

“People Are Either Too Fake or Too Real”: Opportunities and Challenges in Tie-Based Anonymity

Xiao Ma
Jacobs Institute, Cornell Tech
New York, NY, USA
xiao@jacobs.cornell.edu

Nazanin Andalibi
College of Computing and
Informatics, Drexel University
Philadelphia, PA, USA
naz@drexel.edu

Louise Barkhuus
The IT University of
Copenhagen
Copenhagen, Denmark
barkhuus@itu.dk

Mor Naaman
Jacobs Institute, Cornell Tech
New York, NY, USA
mor@jacobs.cornell.edu

ABSTRACT

In recent years, several mobile applications allowed individuals to anonymously share information with friends and contacts, without any persistent identity marker. The functions of these “tie-based” anonymity services may be notably different than other social media services. We use semi-structured interviews to qualitatively examine motivations, practices and perceptions in two tie-based anonymity apps: Secret (now defunct, in the US) and Mimi (in China). Among the findings, we show that: (1) while users are more comfortable in self-disclosure, they still have specific practices and strategies to avoid or allow identification; (2) attempts for de-identification of others are prevalent and often elaborate; and (3) participants come to expect both negativity and support in response to posts. Our findings highlight unique opportunities and potential benefits for tie-based anonymity apps, including serving disclosure needs and social probing. Still, challenges for making such applications successful, for example the prevalence of negativity and bullying, are substantial.

Author Keywords

Anonymity; social media; self-disclosure; CMC; Secret; Wumii; Mimi

ACM Classification Keywords

K.4.3 Computer-supported collaborative work

INTRODUCTION

From masquerade balls to online forums and popular systems like Reddit and 4chan, the design of anonymity functions has

affected social interactions and dynamics in offline and on-line environments [11, 27]. Anonymity allows people to feel less constrained by the expectations of their everyday identities [3, 25], which in turn allows more candid and honest self-expressions [1, 43]. Anonymity also leads users to disclose more information [52], which may result in positive emotional outcomes for the person disclosing [45]. On the other hand, the lack of attribution and responsibility in anonymous environments can encourage malicious behaviors such as trolling [33] and cyberbullying [5, 6, 34]. It is no wonder that society in general, and HCI researchers in particular, have been fascinated with the topic of anonymity.

While early Internet platforms traditionally defaulted to anonymity or pseudonymity, the last decade has witnessed a “real-name”, complete-identity push by some social media platforms like Facebook, Google+, and LinkedIn. Over the last few years, though, a number of new systems have emerged that adopt a new model of anonymity that we call *relation-based* anonymity. These applications often build on mobile device affordances such as persistent user identification and metadata such as location, as well as information about one’s social ties through linkage to real-name social media platforms.

Most notably, *relation-based* anonymous applications such as Secret, Mimi, Whisper, Rumr and Yik Yak all provide mobile services to anonymously share content with people related to them, for example by physical proximity (*proximity-based*) or social affinity (*tie-based*) [37]. Secret, now defunct, was a *tie-based* mobile application when first launched in 2013: it allowed users to “share with friends, anonymously”. Yik Yak, on the other hand, adopted the *proximity-based* model that allows users to see anonymous posts within a 10-mile radius or within a given university campus. Notably, in all these services, individual posts by the same contributor cannot be linked to each other¹. The combination of relationship and

¹In March 2016 (two years after launch), Yik Yak introduced “handles”—permanent pseudonyms that users can choose to post with.

lack of persistent identity resulted in this new and interesting model of anonymity, and presents a wide set of new research questions within HCI.

While HCI research has extensively studied the topic of anonymity, we here extend that literature by providing a qualitative examination of *tie-based* anonymous applications. Past HCI research includes broad examinations of anonymity perceptions and practices across multiple services and systems (e.g., [28, 44]), as well as studies of anonymity on specific platforms [29]. In the latter category, HCI research has examined various platforms, including for instance anonymity on social media [12, 28]; fully anonymous persistent-identity topic-based forums [26, 40, 47, 59]; community ad-hoc discussion threads [1, 30, 31, 35]; proximity-based anonymity services [29]; platforms that allow selective anonymity such as Ask.fm [18]; and anonymity in open collaboration projects such as Wikipedia [20]. Most of these platforms feature anonymity models significantly different from our model of interest here, likely resulting in different dynamics and affordances than tie-based anonymity services. The fact that contributors are aware that their readers are within their friend (and friend-of-friend) circle, provides a different foundation than, for example, pseudonymous social media.

To explore people’s practices and perceptions around this new anonymity model, we posed the following research questions:

RQ1. What type of interactions do tie-based anonymous social networks afford according to regular users?

RQ2. What benefits and affordances do users of tie-based anonymous social networks experience?

RQ3. What are the perceived drawbacks and negative experiences by users of tie-based anonymous social networks?

To this end, we conducted a qualitative interview study with users of two apps: the US-based (and now defunct) Secret, and the China-based Mimi (now known as Wumii²). We conducted semi-structured interviews with users of these applications, and followed up with interviewees seven months later, after a period of significant change in both apps. In particular, during that period, Secret had changed from a tie-based to a proximity-based anonymous sharing model, and subsequently (after we concluded our research) shut down.

We present our results organized around five key themes. First, we present participants’ motivations and general practices in using these services. Second, we illustrate how identity management *does* play a role, even within these anonymous, no-persistent-identity applications. Third, we discuss participants’ widely used de-anonymization activities where they attempted to expose the identity of others on these services. Fourth, we show how the services provided some degree of intimacy and social support. In a fifth theme we consider the negativity and bullying aspects of these services. In addition to these themes, we consider the much-reduced use of the services after a period of six to seven months and the

reasons for the usage drop-off. Finally, the discussion section addresses a key goal of this work: providing an overview of the opportunities and challenges of tie-based applications exposed by this research, to inform the design of future platforms and services that support this model.

BACKGROUND AND RELATED WORK

The HCI community has long been interested in anonymous platforms and services, and how identity practices affect community and social dynamics. In particular, we discuss various types or models of anonymity that were studied in the field, including persistent identity systems, anonymity in ad-hoc discussion threads and mixed anonymity models. Before we expand on the specific types of anonymity, we briefly review literature on self-disclosure.

Self-disclosure in social network environments can be seen as driven by an extended functional process [10] that takes into account audience perceptions and identity representation, balancing benefits and risks. People try to maximize the rewards brought by self-disclosure (e.g., social validation and self-expression) [9, 14] and minimize the risks (e.g., concerns about sensitive information shared with third parties) [10, 16]. Under the condition of anonymity, people have been found to disclose personal information more freely, potentially because they feel less constrained by the expectations of others or perceive less associated risks of social sanction [3, 25]. Anonymity also allows people to share sensitive intimate narratives, and seek and provide support in stigmatized contexts such as sexual abuse or mental illness [1, 2], and explore aspects of their identity that otherwise would have been impossible [12, 56]. Most recently, an interview study of adolescents who use Ask.fm—a platform that allows selective anonymity—found that by strategically using selective anonymity, young people can build from social processes and work towards key developmental goals [18]. Particularly, this study argues that selective anonymity allows for higher levels of perceived authenticity, bypassing social expectations, learning about one’s self, managing identity, and initiating and growing relationships [18].

