

Balancing Audience and Privacy Tensions on Social Network Sites

JESSICA VITAK¹

University of Maryland, USA

STACY BLASIOLA

University of Illinois at Chicago, USA

SAMEER PATIL

New York University, USA

EDEN LITT

Northwestern University, USA

As social network sites grow and diversify in both users and content, tensions between users' audience composition and their disclosure practices become more prevalent. Users must navigate these spaces carefully to reap relational benefits while ensuring content is not shared with unintended audiences. Through a qualitative study of highly engaged Facebook users, this study provides insight into how people conceptualize friendship online as well as how perceived audience affects privacy concerns and privacy management strategies. Findings suggest an increasingly complex relationship between these variables, fueled by collapsing contexts and invisible audiences. Although a diverse range of strategies are available to manage privacy, most participants in this sample still engaged in some degree of self-censorship.

Keywords: privacy, boundary regulation, social network sites, context collapse, impression management, Facebook

Once considered a social space purely for teens and young adults, social network sites (SNSs) have become deeply ingrained in Americans' daily lives, with 73% of online adults now using a SNS, including 65% of those aged 50 to 64 and 46% of those 65 or older (Smith, 2014). This diversification of the user base raises questions regarding how users' network composition affects site use. For example, Facebook maintains a wide range of privacy settings to help users control access to content, while Google+ encourages users to place each connection into a "circle" and designate content distribution accordingly. These network management features highlight the important relationship between audience,

Jessica Vitak: jvitak@umd.edu

Stacy Blasiola: sblas2@uic.edu

Sameer Patil: sameer@sameerpatil.net

Eden Litt: eden.litt@u.northwestern.edu

¹ This research was funded by an internal grant from the College of Communication Arts and Sciences at Michigan State University.

privacy, and disclosures, which has been previously examined in both off-line (e.g., Goffman, 1959; Leary, 1995) and online (e.g., Hogan, 2010; Vitak, 2012) environments.

In this article, we extend previous research in computer-mediated communication by unpacking SNS users' perceptions of their audience and how these perceptions influence their privacy attitudes and the privacy management strategies they employ. Using qualitative methods, we explore users' mental construction of audience to understand how the "imagined audience" (Litt, 2012; Marwick & boyd, 2011) corresponds to reality and how conscious users are of their audience when sharing content on the site. This article is driven by the following question: How cognizant are users of their audience when they use social network sites, and what impact, if any, do users' audiences have on their privacy attitudes and behaviors?

Through an analysis of 26 in-depth interviews, we explore the complex thought process users employ to understand and negotiate large and heterogeneous audiences on Facebook as well as how their perceptions of audience influence their privacy attitudes and behaviors. Specifically, we identify three categories of audience management strategies that help users balance the tension between privacy values and the desire to engage with others in their networks. Finally, we discuss the implications of these findings in deepening our understanding of how privacy is evolving with new communication technologies.

Managing Audience on Social Network Sites

In off-line settings, individuals use knowledge of their audience to help construct the identity they perform (Goffman, 1959). Visual, verbal, and nonverbal cues help them tailor their behaviors in a manner appropriate for a given audience; thus, one's audience provides context for the "performance" of self (Goffman, 1959; Papacharissi, 2012). In online settings, particularly SNSs, some of the cues available in physical settings are removed or limited, but the need for understanding the audience remains; interactants are forced to work with the cues available to them in a given communication environment (Walther, 1992). Compounding this problem is the manner in which SNSs converge multiple distinct audiences into a single homogeneous unit (e.g., Facebook friends)—a process known as context collapse (boyd, 2008; Marwick & boyd, 2011; Vitak, 2012). A Facebook user, for example, may have her parents, coworkers, classmates, and peers as Facebook friends. Considering context collapse within Goffman's (1959) framework, self-disclosure on SNSs becomes problematic when individuals can no longer distinguish the audience for whom they are performing or when they cannot easily alter those performances for different audiences.

Imagining the Audience

In the face of these challenges, SNS users must rely on their imagination to construct the audience for whom their performance is directed (Litt, 2012; Marwick & boyd, 2011). This imagined audience, or "a person's mental conceptualization of the people with whom he or she is communicating" (Litt, 2012, p. 330), serves as a guide on how to perform. Several factors influence the imagined audience, from environmental factors such as SNSs' social norms to individual factors such as a person's motivations for using a site (Litt, 2012). In an empirical study of how the imagined audience plays out on

Twitter, Marwick and boyd (2011) found that some users first considered their purpose for using a site (e.g., self-promotion, keeping in touch with friends), then crafted messages and the corresponding concept of the audience based on that initial purpose. On SNSs, one tactic for imagining the audience is to predict who is likely to see the content. However, predicting the audience is difficult, even when the audience appears bounded. When users craft posts with an audience in mind, they may neglect less visible members, instead focusing on those with whom they have the most contact—typically gauged by comments, likes, and shares (Bernstein, Bakshy, Burke, & Karrer, 2013; Litt, 2012). Recent research examining both self-reported and log data has found that the vast majority of Facebook users cannot accurately assess the full audience for a piece of shared content, with most underestimating the reach of their disclosures (Bernstein et al., 2013).

Facebook researchers have traditionally measured audience through a straightforward count of total friends, but more recent empirical work has attempted to distinguish between the various ways users conceptualize audience. For example, in two studies—one with college students (Ellison, Steinfield, & Lampe, 2011) and one with adults (Ellison, Vitak, Gray, & Lampe, 2014)—Ellison and colleagues established that Facebook users make a mental distinction between the friends formally articulated through the site and the subset of Facebook friends they consider as “actual” friends. Furthermore, in both studies, the researchers established that the number of actual friends significantly predicted perceptions of social capital to a greater extent than the total number of Facebook friends. From this, they suggest users make distinctions within their networks when sharing content. However, Ellison and colleagues did not define actual friends, leaving the interpretation to the individual. Therefore, one goal of this research is to gain a deeper understanding into how Facebook users conceptualize their audience on Facebook, especially in terms of the defining characteristics of a Facebook friend versus an actual friend.

