THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Self Curation, Social Partitioning, Escaping from Prejudice and Harassment: the Many Dimensions of Lying Online

Citation for published version:

Kleek, MV, Smith, DA, Shadbolt, NR, Murray-Rust, D & Guy, A 2015, Self Curation, Social Partitioning, Escaping from Prejudice and Harassment: the Many Dimensions of Lying Online. in *WWW 2015 Companion*. pp. 371-372. <https://doi.org/10.1145/2740908.2745940>

Digital Object Identifier (DOI):

[10.1145/2740908.2745940](https://doi.org/10.1145/2740908.2745940)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

WWW 2015 Companion

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Self Curation, Social Partitioning, Escaping from Prejudice and Harassment: the Many Dimensions of Lying Online

Max Van Kleek
Web and Internet Science
University of Southampton, UK
emax@ecs.soton.ac.uk

Dave Murray-Rust
School of Informatics
University of Edinburgh, UK
d.murray-rust@ed.ac.uk

Amy Guy
School of Informatics
University of Edinburgh, UK
Amy.Guy@ed.ac.uk

Daniel A. Smith
Web and Internet Science
University of Southampton, UK
ds@ecs.soton.ac.uk

Nigel R. Shadbolt
Web and Internet Science
University of Southampton, UK
nrs@ecs.soton.ac.uk

ABSTRACT

Portraying matters as other than they truly are is an important part of everyday human communication. In this paper, we use a survey to examine ways in which people fabricate, omit or alter the truth online. Many reasons are found, including creative expression, hiding sensitive information, role-playing, and avoiding harassment or discrimination. The results suggest lying is often used for benign purposes, and we conclude that its use may be essential to maintaining a humane online society.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous;
J.4 [Social and Behavioral Sciences]: Sociology

Keywords

Lying online; privacy; digital identity; online communities

1. INTRODUCTION

People avoid telling the “full, open, and honest truth” in many situations, whether it involves simply the omission or falsification of information, to more substantial forms of deception and lying. Such behaviours have been shown to amount to, by some accounts, nearly a third of offline interpersonal communications [14, 1]. As individuals increasingly manage multiple social contexts of growing complexity in their daily lives, techniques are required for navigating the interlocking and often antagonistic demands placed on them. Examination of deceptive practices has shown that they often serve as coping strategies for managing and mitigating these complex social situations. Examples of such reasons include protecting one or another’s reputation or identity, such as to preserve particular relationships or ties, avoiding confrontation, showing solidarity with another, covering up accidental transgressions, among others [20, 9].

As the prevalence of computationally mediated socialisation increases, so does the need to understand the role and use of lying and deception in online interaction, and to uncover the kinds of social tensions and attendant complexities that arise from the new social affordances that the Web provides [21]. People now conduct their interactions and curate their identities across a large number of online spaces at a much larger scale, whilst attempting to maintain their privacy, reputation and roles throughout. Deception is a tool used to cope with this complexity, and a lens through which their difficulties and needs can be observed [20, 9].

In this paper, we present a summary of a survey-based study in which we sought to characterise the spectrum of lying and deception practices routinely used online. We are concerned with the ways in which people carry out the use of such behaviours, the reasons for them, and the manners in which such practices arise or are used differently across contexts, situations, and spaces. We aimed this initial study at prolific internet users, who spend, if not a majority, a substantial part of their daily lives in social encounters online, as being most socially engaged online and likely to the widest variety of such behaviours. We are interested in the intent behind the deception, but we do not examine the moral or ethical dimensions of such practices, as these can be highly subjective and grounded in particular personal philosophies.

As described in the following sections, our analysis found that while there are a wide range of reasons people use deception or identity protection online, few reasons for doing so are malicious (or comprised of “dark lies”); in fact, a majority of the reasons pertain to impression management, conflict avoidance, and in order to fit in to groups. Prior to describing these findings, in the following sections, we first briefly review the literature on online deception and lying. Second, we briefly describe our methodology and present our results. In the final section, we discuss the implications of our findings for the development of the Web, in particular within the context of the increasing centralisation of social networks and rise of web identity assurance providers.

2. BACKGROUND

While several detailed frameworks for characterising different kinds of deception have been proposed in the social psychology literature, the most common approach divides lies into two general classes: *everyday lies* of “little consequence or regret”, and *serious lies*, used typically to “hide transgressions”, ranging from misdeeds to betrayals of trust [14]. These notions correspond to

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WWW2015 Florence, Italy

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

the common distinction between harmless or beneficial *white lies*, and socially or morally reproachable lies. A third category of *grey lies* exists in between [10]. In order to characterise the motivations behind lying, Buller and Burgoon derive three categories: *instrumental*, *relational*, pertaining to interpersonal relations, and *identity-related*, pertaining to reputation management [8].

A number of studies have compared online to offline lying practices. When looking at the *richness* of communication environments — the amount and immediacy of feedback for participants — two competing theories arise. *Media richness theory* postulates that lying is more likely in rich environments, while *social distance theory* predicts more lying in impoverished settings, with fewer social cues, such as via e-mail. Hancock et al. [21] analyse this competition, and find support for the latter in their experiments. Belkin et al. compared lying in e-mail to paper, finding that people were likely to lie on e-mail significantly more than in personal letters [31].

