Everyday Online Sharing

Submitted in partial fulfillment for the requirements for
the degree of
Doctor of Philosophy
in
Societal Computing

Manya Sleeper

Carnegie Mellon University
Pittsburgh, PA

May 2016

This research was supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. 0946825, as well as the National Science Foundation under grants CNS-1012763 and DGE-0903659. It was also supported in part by the ARCS Foundation, Google under a Focused Research Award on Privacy Nudges, and IWT.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation or any other sponsors.

Keywords: usability, privacy, social networking sites, access control

Abstract

People make a range of everyday decisions about how and whether to share content with different people, across different platforms and services, during a variety of tasks. These sharing decisions can encompass complex preferences and a variety of access-control dimensions. In this thesis I examine potential methods for improving sharing mechanisms by better understanding the everyday online sharing environment and evaluating a potential sharing tool.

I first present two studies that explore how current sharing mechanisms may fall short on social-networking sites, leading to suboptimal outcomes such as regret or self censorship. I discuss the implications of these suboptimal outcomes for the design of behavioral nudging tools and the potential for improving the design of selective-sharing mechanisms. I then draw on a third study to explore the broader "ecosystem" of available channels created by the services and platforms people move between and combine to share content in everyday contexts. I examine the role of selective-sharing features in the broader audience-driven and task-driven dynamics that drive sharing decisions in this environment. I discuss the implications of channel choice and dynamics for the design of selective-sharing mechanisms.

Using insights from current shortfalls and ecosystem-level dynamics I then present a fourth study examining the potential for adding topic-driven sharing mechanisms to Facebook. I use design mockups and a lab-based interview to explore participants' hypothetical use cases for such mechanisms. I find that these mechanisms could potentially be useful in a variety of situations, but successful implementation would require accounting for privacy requirements and users' sharing strategies.

Contents

List of Figures

List of Tables

1 | Introduction

People make decisions about how and whether to share content online. Outside of organizational or work environments, in everyday contexts, people may share with different individuals or groups, across varied devices, and drawing on different services. These sharing decisions can range from a choice of whether to post a potentially controversial status update on Facebook while waiting for the bus, to deciding which settings and services to use to best share photos with family members, to a group decision around how to share documents with collaborators for editing.

Each of these sharing decisions may encompass a range of dimensions, including available access-control channels, settings, the affordances of different mediums, people or relationships involved in the process, context, type or subject of content, need for verification or security, and underlying preferences. Preferences are rarely based on only a desire to let others view a piece of content, or prevent others from viewing content, at a single point in time. Instead preferences can range from wanting to allow mixed access, to wanting to provide push/pull-based access, to considering access amongst other tasks-at-hand or access based on attributes or features of one's audiences [43, 50, 56, 57, 78, 94].

In some situations, people are able to share content in a manner that matches their sharing preferences. They may do using a single service. Or, they may draw on features or audiences available on multiple platforms, in combination, to meet sharing needs [78]. In other situations, however, on-the-ground decisions may fall short of actual preferences at the time of sharing or at a later date [8, 94]. This gap between user preferences and actual sharing decisions can occur for a variety of reasons. Sometimes technology may fall short. Access-control settings or other mechanisms may not be available to meet users' desired preferences, or these tools may not be usable enough for regular use.

In other cases users may have underlying goals that come into conflict with, or negatively impact, their sharing preferences. For example, a user may wish to garner attention but may also wish to manage their self-presentation or identity [76, 92]. In still other cases users may feel that they are acting in a way that meets their sharing preferences at the time of sharing but, due to changes in life state, knowledge, or attitude they may later regret their sharing [6, 77, 96]

When sharing behaviors fail to meet users' conscious or unconscious preferences, a range of suboptimal outcomes can occur, including coping behaviors to address technological shortfalls [46, 92], regret at the time of posting or in the future [77, 96], threats to identity or presentation of self, undersharing or self censorship [22, 29, 76, 92], and inefficient use of services [94]. Sharing mechanisms should be designed to consider the full use context, as well as these potential suboptimal outcomes.

In this thesis I explore current everyday online sharing decisions, with a focus on attributes that drive these decisions, as well as current shortfalls; I apply knowledge of these behaviors to explore a method for interest-based sharing on Facebook

I focus on three general contributions:

Understanding shortfalls of current access-control mechanisms I draw on two user studies (performed with co-authors) to understand when current access-control mechanisms may fall short for different platforms and types of content. I explore the types of content, people, and sharing decisions for which current access-control mechanisms may not meet user needs. I focus on exploring self censorship and regret as potential suboptimal outcomes.

Understanding users' everyday online sharing decisions Through an interview- and diary-based user study I explore the range of peoples' everyday online sharing decisions across platforms, audiences, and types of content. I establish a baseline understanding of the range of sharing decisions people face throughout the day, as well as the task, audience, and channel-based-feature-driven factors that drive channel choices. I discuss the implications of these dynamics for the design of sharing mechanisms.