RQ1: How do Facebook users conceptualize their audience on the site?

RQ1a: How do users define a Facebook friend?

RQ1b: How does this definition compare to that of an actual friend?

The Role of Privacy Concerns

Individuals attain privacy through regulating social interactions and selectively controlling access to personal information (Altman, 1975). Decisions related to both *what* and *to whom* one discloses pieces of personal information can lead to increased tensions for users. For example, as boyd (2014) demonstrates, teen SNS users typically want to share a lot of personal information, but they only want to do so with a select group of trusted individuals. However, uncertainty about the ability to control access to one’s personal information may increase privacy concerns (Petronio, 2002). The extent to which these concerns influence the type of content posted by SNS users has been studied extensively. Early studies of social media found that privacy concerns prevented some people from adopting a site (Acquisti & Gross, 2006), but once they decided to use a site, there was no relationship between their stated privacy concerns and the type of content they posted (Acquisti & Gross, 2006; Debatin, Lovejoy, Horn, & Hughes, 2009). SNSs, by virtue of their respective structures, encourage various forms of disclosure. Papacharissi (2012) found that the structure of Twitter invites improvisation and play, particularly where tweets stream

into hashtag topics. Comparatively, Facebook's structure encourages users to share personal details such as full name, birthday, hometown, and phone number; research by Lampe, Ellison, and Steinfield (2007) found that sharing information in these fields helped users establish common ground and could lead to increased social connections on the site.

More recent research suggests that privacy concerns align with posting behaviors (e.g., Stutzman, Capra, & Thompson, 2011; Stutzman, Gross, & Acquisti, 2012; Vitak, 2012); however, the reasons for this alignment are unclear. Facebook's many changes to its privacy policies are often accompanied by increasingly complex and granular privacy settings (Goel, 2013). Qualitative work by Vitak and Ellison (2013) found that some older users expressed concerns related to sharing information on the site because of the complexity of privacy settings and a general lack of efficacy navigating the site. On the other hand, privacy concerns are also likely to be influenced by the composition of users' audience on the site; for example, Page, Kobsa, and Knijnenburg (2012) found that users expressed a high degree of concern about how sharing location information in SNS updates might change their relationship with network members. Therefore, we propose the following question to unpack the relationship between audience and privacy concerns:

RQ2: How do Facebook users' audience perceptions impact their privacy concerns?

Privacy Strategies for Managing Audience

Nissenbaum (2010) has argued that privacy management concerns not only controlling access to information but the ability to ensure that information flows appropriately—that is, in the right context and to the right people, a concept she calls "contextual integrity." SNS users engage in a number of social and technical strategies to protect their information. Technical strategies are those that require manipulation of site features to limit the audience for content. These include actions such as changing privacy settings, deleting or blocking audience members from one's network, and untagging or deleting content (Litt, 2013; Young & Quan-Haase, 2009). In both off-line (Brandimarte, Acquisti, & Loewenstein, 2013) and online settings (Stutzman et al., 2011; Vitak, 2012), people are more likely to disclose personal information when they have access to controls that limit its ability to spread. On Twitter, users can set their profiles to be public or private. On Facebook, users have a range of technical features and settings available to them. For example, users may create lists to segment audiences into multiple groups. These strategies are preventative in that they aim to block certain groups from accessing content (Lampinen, Lehtinen, Lehmuskallio, & Tamminen, 2011). Although previous studies have demonstrated users' willingness to create lists (Debatin et al., 2009; Young & Quan-Haase, 2009), actually targeting content to those lists seems to vary across users and sites. Research on Google+ has found that "active users" create multiple circles to manage their content (Kairam, Brzozowski, Huffaker, & Chi, 2012). In the case of Facebook, some research indicates a lack of follow-through when it comes to targeting content to lists (Marder, Joinson, & Shankar, 2012), while other work suggests that people with larger and more diverse networks are more likely to make use of friend lists than those with smaller networks of fewer distinct audiences (Vitak, 2012).

Whereas the technical strategies for audience control rely on manipulating a site's settings or features, social strategies involve social decisions and rely mainly on users' social skills for execution. For example, users may engage in various forms of self-censorship in which they consciously decide not to share something (Das & Kramer, 2013; Sleeper et al., 2013); alternatively, taking a "lowest common denominator" approach, users limit the content of their posts to messages that are appropriate for everyone in the audience (Hogan, 2010).

This study focuses specifically on how a user's Facebook audience influences the strategies he or she employs. We move the conversation beyond a strict categorization of possible strategies users employ to a more detailed discussion of the interaction between audience composition and privacy practices on these sites.

RQ3: What privacy strategies do Facebook users employ to manage their audiences?

Method

Procedure

In April 2011, a random sample of 2,000 graduate students at a large Midwestern U.S. university were invited to complete a survey about their use of online communication tools. At the conclusion of the survey, those who used Facebook were invited to enter their e-mail address if they were interested in participating in a follow-up interview. The survey remained open for two weeks and had a response rate of 25% ($N = 486$); of these, 386 participants reported having an active Facebook account, and 169 provided an e-mail address to be considered for a follow-up interview. The primary goal of conducting these interviews was to gain deeper insight into users' privacy management and self-presentation practices on SNSs; therefore, a type of purposive sampling known as criterion sampling (Patton, 2002) was employed to select participants. Criteria for selection included diversity of networks, potential power differentials in network relationships (e.g., teacher-student), creation of multiple profiles within a single SNS, and high engagement with SNS privacy settings, especially friend lists.

Based on these criteria, 26 people participated in an in-person interview (see Table 1 for descriptive data for each interview participant). The first author compiled each participant's survey responses related to the research questions and conducted semistructured interviews, lasting 30 to 95 minutes, during April and May 2011. Participants were asked questions about their SNS use, actual and perceived audiences, privacy attitudes and behaviors, and impression management strategies. Participants also installed an application that visualized their friend network and described the various groups. All participants received a \$15 Amazon gift card. Among interview participants, 16 (62%) were women, with an average age of 29 (range: 22-53, $SD = 6$).