On the other hand, the lack of attribution and responsibility associated with anonymity can also encourage malicious behaviors such as trolling, hostile commenting [33], deception [24], and cyberbullying [5, 6, 34]. Research suggests, for example, that lack of eye contact is a contributing factor to disinhibition and one of the reasons for online negativity [33], where other research points out that positive attitudes towards cyber-bullying increase later tendencies of this behavior in anonymous fora [6].

Next, we briefly report on the different models of anonymous platforms and services that have been studied in HCI, including recent research on relation-based anonymous platforms.

Models and Dynamics of Anonymous Services

One model of anonymity is persistent identity systems—services in which the identity of the poster may be unknown to others but still consistent from one post to another. This includes network-based systems like Twitter, where the posters can choose whether to use their “real name” or not, as well as

²<https://www.wumii.org/>

topic-based systems (e.g., online forums, discussion groups, etc.) where participants have a persistent identity that may be pseudonymous. Previous studies on topic-based anonymous platforms for example, looked into how identity design impact communication dynamics on the platform [17]. Research has also addressed social support, most often through topic-based anonymous forums, from early studies of interaction practices on forums [47] to more recent studies of limitations to online medical support [26, 59]. Many of these studies illustrate that the anonymity within these forums affords more self-disclosure and emotional support [27]. At the same time, the posting dynamics there express higher accountability and somewhat more risk than the more ephemeral anonymity services.

Most systems with non-persistent identities are content-based, for example 4chan, Reddit and others. As one 4chan study showed [11], 90% of the posts made in the network are anonymous and cannot be linked to other submissions by the same individual. Such submissions arguably allow individuals to express themselves more freely, but result in increased negativity or need for moderation [32]. A significant issue in anonymous systems with non-persistent identities is the issue of source evaluation and credibility; in 4chan, the /b/ community uses shared context expressed in language and images that helps established credibility [11]. Trammell goes further to argue that the validation and approval practices on 4chan make it equivalent in some ways to a persistent user identity community [55].

A few recent examples mixed anonymity models by bridging persistent- and ephemeral-identity platforms. Recent research points out how “throwaway accounts” on otherwise single-identity pseudonymous social media forums enable users to switch identity temporarily [35] to be able to disclose particular information without repercussion. For instance, Andalibi et al. discuss the nuances of anonymity and suggest that sexual abuse survivors are more likely to use the anonymity afforded by throwaway accounts to disclose their experiences and seek support, and men are more likely to do so due to more stigma in the context of abuse [1]. Another recent study [12] looked at the support seeking and response behavior in a mixed anonymous/real-name forum: Facebook confessions boards. On these systems, people post anonymously (via a mediator) but receive replies by identifiable Facebook users. This dynamic lead to little negativity and an appropriate forum for discussing taboo topics and stigmatized identities [12].

As services that support a relation-based anonymity model are quite new, only recently research had started examining these platforms. Kang et al. conducted a qualitative study of users of these new services and emphasized how important it was for their participants to get social validation even from anonymous strangers, echoing early research on Usenet [29]. However, only four of their eighteen participants were Secret users, and the remaining were Whisper and Yik Yak users, services that rely on a proximity relation rather than a set of existing social connections (i.e., being connected through other social networks) [29]. Other recent studies have quan-

titatively examined data from the proximity-based anonymity service Whisper: One study compared Whisper’s and Twitter’s different anonymity models and found that users use anonymity to express their wants, needs, and wishes [15]. Another empirical study using 24 million posts on Whisper looked into structures of user interaction, user engagement, and network stickiness [58].

To summarize, in this paper we investigate tie-based anonymity services, where the perceived affordances, benefits and risks of participation are different than those of more traditional social media services, including anonymous ones.

TIE-BASED ANONYMOUS SOCIAL APPS

The tie-based anonymous social applications have several characteristics that set them apart from previous anonymous chat rooms and forums. Where traditional online forums and chat rooms are most often *topic-based* and support forum-wide identity, ephemeral identity systems make use of some sort of relationship between poster and reader, such as proximity or existing digital social ties. In this work we focus on the unique category of tie-based anonymous applications, where the anonymous application uses data from another social network where social ties are articulated, for example Facebook, or the phone’s address book. With this social network information available to the application, it can mark content as posted by “friend” or “friend of friend”—based on a match done by the application but not exposed to the user.

Secret, available from December 2013 to April 2015, was a mobile application that allowed users to share information anonymously with Facebook friends and friends listed in the user’s telephone contacts. Based on the friends list, the app displayed a stream of posts shared by the user’s friends, friends of friends, as well as popular posts by non-friends. Users would see a list of reverse chronologically ordered posts that include text on top of a solid background or attached picture, with the source identified as “friend” or “friend of friend” when appropriate. In addition to these posts, users could also view popular posts using the “Explore” tab on the application, providing a list of popular and nearby posts by non-friends.

In this version of Secret, users could “like” any post by clicking on the heart-shaped button but only comment on posts from their friends or friends of friends. Comments were anonymous as well, and commenting users were assigned a colored icon consistent within the context of a single post. As a result, there was no built-in mechanism to identify posts or comments made by the same user beyond a single post and discussion thread.

The Secret app changed functionality and look during its lifetime and our study. In December 2014, the app transitioned from tie-based to proximity-based model. Like the popular app Yik Yak, Secret showed anonymous posts from people nearby, rather than from friends. In April 2015, shortly after our last follow-up interview, Secret had shut down.

A Chinese app called Mimi (“secret” in Chinese) was launched in late March 2014. The Mimi app is very similar to the earlier version of Secret, to the point of being considered

a clone. Despite this similarity, Mimi did include additional functions such as private messaging, allowing users to send direct messages to the original author of a post. At the time of writing, Mimi is still available and active under the name Wumii. It is important to note that we included Mimi in the study to increase the generalizability of our findings, rather than to explicitly study cultural differences.

METHOD

Our study was based on in-depth, semi-structured interviews with users of Secret and Mimi, conducted in August and September 2014, with follow-up interviews six to seven months later. Our aim was to investigate uses and practices around the applications, as well as attitudes towards content and anonymity in these new tie-based platforms. We asked questions about participants' current experiences within and around these applications, with the goal of documenting examples of posts and reactions to these posts. We asked about types of content and types of interactions taking place as well as how they thought anonymity affected their interactions on the platform. Seven months after the first round of interviews, we followed up with a subset of participants via short interviews or email correspondence (we approached all but only a subset responded). By then, Secret had changed the posting mechanism to focus on showing nearby posts ("proximity-based"). The second-round interviews focused on the changes to the participants' practices through the past seven months, in particular given the changes to the app in the case of Secret.

Participants

We interviewed 15 participants: six users of the Chinese app Mimi (three men, three women) and nine users (six men, three women) of the US-based app Secret. Participants' ages ranged from 21 to 27 for Mimi, and 22 to 37 for Secret. Participants held a diverse set of jobs, from pre-school teacher and digital manager to advertisement agent and public relations. The Chinese participants all lived in China and Secret users were located in the US. To recruit the participants, we posted fliers at coffee shops in a major American metropolitan area, performed snowball sampling through social networking sites, and approached users who posted about these apps on Twitter and the Chinese-language weChat (through direct messaging and forwarding recruiting messages). The interviews with the Chinese participants were conducted with a Chinese interviewer (one of the authors, a native Chinese speaker) in Chinese and were translated to English by the same person. Participants had to have used the applications for at least two weeks at the time of the interview; most had used it for several months when we conducted the first round of interviews. The first round of interviews lasted 30–60 minutes. The participants were compensated \$20 for their time. All interviews were conducted over the phone due to geographical distance.