Explore the potential of interest-based sharing for Facebook Based on the user studies I find that people tend to base some desired content-sharing decisions around their audiences' interests in topics, a need that is currently unmet by some SNS mechanisms. I draw on a lab-based study to explore potential benefits of a mechanism for interest-based sharing on Facebook.

Thesis outline I begin by outlining high-level background and related work for the thesis in Chapter 2. I then describe two user studies that focus on understanding shortfalls in current access-control mechanisms. Chapter 3 describes a study that focused on understanding regretted messages on Twitter by comparing regretted Twitter posts to regrets in conversation. Chapter 4 describes a study that focused on the types of content participants chose to self-censor rather than post to Facebook as well as the portion, and type, of currently self-censored content, participants might have posted given ideal access-control

mechanisms. In Chapter 5 I then describe an interview- and diary-based user study that explores the factors driving everyday online sharing decisions. This study focused on the audience, task, and channel-feature-related dynamics that drive personal content sharing channel choices as well as the participants' uses of multi-channel strategies. In Chapter 6 I then describe a study that explored the potential impact of adding topic-based sharing mechanisms to Facebook. Finally, I discuss overarching conclusions in Chapter 7.

2 | Background and Related Work

In everyday situations people share different types of content with others using varied platforms and services. Prior work has addressed the importance and challenges of online access control, the impact of the shortfalls of current online access-control mechanisms, as well as potential improvements to access-control mechanisms both for general use and specific to social-networking sites (SNSs).

2.0.1 Online access control is necessary but difficult

Everyday content management, including file management, general communications and sharing, and distributing media like photos and videos, is increasingly moving online. Thus, users with a range of expertise levels increasingly rely on online platforms to manage and share content with others. However, prior work has found that people, even in everyday scenarios, can have complex privacy preferences. Providing usable access-control options for these preferences is challenging; however, when access-control needs are unmet, suboptimal access-control decisions can have negative results.

Online sharing presents a challenging environment

During everyday online interactions users often want to share content (e.g., publish photos, organize events, edit documents, etc.) online, sometimes with one or more people, often to varying degrees. They draw on a range of services to share content during these interactions, including email, SNSs, cloud services, or photo-sharing sites [28, 57, 66]. However, in this context, desired access-control policies can be complex, can include exceptions, and can be difficult or impossible to implement using existing tools and settings [9, 10]. Thus, everyday sharing presents a challenging environment for designing sharing tools that include usable access-control mechanisms.

The desire to selectively share also typically occurs as one component of an end goal or task. For example, a user may want to control access to who can view photos, but their end goal may be to publish the photos to their desired audience. This may necessitate that tools designed for broader use include access-control or selective-sharing components. Sharing

may take place during a wide variety of tasks on social-networking sites or other services, including maintaining contact with friends or family, sharing resources or information, providing or getting social support, having discussions (e.g., about politics), promotional activities, planning, having a conversation, sharing photos or documenting an event, or collaborating [37, 74, 78, 91].

The desire for access control may also, therefore, be further complicated by users who have benefit-oriented goals, for example publicizing or getting attention for a photo, that may be contradictory to access-control or privacy-focused preferences [38, 92].

Online platforms, such as SNSs, also present an environment that can complicate communication dynamics and lead to social, or self-presentation-related, risks that may not be present offline. These platforms may provide mechanisms that requires users to communicate with different types of audiences in a single environment, causing “context collapse" [55]. Platforms can create “group co-presence" by combining different social groups on one application when those groups might not otherwise typically be combined in users’ offline or professional lives [46, 47, 99].

Context collapse and the co-presence of different social groups can create challenges. People may share content intended for an “imagined audience" but may not understand the actual audience who can access the content they share. Users may also struggle to create access-control policies to distinguish between the varied groups who have access to content on the platform [2, 12, 49, 55].

Social media platforms may also develop platform-specific communications norms in accordance with which users may seek to share. For example, McLaughlin and Vitak found that Facebook users observed site norms from other users. Users tended to try to stay within site norms by not sharing too much, sharing highly emotional posts, having fights or private discussions in public, or by posting pictures that might reflect badly on others [59]. Similarly, work found that social media users might unfollow or unfriend others in response to these types of norm violations including sharing too much, sharing inappropriately, or sharing uninteresting content [45, 75].

These sharing and access-control preferences can be further complicated, because online privacy may also depend on others’ actions [52]. For example, on platforms with photo sharing and tagging functionality users may depend on others not to post potentially compromising photos, and rely on social norms and offline strategies to address unwanted photos that are posted or tagged [13, 47].

Sharing in everyday environments may be further complicated because users’ expertise levels, available devices and platforms, demographics, and technical backgrounds may vary more than they might in professional environments. Prior work has found that people use various platforms, including Facebook and email, for different purposes [28], and

users with different demographic or childhood technical backgrounds, tech savviness levels, knowledge or experience with systems, general personal network features, and access to systems, may vary in their use of both access-control settings as well as general communication channels [16, 20, 34, 35, 50, 87].

