

The Imagined Audience and Privacy Concern on Facebook: Differences Between Producers and Consumers

Social Media + Society
April-June 2016: 1–15
© The Author(s) 2016
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/2056305116644615
sms.sagepub.com


Yumi Jung and Emilee Rader

Abstract

Facebook users share information with others by creating posts and specifying who should be able to see each post. Once a user creates a post, those who see it have the ability to copy and re-share the information. But, if the reader has a different understanding of the information in the post than the creator intended, he or she may use the information in ways that are contrary to the intentions of the original creator. This study examined whether post creators (Producers) and readers (Consumers) who are Facebook Friends had similar levels of privacy concern regarding how others might use the information in specific posts, and how their privacy concern about the post varied by whether the imagined audience consisted of Friends, Friends of Friends, or the general Public. The results showed that both Producers and Consumers had similar levels of privacy concern about a post shared with an imagined audience of Friends versus Friends of Friends. However, Consumers believed posts were more private than the Producers themselves did, and showed more privacy concern. This shows that post Consumers care about Producers' privacy, perceive that they are co-owners of the information, and engage in boundary management with Producers.

Keywords

privacy, network distance, imagined audience, visibility, boundary management, social networks

Introduction

Privacy is the ability to control personal information (Westin, 2003). However, when a user shares personal information in a post to an online social network like Facebook she could lose control of her information, resulting in an unwanted disclosure. The user might unintentionally allow the information to be disclosed to someone that she does not intend to receive it by mis-managing privacy settings or being unaware of who has the ability to see her posts. Loss of control leading to a privacy violation can happen because an online social network “flatten[s] social relationships and eliminate[s] context” (Marwick, 2012), making it difficult for the user to keep the different relationships separate. Or, the user's Facebook friends who see the post could re-share that information with others against the user's wishes. Anyone who sees the post can quickly distribute a perfect copy, as well as being able to store it without the user's knowledge. Once information has been shared on a social network site, users must trust their network connections—the people they are connected to on the site directly or through others and who can see their posts—to be considerate of their often unstated

intentions for the spread of their information (Lampinen, Lehtinen, Lehmuskallio, & Tamminen, 2011).

Context collapse and re-sharing are both examples of difficulties that can arise regarding privacy boundary management online. Each person who can see the user's posts becomes a co-owner of the information in the post (Petronio, 2002) and can then use that information however they see fit. The recipients of the information become responsible for protecting it, and enforcing boundaries between who should be able to see it and who should not. In other words, Consumers (the recipients of posts) play a key role in protecting the privacy of the Producers (creators of posts).

Previous studies have investigated privacy management in online social networks solely from the perspective of the

Michigan State University, USA

Corresponding Author:

Emilee Rader, Department of Media and Information, Michigan State University, 404 Wilson Rd, East Lansing, MI 48824, USA.
Email: emilee@msu.edu



Producer of the information (Christofides, Muise, & Desmarais, 2009; Dwyer, Hiltz, & Passerini, 2007). However, it is also important to understand the Consumers' perspectives because of the power they have as co-owners of the Producer's information. Very little is known about discrepancies that might exist between Consumers and Producers regarding their perceptions of contextual factors related to privacy boundaries in social network sites, such as how much a user feels like she can trust her Facebook friends, how sensitive or personal the information in the post is, or who else in the network might be able to see the post. Any discrepancies could play an important role in the occurrence of unwanted disclosures.

To understand differences in perceptions between Producers and Consumers, we recruited pairs of Facebook friends to participate in an online survey. The purpose of the survey was to investigate both Producers' and Consumers' impressions of privacy-related contextual factors related to posts that had been created by one partner and read by the other. The survey asked about their perceptions of how concerned they thought the post creator should be about the privacy of specific Facebook posts if those posts were to be seen by other Facebook users, and how personal the information in the posts was.

Our results showed that Consumers believed Producers would be more concerned about post privacy than they actually were, and also rated the posts as more private. This indicates that Consumers viewed the posts more conservatively than Producers did, and were aware of their shared responsibility protecting the privacy of the Producers' information. It also suggests that while discrepancies exist regarding where the privacy boundary should be, the Consumer perspective leans more toward protecting the Producer's information than disclosing it to others. However, the results also highlight conditions under which boundary turbulence is most likely to occur. This study adds to our understanding of privacy boundary management in social network sites by focusing on the role of the recipients of the information in Facebook posts, and the differences between their perceptions and those of the post creators.

Related Work

Privacy in Online Social Networks

There are positive benefits to sharing information widely on Facebook, including an increase in bridging social capital (Ellison, Vitak, Gray, & Lampe, 2014) and perceived social support (Lai & Yang, 2015). Despite the benefits, however, Facebook users also worry about unwanted disclosures happening if information meant for only some users within their network is revealed more widely than they intended (Fox & Moreland, 2015). One reason why unwanted disclosures occur on Facebook is because of what Marwick and boyd (2011) refer to as "context collapse." When people with

whom a Facebook user has different kinds of relationships in real life all become part of a single friend list, it can be very difficult for the user to tailor what he or she posts appropriately for many different relationship contexts at the same time. The larger the user's friend network is, the more people can see a user's posts, and the more likely unwanted disclosures due to context collapse are to occur (Thon & Jucks, 2014).

When creating a post, Facebook users first consider their ideas and expectations about who will see the post. Litt (2012) calls this the *imagined audience*: the user's "mental conceptualization" of the people with whom she is communicating when she creates a post. According to Litt, the greater uncertainty there is about the characteristics of a user's audience, the more the user relies on what she imagines those characteristics to be. However, the user's imagined audience is rarely an accurate representation. For example, users typically do not expect weak ties, or people they do not feel very close to, as part of their imagined audience (Wang et al., 2011). There is some evidence from previous work that when users have a more specific idea of who is in the audience for their Facebook posts, they are more careful when creating posts, and more likely to alter their posts before submitting them (Wang et al., 2013).

In addition, Facebook users' perceptions of how many people can see the posts they create significantly underestimate the reach of the posts and the size of the actual audience (Bernstein, Bakshy, Burke, & Karrer, 2013). This happens because it is difficult for users to reason about people they are connected to in a social network through other people. The people in a user's friend list are the user's "friends" or first degree connections, and are one step away from the user in the network. However, each of the user's friends has a friend list and first degree connections of their own. These are the user's second degree connections or "friends of friends," and are two steps away from the user in the network. People who are three steps away in the network are "friends of friends of friends," and so on. The number of steps in the social network between two users is the *network distance* between them.

The sheer number of other users that are two or more steps away in the Facebook network is hard for post Producers to imagine (Bernstein et al., 2013). For example, a 2011 paper describing the structure of the Facebook network stated that a user with 100 friends on the social network site had "27,500 unique friends-of-friends and 40,300 non-unique friends-of-friends (Ugander, Karrer, Backstrom, & Marlow, 2011)." Gilbert (2012) proposed that the structure of online social networks makes it very difficult for users to be aware of who they are connected to through other people, and how information flows through the network. Depending on the user's privacy settings, posts she creates can be visible to tens of thousands of other users, making it practically impossible for her to gauge the scope of visibility of the posts throughout the social network. And the more steps away the

audience is from the Producer in the network, the less aware the post author is about who their audience is and who has access to the post.