When approaching the participants again seven months later, we were able to conduct follow-up interviews with five of the Secret users, and received five written responses from the Mimi users. For the remaining four Secret users, three did not respond and one declined to be interviewed again. The

second round interviews with Secret users lasted 10–40 minutes depending on if they still used the application or not; from the participants who responded, only one was still using the application, four were not, but all reported using the app for months after the first interview. As the Mimi app did not change as significantly, we limited the interaction with the Mimi participants to an email inquiry where we asked if and how they were still using the application. In Mimi's case, two participants were still using the application, but four were not, citing that they have used it for a number of months after the first round of interviews.

The recruitment process and study protocol were reviewed and approved for the inclusion of human subjects by an Academic Institutional Review Board (Protocol #1406004763). All the names of our participants were anonymized for analysis. In the results we identify quotes with a code indicating the app used, gender and age of the participant. For example, SF36 is a Secret user, female, who was 36 years old.

Analysis

The analysis of the data was conducted through qualitative categorization and thematic analysis. After the interviews were transcribed and translated into English (when needed), one researcher independently coded all interviews for major categories. Then a second researcher conducted more detailed analysis using sub-categorization and constant comparison [53]. Based on our research questions, we then highlighted relevant mentions in the transcripts, and grouped these categories into major sub-categories iteratively. Comparing the original categories and the sub-categories to each research question, we determined the most important findings that we elaborate on in the following sections. The follow-up interviews were coded and analyzed later but in close connection to the corresponding original interview.

Due to the similarities in terms of functionality of the Chinese app Mimi and the majorly US-based app Secret, we analyzed the two sets of interviews collectively. Our interest was not in cultural differences per se, but rather to sample across two cultures. We closely watched out for potential cultural differences emerging from the interviews. In the end, we only observed minor differences. We therefore continue to describe our findings collectively, and only point out the few subtle cultural differences in the later part of the paper.

RESULTS

We now discuss the results from the first set of interviews. After briefly describing the general motivations and uses of the two services, we explore five themes: Motivations and general uses; the need for identity management in this type of anonymous social media; the de-anonymization that in fact took place; how the app fostered intimacy, support and encouragement; and finally, how users dealt with conflicts inside the app.

Motivations and Uses

This first set of findings echos and enforces the findings of Kang et al.'s investigation of relation-based anonymity apps, which focused on location-based platform but included some

users of a tie-based platform (Secret) [29]. We report on these findings briefly below, to focus on angles and findings specific to tie-based platforms which were not highlighted in earlier research.

One major motivating factor participants cited for using the apps was entertainment: five participants mentioned checking the app when bored as a kind of distraction and mostly glanced over it for locating humorous and funny posts. One user reported that *“it’s sort of my morning routine. I wake up and check my phone; I look at my email, check Twitter, and then I’d go to Secret”* [SM24]. Despite the very different model for anonymity, interviewees sometimes described their motivations in relation to other anonymous social media such as Reddit. One participant talked about Reddit being *“completely different”* in terms of the motivation for using the platform: *“Secret... connects you with friends, and it’s more personalized type of secrets that people are uploading”* [SF36].

A key motivation was illustrated through participants describing the platform as a place where people can “let it out” without having their names attached. *“It allows and enables you just say that stuff that’s on your mind. It may not make everyone happy”* [SM27] or *“things that I would like to tweet, but I don’t want my name attached to it because people might have judgments about that”* [SM24]. One interviewee specifically mentioned anonymous apps as a platform for self-disclosure that avoids social sanction: *“I think everyone has the desire to self-disclose. They want to express the most indecent side of themselves, things that your society doesn’t recognize or value [...] Indicating that you are good at sex is a positive image for you; but you cannot say, I am asexual (on social media) [...] These things can only be posted on anonymous communities”* [MM21]. Other participants reported being able to post untrue content. When asked the reason behind that, one participant responded: *“to see what would happen... and what kind of responses I would get, or what kind of comments or what kind of reaction I could elicit”* [SF36]. The curiosity of the reactions to extreme posts sometimes trumped the social norm of posting “real” secrets and made this type of social probing fairly common.

Several of the participants indicated that posting to these apps allowed them to avoid complications from their friends while still gaining the benefits from social interaction and validation. One participant reported posting a secret that received around ten likes: *“I met my ex boyfriend a few days ago. I got an email from him just now. I deleted it, without reading”* [MF24]. One of the reasons she did not post a particular post on her real-name social networks was that she only wanted to post to release her feelings but did not want anyone to know it was her. Posting under real-name would have led to her friends asking what happened and to “overthink”. These statements illustrate previous findings on the benefits of self-disclosure [16, 45] and earlier findings in the study of Kang et al. [29]. We return to this theme in the discussion.

Motivations for recurring use grew together with the size of the participant’s Secret network. Some participants commented on how the size of their network determined the visibility and readership of posts, and hence significantly affected

their experience on the platform (and the motivation for using the app). As one participant put it, *“If you don’t have a lot of friends, your post is only getting read by however many people in that circle...”* [SF36]. The enjoyment of consumption was also limited with a smaller network, for example: *“[In the beginning] I would go on maybe once a week just because there wasn’t very many updated posts for me to see, so there wasn’t much action happening”* [SM24]. While the value of the network grew with the size of it, it was not simple to “recruit” friends to use the app, precisely due to the app’s anonymous nature, but participants still reported occasionally asking their friends if they were using the app or not.

Identity Management

Six of the participants explicitly stated that the reduced need for identity management was something that made the application in question different. One person for example reported that *“on social networks where people know your identity, [...] you will have a drive to construct your self-image, and intentionally post things that are beneficial [...], that adds to your high-end image construction”* [MM21]. Other participants differentiated the identity aspects of Secret from other anonymous services: *“on [other anonymous social media], you have an identity on the entire site, or multiple identities if you’re logging in with multiple accounts. But no matter what you do [there], no matter what the thread is, it’s the same identity”* [SF22]. Participants reported that the ephemeral identity on the Secret app helped them to not be concerned about identity management since there was no pressure to live up to their “real self” as represented on other social media. This freedom did not necessarily mean that the users presented a false image. SF24, for example, explained how she felt she was representing her true self on Secret: *“I think I’ve never pretended to [...] be anything that I wasn’t or anyone that I wasn’t, and I feel like I guess I’m the most myself on Secret”*.

Interestingly, identity considerations still applied in some cases as participants used the text of the post to hint at, or further obscure, their identity. Interviewees reported crafting their posts carefully, based on their desired audiences: sometimes explicitly targeting their friends or expecting friends to be able to identify them, and at other times making sure to avoid identification by friends. One participant for example explained *“when I’m feeling depressed [...] I don’t want people to know who I am, so I intentionally hide my speech habits. For [other types] I don’t really care so I don’t change anything in particular”* [MF22]. In doing that, these participants took full advantage of the non-persistent nature of identities where different posts cannot be associated with the same author: *“If you present yourself as the same in your daily life, like funny or hilarious, [...] it’s very easy for friends to guess who this is; but if you post something that no one knows about, they couldn’t guess who that is, even if you post a lot of secrets”* [MF24]. In summary, participants were articulate about identity management still being important to a certain degree, perhaps because they were aware of the possible process of de-anonymization, a practice we discuss next.