Features of the content being shared, such as data size, as well as system affordances including perceived convenience, reliability, privacy/security, use for archival and search, and features such as alerts and the ability to comment may also impact choice of communication channel for content sharing [11, 20, 33, 70, 98]. These differences may be more controlled in work environments in which devices, applications, formal access-control policies and hierarchies, are provided. However, in more informal, everyday settings, people often use a wide range of devices or tools, resulting in variety of strategies and ad hoc policies [57].

Thus, users cope with a challenging environment for everyday online sharing. In this context, users also have complex preferences, complicating the need for usable, transparent access-control and sharing options.

2.0.2 Current access-control decisions can be suboptimal

In this challenging environment, faced with complex preferences for sharing and access control, sharing decisions can result in suboptimal outcomes. In the absence of usable sharing or access-control options, people can fall back on coping strategies to try to achieve access control, which can result in suboptimal or inefficient use of services. Alternatively, users sometimes share content that they come to regret at the time of sharing or in the future [77, 96].

Strategies for audience targeting

Users may sometimes limit the audience to whom content is available. They may use formal access-control mechanisms provided by platforms. However, available online access-control tools can be unwieldly, time-consuming, difficult to understand, untrustworthy, or may not provide options that meet all users' needs.

When tools lack usable access-control options people may fall back on informal or ad hoc methods to try to limit the audience for content. For example, SNS users rely on a variety of informal strategies to limit access to the content they post including using multiple profiles to post content for different groups of people, trusting their friend group to maintain the groups' privacy norms, creating stated group rules around sharing, mentally targeting particular audience when sharing, or deciding not to think about posting in a specific manner [46, 47, 55, 83, 100].

Users may sometimes use different services for different purposes, for example using Facebook and email for different types of social interactions [28]. They may use different accounts and services to separate and manage different aspects of their online identifies. For example, users may draw on the different features and levels of control provided by different services to "facet" their identities for privacy and access-control purposes, identity management reasons, or to separate different types of tasks [28, 30, 41, 84, 93]. Similarly, users may limit their friends on a given service to control access to the content they post [90].

On SNSs, users may also cope with a lack of usable access-control options by deciding not to post, or to self censor, some subset of content they might consider sharing. Users choose not to post some types of content because the content might not be appropriate for all audiences who might view it, rather than relying on access-control tools to target specific audiences [22, 38, 46, 47, 55, 76, 90].

Coping strategies may sometimes be successful for preventing undesired access, but may simultaneously create varied inefficiencies. Self censorship can, for example, result in suboptimal use of a service. In a study of Facebook users' self-censored posts, we found that participants would have posted about half of self-censored posts if they had access-control tools that let them do so optimally [76]. Coping strategies may also be unsuccessful at preventing undesired audiences from viewing content. This can occur when users' mental models of access control do not match actual access control [40]. For example, access-control transference when users comment on or re-share SNS posts can be non-transparent to users or misunderstood, which can result in inadvertent or unclear sharing [97].

On SNSs users may also wish to ensure that particular audiences see their content, rather than only focusing on limiting audiences. Litt and Hargittai describe several "audience-reaching" strategies Facebook users may employ to target desired audiences when broadcasting content, including adapting the wording of the post, using features of the site to call attention to the post (e.g., tagging), and timing the post to target particular audiences. Using these strategies may allow users to, for example, reach broader "peripheral audiences," avoid the need to explicitly limit their audiences, or rely on the audience to self select if the content applies to them [51].

Regret

Suboptimal sharing can result in regret, either at the time of sharing or at a later time. Beyond online sharing, work has focused on regretted messages during offline conversation, as a subset of failure events, with an emphasis on types, causes, and efforts to ameliorate

conversational regrets [44, 60, 61, 62].

Prior work found that people tended to regret a variety of types of messages during conversation including blunders, attacks or criticism, making stereotypical references, expressive or cathartic messages or otherwise revealing too much, lies, or telling someone to behave in a specific way [44]. People also tended to associate highly emotional negative states as well as “having a lot on their mind" with saying regretful things [62]. To try to repair these types of conversational regrets, participants often apologized [61].

On Facebook, users have been found to regret posting content related to potentially sensitive topics like alcohol or drugs, sex, profanity, religion, and polticis, as well as negative or argumentative content [96]. Similarly, participants in our study of Twitter regrets tended to regret critical statements, blunders, and tweets that revealed too much. About half of these participants were able to successfully repair their regret, often deleting the tweet, and/or apologizing. However, compared to offline regrets, participants with regrets on Twitter took longer to realize they should regret statements and to repair the regret [77].

Thus, suboptimal online sharing can result in a range of types of regret. Compared to offline contexts, online scenarios present a challenging environment for preventing and addressing such regrets.

2.0.3 Varied modalities and attributes could improve online access control

Prior work has also focused on increasing the usability of online access control, both generally, and with a focus on selective sharing for SNSs. Work has focused on modalities around which access-control decisions can be based, for example changing the timing of access-control decisions or making access-control decisions attribute-based. Previous work has also focused on the attributes around which access control could be based, for example, different types of relationships or measures of relationship closeness. In line with these types of modalities and attributes, prior work has also focused on exploring, improving, and automating grouping tools for selectively sharing content.