People cope with worry about unwanted disclosures on Facebook in various ways. When users feel that they cannot fully control their information, their anxiety about controlling personal information (Dinev & Hart, 2006) increases. Uncertainty about the audience due to context collapse causes users to focus more on audience members to whom unwanted disclosures would be the most damaging for the user, like “parents, partners, and bosses” (Marwick & boyd, 2011), and they imagine how those groups of people might view the information (Marwick, 2012). Users practice self-censorship, limiting what they post (Sleeper, Balebako, & Das, 2013), and they use the built-in technical mechanisms to restrict access (Stutzman & Kramer-Duffield, 2010). However, successfully using the technical mechanisms available in the system to control access to their information is difficult for users, because they often do not understand the correct privacy setting configuration for the level of disclosure they are trying to achieve (Wang et al., 2011). And often, the privacy settings do not capture well the boundaries of the audiences the users want to share with (Stutzman & Kramer-Duffield, 2010).

Boundary Coordination and Turbulence

When Producers begin receiving reactions to their posts from Consumers—especially from Consumers they were unaware of—they may begin to feel regret toward posting what they did (Wang et al., 2011). This is an example of privacy boundary turbulence, which occurs when privacy rules did not work, and misunderstandings, unwanted disclosures or violations occur (Petronio & Durham, 2015). Boundary coordination is the process by which users work to mitigate the consequences of turbulence, or come to agreements with recipients about rules for ownership and disclosure of the information they share (Petronio, 2002; Stutzman & Kramer-Duffield, 2010).

Developing privacy rules and coordinating with recipients in social network sites is complicated because users often have inaccurate or unclear expectations for who is in the imagined audience (Litt & Hargittai, 2014) so they are unaware of who they should be coordinating with, and because privacy management tools do not support boundary coordination very well. In the absence of effective tools, participants in one study reported that they relied on their trust in their network connections to “know how to behave” and to “be considerate” of their wishes and intentions for their information (Lampinen et al., 2011). However, participants in that study also expressed worry about whether their social network connections, and even they themselves, could live up to those expectations.

Producers and Consumers are two sides of the privacy boundary coordination that takes place around information posted to Facebook. Considering how concerned Producers might be about disclosure of their information is the kind of

perspective-taking that co-owners of the information would engage in. Differences in perceptions between Producers and Consumers about how worried one should be about unwanted disclosures, about how private the information in a post is, or about the composition of the imagined audience, could create boundary turbulence and make coordination more difficult. However, most research about privacy on Facebook has focused exclusively on either the Producer or Consumer perspective, independently, and has not explicitly addressed audiences at different network distances away from the Producer. Asking the Producer and Consumer about their perceptions of the same Facebook post and about different audience boundaries could identify discrepancies between their perceptions that can lead to misunderstandings about appropriate disclosure of the information.

Trust and Privacy Concern

Trust is defined as confidence in sharing information with other people or organizations (Gefen & Straub, 2004), and is an important component of privacy boundary coordination (Joinson, Reips, Buchanan, & Schofield, 2010). In order to feel comfortable disclosing information to someone else, one must trust them. However, people often do not have sufficient information about how their disclosures will be accepted by others or how the reactions of others will affect them (Petronio, 2002). Trust can compensate for this uncertainty, and help individuals choose when to disclose information to others (Gefen & Straub, 2004). People tend to trust others by default, expecting them to adhere to privacy rules and norms. But if a betrayal occurs, they are less likely to trust those others again (Petronio, 2002).

Privacy concern and trust do not have a consistent effect on information disclosure. While some studies have found that trust is correlated with privacy concern (e.g., Smith, Dinev, & Xu, 2011), others have found that online users disclose a great deal of information even when they have high levels of privacy concern or low levels of trust (Norberg, Horne, & Horne, 2007). Mothersbaugh (2011) argues that privacy concern affects information disclosure primarily when individuals disclose sensitive information and their trust level is low.

Research Questions

The Relationship Among Role, Network Distance, and Trust. It is harder to trust in situations of uncertainty (Saeri, Ogilvie, La Macchia, Smith, & Louis, 2014). If Producers and Consumers are more uncertain about audiences at greater network distances (friends, friends of friends, etc.), greater network distance should be correlated with less trust that those audiences would protect the information as the Producer would want them to. Also, if the Consumer’s perceptions of the Producer’s level of trust do not match the Producer’s actual level of trust in the imagined audience at different network

distances, this could be a source of boundary turbulence. This is important because of the Consumer's role as a gatekeeper of the information posted by the Producer—in creating the post the Producer delegates control to each person who sees it. As a co-owner of the information, the Producer expects the Consumer to protect the information as he would.

Research Question 1 (RQ1). How are the network distance of the imagined audience away from the Producer and the role of the participant related to *trust* that the audience will protect the information in the post?

The Relationship Among Role, Audience Type, and Post Privateness. Post privateness is another contextual factor about each post that might be important for understanding privacy boundary turbulence and management. Privateness, or how “personal” the information is in the post, is related to the sensitivity of the information and users' willingness or unwillingness to disclose the information broadly (Malhotra, Kim, & Agarwal, 2004). People typically prefer to share information that is more private or personal with others they feel closer to or have a more intimate relationship with (Derlega, Winstead, Mathews, & Braitman, 2008). By asking Producers and Consumers about how private they think each post is, we can indirectly ask them about what their beliefs are about how broadly the information should be shared. If Consumers feel posts are less private than Producers do, they might be more willing to share the information in the posts with a wider audience than the Producer would feel comfortable with. A disagreement like this about the shared privacy boundary could result in unwanted disclosures.

In addition to perceived privateness, which is a property of the post content, each post has an *expected audience* associated with it. The expected audience is the category of people the Producer wants the post to be visible to, and is specified by the Producer using the audience selector tool provided by Facebook. The expected audience that the Producer has chosen for the post (friends, friends of friends, public) might be related to the perceived privateness of the information in the post, either as an expression of the Producer's preferences for the post or as a signal to Consumers about the privateness of the information.

Research Question 2 (RQ2). How are expected audience category and role related to the perceived *privateness* of the information in the post?

The Relationship Among Role, Network Distance, Expected Audience, and Post Privacy Concern. Within an online social network context, the relationship between privacy concern and user behavior is not consistent (McKnight, Lankton, & Tripp, 2011). Greater privacy concern is associated with lower levels of disclosure on Facebook; however, the relationship between concern and use of privacy controls and

settings like the audience selector tool is unclear (Stutzman, Capra, & Thompson, 2011).

It is important to understand Producers' privacy concern regarding information in specific posts that they create, because this concern affects what they choose to post and with whom they share the information. It is even more important to understand Consumers' privacy concern about specific posts created by their friends, because this might affect the degree to which Consumers respect Producers' intentions for the privacy boundaries for their posts. In addition, any differences in privacy concern between Producers and Consumers about specific posts when considering imagined audiences at different network distances might also create boundary turbulence. Turbulence might occur when the Producer's expected audience is more restrictive than the scope of the audience participants were instructed to imagine. Finally, the user's general level of privacy concern (Sheehan & Hoy, 2000) might also impact how concerned they are about the privacy of a post, because there is a relationship between general privacy concern online and information disclosure behaviors.