De-Anonymization

We identified two types of “identity hacking” that participants were aware of. We name them “hard-hacking” (i.e., using technical approaches) and “soft-hacking” (de-anonymization via social engineering and manual procedures). Most participants had no concern for hard-hacking, saying for example: *“I’ve read their articles about how they sort of anonymize things, and so I feel very comfortable with sort of the technical aspects behind things, and how they set up their system”* [SM24]. However, eight participants reported concerns or at least an awareness of “soft hacking”, for example saying that *“some friends could easily figure things out by thinking about it for more than half a second”* [SM24].

Actual practices of soft-hacking by participants were prevalent as well. Seven of the participants expressed how it was “half the fun” or even a game to figure out who was the poster of an item. They used guessing, their knowledge of their friends, or language and speech cues to do so. A participant said he *“knew exactly who it was because six months prior. . . , my friend said the exact same quote to me. . . ”* [SM24]. Participants also practiced de-anonymization, with fairly sophisticated practical methods. Participants reported, for example: (1) trying to figure out who in their contact list was using the app; (2) sharing screenshots to others via text messaging or other social media; (3) comparing (and triangulating) friend lists with other friends who used the app. In contrast, several participants were not motivated, saying that *“while it’s not very difficult to identify who posted the secret, no one has the extra time to actually do this, and nobody would believe it anyway”* [MM21]. When asked if there were any differences between a post from a “friend” or a “friend of a friend”, MF24 responded: *“Of course, I definitely care more about my friends’ posts [. . .] If it’s something shocking from a friend, I am definitely interested in figuring out who that was.”*

Intimacy, Support and Encouragement

As noted above, participants certainly observed and reported posting more intimate content, due to the disinhibition effect of anonymous platforms [54]. Moreover, participants understood (and often reported witnessing) the benefits of such disclosures. One participant for example commented: *“Those things that people are too shy to discuss in real life but truly want to know more about, not involving attacking others, are actually beneficial”* [MF23]. Others said that they would not share the posted content with friends in person if they were worried about being judged. Relatedly, one user described the anonymous sharing as *“almost like gossiping without the negative effects of it landing specifically on someone. . . ”* [SF36].

Ten of the participants described finding encouragement and validation for the sentiment shared on Secret or Mimi, for themselves or for others. One participant for example reported that she posted on Secret about being afraid to go to the gynecologist because of a past sexual assault, something she could not talk openly about. She received several comments and said: *“I think most of them were friends, and they said something encouraging along the lines of ‘it’s important for your health, and I’m really sorry that that all happened. We’re here for you if you reach out to somebody to go with*

you.’ Just trying to help, ease my fear, more or less.” She also received more likes than comments, and attributed that to *“they don’t really know what to say, but they’re just trying to show support or affirmation of some kind”* [SF24].

A more subtle benefit of the tie-based anonymity was a better understanding of attitudes and opinions among one’s social ties, or at least a perception of such understanding. In eight of the interviews, participants reported that they were able to gain a better understanding of their collective group of friends, which in effect influenced their face-to-face social behavior. One participant told us how he had no idea that some of his friends were in fact questioning their sexual orientation: *“For example, [. . .] I had no idea that some people in my friend network questioned their sexual orientation. And they’re afraid to talk about it in person, but they’re okay with talking about it in sort of this anonymous aspect. [. . .] so I think it’s great for me to understand my group of friends better as a whole and maybe next time when I’m in a social gathering, I won’t make any comments about people’s sexual orientation or anything like that. Big [. . .] friendship strategy decisions on my end”* [SM24].

Sometimes, the experienced benefit was to learn about an agreement with others within the anonymous tie-based network. Some very popular posts had a shared sentiment of surprising recognition, for example as described by one participant: *“[. . .] you can find someone who says some random things, and you’ll realize [. . .] (s)he feels the same or is also like this! This is very interesting and very unique”* [MM27]. Another participant reported that one of his posts got three thousand likes and four hundred comments in a few days, saying it really must have *“struck a chord”* [SM27]; he reported receiving comments like *“yep, I do the same thing.”*

People reported being motivated to explicitly provide support. One participant for example said that *“if it can help others, giving one or two sentences and be the audience of that person’s thoughts wouldn’t be a bad thing. After all, they are friends, or your friends of friends. If someone has problems, it’s nice if you can comfort or cheer them up a bit”* [MM27]. On the other hand, this knowledge of ties, even remote, can make negative comments even more hurtful. One user reported stopping posting secrets after receiving very negative responses from several of her quite emotional posts, a topic we now turn to.

Debates, Conflicts and Bullying

As may be expected, postings and comments on these services often lead to debate and conflict. Participants reported that highly contentious topics could lead to non-supportive and judging comments: *“[. . .] these battles [can be really] drawn out and obnoxious to a certain point because both of them are just on such extremes that no one was really having a really worthwhile debate”* [SM27]. Another participant described: *“When there are two different opinions on certain things, and some people are just not very tactful in the way they say things. So if somebody posts something about being cheated on, then there will be commenters who will come in and say, ‘oh you always deserve a second chance’, and some others would say ‘no, once a cheater, always a cheater’. And*

then some of the commenters will fight back and forth about their opinions.”

Although participants did not admit to bullying or directly contributing to conflicts, as expected with a fairly self-selected set of participants, some confessed commenting in an impulsive way, saying things they did not really mean. One participant explicitly admitted commenting un-constructively when he was not in a good mood. Overall, though, participants often claimed that they defended attacked posters from their attackers, and that they reported or called out offenders when inappropriate content was posted.

Participants were exposed to and well aware of negativity in these apps, and were often unsurprised when it occurred, recognizing the anonymity dynamics. *“Whenever you don’t have to be held accountable, I think, unfortunately that there are other people who will use that to hurt people. [...] You tend to lean towards doing something that you probably wouldn’t if you could be held accountable”* [SM37]. This behavior is often referred to as “trolling” behavior, such as degrading and objectionable comments in response to others’ content. Participants were all well aware of trolling on these services, with thirteen of them mentioning the negative characteristics of some of the posts. One participant for example told us about comments that *“say really nasty, negative things”* [SF36]. Another participant specifically mentioned threads bashing an identified person on Secret: *“It’s like a train wreck... it’s a disaster when you look because you just can’t believe that people are this terrible”* [SF22]. Participants reported flagging and reporting posts and comments, as well as calling out people within the social network.

Finally, the negativity on the platforms can lead to decreased engagement and disclosure through the platform. According to one participant, *“Most of the times there were rumors, scandals, made up stories and especially this abusing others or something like that. We say in Chinese ‘negative energy’. It makes people frustrated and disappointed sometimes. [...] So yes after one or two uncomfortable experience I chose not to post on Mimi”* [MF21].

Revisiting: Persistence of Use and Changes in App Functionality

As we mentioned earlier, we contacted participants for a second interview during a period between six to seven months after the first interview. We were primarily interested in changes to the the participants’ practices during the past months, particularly in respect to the changes in the applications. This allowed us to get a longer term perspective of use of tie-based anonymous applications. Between the first and the second interview, the Secret app had gone through significant changes: it transformed into a primarily proximity-based app (much like Yik Yak). Mimi, on the other hand, kept the original tie-based concept, but had added a few new features, such as eliminating the option to see if a post came from friends or friends of friends.