Sherry Turkle’s seminal work on early digital spaces, *Life on the Screen* [36], discusses the ways in which identity construction is a creative act, and linked to self development. This desire to choose how one presents has been confirmed by many studies of online behaviour. For example, the frequency with which gamers assume the opposite gender in online play [24, 32, 37], assume widely different appearances in online virtual environments and massively multiplayer games [42], falsify age, height, weight and appearances in online dating sites [22, 17], and experiment with self-representation in online chat and discussion spaces [16, 39, 38].

Most similar to the work presented in this paper, Caspi et al. [11] conducted a Web survey of Israeli internet users and found that, although a majority of participants believed that deception on the Internet was widespread, only a third ever reported using deception online themselves. Instead of open-ended responses, however, their survey asked participants to select among a fixed set of possible personal attributes (age, sex, marital status, height, weight, etc.) and the reasons for doing so. Our study complements theirs by providing a significantly larger set of ways people use deception, omission and false attributes, and justifications for doing each.

3. DESIGN AND METHOD

In this section, we describe the process used to design the instrument we selected to elicit the online behaviours we wished to study. We follow this discussion with a description of the participant recruitment method, and method of data analysis conducted.

3.1 Survey Design

Early on in the process, we postulated that eliciting the breadth of behaviours we sought to understand – ranging from trivial information omission and fabrication, to deliberate deception and beyond – might be challenging for multiple reasons. On the one hand, behaviours concerned with privacy and safety are often carried out routinely, with benevolent (non-malicious) intention. These kinds of well-ingrained practices might not be reported as lying or deception, due to both their familiarity and their motivation. At the opposite end of the scale, morally reproachable, subversive or malicious practices may not be reported because the person is not comfortable sharing, whether to protect their reputation, their self image or for fear of repercussions.

We addressed these risks by, first, looking for precedent in previous surveys (e.g., the *Questionnaire on Academic Excuses*

Q4a*	Have you ever told lies / “untruths” online? Why?
Q4b	How often do you tell lies / “untruths” on social media?
Q4c	How often do you think your friends lie on social media compared to you?
Q5a*	Do you use any pseudonyms online? Why?
Q5b*	Have you created any fictional personas? Why?

Figure 1: List of questions corresponding to subset of survey discussed in this paper; questions with stars indicate free response; the rest are 5-level Likert responses.

for student lying behaviour [33], elicitation method for daily lying studies [14]) and second, iterating on the survey design consulting other colleagues as experts in the process to shape the specific foci and wordings of our questions. Since we could find no survey that we could directly appropriate, we resorted to designing our own¹. In order to characterise the broad class of behaviours we wanted to examine, we first showed a list of candidate terms including terms such as *deceptions*, *lies*, *falsifications*, *omissions* and *untruths* to several experts, alongside a small but diverse list of example behaviours we wished to seek. Our colleagues, comprising 2 Web Science doctoral students and 3 postdoctoral researchers, gave us feedback about which term(s) they considered most appropriate, and then discussed the range of behaviours we were seeking to elicit. The outcome of this process was to break our three distinct questions: one pertaining to the use of untruths, one pertaining to the use of pseudonyms, and finally to the use of fictional personas, which are identities for characters that were entirely fabricated.

The survey was delivered via the web, and comprised 12 sets of questions including 1 set of demographic questions, and 8 open-answer free responses. In this paper, we focus on the subset of the questionnaire delineated in Figure 1.

3.2 Participant Recruitment

The survey was published online, with no restrictions on participation. General recruitment was carried out by handing out flyers with the URL, and the researcher’s social media presences (primarily Facebook and Twitter). This was augmented by enlisting two people with popular twitter accounts (@TheTomSka, 191k followers, and @DameWendyDBE, 4k followers) to promote the survey. In order to ensure a good selection of passionate internet users – people who live a lot of their life on the web, and care about their online presence – additional recruitment was carried out in person at two events in London: *ComicCon* and the *WebWeWant Festival*.

3.3 Analysis

Analysis of free-response questions was done using a *grounded theory* [35] approach; themes were identified across responses through a process starting with open coding process by each of three researchers separately, followed by a discussion process where themes were refined and combined. Multiple themes were permitted per entry. Once consensus was achieved on themes, all responses for a given question were re-coded against the final set.

4. RESULTS

¹We initially considered several methods besides survey, including semi-structured interviews, and artefact examinations, but fell back to a web-based survey to be the most appropriate for getting a wide sample from a large number of our target population.

Figure 2: Self-reported use of social media, from 1=Never to 5=Often times a day. Medians: facebook = 3, twitter = 5, YouTube = 5, Tumblr = 3, Instagram=2, Vine=1

Out of the $N=500$ survey responses, 39% ($N=198$) provided a gender; 50.2% responded female, 49.8% male, and 1% transgender. With respect to age, 59% responded, 91% were between 18–25, 7% between 26–35, and 2% 36+. The age distribution skew was reflective of, and likely due primarily to, the predominantly young audiences at the two festivals; we discuss the potential implications of this distribution at the end of the paper.