Research Question 3 (RQ3). How are network distance, role, and expected audience related to *post privacy concern*?

The Relationship Among Network Distance, Trust, Post Privateness, and Differences in Post Privacy Concern Between Producers and Consumers. There is some evidence that the same person estimates the likelihood of an unwanted disclosure differently depending on whether they are thinking about themselves and their own information, or others' information. People tend to believe privacy infringement (Baek, Kim, & Bae, 2014) or exposure to privacy risks (Debatin, Lovejoy, Horn, & Hughes, 2009) occur more for others than themselves. However, this belief does not increase privacy protection behaviors or motivate changes to social network privacy settings, unless people had already experienced privacy violations. This suggests that the magnitude and direction of disagreements between Producers and Consumers in their interpretations of trust in the imagined audience and privateness of the same post might be associated with the difference in privacy concern about the unwanted disclosure of information in a Facebook post.

Research Question 4 (RQ4). How are network distance, trust in the imagined audience, and the perceived privateness of the post related to *differences in post privacy concern* between Producers and Consumers?

Method

Participants and Recruiting

We used the Qualtrics platform to conduct an online survey for which we recruited pairs of Facebook friends who had

seen each others' News Feed posts. To recruit participants, we used snowball sampling starting from the authors' Facebook friends between November 2013 and February 2014. Snowball sampling has two advantages for this study over other forms of recruiting. A public snowball sampling message on Facebook can potentially be seen by many more people than other recruiting methods; according to Ugander et al. (2011), the average Facebook user has 214 friends, and an estimated 60,000 friends of friends. Also, snowball sampling is an effective way to recruit when the research requires participants who already know each other. In our sample, only 6 participants (7%) were first degree Facebook friends of the first author. The rest of the participants were second degree or higher connections.

To participate in the study, participants were each required to have recently created at least two Facebook posts that their partner either commented on or Liked before participating in the study. This ensured that the pairs had seen each others' posts as part of their normal use of Facebook. Our institutional review board (IRB) did not permit us to obtain information about the second person in the pair from the person who originally responded to the recruiting advertisement. Therefore, when a potential participant responded to the study advertisement, we explained the study and then asked the potential participant to reach out to Facebook friends that met the recruiting criteria and recruit a partner to participate in the study with them. Both partners answered survey questions in the role of the Producer (creator) of their own two posts, and the role of the Consumer (recipient) of their partner's posts, for a total of four posts shared between them.

We recruited 41 pairs of Facebook friends, for a total of 82 individual users of Facebook. Each pair answered questions about 4 Facebook posts, so there were 164 posts in the sample. A total of 21 pairs were both women, 8 pairs were both men, and 12 were mixed gender. We used the "Inclusion of Other in the Self Scale" (Aron, Aron, & Smollan, 1992) to measure the closeness of the relationship between the participant and his or her partner. It is a 7-point scale measuring emotional closeness that uses images of circles that do not overlap on the low end (coded as 1) and move closer to each other until they almost completely overlap on the high end (coded as 7). The mean closeness was 4.33 (standard deviation [SD] = 1.87), with a maximum of 7 and a minimum of 1. The distribution of closeness ratings is presented in Figure 1.

Participants were predominantly women (66%), and White (51%) or Asian (43%). The average age was 30 (SD = 7). Most participants were heavy Facebook users; 63% reported using Facebook several times per day and 49% said they create posts at least once a week. According to a January 2015 report from the Pew Research Center about the social media use of online adults in the United States, more women (77%) use Facebook than men (66%). In addition, 70% of adult Facebook users visit Facebook daily, and 65% frequently or sometimes share, post, or comment (Pew Research Center, 2014).




Figure 1. Histogram of the Producers' rating of the closeness of the relationship with the Consumer.

Study Procedure

After obtaining consent from both partners in a pair of Facebook friends, we sent a link to the survey to Partner A, who was the member of the pair who initially responded to the study advertisement. Partner A selected the two most recent posts she had created, that Partner B had commented on or liked on Facebook. She entered information into the survey about the two posts, such as the date and the type of post (photo, video, etc.) and then copied and pasted the text of the post into the survey form. She then selected an additional two posts, recently created by Partner B that she had commented on or liked on Facebook and entered the same information. After selecting these four posts, she completed the remaining questions in the survey, as the Producer of her own posts and the Consumer of Partner B's posts.

The survey was configured to automatically email Partner B with a link to the survey when Partner A was finished selecting posts. Partner B's version of the survey was already populated with the information Partner A had entered about the four Facebook posts she had selected. This ensured that both partners answered the survey questions about the same four posts. Partner B completed the survey as the Producer of the two posts that she had created, and the Consumer of the posts Partner A had created. A total of 30% of the posts selected for the study were created within the week prior to the participants completing the survey; 34% were between 1 week and 1 month old; 21% were 1–3 months old; and the remaining 15% of posts were older than 3 months.

After the post selection was completed by Partner A, the remainder of the questions were asked of both partners. The




Figure 2. A diagram illustrating network connections between the Producer and Consumer, and the imagined audience at different network distances from the Producer.

survey asked questions about participants' overall Facebook use and their closeness to their partner in the study. Then, it asked a series of questions about each of the four posts, including the following:

- The privacy setting chosen for the post by the Producer (expected audience);
- Three trust questions, about trusting other users at different network distances not to use the information in the post for other purposes (trust);
- How private the information in the post was (privateness);
- Three privacy concern questions, about the content of the post being shared with other Facebook users at different network distances away from the Producer of the post (concern).

The survey ended by asking a few demographic questions. Participants received a US\$5 Amazon gift card after completing the survey. In addition, Partner A received another US\$5 Amazon gift card when Partner B completed the survey, as a thank-you for their efforts to recruit a partner to participate.

We did several rounds of piloting with people who did not participate in the study and received verbal feedback about question wording and interpretation after each round. We then pre-tested the survey with a final group of people to check for any remaining misunderstandings between what we wanted to ask and interpretations of the questions. One issue uncovered during piloting was how to ask questions about the imagined audience at two degrees (friend of a friend, labeled as "FoF" in Figure 2) or three degrees (friend of a friend of a friend, labeled in Figure 2 as "Stranger (Public)") of separation from the Producer. We decided to use "Partner B's Facebook friends" (where 'Partner B' was replaced with the partner's first name) to represent second degree connections, and "others" to

represent third degree connections because it was easier for participants to understand.

Imagined Audience Manipulation

We were interested in how the network distance of the imagined audience for a Facebook post away from the Producer might affect trust in the audience and privacy concern about the information in the post. The survey asked both members of each pair of participants to imagine that the post was shared with audiences at three different network distances from the Producer. Producers then reported how concerned they would be about the privacy of the post if it were shown to those others and how much they would trust those audience members to protect the information in the post. Consumers answered how concerned they thought the Producer should be about the privacy of the post in those situations. We manipulated network distance of the imagined audience from the Producer in his or her Facebook network by asking the same question about each post three times, phrased differently for audiences at three network distances:

- *First degree.* The Producer's friends, one step in the network away from the Producer, phrased as "my Facebook friends."
- *Second degree.* Friends of friends, two steps in the network away from the Producer, phrased as "[Consumer name]'s Facebook friends."
- *Third degree.* Friends of friends of friends, three steps in the network away from the Producer, phrased as "friend of [partner's name] share[ing] the post with others." Note that third degree connections are far enough away from the Producer that they can effectively be considered as the general "Public."