Table 1. Descriptive Data for Interview Participants.

Name	Sex	Age	Race	Total Facebook friends	Total actual friends	Facebook friend diversity^a	Privacy concerns^b	Use friend lists?^c
James	M	27	White	525	50	8	2.33	Yes
Bambi	F	33	White	280	280	12	3.33	Yes
Tina	F	25	Black	750	50	11	2.00	Yes
Julie	F	28	White	477	450	7	3.00	Yes
Susan	F	28	White	700	200	10	3.00	Yes
Arnold	M	33	White	400	25	12	3.33	Yes
Lela	F	22	White	1,400	20	8	2.67	Yes
Helen	F	31	White	165	100	7	4.00	Yes
Julia	F	22	Latina	1,600	1,000	5	3.67	Yes
Kelly	F	28	White	315	250	11	2.00	Yes
Amy	F	23	White	500	250	8	4.00	Yes
Kyle	M	27	White	251	70	10	2.67	No
Brian	M	29	White	700	500	13	2.67	No
Carl	M	35	White	830	100	6	4.00	Yes
Jonathan	M	35	White	337	200	10	1.33	Yes
Tara	F	26	White	368	200	7	3.67	Yes
Bonnie	F	25	White	548	30	8	4.00	Yes
Mary	F	25	White	302	50	8	3.00	Yes
Jessie	F	28	White	82	32	11	4.67	Yes
Anne	F	28	White	120	80	8	4.67	Yes
Ashley	F	27	Black	370	10	8	2.00	Yes
Andy	M	28	White	465	20	10	3.33	Yes
Dave	M	29	Multi	120	45	5	1.67	Yes
Colin	M	28	White	500	100	10	3.00	Yes
Nancy	F	53	White	200	150	7	2.00	No
Ben	M	31	White	701	10	11	2.67	No
Sum/								
Average	62%F	29	85%W	500	164	8.88	3.03	85%Y

^a Participants were asked to indicate whether they had Facebook friends in 12 relational categories, plus an "other" option. Scores were computed by adding up each category they indicated being present in their network (range: 5–13).

^b Three-item scale of users' concerns related to sharing information through the site, measured on a Likert-type scale ranging from 1 (*strongly disagree*) to 5 (*strongly agree*). See Vitak (2012) for items.

^c Question: "Have you created 'Friend Lists' so you can post updates just to a subset of your Facebook friends?"

Following transcription, the authors developed a codebook reflecting expected themes related to the research questions. The authors conducted content analysis in Dedoose, an online qualitative software program, whereby individual participants' data were used to refine themes as they emerged (Lincoln & Guba, 1985). Line-by-line coding of each transcript was employed using *complete thoughts* as the unit of analysis (Strauss & Corbin, 1998). First, each author coded two transcripts; then they met and refined the codebook. Next, each transcript went through two rounds of coding by the first three authors to ensure consistency of applied themes. Finally, excerpts were exported into Excel, and several metamatrices were created (Miles & Huberman, 1994) to identify patterns across participants, establish representativeness of findings, and detect negative cases.

Findings

RQ1: Conceptualizing Audience: Facebook Friends Versus Actual Friends

The first research question broadly addressed the issue of audience conceptualization on Facebook; in other words, although sites like Facebook have a *technical* designation for users' collection of social ties (e.g., Facebook friends), users may view and articulate these connections in various ways. In line with research by Ellison and colleagues (2011, 2014) distinguishing users' total Facebook friends from those they considered actual friends, participants were asked in the survey how many people fell into each category; then, in the interview, they were reminded of their responses and were asked to provide conceptual definitions.

Participants reported a diverse size range of Facebook friend networks from 82 to 1,600 connections.² In our attempt to answer RQ1a, a consistent theme emerged across participants when defining the principal characteristic: A Facebook friend represented a *known* social connection, be it current, former or, in rare cases, future (e.g., Amy³ friended her future classmates prior to starting medical school the following semester). Many participants explicitly signaled that, at some level, they knew every one of their Facebook friends even if they had not communicated since initially connecting. There were notable exceptions; for example, Brian, a musician, said he was friends with other musicians, although he tended to be "picky" about the ones he connected with on the site and "the ones I'm friends with on Facebook are mostly people that I got to know on Twitter first."

In line with prior research (Dunbar, 2011; Facebook Data Team, 2009), participants noted that, although their Facebook network included a wide variety of connections, most of their friends were acquaintances. Some participants even framed their Facebook friends in business terms, describing the site as a "virtual rolodex" used to manage a diverse set of contacts. Arnold said Facebook gave him "a way to store information for the people you don't really talk with and it's a reminder when you see them." Dave echoed this, saying, "I also consider Facebook friends more like a contact book or a contact list of people that I may need to contact in the future." Finally, some participants were quick to highlight the misnomer of Facebook's "friend" label; for example, Andy said that if he could "do it all over again," he'd

² Survey and interview data were linked prior to analysis.

³ Pseudonyms were created to protect participants' identities.

be much more selective in his friending decisions, because "a Facebook friend, to me, doesn't mean a real friend in real life." Andy, who had 465 Facebook friends, preferred a smaller audience to share his updates and "to keep the integrity of the word *friend* intact."

To help us answer RQ1b, participants discussed how their conceptualization of actual friends differed from that of Facebook friends. Participants reported 10 to 1,000 actual friends, accounting for 36% of the network on average, which is consistent with research using a nonstudent sample (Ellison et al., 2014) and higher than research with an undergraduate sample (Ellison et al., 2011).