Nearly all respondents were very active social media users, although use of particular platforms varied significantly. Figure 2 shows the self-reported Likert scores per platform for 6 social media platforms. The popularity of YouTube and Twitter for respondents was likely influenced by the method of recruitment (via Twitter), and the fact that one of the popular Twitter users who disseminated news of the survey is a popular YouTube. The other platforms, meanwhile, were more divided, with Tumblr being the most divided between highly active ($N=125$, 27%) and those that never used it ($N=144$, 32%). Vine was reported used the least overall with ($N=422$, 91%) reporting having either never or rarely used it.

4.1 Self-reported frequency of deception/lying

In terms of frequency of lying, 77% of participants ($N=387$) responded to *Q4b*, *How often do you lie on social media?* the distribution of answers is displayed in Figure 3a. The median response was 2, with a majority ($N=330$, 85%) of responses answering either a 1 or 2.

Question 4c asked *How often do you think your friends lie on social media compared to you?*, and 77% ($N=386$) again responded overall (Figure 3b). The median value was 3, with

Figure 3: Responses to *Q4b* and *Q4c* on Likert scales a) *How often do you tell lies or untruths on social media?* (1=Never to 5=Often) b) *How often do you think your friends lie on social media compared to you?* (1=Vastly less to 5=Vastly more)

($N=87$, 22%) responding with a value that their friends lie less than they do (e.g. 1 or 2), while ($N=119$, 30%) responded that their friends lie more (e.g. 4 or 5).

4.2 Reasons for Deception

A total of $N=134$ responses were received for *q4*, which asked people to explain whether they remembered telling lies (or “untruths”) online and to explain the circumstances. Out of the total respondents a quarter ($N=34$, 25%) answered that they had or did not lie or use any form of deception online. The rest of the respondents admitted to having, or performing some form of deception regularly. Thematic coding of the remainder of the responses revealed 10 themes, listed in Figure 4, plus an extra for *yes*, a category standing for responses admitting participating in deception with no explanation, and *no* for responses that denied using deception on social media. Tagging each response with the themes yielded $M=1.08$ themes per category, with a minimum of 1 and maximum of 3 ($\sigma^2=0.3$). The number of responses falling into each of the themes is visible in Figure 4 (again, a response may be assigned to more than one theme).

The most prominent theme was **playup** ($N=35$), which corresponded to the rationale of wanting to be more appealing, interesting or attractive to others. There were several subtypes of this activity, starting with simply falsifying personal attributes (height, weight, age) towards what s/he perceived would make him or her more attractive, to *exaggerating* details of stories, to making things “seem relatable”. Four mentioned aspects relating to making one’s self seem popular or important by filling his or her social calendar to appear busy, while two discussed fabricating stories, such as of having met celebrities. Contexts ranged from online dating to social interaction with strangers.

Far less common ($N=9$, 7%) was the opposite reason, in which participants reported deliberately distorting or omitting information in order to not *attract* attention or in many cases to prevent disclosure of illness or situation to protect one’s reputation. This theme, coded as **playdown**, included the following responses:

(R354) Lied about my mental health countless times, denied depression and suicidal thoughts.

(R461) I very selectively curate my online personae, particularly on Facebook, where I am careful to hide my mental illness, my frustrations, and my negative emotions.

(R49) I tend to lie about how sick I am so people don’t worry/employers don’t get anxious.

The second most prominent theme after **playup** was **privacy**, a theme used to encompass a variety of privacy related

Figure 4: Tags and counts for responses to Question 4: Have you ever told “untruths” on social media, given fictitious info, omitted or distorted the truth online?.

concerns. Respondents reported explicitly withholding information often, and, where information was required, providing false values about themselves. The attributes most often mentioned were age ($N=17$), real name ($N=13$), physical location ($N=6$), gender ($N=3$) and birth date ($N=2$) to web sites that they did not trust. Four mentioned that this in order to prevent identity linkage to their real-world identities, e.g.:

(R461) On fetish sites, I will lie about my birthday (displacing my age by a few months to a year in the process) and my hometown, making my identity there harder to connect to my real identity.

Others expressed that they adopted the strategy of falsifying attributes when social networks asked for information that were unnecessary, for example:

(R500) Whenever a social media asks me to provide personal details which are not directly necessary for them to deliver the service (e.g. Facebook asking for my workplace), I constantly feed them wrong information. First and foremost to stop them asking me for such information while at the same time keeping my personal data private.

A different reason given for falsified attributes was coded as **conform**, when falsification was done in order to fit in, in particular to avoid harassment and discrimination. Such behaviour including avoiding potential confrontation surrounding personal beliefs (pertaining to religion or politics), or to personal attributes including gender, age, race or sexual orientation. One participant described her choice of declaring herself as male improved her position in debates online which were often predicated with *ad feminam* attacks on her gender:

(R301) The major untruth I tell is pretending to be a man rather than a woman on YouTube - I know it's bad and not helping the cause, but I know that if I want to convince someone of a particular point, if I pretend to be a man my sayings won't be regarded through the bias of my gender, while if I say opinions (completely disconnected from gender issues) as a woman, it will probably be the 1st thing my opponents will use in a debate.