See Figure 2 for a sketch of what the different network distances look like on a hypothetical network graph. The imagined audience is typically defined as the "mental conceptualization of the people with whom we are communicating" (Litt, 2012). Our imagined audience manipulation places a constraint on this mental conceptualization by instructing participants to imagine an audience with a particular characteristic—in this case, the network distance away from the Producer. Within that constraint, participants were free to rely on their imagination about who that audience might consist of.

Measures

Trust. To measure trust in the imagined audience on Facebook, the survey asked a modified version of a question from Buchanan, Paine, Joinson, and Reips (2007): "I trusted that my Facebook friends would not use my information for any other purpose (such as sharing with others without my permission or using for advertisement)." Consumers were

Table 1. Descriptives for Trust, Post Privacy Concern, and Post Privateness.

	Imagined audience	Producer		Consumer	
		Mean	SD	Mean	SD
Trust	Friend	5.12	(1.54)	4.91	(1.45)
	FoF	4.49	(1.63)	4.73	(1.47)
	Public	4.30	(1.61)	4.36	(1.46)
Privateness		3.64	(1.76)	4.07	(1.81)
Post Privacy Concern	Friend	3.32	(1.86)	3.95	(1.72)
	FoF	3.53	(1.70)	4.09	(1.56)
	Public	4.20	(1.81)	4.46	(1.60)

SD: standard deviation.

Descriptives for trust, post privacy concern, and post privateness for each post by participant role (Producer or Consumer) and levels of imagined audience: Friend (first degree), Friend of Friend (FoF, second degree), and Public (third degree). For example, the mean level of Trust Producers felt for an Imagined Audience of Friends was 5.12. All variables used a 7-point Likert scale, with 1 meaning *low* and 7 meaning *high*.

asked about their perceptions of the Producer's trust toward the imagined audience: "[Producer's name] trusts that his or her Facebook friends would not use the information for any other purpose." Trust was measured using a 7-point Likert scale that ranged from *Strongly Disagree* (1) to *Strongly Agree* (7). The mean of trust for both Producers and Consumers was 4.65 ($SD=1.53$); the descriptives for this question and other measures, including means broken out by whether the Producer or Consumer was answering the question, are available in Table 1. We asked this question three times of each participant, once for the imagined audience at each network distance (Friends, Friends of Friends, and Public).

Post Privateness. We asked both Producers and Consumers to rate each post according to how private they thought it was. The question, "How personal is the information in the post?" used a 7-point semantic differential scale, ranging from *Not personal at all* (1) to *Very personal* (7). The mean for both Producers and Consumers was 3.63 ($SD=1.76$). This measure allowed us to identify any differences that might exist between Producers' and Consumers' perceptions of how private the information was in the same post.

Privacy Concern. The question measuring post privacy concern was also adapted from Buchanan et al. (2007), who created and validated an instrument to measure privacy concern attitudes in a number of different contexts (e.g., medical records, credit cards, email). We tailored the wording of the question differently for Producers and Consumers, who were asked the question three times, once for each network distance of the imagined audience (Friends, Friends of Friends, and Public), about each of the four posts. The Producer question was, "I was concerned about my privacy when I created my post"; the Consumer question was, "[Producer's name] would be concerned about his or her privacy when he or she created the post." We used a 7-point Likert scale that ranged from *Strongly Disagree* (1) to *Strongly Agree* (7). The mean

for post privacy concern statements from both Producers and Consumers was 3.92 ($SD=1.75$). Finally, we also included a participant-level question about general privacy concern while using Facebook, that was answered once by each participant about themselves: "I am concerned about my privacy on Facebook." This question used a 7-point Likert agreement scale ($M=5.54$, $SD=1.24$).

Post Expected Audience. When the survey was conducted, Facebook allowed users to select categories of people to whom they wanted to restrict the visibility of their posts, via the audience selector tool. This is a representation of the user's explicit disclosure boundary for the post, specified using the mechanism Facebook provides for the privacy settings for each post. This is different from the imagined audience manipulation, in that it reflects the actual visibility of the post on Facebook, as specified by the Producer. We included a question about this in the survey as a way to measure the Producer's *expected audience*: "Who did you want to see this post when you created it?" Producers selected Friends (68.3%) and Public (16.5%) more than Friends of Friends (11.6%) and Specific Group (3.7%). Because there were only a few instances where the Producer expected a specific group of people to see the post, we combined the Friends category with the Specific Group category for the analysis.

Results

Trust Decreases as Distance Increases

We conducted a mixed effects linear regression with trust as the dependent variable, to investigate how the network distance of the imagined audience away from the Producer and the role of the user with respect to the information in the post affect trust that the audience will protect the information in the post (RQ1). The predictors were network distance of the imagined audience and role, and we also included random effects terms for post, individual, and pair in the model. We

included an interaction between role and network distance in the model in order to more accurately represent the 2 (role) \times 3 (network distance) categorical structure of the data. The unit of analysis was the post; because there were 41 pairs and 4 posts per pair, the total number of observations for this model and all subsequent analyses is 164. The results of the *Trust* model are presented in Table 2. The intercept ($\beta=4.92$, standard error [SE]=.15) represents the Consumer's perception of the Producer's level of trust for an imagined audience of Friends. The model shows that overall, trust decreased with greater network distance. In other words, the farther away in the network the imagined audience was from the Producer, the less both Producers and Consumers trusted that the audience would treat the information as the Producer intended.

There was also an interaction between role and network distance of the imagined audience, indicated by the negative and statistically significant coefficient in the model for Distance: Friends of Friends by Role: Producer ($\beta=-0.45$, $SE=.17$, $p<.01$). It is easiest to see this interaction in Figure 3, the graph of predicted values from the model (Gelman & Hill, 2006). The values for second degree (Friends of Friends) and third degree (Public) connections are very similar for Producers, and lower than for first degree (Friends). However, while the predicted value of trust that Producers have in their first degree connections is higher than Consumers think they have, Consumers believe Producers have more trust in their second degree connections than the Producers themselves report. In other words, Producers trust their own friends more and trust their friends of friends less than Consumers think they do. Producers' friends of friends may seem more trustworthy to Consumers than Producers if they are connected to the Producer through the Consumer. For example, all first degree connections (Friends) of the Consumer are automatically second degree connections (Friends of Friends) of the Producer, unless they happen to be a mutual Friend of both the Producer and the Consumer. An example of what this network structure looks like is illustrated in the diagram in Figure 2. If friends of friends of the Producer are connected to the Producer through the Consumer, they are likely to be more familiar to the Consumers than the Producers. Therefore, it is plausible that Consumers might trust their own friends more than Producers trust their friends of friends.

Finally, when imagining what the Public (third degree connections) might do with the information in the post, the predicted level of trust shown in Figure 2 is about the same for Producers and Consumers. Overall, the results of this model show that as Facebook users try to imagine who might see posts beyond the people in their own Friend lists that they are directly familiar with, they become less trusting of what those people might do with the information they post.