For the second round of interviews, all of the Secret participants reported a radical change of content over the last

months where bullying and name calling had become prominent alongside adult content, reinforcing the findings of Kang et al. [29] with longitudinal evidence. One participant mused that it was *“a sad commentary about humanity that as soon as we get to be anonymous we abuse it”* [SM37]. Participants reported increased abusive and bullying behavior on the app, and one recalled a heated conflict within her own circle of friends that had taken place between the two interviews. Several participants reported incidents involving posters identifying others *within* the posts. Only one of the five participants we talked to was still using Secret; the other four particularly reported the reason for having left to be the negative and bullying postings. One participant described how she kept seeing her friends post unpleasant content, starting to specifically name others and how, even if she called them out on it, they continued. She said she gradually stopped using Secret until one day she just simply deleted the application from her phone.

Two of the participants mentioned they were initially enthusiastic to have the app change its feed logic to not only proximity-based but also letting them peek at other cities. A participant who were still using the app, for example said *“I like that there is more on the app, there weren’t much, [I] got bored before. Now it is much better and allows us to chat directly, it makes it a lot more interesting. I just think the entire concept behind Secret... now I can see what people are thinking”* [SF36]. She felt that the New York and the San Francisco feed were more interesting than her local Miami feed, which seemed a bit “childish.”

One of the main changes in the Mimi app was that the app no longer showed explicitly whether posters were friends, or just friends of friends, thus making de-identification more difficult. All but two of our six participants reported having given up using the app with different reasons: Two felt the application did not provide enough content and entertainment. Two others tied their disengagement to the anonymity change: not being able to distinguish between friends and friends of friends, the de-anonymization was limited and the app was not as fun.

DISCUSSION

Before our discussion, we note a few cultural differences in relation to tie-based anonymity systems—differences that were not readily apparent from our study. As an investigation into this topic, we included Chinese and US participants in order to be comprehensive rather than to be comparative. Indeed, we analyzed the data with an eye for differences in how the themes were expressed by participants of different cultures but no such differences emerged. Of course, the fact that the Mimi app is still active, versus the disappearance (and preceding adaptation) of Secret from the US market, may by itself suggest cultural differences, but these are a matter of another, more targeted investigation.

We now turn to our discussion where we highlight opportunities and challenges for tie-based anonymity applications. We use these opportunities and challenges to motivate several ideas and implications for future tie-based anonymity platforms. We acknowledge that some of the opportunities we

identified *can* also be problematic depending on the specific implementation, and that finding ways to design for the trade-offs between the benefits and risks of tie-based anonymity is a fruitful area for future work.

Opportunity: Self-disclosure

Based on our findings, we believe the self-disclosure opportunities with tie-based anonymity goes beyond what is afforded by other anonymous services, particularly due to the connection with friends.

As recognized by past research on self-disclosure, the disclosure itself may provide significant benefits [16, 45]. As is also known, under the condition of anonymity, users tend to disclose more, a phenomena known as the disinhibition effect [54], which we also identified in our participants' remarks. Furthermore, our interviews suggest that free from identification concerns, participants perceive less associated risks of social sanction or issues of self-representation, similar to earlier findings [3, 25]. As is the case for other anonymous systems [12, 29], tie-based systems thus allow greater self-disclosure than other social applications, as well as provide grounds for self-expression and relief.

The core value, we believe, lies in the potential reward for self-disclosure in tie-based applications, which is potentially more substantial than in other anonymous applications. Self-disclosure is a function of both risk and reward [10, 41] where rewards may include, for example, social validation and self-expression [14]. As mentioned above, the opportunity for self-disclosure in tie-based anonymity applications is strengthened through the reduced risks of being identified. Importantly, it is also strengthened through the heightened reward (compared to other anonymous platforms) given that the post is likely to be seen by a subset of one's social ties; despite trends of negativity and mean behavior, participants reported being able to receive support and encouragement based on their disclosure. In other words, while anonymity in general allows people to explore aspects of their identity that otherwise would have been impossible [40, 56], tie-based systems allow one to do that *in the context of their social ties*. Or, in Goffman's terms [22], tie-based anonymity may allow backstage-like performance in a front stage-like context.

Opportunity: Extending Social Behavior

Beyond self-disclosure, our findings highlight an opportunity for extending social behavior through tie-based anonymity services, providing social affordances that go beyond identifiable media and other types of anonymous systems. It was clear from our participants that they, or others on these tie-based services, were engaging in social behavior and extracted social insights about their social ties in a way that is not readily attainable from other services or classic social networks.

First, the opportunity to perform social probing, with little risk, and in the context of friends, is a powerful and unique affordance of tie-based anonymous services. Our findings suggest that users can go beyond self-disclosure to post messages, whether true or false, that are intended to elicit reactions or "test the waters" with their group of social ties.

This opportunity is available in anonymized settings, where an individual is not at risk by posting the information, and by replying to it with honest attitudes. Furthermore, only in tie-based anonymized settings people can post such message to get feedback from and exposure to friends and social ties. Prior work on identity transition has similarly found a need for people to be able to test out aspects of their identities that they may feel vulnerable about in online spaces [23]. In doing such posting, it may fill a social probing role similar to one of humor [19]; relatedly, tie-based applications allow radical ideas to be tested in social settings. Even here, the expression can take a slant that is sometimes playful and more tentative.

Second, and relatedly, is the opportunity for an individual user of tie-based anonymous services to develop a better understanding of their group of friends and social ties using information shared on the application, even without explicit probing activities. Some participants in our study hinted at a potential for a better window into their friends' attitudes, opinions, and perhaps even actions due to the properties of the tie-based services, aligned with the fundamental human need to belong [7]. Note that these impressions may or may not be warranted (e.g., as only a minority of most participants' friends were active on these services), and the understanding may not even be accurate (e.g., as participants themselves noted, information may be untrue). Nevertheless, the opportunity to be—or merely to feel—more informed about one's social ties is significant. We emphasize that this affordance is unique to tie-based systems. Social information processing theory [57] has long considered how people slowly develop mutual impressions and understanding over computer-mediated communication. Here, the learning is not about a specific connection, and definitely not an identified one, but rather about a set of social ties. A critical difference from other types of anonymous groups is the tie-based setup, where each individual is connected to their own social ties, as opposed to the topic- and content-based, or location-based anonymous groups or communities in other anonymity models.

Opportunity: Anonymity at Will

Finally, an opportunity that has yet to be widely utilized within social networks is the new mechanism for "anonymity at will". We saw from our interviews the users' ability to construct their posts in relation to who would be reading them, thus providing the level of anonymity that they wanted each post to have. Sometimes they were able to be completely anonymous, but they felt they could craft posts to signal identity to some subset of readers or all. Ellison et al. suggest that particularly for adolescents, being able to selectively engage in anonymous interactions with a group of known peers has, key implications for identity and social development [18].

Social steganography [39] refers to a practice of hiding secret messages in plain sight, where unwanted parties are not aware of their existence, or are unable to penetrate their actual meaning. This sharing practice is not uncommon in other social media, where for example location is revealed but only to the level of detail that an inside circle of friends would under-

stand [4, 36], or posts are “coded” through culturally specific references, inside jokes and other linguistic tools [13].