Another set of responses ($N=6$) involve tricking the system in some way (**system**), predominantly falsification of age in order to circumvent controls on age-restricted websites.

A smaller category ($N=6$) corresponded to deception or lies told for fun, humour, or “just messing about”. The tag **creative** was used for this group, which included examples such as pretending to have a twin, pretending to have met someone famous, or permuting another person's words.

Lies used to diffuse, or bring an end to, unwanted social situations we called **mitigate**. This class ($N=7$) was a superset of butler lies, while butler lies serve primarily to terminate and divert unwanted social interactions, the lies in this category included those told to be polite, such as agreeing with a person to avoid an argument. Meanwhile, **safety** ($N=3$) corresponded to the responses describing omission or falsification to avoid compromising one's physical safety, or from potential litigation for potentially illegal activities.

Some users described the use of deception in order to deceive, trick or manipulate situations to the individual's advantage; such reasons were coded **soceng** ($N=2$) because it reflected the common notion of “social engineering”. These responses described falsification of academic credentials for jobs and posing as another person online and attempting to attract her partner's attentions as this fake identity in order to test her partner's loyalty.

Finally, **explore**, and **coherence** each had two responses. The first, **explore**, pertained to responses that discussed experimenting with aspects of their identity, in particular to “test the reactions of others”, a category corresponding to the tag **experiment** for $q5$, as described in the next section. Meanwhile, **coherence** was the act of lying in order to maintain consistency with other lies told elsewhere to prevent lies from being discovered.

4.3 Pseudonyms

A total of $N=286$ responses were received for $q5a$, which asked for information about whether participants had used pseudonyms, and why. A group ($N=82, 27\%$) claimed not to use pseudonyms online, and a further group ($N=5, 2\%$) gave answers which were unclear. This left 70% of respondents claiming to have used an online pseudonym.

The most common reason for pseudonym use was tagged as **separation** ($N=63, 22\%$). This covers several different lines of division. The three most prevalent reasons were i) separating online and offline lives; ii) separating personal and professional identities; and iii) maintaining distinction between groups of friends or family:

(R266) ... It was mainly done to slightly separate my identity from reality and the internet.

(R79) ... I also do not want future employers and such to be able to find all of my social media straight away and making judgements based on it.

(R150) ... I used to have a nerdy YouTube channel which I did not want my peers finding out about, so almost all of my online activity connected to that was under a different (screen) name

Figure 5: Pseudonyms — Tags and counts ($N = 286$) for responses to Question 5a: *Do you use pseudonyms on any social media platforms?* Responses which were in the negative or unclear have been removed.

Related to **separation**, several people used pseudonyms to **hide** ($N=8$) their activities online. This is distinct as it covers activities that they would like *no-one* to know about, rather than seeking to separate different identities. Most commonly, this had to do with pornography:

(R183) Yes, especially when using pornographic sites such as Chaturbate.

However, there were also examples of more general hiding:

(R398) I have to do things that people don't need to know about but I don't hide my real persona

Ignoring the 'yes' answers, the next most common reasons were **privacy** and **safety**; while these codes are related, there are some distinctions in the meanings we found. **safety** ($N=22$) related to a fear of repercussions spilling out of that particular online world. Some of these threats were specific ideas of violence:

(R168) As a person on the internet with (rather unpopular) opinions I find myself constantly subjected to pretty severe harassment such as very graphic rape and death threats, so I feel it would be safer to reveal little to no identifying information on certain platforms.

Many people were concerned about the idea of being stalked, of what might happen if people could find them 'in real life', while others had a general sense that one should be safe or

careful online:

(R184) Yes, I do, because I am concerned that people might stalk me if they know my real name.

(R169) ...tends to involve a lot of total strangers, so I feel I need to be more careful.

This is distinct from the responses concerned with a more general notion of **privacy** ($N=36$). This code was used for responses which simply mentioned privacy, or a desire for one's data not to be shared. This ranged from a passive sense of not wanting to share more than necessary to an active, explicit desire to maintain privacy:

(R151) ...I just don't feel the need to have that info on there at the moment ...

(R241) I use pseudonyms to maintain privacy and also to [The response cuts off here. The authors are intrigued as to what was coming next.]

Some users were also change names in order to reduce the ability of systems to **spy** on them, or share their data unnecessarily ($N=4$).

Not all uses of pseudonyms related to hiding or privacy. A significant number of people (**nohide**, $N=15$) explicitly stated that they were not using a pseudonym in an attempt to hide, while several carried on using pseudonyms out of **habit** ($N=10$).

(R383) I use pseudonyms because they're fun, I don't use them to hide my identity, I'm not batman.