Consumers Believe Posts Are More Private

To investigate how differently Producers and Consumers perceive the privateness of posts the Producers shared with

Table 2. Results From the Regressions for Trust and Post Privateness.

	RQ1: Trust		RQ2: Privateness	
Intercept	4.92***	(.15)	4.21***	(.17)
Role: Producer	0.21*	(.12)	-0.52***	(.06)
<i>Imagined or expected audience</i>				
Friends of Friends	-0.18	(.12)	0.44*	(.17)
Public	-0.56***	(.12)	-1.02***	(.15)
FoF * Producer	-0.45**	(.17)	-0.18	(.23)
Public * Producer	-0.27	(.17)	0.49**	(.17)
<i>Random effects</i>				
	(SD)		(SD)	
Level 1 (Posts)	0.32		1.22	
Level 2 (Individuals)	0.98		0.83	
Level 3 (Pairs)	0.32		0.65	

SD: standard deviation.

"F of F" stands for "Friends of Friends." Audience for the Trust model is the imagined audience; for the Privateness model, it is the expected audience. The reference categories for both models were Consumer (role) and Friends (imagined audience).

* $p<.10$; ** $p<.05$; *** $p<.01$; **** $p<.001$.




Figure 3. Predicted values from the mixed effects regression for Trust (RQ1).

different expected audiences specified by the Producer of the post using the privacy settings mechanism (RQ2), we conducted a second mixed effects linear regression. The privateness of the information in the post was the dependent variable of this model; the predictors were an interaction between role of the person viewing the post (Producer or Consumer) and the expected audience type that the Producer specified (Friend, Friend of Friend, Public). We included the interaction in the model because all combinations of the role and expected audience categories were represented in the data, and including the interaction more accurately represents this in the model. This model also included random effects for post, individual, and pair. The intercept ($\beta=4.21$, $SE=.17$) represents how private posts that the Producer intended for

friends were, as evaluated by the Consumer of the post, for an expected audience of Friends. The results of this model are presented in Table 2.

The results show that there was a statistically significant difference in how Producers and Consumers view the privateness of the same post ($\beta = -0.52$, $SE = .06$, $p < .001$). Consumers evaluated posts intended for expected audiences who were first and second degree connections of the Producer to be more private than Producers did. In addition, posts shared with an expected audience that included second degree connections (Friends of Friends) were rated by both Producers and Consumers to be more private than posts shared with first degree connections ($\beta = 0.44$, $SE = .17$, $p < .05$). Posts shared with third degree connections (Public) were the least private ($\beta = -1.02$, $SE = .15$, $p < .001$).

There was an interaction between role and expected audience specified by the Producer, illustrated in Figure 4, which shows the predicted values generated from this model. Both Producers and Consumers rated posts intended for a Public expected audience as being the least private, and the predicted values were almost identical. This means that both Producers and Consumers were in agreement that public posts are not very private. However, they disagreed about how private the posts shared with the other audience types were; Consumers actually believed the posts were more private than Producers did. It is important to note that just because Consumers and Producers were first degree connections with each other, this does not mean they were strong ties; the mean closeness rating was 4.33 out of 7, and the values were spread fairly evenly from 2 to 7 on the scale (see Figure 1 for the distribution). This result is interesting and cause for optimism about re-sharing of posts. Consumers who think posts are more private than the Producers themselves do might be less likely to use the information in the posts for other purposes. In addition, this might signal that Consumers could show more privacy concern than Producers about these posts.

Producers Are Less Concerned About Post Privacy Than Consumers

To identify how concern for the privacy of the Producer's posts differs by role and distance of the imagined audience away from the Producer (RQ3), we conducted a third mixed effects linear regression model with privacy concern as the dependent variable. The predictors were role, network distance of the imagined audience, and expected audience (the privacy setting Producers chose for their posts). We included a role by network distance interaction in the model, and trust, post privateness, and general concern about privacy on Facebook as fixed effect controls. We used random effects controls in this model only for posts and individuals, because the variance for pair was marginal. The reference categories for the model were Consumer (role), Friends (imagined audience), and Friends (expected audience); trust, post privateness, and general Facebook privacy concern were




Figure 4. Predicted values from the mixed effects regression for Post Privateness (RQ2).

centered at their means. The results of this model, presented in Table 3, show a statistically significant association between role and post privacy concern.

Producers exhibited less post privacy concern than Consumers thought they should about the same post ($\beta = -0.52$, $SE = .13$, $p < .001$), even after controlling for trust, post privateness, and the Producer's expected audience for the post. Trust was negatively associated with post privacy concern: a one unit increase in trust meant a decrease in concern of -0.18 points ($SE = .03$, $p < .001$). This indicates that as trust increases, post privacy concern decreases. In contrast, the privateness of the post is positively associated with post privacy concern: a one unit increase in post privateness means a 0.19-point increase in post privacy concern ($SE = .03$, $p < .001$).

The graph of predicted values (Figure 5) shows the level of post privacy concern with trust, post privateness, and general Facebook privacy concern held at their means. Lower values on the y-axis represent less post privacy concern. This graph illustrates that post privacy concern was higher for an imagined audience of second degree connections than first degree, and higher still for third degree (Public) connections.

Expected Audience and Imagined Audience

It is reasonable to expect that if the expected audience is first degree connections (Friends), but the network distance of the imagined audience is farther away from the Producer in the network, post privacy concern would be higher. This is in fact what the results show. Predicted values of post privacy concern are highest for both Producers and Consumers when the network distance of the imagined audience of the post is "Public" and the expected audience as specified by the Producer is "Friends." Figure 5 shows that the predicted value of post privacy concern for "Friends" as both expected and imagined audience is right


Table 3. Results From the Regression for Post Privacy Concern.

	RQ3: Post Privacy Concern	
Intercept	4.06 ^{***}	(.14)
Role: Producer	-0.52 ^{***}	(.13)
<i>Imagined audience</i>		
Friends of Friends (FoF)	0.11	(.13)
Public	0.42 ^{**}	(.13)
<i>Expected audience</i>		
Friends of Friends (FoF)	-0.06	(.18)
Public	0.56 ^{***}	(.16)
Privateness	0.19 ^{***}	(.03)
Trust	-0.18 ^{***}	(.03)
General FB Privacy Concern	0.28 ^{**}	(.09)
Imagined FoF * Role	-0.01	(.19)
Imagined Public * Role	0.31 [*]	(.19)
<i>Random effects</i>		
	(SD)	
Level 1 (Posts)	0.44	
Level 2 (Individuals)	0.89	

SD: standard deviation; FB: Facebook.

The reference categories for the model were Consumer (role), Friends (imagined audience), and Friends (expected audience).

* $p < .10$; ** $p < .05$; *** $p < .001$.

**Figure 5.** Predicted values from the mixed effects regression for Post Privacy Concern (RQ3).

around 4, which was the center of the scale. Post privacy concern increases slightly for an imagined audience of second degree connections, and even more for third degree connections. This means that when imagining posts that were more private being seen by the most ambiguous audience, both post creators and recipients reported the most concern about the privacy of the post. Finally, there were only very small differences in the predicted values for post privacy concern between the expected audiences of Friends (first degree) versus Friends of Friends (second degree). This indicates that participants saw little difference between these two audiences regarding post privacy concern.