In a slightly different way, tie-based anonymity services create new avenues for people to transmit sensitive messages to friends: here the content is exposed but the *identity* of the poster may be only known to some specific target set of individuals through the crafting of the message. We term this practice “identity steganography”. As such, any new avenue for communication could serve an important social role [13]. For example, users can use identity steganography to seek social support while preserving privacy, especially in tie-based anonymity as friends who have the relevant context may know the identity of the original poster and reach out privately.

We discussed opportunities and potential benefits, but of course, there are darker sides to tie-based anonymity. Next, we discuss some of the key challenges.

Challenge: Negativity and Bullying

As is the case in other anonymous platforms, negativity and bullying are a significant challenge. Throughout the life of anonymous or pseudo-anonymous networks (i.e., site-wide user names, nick-name based networks), “trolling” can significantly curb participation and constructive contributions [33, 34]. Participants in our study also described widespread negativity, getting even worse later on in the lifetime of the Secret application. This negativity increase could potentially be due to the fact that anonymity accentuate the interchangeability of group members, making trolling behavior “catch on” and converge to become the social norm [46]. While our results do not explicitly hint at that, there is a further risk of particular forms of abuse when the negativity is expressed by a friend, from the comfort of their anonymity cover. As shown in our findings and by Kang et al. [29], this negative behavior is likely the reason for a decline of the use of these applications over time.

While participants reported “flagging” certain posts and otherwise trying to prevent abuse in the platform, this issue remains a challenge for services that would require significant resources to address. A sliver of hope lies in the fact that in mobile-based applications (which relation- and tie-based anonymity services tend to be), a robust user identifier is available to the service (i.e., an identifier of the mobile device used to access the app) and it is hard for users to create a new access avenue. This feature makes banning users, at least theoretically, more plausible than for other anonymous services, for example on the Web. However, given the amount of potential interactions, moderating, flagging and banning content is likely to remain a challenge for any popular tie-based anonymity platform.

Challenge: Rumors and Gossip

Beyond negativity, the lack of accountability also transforms tie-based anonymous services into places where rumor and gossip can be spread easily. Through our interviews we also uncovered a general perception by our participants that the two secret apps were rife with rumors and gossip. It has been long argued that gossip has a social function [21, 42]; gossip

provides ways of making sense to help us cope with uncertainty and anxiety [49, 50], and can serve to convey information about the society and culture [8, 60]. Interestingly, “gossip tends to have an ‘inner-circleness’ about it, in that it is customarily passed between people who have a common history or shared interests” [51], aligning gossip outcomes well with tie-based anonymous applications.

On the other hand, social psychology defines rumors as “a statement whose truth-value is unverifiable or deliberately false” [48]. Rumors thrive in secret environments, and often has significant negative outcomes [50]. The challenge of controlling the spread of rumors on social media platforms has been substantial. Even with real-name or persistent identity on social media, where one can verify or identify the source of information, rumor practices is an open issue and an active research area [48]. In tie-based anonymous applications, this evaluation is much harder or impossible. Even worse, individuals run little risk in spreading rumors on these platforms. Participants felt that to be engaging, the content needed to be believable but as we pointed out above, the tie-based services often offered content that was “too fake or too real”. The uncertainty, and inability to evaluate content, could be a challenge for adoption and continued use.

Adding to the challenge is the fact that due to the unverifiable nature of rumors, hurtful rumors and gossiping may be harder to moderate than negativity and bullying, which are more often immediately identifiable.

Challenge: Critical Mass

One challenge for any social media application is critical mass and adoption [38]. Many social media services rely on known ties and existing social networks for diffusion, for example via mechanisms to invite friends to join. To have enough compelling content from one’s friend, tie-based anonymity services also need to reach scale in adoption. This presents a significant challenge for these systems, as they cannot rely on traditional diffusion techniques: an individual cannot reveal their own identity to invite people to join the service. Of course, due to the anonymity and lack of persistent identity, one cannot create ties directly on these platforms either, as one would for example on Facebook or Twitter. As a result, tie-based anonymous platforms could find it hard to grow the network, and hence limit the available tie-based content and audience for any single user. This challenge had direct impact on the participants in our study, whom could not get enough content to keep them entertained and engaged in the apps. While proximity-based anonymous services could simply expand content by extending the radius of relevant posts, in tie-based services it is unclear whether the context (i.e., social radius) can be extended meaningfully (one might not care to read updates from a friend-of-friend-of-friend). All in all, without the ability to directly invite one’s friends or create new ties on a tie-based anonymity service, the availability of content on the platform would be quite sparse.

Design Implications

We provide several design implications and considerations for providing tie-based anonymity services. Importantly,

some of these suggestions are not for stand-alone tie-based services, but instead argue for inclusion of similar features in other social media platforms.

Anonymity flashes. One way to allow for the benefits listed above while perhaps circumventing some of the challenges is to directly build tie-based anonymity features into the experience of identity-based social network services like Twitter or Facebook. For example, Facebook could allow users to post one anonymous post per day (an “anonymous flash”) that would be broadcast—without attribution—to the user’s network. Deep integration into such social network service can allow comments and replies to the post to be identified (for instance as described in [12]) or similarly anonymous, each option provides different benefits.

Guessing game. Given the magnitude of de-anonymization attempts, tie-based services or functions may consider including a formal mechanism where, via private messages, individuals can be queried and admit (or perhaps prove) that they posted any specific content item (“Alice, is that you?” queries Bob). In “pure” tie-based anonymity services even the fact that individuals participate in the system is obscure, which would make such mechanism difficult to implement. However, such mechanism would be easier to implement if the tie-based system is integrated into a persistent identity platform like Twitter or Facebook.

Minimize de-identification risks. To make it harder for users to engage in “contact list hacking”, future applications should make permanent the app content shown to the user, and not dependent on the user’s current set of followees. In other words, once a user sees a post, it always remains on their list, preventing the option of hacking your contact list by removing contacts one-by-one to figure out who made one specific post.

Curbing negativity. Given that anonymity and loose social norms can quickly cause tie-based anonymity services to degenerate into a hostile and negative environment, it is critical to offer more comprehensive support to battling this type of content. Immediate flagged content removal (delete-then-verify) and other more proactive techniques may be a better approach. One can perhaps consider that users posting negative comments slowly lose their privileges (e.g., their anonymity) on the service, essentially circumventing outright banning but limiting the opportunity for negative postings by any one user. Note that the main enablers of these tie-based systems—the fact that they are tied to a specific ID, often a mobile phone identifier, and even when it is not exposed in the service—should be useful for banning and restricting users, compared to other anonymous services like news commenting, or online forums.

CONCLUSION

In this paper we provided a qualitative in-depth examination of user practices in tie-based anonymity services, and showed why the affordances provided by these services form a new model of anonymity in computer-mediated communication. Such tie-based systems offer opportunities for self-disclosure that are not available in “real-name” networks or persistent-

identity anonymous networks. Furthermore, these services can provide an environment where people can test the waters with their friends, exploring attitudes and possible reactions for complicated situations in their lives. We also highlighted a set of challenges for tie-based anonymity systems. Without having to rule whether tie-based anonymity services have any merit or potential for broad appeal and mass adoption as stand-alone applications, we believe our design considerations may be useful for providing the benefits of such services in the future in various settings.