Tag or attribute-based access control

Prior work has focused on improving the usability of online access control by allowing users to control access to content by defining access-control rules using user-created tags or other attributes.

Klemperer et al. found that tags created naturally by users for photos were viable for access control, and that users found hypothetical tag-based policies usable [43]. In online content-sharing systems, there are a range of potential attributes that can be drawn on

for such systems, ranging from user-defined tags to system-defined metadata that users commonly draw on for search and recall like file location, type or format, time of last usage, keywords, or events associated with the content [15].

Several systems have been proposed that offer access-control decisions using tag or attribute-based policies. For example, Au Yeung et al. created a prototype system for creating access-control policies for Flickr using descriptive tags and linked data for photos [101]. Hart et al. created a tag-based system to provide access control for Wordpress, an online blogging system. They found that users were able to create policies more quickly, and with equal accuracy, using the tag-based tool as compared to traditional tools [32]. More broadly, Mazurek et al. proposed a distributed, attribute-based file access-control system that was able to express policies for personas drawn from user studies with low overhead [58].

Grouping tools

Previous work has also focused on exploring the groups of people with whom users may share content and creating improved tools and interfaces for providing users with automatically-created groups of people with whom to share.

Several SNSs, including Facebook and Google+, provide manual or partially-automated grouping mechanisms, like Facebook's Friend Lists or Circles on Google+. Researchers examined the types of groups that emerge on these sites, how well these mechanisms capture the types of groups that emerge from friends present on these sites, and how these mechanisms are typically used.

In a field study of the Circles created by early Google+ users, Kairam et al. found that participants tended to create Circles that reflected either "life facets" like work or school or strong or weak tie strength [38]. Work also found that users used these Google+ Circles for a range of purposes, beyond privacy-based selective sharing, including directing content to appropriate audiences, for appropriateness, relevance, and to try to maximize audience [38, 97].

Prior work has also examined both how people use Facebook's Friend Lists feature, as well as how well the feature matches types of groups people naturally create from their friends in different scenarios. For example, Facebook users with larger and more diverse networks were found to more frequently use Friend Lists, often to target specific audiences to "recreate some of their offline contexts" or to target relevant people in their networks [89, 92]

Several studies have also found that, when asked to group their Facebook friends, people typically group them into groups that correspond to life-stage or contextual re-

lationships. Kelley et al. found that, when asked to group their Facebook friends using various lab-based methods, groups participants created tended to correspond to school, family, specific locations, and people the participants couldn't identify [42]. De Wolf et al. had similar results for a study of young adults in Belgium, finding that participants tended to categorize their friends according to interest-based categories, geographic-community-based categories, people who knew each other, mutual friends, types of contacts, or personality traits [23]. In a study of categories of SNS-friends in Singapore, Zhang et al. also found that participants tended to describe school, work, interest, and family-based groups, but found some variation by gender, ethnicity, and age [102].

While people were able to describe a number of groups into which they could sort their existing friends, static groups were not found to perform well in supporting real-time sharing decisions, which suggested that grouping tools might need to be dynamic to be effective on an ongoing basis [42]. As an alternative to specific static groups, Wiese et al. found that self-reported closeness was the strongest predictor of willingness to share various types of information [99].

Prior work has also examined creating machine-learning-driven automated grouping tools that use various attributes to algorithmically predict groups for sharing or privacy settings. Although prior work has found that permanent access-control list membership can be difficult to predict using traditional algorithms [25, 63], several systems have used automated or partially-automated interfaces to provide more usable options than traditional interfaces.

For example, Amershi et al. created an interactive automated grouping tool called ReGroup for sharing content on Facebook using seventeen demographic, life stage, interest, and social features. It iteratively learned, and presented, groups to users and suggested additional members and characteristics for filtering the groups. Participants found it more effective for creating large and varied groups than the traditional grouping models [4]. Similarly, Fang and LeFevre designed a "privacy wizard" that created privacy policies for Facebook friends based on user input for a subset of their friends, gathered over iterative rounds [27].

Timing of access control

Another method for improving the effectiveness of access-control tools is by changing when users can set or adjust access-control policies.

Mazurek et al. looked at the usability and utility of allowing users to make reactive, rather than a priori access-control decisions through an experience-sampling study. They found that reactive policies can facilitate policies that are contextually-dependent but

difficult to define using traditional models. Additionally, many participants preferred reactive, or partially-reactive, systems to traditional systems [57]. Similarly, Bauer et al. looked at a smartphone-based door-unlocking system and found that the ability to reactively provide permission rather than distribute keys a priori helped create policies that better matched user preferences [9].

Work has also looked at the possibility of sharing impermanent content. Ayalon and Toch found that sharing preferences for Facebook content faded over time and changed based on life events, suggesting the potential for impermanent sharing mechanisms on Facebook [6]. However, Bauer et al. found that while some participants wanted the visibility of posts to change over time, participants tended not to be able to accurately predict which posts they would prefer to become more private, indicating difficulties in creating a priori fading mechanisms [8].