Disagreements About Trust and Privateness

The nature of the relationship between privacy concern, trust, and post privateness is more clearly illustrated by the analysis for RQ4. We asked how trust in the imagined audience and the privateness of the post are related to the differences in post privacy concern between Producers and Consumers. We used an ordered multinomial mixed effects regression, with three predictors: the network distance of the imagined audience for the post, the difference between Producer and Consumer trust in the imagined audience, and the difference between their views of the privateness of the post. This model also has a random effect control for individuals.

We calculated the difference values for each variable by subtracting Consumer values from Producer values for the same post and then transformed them into categorical variables, with Consumers as the reference category. This means that in the table of results, Table 4, “More Concern” means “Consumers reported more concern about post privacy than Producers,” “Same Audience Trust” means the difference between Producer and Consumer trust was zero, and “Less Private” means Consumers thought the information in the post was less personal than Producers did.

The dependent variable of this model has three categories: Consumers reporting more post privacy concern, the Consumers and Producers reporting the same amount of concern, or Consumers reporting less concern than Producers. The reference categories for the model are Friends (network distance), More Audience Trust (trust), and More Private (privateness). The predictors in the model indicate the likelihood of moving from one category of the dependent variable to the next. For example, the large, negative coefficient for “Trust Audience Less” in Table 4 ($\beta = -0.99$, $SE = .29$, $p < .001$) indicates that when the Consumer trusts the imagined audience at a given network distance less than the Producer does regarding a specific post, she is unlikely to also be less concerned about post privacy than Producers.

The graph of predicted probabilities in Figure 6 shows the likelihood that the Consumer is less concerned than the Producer about the privacy of the post the Producer created, given combinations of trust, post privateness, and the network distance of the imagined audience. Clearly, this is not very likely; the predicted probabilities range from 8% to 41%. Overall, Consumers are most likely to be more concerned about post privacy than Producers, not less. However, it is still important to focus on instances where Consumers are less concerned than Producers about the privacy of a specific post, because this is where boundary turbulence and unwanted disclosure would be the most damaging for the Producer. The results show that the greatest chance of boundary turbulence occurs when the imagined audience is the farthest away from the Producer, the Consumer thinks the post is less private than the Producer does, and also trusts the imagined audience more.

Table 4. Results From the Regression for the Difference in Post Privacy Concern Between Producers and Consumers.

	RQ4: Concern Difference	
More Concern Same Concern	−0.07	(.33)
Same Concern Less Concern	1.42 ^{***}	(.34)
<i>Imagined audience</i>		
Friends of Friends	0.08	(.24)
Public	0.46 [*]	(.24)
<i>Trust</i>		
Same Audience Trust	−0.35	(.32)
Trust Audience Less	−0.99 ^{***}	(.29)
<i>Privateness</i>		
Equally Private	0.42	(.32)
Less Private	0.60 [*]	(.35)
<i>Random effects (SD)</i>		
Level 2 (Individuals)	1.62	

SD: standard deviation.

The reference categories for the model are Friends (imagined audience), More Audience Trust (trust), and More Private (privateness). Difference values were calculated by subtracting the Consumer's rating from the Producer's rating for the same post and then transforming them to three categories: Consumer > Producer, Consumer = Producer, Consumer < Producer.

* $p < .10$; ^{*} $p < .05$; ^{**} $p < .01$; ^{***} $p < .001$.

Limitations

We asked Consumers about privacy concern, but we did not ask about potential or actual re-sharing or disclosure of the information, or perceptions of actual boundary turbulence related to the posts participants chose for use in the survey. So, we cannot tell from this study whether Consumers are willing to re-share Producers' posts anyway, despite feeling like Producers should be concerned.

Another limitation is that participants were all first degree connections of each other. It is possible that had second or third degree connections participated, they might have had different understandings of the posts, and of the perspective of the Producer regarding privacy boundaries. For example, Consumers might more readily share or re-use someone's information if they do not know the Producer directly. In addition, we asked participants to recruit their own partners, and they may have been more likely to ask strong ties to participate along with them. We used emotional closeness to operationalize tie strength, following studies of Facebook tie strength and social capital, like Gray, Ellison, Vitak, and Lampe (2013), and found that participants reported a wide range of relationship closeness with their partners (see Figure 1). However, we did not measure tie strength based on frequency of interaction or other behavioral proxy metric. While there was diversity in the emotional closeness ratings, it is possible that all partners in this study might be classified as strong ties based on interaction metrics, and if so these results may not generalize to friends who do not interact with each other often on Facebook.




Figure 6. Predicted probability of the Consumer being less concerned than the Producer about the privacy of a post, from the ordered multinomial mixed effects regression (RQ4). Each panel shows one combination of the levels of imagined audience (Friends, Friends of Friends, Public) by the difference in how personal the information in the post is as perceived by the Producer and Consumer (Consumer's perception is more, equally, or less private).

Finally, responses to the survey questions were self-report, and as such might be limited by social desirability bias, which occurs when a study involves socially sensitive issues (Grimm, 2010). Experimenter effects are also possible, which occur when participants are aware of the purpose of the study. Privacy can be a sensitive issue, so it is possible that Consumers rated their concern higher to show that they respect the Producer's privacy, or because they realized the study is about privacy. Also, while the consent form and instructions specified that only the researchers would have access to the survey responses, if participants believed their partner would see their responses this may have also caused Consumers to respond in a biased fashion.

Discussion

Boundary coordination is difficult under the best of circumstances; on Facebook it is complicated by misunderstandings about one's audience and ownership of the information. Boundary turbulence occurs when people have difficulty coordinating to develop and enact rules for when and with whom private information should be shared (Petronio, 2002); this can happen on Facebook when the Consumer of the information in a post believes that there will be no

consequences if they reveal the information to others, but the Producer would rather the information is not shared.

In this study we found that trust, post privateness, and post privacy concern are all factors related to the management of boundaries and co-ownership of private information, and that Consumers do not have the same understanding as Producers about the same posts. The results show that Consumers took a more protective stance toward Producers' posts than the Producers themselves did. This result agrees with "comparative optimism" discussed in Baek et al. (2014): people generally think an unwanted disclosure is more likely for others than for themselves. This could explain why Consumers' concern was greater than Producers'—because they may believe a privacy violation is more likely for the Producer than for themselves.

Trust that the imagined audience would not use the information they receive via Facebook posts contrary to the wishes of the Producer decreased with increased network distance of the imagined audience for both Producers and Consumers. This indicates that the confidence that both Producers and Consumers had that the other users at different network distances from the Producer would respect privacy boundaries decreased as ambiguity about the imagined audience increased. This may indicate that Facebook users do not extend trust in their first degree connections (Friends) to the people they are connected to through their friends. If so, this means that automated tools designed to provide personalized recommendations and advice about privacy settings (e.g., Dong, Jin, & Knijnenburg, 2015) and thereby help users manage their privacy should not assume that second degree connections are more worthy of trust than strangers just because they are closer in the network.