ACKNOWLEDGMENTS

This research was funded by AOL through the Connected Experience Lab. We would like to thank Raz Schwartz and Funda Kivran-Swaine for conducting interviews with Secret users, as well as all the participants for their time.

REFERENCES

1. Nazanin Andalibi, Oliver L Haimson, Munmun De Choudhury, and Andrea Forte. 2016. Understanding social media disclosures of sexual abuse through the lenses of support seeking and anonymity. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3906–3918. DOI : <http://dx.doi.org/10.1145/2858036.2858096>
2. Nazanin Andalibi, Pinar Ozturk, and Andrea Forte. 2017. Sensitive self-disclosures, responses, and social support on Instagram: The case of #depression. In *Proceedings of the 20th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '17)*. ACM, New York, NY, USA. DOI : <http://dx.doi.org/10.1145/2998181.2998243>
3. John A Bargh, Katelyn YA McKenna, and Grainne M Fitzsimons. 2002. Can you see the real me? Activation and expression of the “true self” on the Internet. *Journal of Social Issues* 58, 1 (2002), 33–48. DOI : <http://dx.doi.org/10.1111/1540-4560.00247>
4. Louise Barkhuus and Juliana Tashiro. 2010. Student socialization in the age of facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 133–142. DOI : <http://dx.doi.org/10.1145/1753326.1753347>
5. Christopher P Barlett and Douglas A Gentile. 2012. Attacking others online: The formation of cyberbullying in late adolescence. *Psychology of Popular Media Culture* 1, 2 (2012), 123–135. DOI : <http://dx.doi.org/10.1037/a0028113>
6. Christopher P Barlett, Douglas A Gentile, and Chelsea Chew. 2014. Predicting cyberbullying from anonymity. *Psychology of Popular Media Culture* 5, 2 (2014), 171–180. DOI : <http://dx.doi.org/10.1037/ppm0000055>
7. Roy F Baumeister and Mark R Leary. 1995. The need to belong: desire for interpersonal attachments as a

- fundamental human motivation. *Psychological Bulletin* 117, 3 (1995), 497–529. DOI : <http://dx.doi.org/10.1037/0033-2909.117.3.497>
8. Roy F Baumeister, Liqing Zhang, and Kathleen D Vohs. 2004. Gossip as cultural learning. *Review of General Psychology* 8, 2 (2004), 111–121. DOI : <http://dx.doi.org/10.1037/1089-2680.8.2.111>
 9. Natalya N Bazarova. 2012. Public intimacy: Disclosure interpretation and social judgments on Facebook. *Journal of Communication* 62, 5 (2012), 815–832. DOI : <http://dx.doi.org/10.1111/j.1460-2466.2012.01664.x>
 10. Natalya N Bazarova and Yoon Hyung Choi. 2014. Self-disclosure in social media: Extending the functional approach to disclosure motivations and characteristics on social network sites. *Journal of Communication* 64, 4 (2014), 635–657. DOI : <http://dx.doi.org/10.1111/jcom.12106>
 11. Michael S Bernstein, Andrés Monroy-Hernández, Drew Harry, Paul André, Katrina Panovich, and Gregory G Vargas. 4chan and/b: An analysis of anonymity and ephemerality in a large online community. In *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media (ICWSM '11)*.
 12. Jeremy Birnholtz, Nicholas Aaron Ross Merola, and Arindam Paul. 2015. “Is it weird to still be a virgin”: Anonymous, locally targeted questions on Facebook confession boards. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2613–2622. DOI : <http://dx.doi.org/10.1145/2702123.2702410>
 13. danah boyd. 2014. *It's complicated: The social lives of networked teens*. Yale University Press.
 14. Nancy L Collins and Lynn Carol Miller. 1994. Self-disclosure and liking: A meta-analytic review. *Psychological Bulletin* 116, 3 (1994), 457–475. DOI : <http://dx.doi.org/10.1037/0033-2909.116.3.457>
 15. Denzil Correa, Leandro Araújo Silva, Mainack Mondal, Fabricio Benevenuto, and Krishna P Gummadi. The many shades of anonymity: Characterizing anonymous social media content. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media (ICWSM '15)*.
 16. Valerian J Derlega, Barbara A Winstead, and Kathryn Greene. 1997. Self-disclosure and starting a close relationship. *Handbook of Relationship* (1997), 153–174.
 17. Judith S Donath. 1999. Identity and deception in the virtual community. In *Communities in Cyberspace*. 29–59.
 18. Nicole B Ellison, Lindsay Blackwell, Cliff Lampe, and Penny Tieu. 2016. “The question exists, but you don’t exist with it”: Strategic anonymity in the social lives of adolescents. *Social Media + Society* 2, 4 (2016). DOI : <http://dx.doi.org/10.1177/2056305116670673>
 19. Hugh C Foot and May McCreaddie. 1997. Humour and laughter. In *The Handbook of Communication Skills*. Routledge, 259–285.
 20. Andrea Forte, Nazanin Andalibi, and Rachel Greenstadt. 2017. Privacy, anonymity, and perceived risk in open collaboration: A study of Tor users and Wikipedians. In *Proceedings of the 20th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '17)*. ACM, New York, NY, USA. DOI : <http://dx.doi.org/10.1145/2998181.2998273>
 21. Max Gluckman. 1963. Papers in honor of Melville J. Herskovits: Gossip and scandal. *Current Anthropology* 4, 3 (1963), 307–316.
 22. Erving Goffman. 1959. The presentation of self in everyday life. (1959).
 23. Oliver L Haimson, Anne E Bowser, Edward F Melcer, and Elizabeth F Churchill. 2015. Online inspiration and exploration for identity reinvention. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3809–3818. DOI : <http://dx.doi.org/10.1145/2702123.2702270>
 24. Jeffrey T Hancock, Jennifer Thom-Santelli, and Thompson Ritchie. 2004. Deception and design: The impact of communication technology on lying behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 129–134. DOI : <http://dx.doi.org/10.1145/985692.985709>
 25. Erin E Hollenbaugh and Marcia K Everett. 2013. The effects of anonymity on self-disclosure in blogs: An application of the online disinhibition effect. *Journal of Computer-Mediated Communication* 18, 3 (2013), 283–302. DOI : <http://dx.doi.org/10.1111/jcc4.12008>
 26. Jina Huh. 2015. Clinical questions in online health communities: The case of “see your doctor” threads. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 1488–1499. DOI : <http://dx.doi.org/10.1145/2675133.2675259>
 27. Jina Huh, Rupa Patel, and Wanda Pratt. 2012. Tackling dilemmas in supporting ‘the whole person’ in online patient communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 923–926. DOI : <http://dx.doi.org/10.1145/2207676.2208535>
 28. Ruogu Kang, Stephanie Brown, and Sara Kiesler. 2013. Why do people seek anonymity on the Internet?: Informing policy and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2657–2666. DOI : <http://dx.doi.org/10.1145/2470654.2481368>