A slightly larger number ($N=18$) used the pseudonyms to aid in their **discoverability**, by having a common name across several platforms, or to **conform** to the norms of the community ($N=6$). Similarly, several people ($N=10$) **reused** real-world identities, often in order to allow people they know offline to find them. There is often an exclusive component to these responses, that only the desired set of people will be able to find them:

(R223) Normally just a username which is based on my real name because if you know me then you will know it is me otherwise you would not

People also used pseudonyms to support **creative** activities, or simply for amusement ($N=8$). They also allowed the **expression** of parts of their personality without fear of repercussions ($N=9$), sharing of **intimate** content ($N=3$), and a presentation closer to their internal **identity**:

(R44) I really identify as a guy, so I go by a male name. Nobody IRL knows about that though. I do this cause I just want to be... Who I really am inside? Cheesy, but true.

Some people ($N=3$) had a dislike of their civil name, and simply wanted a change, or had a desire to create **bespoke** identities for certain activities ($N=3$).

Finally, a few people used pseudonyms in order to have multiple accounts to manipulate the sociotechnological **system** ($N=3$) – avoiding copyright issues, or tracking who sends spam mails. One person created a pseudonym to escape **discrimination**, one in the pursuit of **sex**, and three people explicitly mentioned Google+'s insistence on real names or merging accounts.

4.4 Personas

A total of $N=267$ responses were received for *q5b*, in which participants were asked if and why they had created any fictional personas for use on social media. 65% reported that they do not

Figure 6: Personas — Tags and counts ($N = 267$) for responses to Question 5b: “Have you created any fictional personas (e.g., characters, *alter-egos*) to use on social media?”. Responses which were in the negative or unclear have been removed.

or never have; 5% responded in an unclear manner or described pseudonyms rather than personas. Of the remaining third, the most common reason was for **creative** purposes ($N=21$), including to entertain themselves or others. Related to this are those who explicitly state they’re role-playing a fictional **character** ($N=11$) and those creating **bespoke** identities ($N=1$).

(R44) I just role-play characters I like to escape from my everyday hell hole.

(R256) I use another persona to have fun telling fictional stories.

(R443) I have a blog that I update in the voice of a character but thats for my own personal use as it’s helping me to write a book

(R449) I have and i did it because i created a fictional character and i wanted to give the illusion that the character was real

(R482) I did so to make fun of some naive friends on a facebook group.

The next most common response ($N=10$) was to experiment, including testing the reactions of others to different ages, genders or political views, or for self-exploration.

(R112) I use to when I was younger on tchat to see How people talk to different kind of people (male, female, younger, older etc...)

(R303) yes. many... i do this to role play different personalities online and sometimes learn more about my actual persona by doing so. i like the act.

(R461) ...I have created two alter-egos. One was a short-lived novelty account that posted in the voice of a fictional character, while the other is a member of a hate group whom I used as a kind of psychological experiment in empathy—by performing as a member of that group, I came to a fuller understanding of what compels their bigotry.

$N = 8$ responses were tagged with **separation**, where respondents created personas to separate work and social lives or posting of different content types.

(R381) Yes, to comment on Youtube, because I don’t want Google+ on my regular upload account.

(R444) i’ve got accounts to post on when i feel annoyed so that friends/family dont see and it doesnt affect their impression of me

(R492) Yes, I have 2 different twitter accounts that I use, one for general Fan base use which I am an overactive mad sloth and one which is for school people to think is my only one

Some users took on pseudonyms for **privacy** ($N=3$) or to aid their self **expression** ($N=3$) finding it gave them the power to give voice to parts of their personality:

(R371) Yes, it helps me be more confident and say things to people that I would otherwise be unable to say.

Social engineering was also a motivation (**soceng**, $N=3$), typically pretending to be someone new to gain trust or find out people’s private opinions:

(R473) I once created a fake persona to ingratiate myself with an online community and see what they were saying about me in private.

Finally there was one respondent with each of the following motivations: resistance to the system **spying** on them, or explicitly fighting the Google+ real names policy (**plus**); force of **habit**; presenting an **identity** closer to their ‘true self’; and for **sex**.

5. DISCUSSION

Examining themes common to all of the questions we analysed, several could be considered reflections of offline social practices. Impression management behaviours such as **playup**, **conform** and **mitigate** commonly occur in day-to-day life. The online performances aimed at impressing friends and attempting to diffuse awkward social encounters seemed largely analogous with their face-to-face equivalents.

Similarly, a number of participants attempted to **playdown** or not disclose problems they were having – they described their motivations as not wanting to worry others, or not wanting employers to find out. These participants are essentially using lying to manage how others perceive them, effectively giving them more control over their illnesses, rather than being forced to disclose them, and having to deal with potential consequences of that disclosure. This particular use goes beyond the *butler lies* phenomenon discussed previously, and instead enables control of psychological projection and public perception of self online.

Pseudonyms and personas, meanwhile, were commonly used as mechanisms for preventing *context collapse* [23, 2, 30], maintaining a *separation* of concerns between different facets of

respondent's lives. Identity was partitioned based on both the content posted and the intended audience. This included having separate Twitter accounts for personal vs. professional posts; 'secret' accounts used to interact with fandom communities away from the judgemental eyes of peers; and pseudonymous Tumblrs which allow the solicitation of advice from strangers regarding their non-parent-friendly **intimate** secrets. Consistent pseudonyms were reported as useful for allowing those *in the know* to track them across platforms (**discoverability**), or link certain aspects of their persona together whilst excluding others, without requiring the sharing of any personal details.