A clear distinction emerged between participants' broad Facebook audience and the smaller subset of that audience participants viewed as actual friends; two primary themes dominated their discussion. First, actual friends were characterized by communication *channel* used and communication *frequency*. In other words, participants said that when they thought about actual friends, they thought of people they saw in person or talked to on the phone at least occasionally. Although participants spoke in general terms, stating that actual friends were people they were likely to hang out with in person, send e-mails to, or otherwise communicate with outside of Facebook, some gave specific examples that highlighted the distinction. Tina, who said most of her Facebook friends were not really friends, described her actual friends as those with whom she would share her more intimate life events: "An actual friend, I think, as I go through planning these life events, as I think about graduation or a baby shower, is someone who I would invite to actually share in those moments with me in person." In this way, actual friends have a very specific meaning.

The second theme to emerge when distinguishing actual friends was that this group contained people who could provide participants with various social, emotional, and instrumental resources encapsulated within the construct of social capital, which is in line with recent quantitative research (Ellison et al., 2014). For example, Helen said, "I want to know what's up with them, and they are people that I feel I could call or contact if I needed information or I needed help." Likewise, when Kyle tried to distinguish this smaller audience from his broader group of Facebook friends, he said, "I would say an actual friend is someone I could call on to help me out if I had a problem . . . someone who would help me move or someone who I could tell something personal about myself." Participants perceived an actual friend as someone who provides bonding resources, which has traditionally been associated with closer ties (Putnam, 2000).

As these descriptions highlight, participants generally viewed their Facebook network as comprising various social groups, or audiences, with at least one subset being those with whom they maintain a more meaningful relationship in off-line spaces and to whom they turn for emotional and instrumental support.

RQ2: Audience and Privacy Concerns

The second research question explored the relationship between users' audiences on Facebook and their privacy attitudes, specifically whether their network composition was related to particular privacy concerns. Overall, participants in this sample—who were selected because of their active engagement in

privacy management on SNSs—framed the relationship between privacy and audience in more proactive terms than would likely be seen in a more general sample. That said, participants voiced significant concerns related to two types of imagined audiences: those with whom they had articulated a formal relationship (i.e., Facebook friends) and those outside of the intended audience for their disclosures (e.g., current or future employers).

Because the people interviewed for this study exhibited high levels of self-monitoring, it was unsurprising that many participants voiced an awareness of how Facebook shares information as well as the audience with whom that information is shared. For example, Tina highlighted the tension she experienced between participating on the site and the potential audience implications of sharing information—with both known and unknown others:

I'm a very private person actually and Facebook kind of challenges that, where if you confirm you're going to an event, then people know where you're going. People know who your friends are. I don't really like that, so I try to keep it minimal in terms of my interactions with others.

Likewise, Helen said, "I'm very, very aware that whatever gets out cannot be brought back. It's kind of a Pandora's Box. So I'm concerned that personal information that I would not want out in public would somehow get released." Both examples highlight the cognitive process users engage in, with special consideration being given to *who* will see the content being shared.

Conversely, some participants expressed little concern that members of their network had access to information in their profile and the content they shared, because, as Andy put it, "I don't think most of them take any interest in me." Some participants also expressed little to no audience-related privacy concerns, but credited the strategies they employed to manage the relationship between their privacy concerns and their audience. These strategies (which will be discussed in detail in the discussion of RQ3), ranged from self-censorship to use of privacy settings to segment their audience and limit content distribution. Among participants employing various privacy-based strategies, however, some still expressed high levels of concern about sharing information, either related to audience composition (size and/or diversity) or a general lack of trust in Facebook itself. For example, Amy, who had many future classmates as Facebook friends, worried about information in her profile being taken out of context:

When I was younger, I thought it was funny, like my friends, we'd post while we were drinking, stupid stuff we're saying . . . looking at med school friends, that's not something I want anybody to see, especially since I don't know these people yet.

On the other hand, Bambi's concerns were largely derived from her lack of trust in how Facebook would treat the information:

I don't really trust the privacy on Facebook. I have not read the terms of service in their entirety. I know enough to be cautious, but I just have an issue . . . I'm sort of like, I don't want to think that privacy is dead, so I'd like to maintain a little bit of it.

Because this sample was composed mostly of employed individuals or people on the job market, privacy concerns related to potential employers viewing their Facebook profiles dominated interviews, even though most participants explicitly stated that they did not share anything they thought was controversial or incriminating. The concerns participants expressed took many forms. Most participants mentioning concerns about employers simply described being more cautious and conscientious in making posting-related decisions. For example, Brian noted, "You just have to be really careful about what you're sending out, what image of yourself you're projecting into the digital world." Other participants described taking explicit actions as a result of their increased privacy concerns. For example, James described a photo "triage" he conducted a few years earlier following a similar activity by some friends who had just completed their degrees. Likewise, Jessie, who admitted she had not viewed her privacy settings page "in ages," said her concerns about potential employers viewing her profile were significant enough to prevent her from sharing content: "[I won't post things to Facebook because] people looking for me as a potential employee or future coworkers might be able to see it and think less of me for it."

RQ3: Audience-Limiting Privacy Strategies

Participants described various practices aimed at limiting the audience for SNS disclosures. We classified these into four groups based on the mechanism used to enact audience control.

1. Network-based control. When considering privacy management, the most basic strategies involve granting or denying access—similar to the role of a bouncer at a club. Among our sample, many participants reported engaging in *network*-based privacy strategies, whereby they used the site's tools to proactively control the audience for their content.

As stated above, participants described their friend networks as comprising only known connections. Jessie, for example, rejected friend requests from people she did not recognize. Knowing a person, however, served as an initial criterion but did not guarantee Facebook friendship. Participants chose to ignore or deny requests from known others for a multitude of reasons. For example, Kelly did not accept friend requests from former high school classmates because she said she had nothing to talk to them about. Helen described denying her father's friend request to avoid dealing "with drama." Another reason participants chose not to accept friend requests was maintaining boundaries across disparate social contexts. This became most apparent when participants described overt power dynamics, such as their relationships with students, professors, and supervisors. Ashley, for instance, did not accept friend requests from previous students, explaining, "It just changes the dynamic when they see you in a professional way and they also see your personal side."