Identifiability when sharing

Access-control tools can also be used in combination with systems that allow the people sharing to be more or less identifiable. For example, some services, like Facebook require that users' shared content be tied to a real name; other services, like Twitter, allow content to be tied to a pseudonym, while other services, like YikYak, allow content to be shared anonymously [82].

Increased anonymity can have both benefits and downsides. Anonymity can lead to increased disinhibition, which can allow users to express or access potentially sensitive or controversial content, such as critiques or feedback. Users may also be more comfortable accessing support groups or asking questions about taboo topics in anonymous forums. For example, on Facebook, people use anonymous "confession boards" to ask questions about topics like sex or drugs [14, 39, 82]. Anonymity can also allow users to improve their ability to control how they manage online-self-presentation boundaries [39, 82].

The increased disinhibition that comes with anonymity can also lead users, however, to become less civil or to use anonymity for illegal or other "socially undesirable activities" [39]. For example, anonymous, or partially anonymous users on Twitter were found to tend to tweet more, and to tweet on sensitive topics such as porn, sexual topics, or drugs [67]. Similarly, anonymity can make content appear untrustworthy if anonymous sharing is not the norm for a platform [82].

3 | Understanding shortfalls: Regrets on Twitter

To create selective-sharing mechanisms that effectively help users share in everyday environments, it is important to understand how current sharing mechanisms may fall short. One metric for measuring how sharing mechanisms might be improved is to examine where people currently share content online and then later regret their decision to share. Improved sharing tools can seek to prevent these types of regrets or help users more rapidly, or successfully, repair regrets when they occur.

In everyday online sharing contexts, people often use social-networking sites, like Twitter, to share messages that they might otherwise share during conversation offline. When looking for opportunities to improve sharing mechanisms for these sites, we chose to compare regrets on Twitter to regrets in conversation. The dynamics of regrets for messages during conversation offline have been studied extensively. Thus, comparing regretted messages on Twitter to regretted messages in conversation allowed us to determine where online regretted messages might differ from offline regrets

These differences provide insight into where it might be possible to add or adjust online sharing mechanisms to help prevent regretted messages or ameliorate their effects, in relation to the types of regrets one might experience in everyday offline life.

In this chapter I describe a survey-based study (performed with colleagues) that explored, and compared, causes of regret and actions to repair regret on Twitter and in conversation. I highlight potential uses of the insights for online mechanisms to reduce and ameliorate the types of regret that may occur when sharing on Twitter. The majority of this chapter previously appeared in the proceedings of CHI 2013 [77].

3.1 Introduction

It is easy to say something you regret, angrily insulting a loved one or inadvertently letting a secret slip. However, Twitter, a social networking service, enables these types of regrettable messages to spread rapidly and broadly, and to remain available for extended periods of time. Twitter's ability to broadcast messages widely and retain them indefinitely potentially alters the dynamics of regretted communications. In extreme cases, Twitter

has enabled highly-publicized instances of regret, like Rep. Anthony Weiner's infamous tweet that led to his resignation [17]. However, everyday Twitter use can lead to more mundane regrets. As in conversation, Twitter users insult others, accidentally reveal private information, and express emotion in heated moments.

Thus it is worthwhile to investigate regret both on Twitter and for in-person conversations. Past studies of in-person regret have identified factors that lead to regret, methods for becoming aware of regret, and strategies for repairing harm [44, 61, 62]. However, Twitter presents different features and limitations than offline conversation. Beyond offering wider audiences and increased message persistence, Twitter lacks face-to-face channels, such as body language, for transmitting apologies or indicating offense.

We explore regretted messages Twitter users posted on Twitter or said during in-person conversations. We aim to improve understanding of regrets on Twitter by comparing them with in-person regrets. By examining these regrets, as well as how people became aware of regrets in person and on Twitter, we also identify preliminary design directions for preventing and ameliorating regrets on Twitter.

Specifically, we examine four research questions:

- Q1: What states of being lead to regret on Twitter and in person?
- Q2: What types of regret occur on Twitter and in person?
- Q3: How do people become aware of regretted messages on Twitter versus in person?
- Q4: What repair strategies do people use to cope with regretted messages on Twitter and in person?

To address these questions, we ran a 1,221-participant online Mechanical Turk survey with two conditions. In one condition, we asked Twitter users to report on one message they regretted saying during an in-person conversation. In the other, we asked parallel questions about a message they regretted posting on Twitter. We collected information on the incident, the participant's emotional state preceding the incident, how the participant became aware of the regret, and any mitigation strategies employed. We used these answers to understand and compare drivers and consequences of regretted messages during in-person conversation and on Twitter.

3.2 Twitter background

Twitter is an online social-networking site where users post tweets, which are text-based messages of 140 characters or less. These messages are broadcast to a user's followers in relationships that are often asymmetric.