Whilst most people are told from a young age not to talk to strangers in the street, the even more uncertain nature of the audience of online interactions seem to make many of our respondents innately wary. Altering or omitting personal details was considered 'the done thing' by many, who either feared for their physical safety or just wanted to avoid nasty comments. Some had a sense that they would be stalked by strangers if they revealed their location, regardless of whether or not they considered their online activities provocative.

Lies in the form of impersonations, parodies, role-playing, or storytelling were used creatively to entertain others and alleviate boredom—just as joking around in person would do. The behaviours reported are extensions of ways in which people construct the multiple facets of their identity offline. This is consistent with findings reported by boyd following ten years of ethnographic studies of social media use by teenagers [3], that the primary attraction of social media to young people is the ability to claim a social space of their own, in which they can 'hang out' when restricted from being physically co-located with their peers. boyd argues that privacy norms have not changed as technology executives like Eric Schmidt and Mark Zuckerberg would have us believe, but rather young people are continuously evolving new ways to maintain much-desired control over social situations [3, 25].

Some of the reported behaviours serve to highlight differences between online and offline practices. While role-playing is used in the real world in order to help people work through difficult or novel situations, the malleability of identity on social networks enables participants a greater control over how they present. This allowed several people to put themselves in the shoes of others, to experience the treatment given to women, or the feeling of being part of a hate group. In these cases, the intention was clearly to deceive others, in order to get a 'realistic' experience, although the deception was carried out for seemingly benign purposes. In addition, a small number of respondents reported being able to project their true selves online in a way that they cannot elsewhere. Others could alter their identity to avoid discrimination, allowing an ease of engagement which was otherwise not available. This illustrates empowering potential of the Web, where the ability to control information about oneself can be a positive force for good.

Another observation from our study relates to how platform restrictions become barriers to the kinds of activities we described. Platforms can limit control over identity accidentally or deliberately, through policy or by constraining the affordances provided. In particular, it is clear that several of the deception strategies described were deployed in order to preserve safety, privacy, or separation of identities in the face of platforms that were designed to thwart such separation and/or anonymous use. Common examples include providing false attributes to platforms that required personal info "it had no business asking for" and creating separate identities where platforms provided no

means of opting out of advertising or tracking. Perhaps the most irksome to the participants of our study was the consolidation of YouTube and Google+ identity namespaces with the introduction of policies requiring the use of real names. Opposition to this policy gathered over 240,000 signatures in a petition in 2013 when the change was made [15], indicating the widespread desire to maintain separate, controllable identities. Examples of careful and deliberate control over public profile information on YouTube are documented by Guy et al. [18], showing that strategies for persona management continue despite attempts by Google to reduce the fluidity of identities of their users.

While it is apparent that many of the deceptions discussed are neither new nor malicious, and complement or mirror pre-Web forms of social mediation, we do not suggest that all such activities were wholesome or innocent. Responses in the **soceng** category included creating fake accounts to stalk an ex-partner or to test the faithfulness of a current one; gaining trust to see what people were saying about them behind their backs; and manipulating social situations for personal advantage. Additionally, several respondents admitted to manipulating technical systems, to gain access to protected resources in games, access age-restricted content and so on.

In addition to the behaviours here, there are many malicious or undesirable behaviours which were not reported in our study, but which are facilitated by the ability to create and alter identities. *Astroturfing* [13, e.g.] has become common online [43], with corporations and governments employing sophisticated identity management software to carry out large scale operations. Possibly the most famous of this is the "50 Cent Party", hired by the government of the People's Republic of China to post favourable comments towards party policy [40]. On a smaller scale, *sock-puppets* – multiple accounts controlled by a single person – are used to skew ideas of consensus and distort discussion in online societies, leading to attempts to automatically identify such accounts [6, 34]. Personas can be constructed for the purpose of *trolling*, whether it is overtly offensive in order to cause outrage or more subtle manipulation to trick people into wasting effort or taking caricatured positions, and correlations have been shown between enjoyment of trolling and everyday *sadism* [7].

Many of these activities are a form of *obfuscation*, in some way hiding the truth, polluting the data pool and diminishing trust. The ethical issues here are complex and contextual, with the viewpoints of different actors having considerable divergence [5]. However, we contend that the great majority of the – admittedly self-reported – lies detailed here were carried out for justifiable, defensible reasons, which contribute to the richness and vitality of the online social fabric.

6. LIMITATIONS

Among the limitations of the study, first, the self-report of lying behaviours may be different from actual practices for several reasons; retrospective bias effects may cause consistent under-reporting (e.g. "I think I am a mostly honest person, therefore I really must not lie that much"). A second reason that self-report is challenging here is that, due to the degree to which lying practices may be ingrained, there may be classes of behaviours that people may not consider, realise or think of as lying or deception at all. Indeed, a major class of butler lies were not even perceived as lies by participants of prior study [19]. In order to mitigate this effect, we iterated on the wording of the survey questions to try to elicit as wide a variety of relevant behaviours as possible. Second, as with all surveys, selection-bias effects may have affected the results; in particular, those that volunteered (or, indeed, took

any notice to begin with), were perhaps more likely than not to have a pre-existing interest in topics. This may have biased results towards those with opinions or thoughts on the topic.