A related behavior reported by many participants—with very different implications—was that of "hiding" a friend; in fact, participants reported more instances of hiding friends than outright unfriending them. This finding is in line with recent quantitative work showing that Facebook users demonstrated a significantly higher intention to hide contacts than to unfriend them (Peña & Brody, 2014). Importantly, hiding a connection has no effect on the hidden user's ability to view content. However, many participants described engaging in this practice—often in lieu of unfriending another user—because it had fewer

potentially negative relational implications. For example, Helen described her rationale for hiding a group of high school friends:

I hid a lot of my high school friends after the reunion because I'm just not really interested in what's going on . . . I feel bad to completely defriend them, but, you know, we were good friends 10 or 15 years ago, but we're not really that close now. So I feel bad to defriend them, but I don't really want to get frequent updates on what's going on in their lives.

Privacy considerations did sometimes lead to defriending and/or blocking people. Life changes, such as breakups, were the most commonly reported reason, while other participants described "friend purges" in which they removed a significant number of friends—sometimes hundreds—from their network.

2. Platform-based control. In addition to curating their list of connections, participants reported leveraging the site's privacy settings. Several participants used the phrase "locked down" when referring to their use of privacy settings to protect personal information from unwanted audiences. As Bonnie noted, "I have it set that only my friends can see my pictures, and if you search for my profile, you can only see my profile picture and some of the limited information." The most common use of privacy settings was to limit access to one's profile only to direct connections (i.e., "friends only").

Within the "friends only" setting, the vast majority of participants in this sample established an additional subset of controls via friend lists.⁴ A common implementation involved separating friends into two distinct groups: one with access to all (or nearly all) shared content and a second group—which was usually, but not always, smaller—that could only view a stripped-down profile or designated posts (i.e., "limited profile"). Many participants acknowledged using limited profiles for specific individuals or groups such as family members, colleagues, and students. Kelly would not allow professors access to all her information, while Susan did the same with her students. To construct even more refined audience segmentation, many participants created multiple friend lists. Participants felt they achieved audience segmentation by targeting particular disclosures to some groups and not others. Lela, for example, vigilantly assigned each friend to a list or lists, thereby creating a different interaction experience for each friend.

3. Content-based control. Despite the level of control offered by lists, nearly all participants described significant self-reflection regarding the audience for their disclosures. Participants reported social strategies such as actively self-censoring and tailoring disclosures because of concerns about context collapse. Certain topics or themes, presented below, were common targets for self-censorship.

Professional considerations were a notable cause of self-censorship due to context collapse. Amy, for example, was concerned about repercussions of past disclosures:

⁴ Friend list use was one criterion used to invite participants to participate in the interview study; this is the reason for such high use of friend lists (86%) among the sample. In the full sample ($N = 386$), 17% of participants reported using friend lists.

Now that I'm moving into a professional part of my life, I don't want things that I did when I was a teenager coming back to bite me. . . . [The posts] may have been when I was 17, but most people don't look for the date.

Although most participants shared similar concerns regarding potential employers, Ben was an exception: "If someone wants to deny me a professional opportunity because of my politics, I don't want to be part of that professional opportunity." For some, intimate or personal details were not considered appropriate for Facebook and were also self-censored. These topics included information about family (especially children), health, sexual matters, address and contact information, and emotions.

A common strategy to deal with context collapse was to ensure content appropriateness for the broadest audience, known as the "lowest common denominator" (Hogan, 2010). Use of this strategy often resulted in less specific disclosures, as reflected in Bonnie's comment: "I keep things fairly generic because I use it both for social networking and also some professional networking." For some, applying this strategy decreased the amount and frequency of disclosures. For example, Andy noted, "The effect of having this many friends makes me say fewer and more general things than many and more specific things." Many participants stayed away from politics and religion to avoid divisive conflicts as well as negative judgment from their network.

Although topic was typically the most salient consideration, participants took into account other aspects, such as posting frequency, relevance, length, and even off-line connections and activities of friends when contemplating disclosures. For example, Tina refrained from certain posts to avoid unwanted "social leakage" off-line:

If I post, "I'm bored right now" at 4:30, and a colleague sees it, they know I was in class at 4:30. And you don't know whose advisor is who, and they're like, "Oh yeah, she posted she was bored while she was in your class," I worry about things like that getting around.

The imagination of an unintended audience gaining access to the disclosure motivated Tina to censor herself.

In some cases, participants enacted self-censorship by taking communications to private messages or chat, or switching to external channels such as phone or face-to-face communication. Andy captured this effectively: "In a private message you can communicate more freely without having to be concerned with the social world looking at what you're saying, and judging you or feeling as though they might be judging you."

4. Multiple profiles as control. Although uncommon, a few participants catered to divergent audiences by separating various facets of their persona into separate SNS accounts. Typically, the accounts were maintained on different SNS platforms (e.g., Twitter, LinkedIn, Blogger). Such separation enabled individuals to not only exploit the differences in interactive affordances and disclosure norms but

target disclosures appropriately based on the differences in audience characteristics and composition across these platforms. As Julia explained:

I have one Facebook, I have two Twitters, but because I do work in the social media, I also have that extended part that I do separate of my own personal things. So I represent my company on that. . . . But then I have the private Twitter and public Twitter, and private Facebook and public LinkedIn.

Some participants reported using multiple accounts on the same SNS platform to separate audiences based on relationship (social vs. professional), disclosure purpose, and other related reasons. Consider Jonathan, who used three Twitter accounts: one for a professional audience, one for a personal audience, and one for frivolous fun.

