message within the pre-specified time window, and the sender would still believe it to be unread and not be upset by a delayed reply. In effect, such features would create new types of ambiguity around who uses particular features, when they are or are not enabled, and whether or not they are being used to share accurately.

Another finding here from the analysis of receiver explanations is that people use many sources of information in assessing the content of messages. With today’s systems, however, it may be difficult for message senders to know what the receiver is aware of (e.g., “Did she see my Facebook post tagged with my Chicago location?”). This can make it more difficult to construct plausible and coherent explanations, as discussed by [22]. We therefore urge designers to consider ways of making message senders in various media aware of information known to be available to a message recipient. In a simple form, this could include a set of checkboxes (i.e., “Receiver may know: __ my location; __ when I read my messages; __ my calendar), with links to the potential source of the information. This will allow senders to craft plausible rationales, even as sources of ambiguity (i.e., who has access to what and what features are enabled) themselves may be ambiguous. To be clear, we are not advocating helping people deceive and possibly hurt others, but we do want to help them manage their availability and avoid communication overload. This means being sensitive not just to the need for ambiguity and crafting plausible explanations, but also sensitivity to factors that influence perceptions and usage of butler lies.

Limitations and Future Work

There are several limitations that should be considered when interpreting these findings, and that motivate important future work. First, we studied students at a university in the United States and these results may not generalize to other populations. However, college-aged young adults are the most active users of text messaging [26], so this sample is useful for understanding everyday text messaging behavior. At the same time, we urge future study of butler lie usage and perception in cultures that may have very different views of deception [18], among other demographic groups that may use SMS differently [26], and in non-social contexts such as work, where behavior may differ as well [2, 15].

Second, our design and recruitment strategy meant that most participants came to the lab with relatively close friends. There was little variance on the dimensions of relationship type and closeness. We can therefore draw few inferences about the role of these factors in availability management. We urge future research that explicitly examines participant interactions with partners from multiple relationship categories (e.g., romantic partners, close family members, distant family, acquaintance, etc.) to more closely examine these phenomena, which we believe likely impact behavior.

Third, our participants were disproportionately female. We ran all analyses controlling for gender and found no gender effects. However, we are missing demographic information due to a database failure for a significant number of participants. We do not know what effects missing data could have on our findings, but we have no reason to expect that data would be significantly different than the demographic data we do have.

Last, our survey method relies on self-reports of lying, and requires participants to be honest about their deceptive behaviors. The method used here, however, in which participants rated and explained actual messages and deceptions, improves upon other self-report deception methods, such as diary studies, that rely on the participants memory for prior conversations. Future studies may improve on this method further by directly collecting messages from participants’ phones rather than having them enter them into a survey. Additionally, while the explanations participants provided in this study gave us insight to multiple perspectives on the same messages, future qualitative work is required to more thoroughly explore this rich area through interviews with people who regularly use text messaging.

ACKNOWLEDGEMENTS

This work is supported in part by the National Science Foundation (Award #IIS 0915081). We acknowledge valuable research assistance from Jon Culver, Mark Thomas, and Jin Zhao.

REFERENCES

1. Aoki, P., & Woodruff, A. (2005). Making space for stories: ambiguity in the design of personal communication systems. *Proc. ACM CHI*, 181-190.
2. Birnholtz, J., Dixon, G., & Hancock, J. (2012). Distance, ambiguity and appropriation: Structures affording impression management in a collocated organization. *Computers in Human Behavior*, 28, 1028-1035.
3. Birnholtz, J., Guillory, J., Hancock, J. T., & Bazarova, N. (2010). “on my way”: Deceptive Texting and Interpersonal Awareness Narratives. *Proc. ACM CSCW*, 1-4.
4. Boehner, K., & Hancock, J.T. (2006). Advancing ambiguity. *Proc. ACM CHI*, 103-107.