When we compared Consumers' and Producers' ratings of the privateness of information in the Facebook posts, we found that Consumers believed the posts were more private than Producers did. Interestingly, Consumers' perceptions of the privateness of Producers' posts were not correlated with how close they felt their relationship was with the Producer, $r(162) = -0.01$, $p = .82$. This indicates that tie strength as operationalized by the emotional closeness measure is not meaningfully related to perceptions of the privateness of post content. One unexpected finding was that both Producers and Consumers rated the content of posts that were created for an expected audience of Friends of Friends to be more private than posts created for an expected audience of Friends, after controlling for individual differences. This may be evidence that Facebook users are not always able to use the privacy mechanisms provided by the system to specify settings that adequately protect their information, which has been found by other researchers (e.g., Wang et al., 2011). In other words, users may choose privacy settings that allow post content they rated as more private to be exposed to a wider audience, which seems to be a failure to use the privacy mechanisms to specify an audience that is appropriate for the privateness of the content.

Both Producers and Consumers were least concerned when they imagined that a post they expected to be seen by the Public was shown only to Friends. However, they were most concerned when they imagined a post that only Friends were expected to see being shown to a Public audience. This might indicate that the expected audience—the privacy setting chosen for the post—is a signal to other users about the sensitivity of the information in the post and therefore where the sharing boundaries should be. Consumers also believed that Producers were more concerned about the privacy of their posts than they actually were. This belief on the part of Consumers is good news for Producers, because it indicates that Consumers' disclosure boundaries for the posts they read on Facebook are more restrictive than the post creators' are. It also suggests that a social network site user's first degree connections will understand and respect their intentions for the information they choose to disclose in their posts.

However, the results also point to conditions under which boundary turbulence, or disagreements about privacy rules that might lead to unwanted disclosures, would be most likely to occur. Consumers trusted Producers' second degree connections, some of whom could be connected to the Producer through the Consumer and therefore are the Consumer's own first degree connections, more than Producers did. This means that if Consumers think of their own Facebook friends when they imagine the audience for the Producer's post, they might be more likely to re-share the information than if they imagined an audience of people unknown to the Producer.

For example, these results suggest that a Consumer may feel more comfortable re-sharing a post if the people they see have commented on the post are mutual friends between the Producer and Consumer. This may bias the Consumer's perception of the imagined audience for the re-shared post toward people who are known to both themselves and the Producer, and away from people that the Producer does not know. In addition, Consumers are most likely to be less concerned than Producers if they imagine a public audience they trust gaining access to a post they think is less private than the Producer does. In other words, when Consumers trust the imagined audience more than Producers, they may be more likely to re-share information with people who are not within the Producer's intended privacy boundary.

These results allow us to begin to more precisely identify why unwanted disclosures might happen. It is not easy to find out who one's second and third degree network connections on Facebook might be. Gilbert (2012) wrote about this problem as one of "triadic awareness": two different friends of a given Facebook user might not be aware of each other, nor whether conditions might occur where they can see each others' posts. This would happen, for example, if a user has a friend from school and a friend from work who have no connection to each other except through the user. That user then serves as a *de facto* gatekeeper for the posts created by these

two friends. Our results show that in situations like this, the network structure itself adds to the ambiguous context of social connections online. Producers in our study were more concerned about a public post (expected audience) being shown to a third degree connection (a highly uncertain imagined audience) than they were about a friends-only post (restrictive privacy boundary) being shown to second degree connections. This indicates that an unknown audience is scarier than an audience consisting of friends of friends, when in reality, most of a user's friends of friends are effectively strangers to them.

The privacy management tools Facebook currently provides, like the "privacy checkup" which helps users become more aware of what audience(s) they are sharing with, support a mental model of privacy as control over access. It is a move in the right direction to help users regain control over unwanted sharing by reminding them of the default audience setting, and removing the "friends of friends" option from the audience selector privacy mechanism. But our study demonstrates that privacy concern is a relational property that changes according to network structure, and varies depending on both the Producer's and Consumer's imagined audience for information. It may be difficult for Consumers to imagine an audience for a Public post created by the Producer, were they to re-share it, which contains many people who are not also friends of the Producer. Because Facebook users have so many second degree connections, the "Public" audience setting on the post might not be a useful cue for protecting privacy boundaries. Reminding users their Public posts can be seen by anyone does not provide concrete, actionable information that they can use to imagine who beyond their immediate, salient friends might be able to see their posts if they were re-shared.

To help alleviate privacy concern, Facebook might create ways to help both Producers and Consumers understand the visibility and reach of posts, especially considering network connections that they both have in common. Although Producers can specify an expected audience for their posts, the practical audience-related implications of a Consumer liking or commenting on a public post and thereby making the post visible to people the Producer is unaware of are hard for users to reason about. A possible way to reduce uncertainty might be to estimate and report the size of the true audience for each post to the Producer, or to show Consumers the amount of overlap between their friend network and the network of the Producer. Users may underestimate the extent of unwanted disclosures that are not explicitly visible, especially if they have not correctly specified their expected audience. To help users more correctly select the audience for a post, the interface might prompt the Producer to rate the privateness of a post before they create it, and then compare that rating against the expected audience setting for the post. Feedback like this might enhance users' ability to coordinate privacy boundaries as properties of relationships in a network, rather than binary access

control decisions. It might even be possible to reveal that information to Consumers in a lightweight way, helping Consumers to understand Producers' intentions for their posts as way to support boundary co-ownership as well as access control. Consumers were more concerned about post privacy than Producers, so if they knew something about the privateness of the post from the Producer's standpoint, they might take this into consideration when commenting on, liking, or tagging someone in a post.

Conclusion

Once a Facebook post is shared with others, the Producer and Consumers of the post share ownership of the information in the post. We found that Consumers who were first degree connections of the Producer of a post reported privacy-friendly attitudes consistent with collective privacy boundaries and co-ownership. However, although Consumers are sensitive to Producers' privacy, it is still possible that they might misunderstand the Producer's original intentions for the post, and disagree with the Producer about how much they can trust the imagined audience. Considering how unwanted disclosures happen often through Consumers, these disagreements should receive more attention. Boundary coordination between the Producer and Consumer might be more effective if these disagreements were reduced.

Acknowledgements

We thank Michelle Rizer for her assistance with piloting and revising the survey questions, the Michigan State University BITLab research group for feedback on the survey and conversations about the results, and the anonymous reviewers for their helpful comments.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported in part by an endowment from AT&T to the Department of Media and Information at Michigan State University.