I have one which is strictly professional, so I connect to other folks who are studying technology or tourism, and they usually connect back to me as well, and that's its sole purpose. It is fully public. I have a second account that is sort of my fun account. It is the one that I'll connect to like Neil Patrick Harris. Or I listen to a lot of stuff off CNET, so their various podcasts, and so I'll connect to them. It's my frivolous feed . . . that one is purely for fun. It is also completely public. And then my third account is my personal account. The only people I'm connected with are personal friends of mine. That's fully private. And it sort of serves as an SMS [short message service] replacement.

Although this audience segmentation strategy likely decreases the chances of content being viewed by unintended audiences, it is also more costly to the individual in resources and time, which may explain why most participants did not employ it.

Discussion

This study provides a rich exploration of the tensions between SNS users' audiences and their disclosure practices on the site, as moderated by their privacy attitudes and practices. Research in this area is evolving from work examining the relationship between a standard set of privacy concerns individuals may have and the profile fields users complete (e.g., Acquisti & Gross, 2006; Stutzman et al., 2011) to more recent work delving into the strategies available to SNS users (Lampinen et al., 2011; Stutzman & Hartzog, 2012). The present work expands on previous research by exploring the evolving conceptualization of friendship in networked spaces and by focusing on the extent to which individuals' imagined audience—both those for whom they are actively thinking of when sharing content as well as those they may consider as a potential viewer of that same piece of content later on—impacts their privacy concerns and privacy practices on the site. As the user base of sites like Facebook continues to grow and diversify (Smith, 2014) and context collapse becomes an increasing impediment to participation for some users (Vitak & Ellison, 2013), untangling these relationships may provide important insights for theory and design.

Social media generally—and Facebook specifically—have significantly impacted how individuals consider relationship maintenance, interpersonal communication, and friendship (boyd, 2008). By labeling a technical connection between two users as a friendship, Facebook adds layers of implicit meaning to a relationship that may differ from how the two individuals may have characterized it. For example, in this study, participants overwhelmingly described a Facebook friend as simply someone they had met—with no indication of relational closeness or categorization going into that conceptualization—whereas an actual friend—a proxy for relational closeness—was defined by communication frequency, perceived social capital, and number of communication channels employed. This finding supports Haythornthwaite's (2005) media multiplexity research describing a positive correlation between tie strength and quantity of communication channels used. Similar to Ellison et al.'s (2014) results, this study provides qualitative support for the notion that, although Facebook friends played an important role in participants' lives, it was the actual friends that they viewed as sources of social, emotional, and instrumental support.

When considering the role of new social technologies in relational processes, it is essential to differentiate these newer forms of communication from more traditional ones, because the features and normative practices can differ widely. "Networked publics," as boyd (2008) has termed them, blur the boundaries between private and public disclosure and collapse unique sets of connections into one audience. This is likely the reason for participants' broad definition of a Facebook friend as well as the struggles they shared regarding friending practices and decisions to self-censor posts. Participants recognized that there were benefits to sharing content with their networks, in line with Burke and Kraut (2013) and Ellison et al. (2014), but they employed a range of social and technical strategies to share the right information with the right audience—the imagined audience (Litt, 2012). In an attempt to reach the imagined audience, their privacy strategies included four types of controls: network-based (e.g., only letting certain people into their networks), platform-based (e.g., using the site's privacy settings), content-based (e.g., deciding to not post an update at all or editing a post to avoid upsetting some subset of their network), and profile-based (e.g., maintaining separate Twitter accounts to interact with different parts of one's network). However, because the social norms around sharing continue to evolve (McLaughlin & Vitak, 2012), participants acknowledged turbulent moments in which their imagined and actual audiences did not align, such as inappropriate photos being posted and arguments occurring through public channels (Litt & Hargittai, 2014; Petronio, 2002). These negative experiences were described as catalyzing moments for changing their sharing and privacy behaviors (Debatin et al., 2009; Litt, 2013; Petronio, 2002).

Context collapse leads to both privacy concerns and employment of audience-limiting privacy strategies. However, because Facebook encourages disclosure and interaction with one's network at large, strategies that involve re-creating off-line social boundaries are not only inconsistent with the stated norms of the site but can be very time-consuming and require a moderate to high level of skill. Prior research suggests that several user characteristics relate to the norms for managing context collapse, including network composition (Vitak, 2012), online skills (e.g., Litt, 2013; Vitak & Ellison, 2013), and demographics (Litt, 2013). This study provides additional factors that may influence users' privacy management strategies—or lack thereof—and extends our understanding of privacy management by focusing on its link to audience composition. For example, participants who used friend lists reported feeling more comfortable sharing content because they were able to restrict access to the audiences of

greatest concern—typically family members, coworkers, professors, or students. At the same time, they expressed frustration or apathy toward using these lists because of the work involved. These factors likely contribute to many users' preference for broad approaches such as self-censorship of their disclosures.

The blurring of private and public in these spaces and the collapsing of users' networks into a single audience influenced privacy perceptions among this sample. Specifically, findings suggest that, when self-presentation is important—as would be the case for users on the job market or those maintaining a diverse network or connections—users consciously think about a particular audience when sharing content. For some, the lowest common denominator became the imagined audience. For example, future employers were part of many participants' imagined audiences, often leading to significant self-censorship. It is worth noting that participants were not always concerned about their current audience, but were in part fixated on a hypothetical/potential future audience. However, it is important to keep in mind these are self-reported data, and the extent to which the thought process for imagining audiences actually impacts participants' behavior is likely to depend on a number of factors—online skills, social skills, current privacy settings, and existing disclosure behaviors, among others (Litt, 2012). Nonetheless, participants' audience and privacy thoughtfulness suggest that people are not sharing content indiscriminately, but rather are thinking about who might see that content or how it might be interpreted by various audiences, especially when they imagine an important or influential audience for that disclosure.

Limitations

This study employed qualitative interviews to gain deeper insights into SNS users' conceptualizations of audience as well as how their audience influenced privacy perceptions and practices. Because the study employed purposive sampling, readers should use caution in interpreting findings, because the small, homogeneous sample of highly engaged U.S. graduate students limits generalizability of findings. To some degree, all Facebook users experience issues related to the unique dynamics of SNSs (e.g., context collapse), but they were likely magnified within this privacy-conscious sample; therefore, the findings presented here may not reflect that of a more heterogeneous population.