Twitter has several conventions that aid in sharing. Users can direct a message to a

handful of specific users by crafting an @-reply. Users indicated by the @-reply will be alerted to the message through email or the Twitter client, but the message itself is public. A direct message (DM) allows a user to send a private message to a single person. A user can also add #hashtags to a tweet to categorize it, better enable searches as part of a trend, or provide contextual information. Tweets are publicly accessible unless an account is protected. Only a user's approved followers can view a protected user's tweets.

3.3 Methodology

Our goal was to analyze regrets that Twitter users had experienced on Twitter and during in-person conversations. We conducted a large-scale online survey from August to September 2012 using Amazon's Mechanical Turk (MTurk). We asked each of 1,221 MTurk Twitter users to describe one thing they had said and then later regretted (the regretted message) either during in-person conversation or on Twitter, depending on the condition to which the participant was assigned. We collected a description of the message, the context, how they became aware of the regret, and how they sought to repair the regret. It took participants 14.5 minutes on average to complete the survey, for which they were paid $0.75 (within the typical pay range for MTurk [19]).

3.3.1 Participant selection and conditions

We screened for US MTurk workers over 18 years old who self-reported English proficiency and relatively frequent Twitter use (having had a Twitter account for at least a month and posting at least monthly). Of the 3,175 MTurk workers who started the survey, 946 did not meet these requirements. The majority (609) were disqualified for posting less than once a month on average.

3.3.2 Survey

Conditions After the initial screening questions, participants were split into two conditions in a round-robin fashion. The first condition was conversational regret, which mirrored previously described work. The second condition asked parallel questions, slightly reworded to focus on Twitter regret. In both conditions, participants were asked to recall a time when they said or tweeted something and then regretted it, with the wording and format of the prompt based on Meyer's work on in-person messaging regrets [61, 62]. Our prompt for **conversational-regret** participants was:

> "Please recall an occasion when you **said** something during an **in-person** conversation and then regretted saying it. This may be something that you

regretted saying immediately or that you regretted saying later."

Our **Twitter-regret** prompt was similar:

> "Please recall an occasion when you **tweeted** something and then regretted tweeting it. This may be something that you regretted tweeting immediately or that you regretted tweeting later."

Survey structure Participants in both conditions who could not recall a regret were directed to an alternate survey that asked them about why they did not have regrets. We do not report the results of this survey, as the goal was only to ensure an equal workload for either positive or negative responses. Of the 1,879 participants who qualified for the study, 601 (456 for Twitter and 145 for conversational regret) could not recall regrets.

Participants who were able to recall regrets completed a survey about the regretted messages they reported in response to the initial prompt. The survey drew heavily on questions and structure from in-person messaging regrets work [44, 61, 62] and included several groups of related questions. We asked participants about the following:

Regretted message description: a series of essay questions that asked the participant to describe the message in detail, including the context, the reason why they said/tweeted it, the intended audience, the audience's reaction, why they regretted the message and any consequences

Circumstances: follow-up questions about their state when they delivered the message

Awareness: free response about how they became aware that they should not have said the message, followed by a multiple choice selection of how quickly after the message they realized they should regret the message

Repair strategies: a description of whether, how, and how successfully they tried to repair any harm caused by the message; participants were also asked to rate the seriousness of the regret before and after repair

Twitter specifics: questions on Twitter usage (e.g., client and device tweeted from, is/was the account protected)

Demographics: basic demographic questions

We based the general survey structure on the format used in previous work on in-person regrettable messaging [61, 62]. Specifically, we used Meyer's format of asking participants to provide one regret and then probing for details. Although this format has several weaknesses, as outlined in Limitations, it has been used repeatedly to examine in-person messaging regrets.

Quality control on Mechanical Turk While MTurk has been shown to produce quality samples and results [19], surveys on MTurk should be designed to encourage quality responses.

We took several quality control measures. First, we only used MTurk workers who had over a 95% approval rating on the site. Second, we front-loaded longer essay questions. By putting these questions earlier in the study, we encouraged lazy or unmotivated participants to drop out early or to enter nonsensical data where it was visible. It also made it easy for honest survey participants to return to the task, without feeling like they still needed to invest large amounts of time. We removed a small number of participants (25) from the dataset who provided nonsensical or non-English answers to the free response fields.

We also removed responses from 32 conversational-regret participants who responded about a regret on Twitter. We believe they did so because they were primed to think about Twitter when recruited as Twitter users. An additional 350 participants were removed for not completing the survey.

Data analysis We surveyed MTurk users who posted on Twitter about a regretted message either said in-person or posted on Twitter. Although the surveys for each condition were designed to be parallel, the fundamentally different contexts preclude statistical comparisons between conditions. To explore characteristics of how regret on Twitter compares with in-person regret, we present the results of the Twitter- and conversational-regret conditions side-by-side. The proportions of participants reporting different answers are only meant to illuminate general themes and trends, not to be compared statistically.

Within a single condition, we perform statistical analyses. We use logistic regression to evaluate the relationship between types of regret and whether the audience was a group or individual, the relationship between awareness mechanisms and whether or not regret was experienced immediately, and the impact of repair strategy on the success level. Demographics were compared between conditions using a Wilcoxon test for numerical data and χ^2 tests for categorical data. All tests use a significance level of $\alpha = .05$.