References

- Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of Personality and Social Psychology*, 63, 596–612.
- Baek, Y. M., Kim, E.-M., & Bae, Y. (2014, February). My privacy is okay, but theirs is endangered: Why comparative optimism matters in online privacy concerns. *Computers in Human Behavior*, 31, 48–56.
- Bernstein, M. S., Bakshy, E., Burke, M., & Karrer, B. (2013). Quantifying the invisible audience in social networks. In

- Proceedings of the CHI 2013 Conference on Human Factors in Computing Systems* (pp. 21–30). New York, NY: ACM Press.
- Buchanan, T., Paine, C., Joinson, A. N., & Reips, U. D. (2007). Development of measures of online privacy concern and protection for use on the Internet. *Journal of the Association for Information Science and Technology*, 58, 157–165.
- Christoffides, E., Muise, A., & Desmarais, S. (2009). Information disclosure and control on Facebook: Are they two sides of the same coin or two different processes? *CyberPsychology & Behavior*, 12, 341–345.
- 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.
- Derlega, V. J., Winstead, B. A., Mathews, A., & Braitman, A. L. (2008, May). Why does someone reveal highly personal information? Attributions for and against self-disclosure in close relationships. *Communication Research Reports*, 25, 115–130.
- Dinev, T., & Hart, P. (2006). Internet privacy concerns and social awareness as determinants of intention to transact. *International Journal of Electronic Commerce*, 10(2), 7–29.
- Dong, C., Jin, H., & Knijnenburg, B. P. (2015). Predicting privacy behavior on online social networks. In *International Conference on Weblogs and Social Media* (pp. 91–100). Palo Alto, CA: Association for the Advancement of Artificial Intelligence.
- Dwyer, C., Hiltz, S., & Passerini, K. (2007). Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace. In *2007 Proceedings of the Americas Conference on Information Systems* (p. 339). Retrieved from <http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1849&context=amcis2007>
- Ellison, N. B., Vitak, J., Gray, R., & Lampe, C. (2014). Cultivating social resources on social network sites: Facebook relationship maintenance behaviors and their role in social capital processes. *Journal of Computer-Mediated Communication*, 19, 855–870.
- Fox, J., & Moreland, J. J. (2015, April). The dark side of social networking sites: An exploration of the relational and psychological stressors associated with Facebook use and affordances. *Computers in Human Behavior*, 45, 168–176.
- Gefen, D., & Straub, D. W. (2004). Consumer trust in b2c e-commerce and the importance of social presence: Experiments in e-products and e-services. *Omega*, 32, 407–424.
- Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press.
- Gilbert, E. (2012). Designing social translucence over social networks. In *Proceedings of the CHI 2012 Conference on Human Factors in Computing Systems* (pp. 2731–2740). New York, NY: ACM Press.
- Gray, R., Ellison, N. B., Vitak, J., & Lampe, C. (2013). Who wants to know? Question-asking and answering practices among Facebook users. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work* (pp. 1213–1224). New York, NY: ACM Press.
- Grimm, P. (2010). Social desirability bias. In J. Sheth & N. Malhotra (Eds.), *Wiley International Encyclopedia of Marketing*. Chichester, UK: John Wiley & Sons.
- Joinson, A., Reips, U.-D., Buchanan, T., & Schofield, C. B. P. (2010). Privacy, trust, and self-disclosure online. *Human-Computer Interaction*, 25(1), 1–24.
- Lai, C. Y., & Yang, H. L. (2015). Determinants of individuals' self-disclosure and instant information sharing behavior in micro-blogging. *New Media & Society*, 17, 1454–1472.
- 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 CHI 2011 Conference on Human Factors in Computing Systems* (pp. 3217–3226). New York, NY: ACM Press.
- Litt, E. (2012, July). Knock, knock. Who's there? The imagined audience. *Journal of Broadcasting & Electronic Media*, 56, 330–345.
- Litt, E., & Hargittai, E. (2014). A bumpy ride on the information superhighway: Exploring turbulence online. *Computers in Human Behavior*, 36, 520–529.
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15, 336–355.
- Marwick, A. E. (2012). The public domain: Social surveillance in everyday life. *Surveillance & Society*, 9, 378–393.
- Marwick, A. E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13, 114–133.
- McKnight, D. H., Lankton, N., & Tripp, J. (2011). Social networking information disclosure and continuance intention: A disconnect. In *2011 44th Hawaii International Conference on System Sciences (HICSS)* (pp. 1–10). New York, NY: IEEE.
- Mothersbaugh, D. L. (2011). Disclosure antecedents in an online service context: The role of sensitivity of information. *Journal of Service Research*, 15, 76–98.
- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41, 100–126.
- Petronio, S. (2002). *Boundaries of privacy: Dialectics of disclosure*. Albany: State University of New York Press.
- Petronio, S., & Durham, W. T. (2015). Communication privacy management theory: Significance for interpersonal communication. In D. O. Braithwaite & P. Schrodt (Eds.), *Engaging theories in interpersonal communication: Multiple perspectives* (pp. 335–347). Thousand Oaks, CA: SAGE.
- Pew Research Center. (2014). *Social media update 2014*. Retrieved from <http://www.pewinternet.org/2015/01/09/social-media-update-2014/>
- Saeri, A. K., Ogilvie, C., La Macchia, S. T., Smith, J. R., & Louis, W. R. (2014). Predicting Facebook users' online privacy protection: Risk, trust, norm focus theory, and the theory of planned behavior. *The Journal of Social Psychology*, 154, 352–369.
- Sheehan, K. B., & Hoy, M. G. (2000). Dimensions of privacy concern among online consumers. *Journal of Public Policy and Marketing*, 19, 62–73.
- Sleeper, M., Balebako, R., & Das, S. (2013). The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the CSCW 2013 Conference on Computer Supported Cooperative Work and Social Computing* (pp. 793–802). New York, NY: ACM Press.
- Smith, H. J., Dinev, T., & Xu, H. (2011, November). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35, 989–1016.

- Stutzman, F., Capra, R., & Thompson, J. (2011, January). Factors mediating disclosure in social network sites. *Computers in Human Behavior*, 27, 590–598.
- Stutzman, F., & Kramer-Duffield, J. (2010). Friends only: Examining a privacy-enhancing behavior in Facebook. In *Proceedings of the CHI 2010 Conference on Human Factors in Computing Systems* (pp. 1553–1562). New York, NY: ACM Press.
- Thon, F. M., & Jucks, R. (2014). Regulating privacy in interpersonal online communication: The role of self-disclosure. *Studies in Communication Sciences*, 14, 3–11.
- Ugander, J., Karrer, B., Backstrom, L., & Marlow, C. (2011). *The anatomy of the Facebook social graph*. Retrieved from <http://arxiv.org/abs/1111.4503>
- Wang, Y., Leon, P. G., Scott, K., Chen, X., Acquisti, A., & Cranor, L. F. (2013). Privacy nudges for social media: An exploratory Facebook study. In *Proceedings of the 22nd International Conference on World Wide Web Companion* (pp. 763–770). New York, NY: ACM Press.
- Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., & Cranor, L. F. (2011). I regretted the minute I pressed share: A qualitative study of regrets on Facebook. In *Proceedings of the 2011 Symposium on Usable Privacy and Security (SOUPS)*. New York, NY: ACM Press.
- Westin, A. F. (2003, April). Social and political dimensions of privacy. *Journal of Social Issues*, 59, 431–453.

Author Biographies

Yumi Jung (MA, Rutgers University) is a PhD candidate in Media and Information Studies at Michigan State University. Her research interests lie in the area of privacy and social networks, focusing on individuals' privacy preferences and the tension between public security and individual privacy.

Emilee Rader (PhD, University of Michigan) is an assistant professor in the Department of Media and Information at Michigan State University. Her research program is focused on understanding the effects of interdependence between human and algorithm behavior in socio-technical systems.