Another limitation of this study is that it is reflective of only one specific demographic that we targeted; young, Western, social media enthusiasts comprising YouTubers and other “web nerds”, as these individuals have been shown to have complex, entangled online social lives [4, 28, 26, 27, 29]. As such, the kinds of concerns and experiences people reported may not be representative of other Web demographics; for example, some demographics may be less likely to maintain separate fictional personae online, or have any need to keep separate their social media fanbases. However, studies of specific online groups, such as gamers on MMORPGs [41] have demonstrated that demographics were considerably more diverse than previously suggested, particularly in specialised online communities [12]. To find out more about other groups, we plan to survey attitudes and practices of other demographics in follow-up work.

Finally, this study is an exploration of the kinds of untruthful practices carried out rather than an attempt to rigorously determine how often they are used. As such, we have not leaned heavily on any quantitative analysis – frequency counts have been used as an organising principle rather than a means of comparison or a fundamental part of our claims. Realistically estimating the frequency of actual lying and deception practice online presents significant methodological challenges, and, as a result, is out of our research priority for the immediate foreseeable future.

7. CONCLUSIONS & FUTURE WORK

In summary, this study found that people self-reported many routine kinds of lying, deception and omission strategies, reflecting a variety of needs and coping strategies for sustaining healthy, safe, and fun social interactions online. Only a small proportion of responses found deliberate attempts to socially manipulate others, while the vast majority corresponded to instances of trying to make oneself look good, maintaining separation among one’s personal, professional and other social roles, fit in with others, avoid harassment, avoid causing others’ worry, and to protect themselves from potentially harmful violations of privacy.

In immediate ongoing work, we are expanding our analysis to the other questions to identify platform- and demographic- differences in lying and deception practices. For example, despite not asking about platforms in q4 or q5, many participants mentioned adopting behaviours for specific platforms, for example, to separate their “intimate” content on Tumblr, or to mitigate potential privacy concerns with trolls on Reddit or YouTube. Longer term, we wish to further develop our taxonomy of lying, omission and deception behaviours in order to translate the needs they imply into implications for the design of platforms for future Web communities.

The fact that users must take active steps to circumvent the default behaviour of systems to maintain their online presence(s) suggests that current social media platforms have some way to go to provide a service that sufficiently affords the complex self-representation needs of users. The variety of benign and positive reasons users had for creating mistruths indicates that these representations should be supported in order to maintain vibrant online spaces.

8. ACKNOWLEDGEMENTS

This project was supported by the *Theory and Practice of Social Machines* project, funded by the EPSRC under grant

EP/J017728/1. We are grateful to all of our participants for volunteering a rich and varied set of responses, and to TomSka and Dame Wendy Hall for help with recruitment.

References

- [1] S. Bok. *Lying: Moral choice in public and private life*. Random House LLC, 2011.
- [2] d. boyd. *Faceted id/entity: Managing representation in a digital world*. PhD thesis, Massachusetts Institute of Technology, 2002.
- [3] d. boyd. *It’s Complicated: The Social Lives of Networked Teens*. Yale University Press, 2014.
- [4] D. Boyd. *It’s complicated: The social lives of networked teens*. Yale University Press, 2014.
- [5] F. Brunton and H. Nissenbaum. Vernacular resistance to data collection and analysis: A political theory of obfuscation. *First Monday*, 16(5):1–16, 2011.
- [6] Z. Bu, Z. Xia, and J. Wang. A sock puppet detection algorithm on virtual spaces. *Knowledge-Based Systems*, 37:366–377, 2013.
- [7] E. E. Buckels, P. D. Trapnell, and D. L. Paulhus. Trolls just want to have fun. *Personality and Individual Differences*, 67:97–102, 2014.
- [8] D. B. Buller, J. K. Burgoon, A. Buslig, and J. Roiger. Testing interpersonal deception theory: The language of interpersonal deception. *Communication theory*, 6(3):268–289, 1996.
- [9] J. K. Burgoon, G. M. Stoner, J. A. Bonito, and N. E. Dunbar. Trust and deception in mediated communication. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences (HICSS’03) - Track1 - Volume 1*, HICSS ’03, pages 44.1–, Washington, DC, USA, 2003. IEEE Computer Society.
- [10] C. Camden, M. T. Motley, and A. Wilson. White lies in interpersonal communication: A taxonomy and preliminary investigation of social motivations. *Western Journal of Speech Communication*, 48(4):309–325, 1984.
- [11] A. Caspi and P. Gorsky. Online deception: Prevalence, motivation, and emotion. *CyberPsychology & Behavior*, 9(1):54–59, 2006.
- [12] W. M. Chan and B. E. Dicianno. Virtual socialization in adults with spina bifida. *PM&R*, 3(3):219–225, 2011.
- [13] C. H. Cho, M. L. Martens, H. Kim, and M. Rodrigue. Astroturfing global warming: It isn’t always greener on the other side of the fence. *Journal of Business Ethics*, 104(4):571–587, 2011.
- [14] B. M. DePaulo, D. A. Kashy, S. E. Kirkendol, M. M. Wyer, and J. A. Epstein. Lying in everyday life. *Journal of personality and social psychology*, 70(5):979, 1996.
- [15] J. Doe. Change the youtube comment section back to its original form., November 2013.
- [16] J. S. Donath et al. Identity and deception in the virtual community. *Communities in cyberspace*, 1996:29–59, 1999.