29. Ruogu Kang, Laura Dabbish, and Katherine Sutton. 2016. Strangers on your phone: Why people use anonymous communication applications. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 359–370. DOI : <http://dx.doi.org/10.1145/2818048.2820081>
30. Lee Knuttila. 2011. User unknown: 4chan, anonymity and contingency. *First Monday* 16, 10 (2011). DOI : <http://dx.doi.org/10.5210/fm.v16i10.3665>
31. Cliff Lampe and Erik Johnston. 2005. Follow the (slash) dot: Effects of feedback on new members in an online community. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work (GROUP '05)*. ACM, New York, NY, USA, 11–20. DOI : <http://dx.doi.org/10.1145/1099203.1099206>
32. Cliff Lampe and Paul Resnick. 2004. Slash(dot) and burn: Distributed moderation in a large online conversation space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 543–550. DOI : <http://dx.doi.org/10.1145/985692.985761>
33. Noam Lapidot-Leffler and Azy Barak. 2012. Effects of anonymity, invisibility, and lack of eye-contact on toxic online disinhibition. *Computers in Human Behavior* 28, 2 (2012), 434–443. <http://dx.doi.org/10.1016/j.chb.2011.10.014>
34. Danielle M Law, Jennifer D Shapka, Shelley Hymel, Brent F Olson, and Terry Waterhouse. 2012. The changing face of bullying: An empirical comparison between traditional and internet bullying and victimization. *Computers in Human Behavior* 28, 1 (2012), 226–232. <http://dx.doi.org/10.1016/j.chb.2011.09.004>
35. Alex Leavitt. 2015. “This is a throwaway account”: Temporary technical identities and perceptions of anonymity in a massive online community. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 317–327. DOI : <http://dx.doi.org/10.1145/2675133.2675175>
36. Jessica Lingel, Aaron Trammell, Joe Sanchez, and Mor Naaman. 2012. Practices of information and secrecy in a punk rock subculture. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)*. ACM, New York, NY, USA, 157–166. DOI : <http://dx.doi.org/10.1145/2145204.2145230>
37. Xiao Ma, Jeffery T Hancock, and Mor Naaman. 2016. Anonymity, intimacy and self-disclosure in social media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3857–3869. DOI : <http://dx.doi.org/10.1145/2858036.2858414>
38. M. Lynne Markus. 1987. Toward a “critical mass” theory of interactive media: Universal access, interdependence and diffusion. *Communication Research* 14, 5 (1987), 491–511. DOI : <http://dx.doi.org/10.1177/009365087014005003>
39. Alice E. Marwick and danah boyd. 2011. The drama! Teen conflict, gossip, and bullying in networked publics. A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society. Available from SSRN. (2011). <http://ssrn.com/abstract=1926349>
40. Mark W. Newman, Debra Lauterbach, Sean A. Munson, Paul Resnick, and Margaret E. Morris. 2011. It’s not that I don’t have problems, I’m just not putting them on Facebook: Challenges and opportunities in using online social networks for health. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. ACM, New York, NY, USA, 341–350. DOI : <http://dx.doi.org/10.1145/1958824.1958876>
41. Julia Omarzu. 2000. A disclosure decision model: Determining how and when individuals will self-disclose. *Personality and Social Psychology Review* 4, 2 (2000), 174–185. DOI : http://dx.doi.org/10.1207/S15327957PSPR0402_05
42. Robert Paine. 1967. What is gossip about? An alternative hypothesis. *Man* 2, 2 (1967), 278–285. DOI : <http://dx.doi.org/10.2307/2799493>
43. W Babnett Pearce and Stewart M Sharp. 1973. Self-disclosing communication. *Journal of Communication* 23, 4 (1973), 409–425. DOI : <http://dx.doi.org/10.1111/j.1460-2466.1973.tb00958.x>
44. Sai Teja Peddinti, Keith W. Ross, and Justin Cappos. 2014. “On the Internet, nobody knows you’re a dog”: A Twitter case study of anonymity in social networks. In *Proceedings of the Second ACM Conference on Online Social Networks (COSN '14)*. ACM, New York, NY, USA, 83–94. DOI : <http://dx.doi.org/10.1145/2660460.2660467>
45. James W Pennebaker. 2012. *Opening up: The healing power of expressing emotions*. Guilford Press.
46. T Postmes, R Spears, and M Lea. 2000. The formation of group norms in computer-mediated communication. *Human Communication Research* 26, 3 (2000), 341–371. DOI : <http://dx.doi.org/10.1111/j.1468-2958.2000.tb00761.x>
47. Jenny Preece. 1999. Empathic communities: Balancing emotional and factual communication. *Interacting with Computers* 12, 1 (1999), 63–77. DOI : [http://dx.doi.org/10.1016/S0953-5438\(98\)00056-3](http://dx.doi.org/10.1016/S0953-5438(98)00056-3)
48. Vahed Qazvinian, Emily Rosengren, Dragomir R Radev, and Qiaozhu Mei. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '11)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 1589–1599.

49. Ralph L Rosnow. 1988. Rumor as communication: A contextualist approach. *Journal of Communication* 38, 1 (1988), 12–28. DOI : <http://dx.doi.org/10.1111/j.1460-2466.1988.tb02033.x>
50. Ralph L Rosnow. 2001. Rumor and gossip in interpersonal interaction and beyond: A social exchange perspective. In *Behaving badly: Aversive behaviors in interpersonal relationships.*, Robin M. Kowalski (Ed.). American Psychological Association, Chapter 8, 203–232. DOI : <http://dx.doi.org/10.1037/10365-008>
51. Ralph L Rosnow and Eric K Foster. 2005. Rumor and gossip research. *Psychological Science Agenda* 19, 4 (2005).
52. William B Stiles. 1987. “I have to talk to somebody”. In *Self-Disclosure: Theory, Research, and Therapy.* Springer US, 257–282. DOI : http://dx.doi.org/10.1007/978-1-4899-3523-6_12
53. Anselm Strauss and Juliet M Corbin. 1997. *Grounded theory in practice.* Sage.
54. John Suler. 2004. The online disinhibition effect. *Cyberpsychology & behavior* 7, 3 (2004), 321–326. DOI : <http://dx.doi.org/10.1089/1094931041291295>
55. Matthew Trammell. 2014. User investment and behavior policing on 4chan. *First Monday* 19, 2 (2014). DOI : <http://dx.doi.org/10.5210/fm.v19i2.4819>
56. Sherry Turkle. 1995. *Life on the screen: Identity in the age of the Internet.* Simon & Schuster Trade.
57. Joseph B Walther. 1996. Computer-mediated communication: Impersonal, interpersonal, and hyperpersonal interaction. *Communication Research* 23, 1 (1996), 3–43. DOI : <http://dx.doi.org/10.1177/009365096023001001>
58. Gang Wang, Bolun Wang, Tianyi Wang, Ana Nika, Haitao Zheng, and Ben Y Zhao. 2014. Whispers in the dark: Analysis of an anonymous social network. In *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14).* ACM, New York, NY, USA, 137–150. DOI : <http://dx.doi.org/10.1145/2663716.2663728>
59. Yi-Chia Wang, Robert Kraut, and John M. Levine. 2012. To stay or leave?: The relationship of emotional and informational support to commitment in online health support groups. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12).* ACM, New York, NY, USA, 833–842. DOI : <http://dx.doi.org/10.1145/2145204.2145329>
60. Sally Yerkovich. 1977. Gossiping as a way of speaking. *Journal of Communication* 27, 1 (1977), 192–196. DOI : <http://dx.doi.org/10.1111/j.1460-2466.1977.tb01817.x>