5. Bond, C. F. and DePaulo, B., M. (2006). Accuracy of Deception Judgments. *Journal of Personality and Social Psychology*, 10 (3), 214-234.
6. Chen, B. X. (2011). *Always on: how the iPhone unlocked the anything - anytime - anywhere future--and locked us in*. Cambridge, MA: Da Capo Press.
7. Clark, H. H. (1996). *Using language*. New York: Cambridge University Press.
8. DePaulo, B.M., Kashy, D.A., Kirkendol, S.E., Wyer, M.M., and Epstein, J.A., (1996). Lying in Everyday Life. *Journal of Personality and Social Psychology*, 70(5), 979-995.
9. Fielder, K. and Walka, I. (1993). Training Lie Detectors to Use Nonverbal Cues Instead of Global Heuristics. *Human Communication Research*, 20(2), 199-223.
10. Frank, M.G. & Ekman, P. (1997). The ability to detect deceit generalizes across different types of high stake lies. *Journal of Personality and Social Psychology*, 72, 1429-1439.
11. Grinter, R., & Eldridge, M. (2003). Wan2tlk?: Everyday Text Messaging. *Proc. ACM CHI*, 441- 448.
12. Hancock, J. T., Birnholtz, J., Bazarova, N., Guillory, J., Perlin, J., & Amos, B. (2009). Butler lies: awareness, deception and design. *Proc. ACM CHI*, 517-526.
13. Hancock, J., Thom-Santelli, J., and Ritchie, T. (2004). Deception and design: the impact of communication technology on lying behavior. *Proc. ACM CHI*, 129-134.
14. Häkkinä, J., and Chatfield, C. (2005). 'It's like if you opened someone else's letter' – User Perceived Privacy and Social Practices with SMS Communication, *Proc. ACM MobileHCI*, 219-222
15. Leonardi, P. M., Treem, J. W., & Jackson, M. H. (2010). The Connectivity Paradox: Using Technology to Both Decrease and Increase Perceptions of Distance in Distributed Work Arrangements. *Journal of Applied Communication Research*, 38(1), 85-105.
16. Levine, T. R., Park, H. S., & McCornack, S. A. (1999). Accuracy in detecting truths and lies: Documenting the “veracity effect.” *Communication Monographs*, 66, 125-144.
17. Levine, T. R., Kim, R. K., Park, H.S., & Hughes, M. (2006). Deception detection accuracy is a predictable linear function of message veracity base-rate: A formal test of Park and Levine's probability model. *Communication Monographs*, 73, 243-260.
18. Ma, F., Xu, F., Heyman, G. D., & Lee, K. (2011). Chinese children's evaluations of white lies: Weighing the consequences for recipients. *Journal of Experimental Child Psychology*, 108, 308-321
19. McCornack, S. A., & Parks, M. R. (1986). Deception detection and relationship development: The other side of trust. In M. L. McLaughlin (Ed.), *Communication yearbook 9* (pp. 377-389). Beverly Hills, CA: Sage.
20. Miller, G. R., Mongeau, P.A., & Sleight, C. (1986). Fudging with Friends and Lying to Lovers: Deceptive Communication in Personal Relationships. *Journal of Social and Personal Relationships*, 3, 495-512.
21. Perry, M., O'Hara, K., Sellen, A., Brown, B., & Harper, R. (2001). Dealing with mobility: understanding access anytime, anywhere. *ACM TOCHI*, 8, 323-347.
22. Read, S. (1992). Constructing accounts: the role of explanatory coherence. *Explaining the self to others*. M. L. McLaughlin, M. J. Cody and S. Read. Hillsdale, NJ, Lawrence Erlbaum Associates: 3-19.
23. Reynolds, L., Gillette, S., Marder, J., Miles, Z., Vodenski, P., Weintraub, A., et al. (2011). Contact stratification and deception. *Proc. ACM CSCW*, 221-224.
24. Salovaara, A., Lindqvist, A., Hasu, T., Häkkinä, J. (2011). The Phone Rings but the User Doesn't Answer: Unavailability in Mobile Communication. *Proc. ACM MobileHCI*, 503-512.
25. Serota, K.B., Levine, T.R., and Boster, F.J. (2010). The Prevalence of Lying in America: Three Studies of Self-Reported Lies. *Human Communication Research*, 36(1), 2-25.
26. Smith, A. (2011). *Americans and Text Messaging*. Washington, DC: Pew Research Center's Internet & American Life Project.
27. SPSS Technical Report. *Linear Mixed Modeling*. http://www.spss.ch/upload/1107355943_LinearMixedEffectsModelling.pdf.
28. Tang, J. C. (2007). Approaching and leave-taking: Negotiating contact in computer-mediated communication. *ACM TOCHI*. 14 (1), Article 5.
29. Turkle, S. (2011). *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books.
30. Vanden Abeele, M., & Roe, K. (2008). White cyberlies: The use of deceptive instant messaging statuses as a social norm. *Presented at Conference of the International Communication Association*.
31. Vrij, A. (2008). *Detecting lies and deceit: pitfalls and opportunities*. Chichester: Wiley.
32. Waidd, W. M., Orne, M. T. (1981). Cognitive, Social, and Personality Processes in the Physiological Detection of Deception, In: Leonard Berkowitz (Ed.), *Advances in Experimental Social Psychology*, Academic Press, 14, 61-106.
33. Weilenmann, A. (2003). "I can't talk now, I'm in a fitting room": formulating availability and location in mobile-phone conversations. *Environment and Planning A*, 35, 1589-1605