- [17] N. Ellison, R. Heino, and J. Gibbs. Managing impressions online: Self-presentation processes in the online dating environment. *Journal of Computer-Mediated Communication*, 11(2):415–441, 2006.
- [18] A. Guy and E. Klein. Constructed identity and social machines: A case study in creative media production. In *WWW 2014 Companion*, Seoul, 2014. ACM.
- [19] J. Hancock, J. Birnholtz, N. Bazarova, J. Guillory, J. Perlin, and B. Amos. Butler lies: awareness, deception and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 517–526. ACM, 2009.
- [20] J. T. Hancock. Digital deception. *Oxford handbook of internet psychology*, pages 289–301, 2007.
- [21] J. T. Hancock, J. Thom-Santelli, and T. Ritchie. Deception and design: The impact of communication technology on lying behavior. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 129–134. ACM, 2004.
- [22] J. T. Hancock, C. Toma, and N. Ellison. The truth about lying in online dating profiles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 449–452, New York, NY, USA, 2007. ACM.
- [23] B. Hogan. The presentation of self in the age of social media: Distinguishing performances and exhibitions online. *Bulletin of Science, Technology & Society*, page 0270467610385893, 2010.
- [24] Z. Hussain and M. D. Griffiths. Gender swapping and socializing in cyberspace: An exploratory study. *CyberPsychology & Behavior*, 11(1):47–53, 2008.
- [25] B. Johnson. Privacy no longer a social norm, says facebook founder, January 2010.
- [26] A. Lenhart, J. Kahne, E. Middaugh, A. R. Macgill, C. Evans, and J. Vitak. Teens, video games, and civics: Teens’ gaming experiences are diverse and include significant social interaction and civic engagement. *Pew Internet & American Life Project*, 2008.
- [27] A. Lenhart, K. Purcell, A. Smith, and K. Zickuhr. Social media & mobile internet use among teens and young adults. millennials. *Pew Internet & American Life Project*, 2010.
- [28] S. Livingstone. Taking risky opportunities in youthful content creation: teenagers’ use of social networking sites for intimacy, privacy and self-expression. *New media & society*, 10(3):393–411, 2008.
- [29] M. Madden, A. Lenhart, S. Cortesi, U. Gasser, M. Duggan, A. Smith, and M. Beaton. Teens, social media, and privacy. *Pew Research Center*, 2013.
- [30] a. E. Marwick and d. boyd. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1):114–133, July 2010.
- [31] C. E. Naquin, T. R. Kurtzberg, and L. Y. Belkin. The finer points of lying online: e-mail versus pen and paper. *Journal of Applied Psychology*, 95(2):387, 2010.
- [32] A. Potts. ‘LOVE YOU GUYS (NO HOMO)’ How gamers and fans play with sexuality, gender, and Minecraft on YouTube. *Critical Discourse Studies*, (ahead-of-print):1–24, 2014.
- [33] M. Roig and M. Caso. Lying and cheating: Fraudulent excuse making, cheating, and plagiarism. *The Journal of Psychology*, 139(6):485–494, 2005.
- [34] T. Solorio, R. Hasan, and M. Mizan. A case study of sockpuppet detection in wikipedia. In *Workshop on Language Analysis in Social Media (LASM) at NAACL HLT*, pages 59–68, 2013.
- [35] A. Strauss and J. M. Corbin. *Basics of qualitative research: Grounded theory procedures and techniques*. Sage Publications, Inc, 1990.
- [36] S. Turkle. *Life on the Screen*. Simon and Schuster, 2011.
- [37] N. Van Doorn and L. Van Zoonen. Theorizing gender and the internet: Past, present, and future. *The Routledge Handbook of Internet Politics*, pages 261–74, 2008.
- [38] M. Whitty and J. Gavin. Age/sex/location: Uncovering the social cues in the development of online relationships. *CyberPsychology & Behavior*, 4(5):623–630, 2001.
- [39] M. T. Whitty. Liar, liar! an examination of how open, supportive and honest people are in chat rooms. *Computers in Human Behavior*, 18(4):343–352, 2002.
- [40] Wikipedia. 50 cent party, 2015.
- [41] N. Yee. The demographics, motivations, and derived experiences of users of massively multi-user online graphical environments. *Presence*, 15(3):309–329, 2006.
- [42] N. Yee, J. N. Bailenson, and N. Ducheneaut. The proteus effect: Implications of transformed digital self-representation on online and offline behavior. *Communication Research*, 2009.
- [43] J. Zhang, D. Carpenter, and M. Ko. Online astroturfing: A theoretical perspective. In *Americas Conference on Information Systems 2013 Proceedings*, 2014.