Conclusion

How, what, and with whom people communicate online provides researchers with a wide range of questions to explore. In this study, we explored a critical piece of this puzzle—how Facebook users conceptualize their audience when sharing content and how they consciously negotiate the tensions between the composition of that audience and their personal disclosures. Through analysis of 26 interviews, we have begun unpacking the complexity of this relationship—how users can alter their behavior to fit their audience when that audience is forever changing or half hidden as well as how some audiences may increase privacy concerns and, consequently, strategies used to manage the visibility of content. Participants in this study made a stark distinction between actual friends and Facebook friends based on communication frequency and channel selection as well as their resource support potentiality. These results showcase how users' perceived audiences impacted their privacy concerns, with this sample placing a heavy emphasis on an imagined audience of current and future employers. Finally, our findings

highlight four audience-limiting privacy strategies people engaged in in an attempt to reach their desired privacy: network-based, platform-based, and content-based control strategies as well as use of multiple profiles as control.

These findings reveal some of the cognitive processes users engage in when conceptualizing their audience on Facebook as well as the perceived impact that audience has on their privacy perceptions and practices. Notably, even among this privacy-conscious sample, participants voiced annoyance and struggles with enacting social and technical privacy strategies on Facebook. This suggests a need for more accessible tools that enable users to feel comfortable when making disclosures on the site. As these sites become more ubiquitous and ingrained into our daily lives, it will become increasingly important for us to address questions posed in this study and to understand the benefits and challenges presented by complex social technologies. Especially with increases in adoption by potentially at-risk populations (e.g., teens and older adults) and networks including a wider range of people with different backgrounds, value systems, and power dynamics, researchers and practitioners must consider solutions to problems of media literacy and skills, usability, and alternatives for sharing personal information.

Finally, even highly engaged and privacy-conscious users have audience-related concerns, which often lead to self-censorship. Social capital researchers have consistently argued that, in order for users to fully reap the relational benefits of social media use, they must be willing to disclose personal information (Vitak, 2012) and interact with network members (Burke, Kraut, & Marlow, 2011; Ellison et al., 2014). Likewise, communication theory and research has demonstrated the important role of such interactions in relationship formation and maintenance, both off-line (Altman & Taylor, 1975; Duck, 1994) and online (Tong & Walther, 2011; Walther, 1992). SNSs have great potential for facilitating communication and relationship development for both strong and weak ties; however, users are less likely to engage with these sites if they have significant privacy concerns. Therefore, studying and understanding the tensions between audience, information disclosure, and privacy—and creating an environment that balances these constructs—is vital to the future of these technologies.

References

- Acquisti, A., & Gross, R. (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. *Lecture Notes in Computer Science*, 4258, 36–58.
- Altman, I. (1975). *The environment and social behavior*. Belmont, CA: Wadsworth.
- Altman, I., & Taylor, D. (1973). *Social penetration: The development of interpersonal relationships*. New York, NY: Holt.
- Bernstein, M., Bakshy, E., Burke, M., & Karrer, B. (2013). Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 21–30). New York, NY: ACM. doi:10.1145/2470654.2470658

- boyd, d. (2008). *Taken out of context: American teen sociality in networked publics* (Doctoral dissertation). University of California, Berkeley, Berkeley, CA.
- boyd, d. (2014). *It's complicated: The social lives of networked teens*. New Haven, CT: Yale University Press.
- Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences, privacy and the control paradox. *Social Psychological and Personality Science, 4*, 340–347. doi:10.1177/1948550612455931
- Burke, M., & Kraut, R. (2013). Using Facebook after losing a job: Differential benefits of strong and weak ties. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 1419–1430). New York, NY: ACM. doi:10.1145/2441776.2441936
- Burke, M., Kraut, R., & Marlow, C. (2011). Social capital on Facebook: Differentiating uses and users. In *Proceedings of the 29th International Conference on Human Factors in Computing Systems* (pp. 571–580). New York, NY: ACM.
- Das, S., & Kramer, A. (2013). Self-censorship on Facebook. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media* (pp. 120–127). Washington, DC: Association for the Advancement of Artificial Intelligence.
- Debatin, B., Lovejoy, J. P., Horn, A. K., & Hughes, B. N. (2009). Facebook and online privacy: Attitudes, behaviors, and unintended consequences. *Journal of Computer-Mediated Communication, 15*, 83–108. doi:10.1111/j.1083-6101.2009.01494.x
- Duck, S. (1994). *Meaningful relationships: Talking, sense, and relating*. Thousand Oaks, CA: SAGE Publications.
- Dunbar, R. I. M. (2011, June). How many "friends" can you really have? *IEEE Spectrum, 81*, 83.
- Ellison, N., Steinfield, C., & Lampe, C. (2007). The benefits of Facebook "friends": Exploring the relationship between college students' use of online social networks and social capital. *Journal of Computer-Mediated Communication, 12*, 1143–1168. doi:10.1111/j.1083-6101.2007.00367.x
- Ellison, N. B., Steinfield, C., & Lampe, C. (2011). Connection strategies: Social capital implications of Facebook-enabled communication practices. *New Media & Society, 3*, 873–892. doi:10.1177/1461444810385389
- Ellison, N., Vitak, J., Gray, R., & Lampe, C. (2014). Cultivating social resources: The relationship between bridging social capital and Facebook use among adults. *Journal of Computer-Mediated Communication, 19*, 855–870. doi:10.1111/jcc4.12078