3.3.3 Participant demographics

After quality-control removals, 1,221 people reported regrets: 747 for conversational regret (72% of those who started) and 474 for Twitter (41%). The mean age was 30.3 (28.2 for Twitter and 31.7 for conversational regrets). Overall, 53% of participants were female and 46% were male (10 preferred not to answer). The gender breakdown was almost identical for the Twitter- and conversational-regret conditions. Of the participants, 26% were students and 10% were unemployed. The remainder were primarily employed in science (9%), service (8%), and art (8%) occupations. There were no significant differences between the Twitter- and conversational-regret participants in age, gender or occupation, nor were there significant demographic differences between participants who did and did not report regrets.

3.4 Analysis and results

3.4.1 States of being leading to regret

People often say things they later regret because of demands on mental capacity that impair thought processes. We found that both Twitter- and conversational-regret participants were often in negative, highly emotional states prior to regret. Meyer outlines several factors that contribute to “cognitive load," “physiological state," and “emotional state," which can potentially lead to regret [62]. We asked participants about these states. Based on Wang et al. [96], we also asked whether they were drunk at the time of the message. We asked participants to rate on a five-point scale how much or how little each factor applied immediately before they tweeted or spoke. A one indicated “Not at all" and a five indicated “Very much so." They rated each of the following: “I was fearful or frightened," “I had a lot on my mind," “I was feeling excited," “I felt ill," “I was worried," “I was nervous or anxious," “I was drunk," “I was angry," “I was stressed," “I was tired/fatigued," “I was happy," “I was hung over," and “I felt frustrated."

Consistent with Wang et al.’s work on Facebook regrets [96], we found that both in person and on Twitter, highly emotional negative states were most common prior to regret. Participants commonly reported a four or a five for stress (46% of Twitter and 50% of conversational participants), anger (51% and 43%), or frustration (58% and 53%) prior to the regrets. Participants also often had something on their minds (54% and 51%). Somewhat less common were positive emotions, including feeling excited (26% and 17%) or happy (22% and 21%).

3.4.2 Types of regret

We also looked at types of regrets participants reported for Twitter and for in-person conversations. In both conditions, participants most commonly reported regretting messages that were critical of others. However, on Twitter, participants more commonly regretted content that was expressive/cathartic and that was intended for groups of people.

Types of regret We coded each regret described by participants into one of Knapp et al.'s categories for types of regretted in-person conversational messages [44], specifically:

- **Blunder:** "not normally perceived by a third-party observer as problematic"; mistakes, factual issues; includes typos or errors during conversation
- **Direct attack:** "critical statements directed at a person, the person's family, or the person's friends [...] general rather than specific"
- **Group reference:** stereotypical references about a group (e.g., ethnic, racial)
- **Direct criticism:** critical statements about "something specific" about a person
- **Reveal/explain too much:** telling "more than the situation calls for"; e.g., undesired personal information or a secret
- **Agreement changed:** agreeing to something, then later changing one's mind
- **Expressive/catharsis:** general "expressions of feeling and emotion"
- **Lie:** "knowingly lying to another person"
- **Implied criticism:** "critical remarks that are implicit" and can be "teasing remarks"
- **Behavioral edict:** telling someone to behave in a certain way

Two coders independently coded all the regrets based on Knapp et al.'s categories. Two coders reached a consensus for any regrets for which there were discrepancies.

Across both conversational and Twitter regrets, participants most commonly regretted critical statements (Table 3.1). Common critical statements included direct attacks and direct criticisms; 29% of conversational and 20% of Twitter regrets were direct criticisms, while 14% of conversational and 13% of Twitter regrets were direct attacks.

Blunders also arose frequently for both conversational and Twitter regrets, although more often for conversational (11% for Twitter, versus 16% for conversational). Although both Twitter- and conversational-regret participants reported some similar blunders, such as saying/posting messages they later found out were false or that had been said/shown to someone who found them offensive, some blunders were unique to Twitter. On Twitter, time-delayed blunders sometimes caused participants to regret messages because of an event or change in context. For example, one participant regretted tweeting about a drive-by

Participant-Reported Types of Regret

	Twitter		Conversation	
Reveal too much	117	25%	105	14%
Direct criticism	96	20%	213	29%
Expressive	64	14%	15	2%
Direct attack	62	13%	108	14%
Blunder	51	11%	120	16%
Implied criticism	34	7%	84	11%
Group reference	13	3%	21	3%
Agreement changed	3	1%	10	1%
Behavior edict	2	0%	28	4%
Lie	1	0%	25	3%
Other	31	7%	18	2%

Table 3.1: Types of regret for Twitter and Conversation

shooting in his friend's hometown when that friend was later killed in a drive-by shooting. Twitter, as an online interface, also allowed blunders caused by typos and broken links, which several participants found embarrassing. For example, one participant reported being "made fun of" for tweeting that he "used a lot of hags on [his] car."