- Facebook Data Team. (2009, March 9). *Maintained relationships on Facebook*. Retrieved from <http://www.facebook.com/notes/facebook-data-team/maintained-relationships-on-facebook/55257228858>
- Goel, V. (2013, August 29). Facebook to update privacy policy, but adjusting settings is no easier. *The New York Times*. Retrieved from <http://bits.blogs.nytimes.com/2013/08/29/facebook-to-update-privacy-policy-but-adjusting-settings-is-no-easier/>
- Goffman, E. (1959). *The presentation of self in everyday life*. New York, NY: Anchor.
- Haythornthwaite, C. (2005). Social networks and Internet connectivity effects. *Information, Communication & Society*, 8, 125–147. doi:10.1080/13691180500146185
- Hogan, B. (2010). The presentation of self in the age of social media: Distinguishing performances and exhibitions online. *Bulletin of Science, Technology and Society*, 30, 377–386. doi:10.1177/0270467610385893
- Kairam, S., Brzozowski, M. J., Huffaker, D., & Chi, E. (2012). Talking in circles: Selective sharing in Google+. In *Proceedings of the 30th International Conference on Human Factors in Computing Systems* (pp. 1065–1074). New York, NY: ACM.
- Lampe, C., Ellison, N., & Steinfield, C. (2007). A familiar Face(book): Profile elements as signals in an online social network. In *Proceedings of the 25th International Conference on Human Factors in Computing Systems* (pp. 435–444). New York, NY: ACM. doi:10.1145/1240624.1240695
- Lampinen, A., Lehtinen, V., Lehmuskallio, A., & Tamminen, S. (2011). We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the 29th International Conference on Human Factors in Computing Systems* (pp. 3217–3226). New York, NY: ACM. doi:10.1145/1978942.1979420
- Leary, M. R. (1995). *Self-presentation: Impression management and interpersonal behavior*. Madison, WI: Brown & Benchmark.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Beverly Hills, CA: SAGE Publications.
- Litt, E. (2012). Knock, knock. Who's there? The imagined audience. *Journal of Broadcasting and Electronic Media*, 56, 330–345. doi:10.1080/08838151.2012.705195
- Litt, E. (2013). Understanding social network site users' privacy tool use. *Computers in Human Behavior*, 29, 1649–1656. doi:10.1016/j.chb.2013.01.049
- Litt, E., & Hargittai, E. (2014). A bumpy ride on the information superhighway: Exploring turbulence online. *Computers in Human Behavior*, 36, 520–529.

- Marder, B., Joinson, A., & Shankar, A. (2012). Every post you make, every pic you take, I'll be watching you: Behind social spheres on Facebook. In *Proceedings of the 45th Annual Hawaii International Conference on System Sciences* (n.p.). New York, NY: Computer Society Press.
doi:10.1109/HICSS.2012.12
- Marwick, A., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society, 13*, 114–133. doi:10.1177/1461444810365313
- McLaughlin, C., & Vitak, J. (2012). Norm evolution and violation on Facebook. *New Media & Society, 14*, 299–315. doi:10.1177/1461444811412712
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An expanded sourcebook* (2nd ed.). Thousand Oaks, CA: SAGE Publications.
- Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Redwood City, CA: Stanford University Press.
- Page, X., Kobsa, A., & Knijnenburg, B. P. (2012). Don't disturb my circles! Boundary preservation is at the center of location-sharing concerns. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (pp. 266–273). Washington, DC: Association for the Advancement of Artificial Intelligence.
- Papacharissi, Z. (2012). Without you, I'm nothing: Performances of the self on Twitter. *International Journal of Communication, 6*, 1989–2006.
- Patton, M. Q. (2002). *Qualitative research and evaluation methods*. Thousand Oaks, CA: SAGE Publications.
- Peña, J., & Brody, N. (2014). Intentions to hide and unfriend Facebook connections based on perceptions of sender attractiveness and status updates. *Computers in Human Behavior, 31*, 143–150.
doi:10.1016/j.chb.2013.10.004
- Petronio, S. (2002). *Boundaries of privacy: Dialectics of disclosure*. Albany, NY: State University of New York Press.
- Putnam, R. (2000). *Bowling alone: The collapse and revival of American community*. New York, NY: Simon & Schuster. doi:10.1145/358916.361990
- Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., & Cranor, L. F. (2013). The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work* (pp. 793–802). New York, NY: ACM.
doi:10.1145/2441776.2441865

- Smith, A. (2014). *Older adults and technology use*. Washington, DC: Pew Internet Project.
- Strauss, A., & Corbin, J. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Thousand Oaks, CA: SAGE Publications.
- Stutzman, F., Capra, R., & Thompson, J. (2011). Factors mediating disclosure in social network sites. *Computers in Human Behavior, 27*, 590–598. doi:10.1016/j.chb.2010.10.017
- Stutzman, F., Gross, R., & Acquisti, A. (2012). Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality 4*, 7–41.
- Stutzman, F., & Hartzog, W. (2012). Boundary regulation in social media. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (pp. 769–778). New York, NY: ACM. doi:10.1145/2145204.2145320
- Tong, S., & Walther, J. B. (2011). Relational maintenance and CMC. In K. B. Wright & L. M. Webb (Eds.), *Computer-mediated communication in personal relationships* (pp. 98–118). New York, NY: Peter Lang.
- Vitak, J. (2012). The impact of context collapse and privacy on social network site disclosures. *Journal of Broadcasting and Electronic Media, 56*, 451–470. doi:10.1080/08838151.2012.732140
- Vitak, J., & Ellison, N. (2013). “There’s a network out there you might as well tap”: Exploring the benefits of and barriers to exchanging informational and support-based resources on Facebook. *New Media & Society, 15*, 243–259. doi:10.1177/1461444812451566
- Walther, J. B. (1992). Interpersonal effects in computer-mediated interaction: A relational perspective. *Communication Research, 19*, 52–90. doi:10.1177/009365092019001003
- Young, A. L., & Quan-Haase A. (2009). Information revelation and Internet privacy concerns on social network sites: A case study of Facebook. In *Proceedings of the Fourth International Conference on Communities and Technologies* (pp. 265–274). New York, NY: ACM. doi:10.1145/1556460.1556499