Participants also regretted expressive or cathartic content more frequently on Twitter than in person (14% versus 2%). These expressive statements were typically tweeted when participants were angry or upset. They often served to vent or express frustration on topics such as work, relationships, or politics. Often, the goal was to allow others to sympathize or "know what [the participant] was going through." Participants tended to regret the message later after re-thinking how it would sound, or after someone who viewed it became upset. For example, one participant described tweeting "Last day of my internship, so excited to be done," because she "was unhappy with how the internship treated [her] and what had happened [...and] wanted [her] friends to see it because they knew [she] was having a rough time." However, she regretted the tweet when her internship coordinators saw it and sent her an email telling her she needed to delete the tweet. In contrast, expressive regrets during in-person conversations tended to be part of arguments or opinions.

Type and audience Participants also specified whether they intended the messages to be seen or heard by individuals, or by multiple people. Twitter-regret participants were more likely to target multiple people (73% of Twitter regrets, versus 24% of conversational), likely because of Twitter's broadcast capabilities.

Certain types of regretted messages were more frequently intended for multiple people, especially on Twitter. When the intended audience comprised multiple people, rather than

an individual, Twitter-regret participants were significantly more likely to report a blunder ($p = 0.008$), content that revealed too much ($p = 0.005$), or expressive/cathartic content ($p = 0.003$). Of Twitter blunders, 82% were intended for multiple people, versus 33% of reported in-person blunders. Twitter-regret participants often said that they wanted to tweet to friends, coworkers, or others interested in a specific topic, but regretted the tweet because they made an error that caused confusion or made them look bad. For example, one participant reported tweeting, “Congratulations to B for being elected ALA Councilor," intending the message for other librarians in South Carolina. She later realized that the individual was actually a candidate for the position, rather than having been elected, and regretted the tweet because “it was embarrassing."

Twitter-regret participants who regretted expressive or cathartic posts also tended to target multiple people rather than an individual (84% of expressive/cathartic regrets). Participants often hoped to share political or negative feelings with the general public or their friends because they “wanted to vent" or express their feelings “to anyone that would listen."

Regretted statements on Twitter that revealed too much also tended to be targeted at multiple people (80%). Many participants tweeted personal information, such as details about their lives or relationships, and then regretted sharing them on Twitter. Several participants also reported having both personal and professional accounts and regretting tweeting personal information on their professional Twitter accounts. For example, one participant said that he regretted tweeting “on my professional twitter account about a night of heavy drinking" because it seemed “unprofessional."

In contrast, conversational-regret participants were significantly more likely to report regrets that were direct attacks ($p = 0.024$) when the intended audiences were individuals (67%) rather than multiple people. Participants were typically angry or arguing with the recipient of the message. For example, one participant “screamed at my father that ‘I hate him’ in an argument" because his father kept him from attending a party. On Twitter, such attacks were commonly focused at groups (68%), and participants reported wanting their anger to be seen. For example, one participant had a conflict with a friend, and wrote “she’s so annoying and whiny," intending “it to be seen by friends."

Unintended audience We also coded for regretted messages having unintended audiences. In conversation, unintended audiences included people overhearing messages (e.g., by walking into a room) or being told about them. On Twitter, most of the tweets reported were public tweets. However, participants still had particular audiences in mind when they tweeted. Unintended audiences occurred because people other than the intended audiences saw or heard about the tweets.

For Twitter regrets, 13% had unintended audiences, compared to 5% of in-person regrets. Unintended audiences occurred most commonly on Twitter for regrets that revealed too much (23% of regrets that revealed too much), often because participants tweeted something private, insulting, or about work, which they later realized they didn't want everyone to know. For example, one participant described how she tweeted "something sexual and my [T]witter at the time was public, so I freaked out when I saw that my brother's screen name popped up on Recommended Twitter."

Level of regret To measure level of regret, we asked participants "In your opinion, how serious of a problem was it that you said the messages, at the time you said it" (or tweeted it), based on a question from [61]. Participants responded from one ("Not at all") to five ("Very much so"). We consider participants who reported a four or a five to have had a high level of seriousness and below a four to have had a low level.

For Twitter, 18% of messages had high levels of seriousness. For conversational regrets, 38% had high levels of seriousness. However, the interpretation of the difference is somewhat ambiguous; the seriousness of regrets across contexts may not be directly comparable. For instance, a serious conversational regret may differ from one on Twitter.

3.4.3 Awareness of regret

Individuals must become aware of regrets to address them. Conversational-regret participants tended to become aware of regret more quickly and relied more on audience actions, such as body-language cues. Twitter participants more often reported realizing regrets themselves or had audience members tell them they should regret the message.

Means of awareness We asked each participant to describe in a free response how they became "aware [they] shouldn't have said the message." Two coders created a set of codes for means of awareness based on types of awareness outlined in Meyer's work on regretted messaging [61] using a set of 100 regrets (Table 3.2). The same two coders then independently coded the regrets based on these codes. A third coder also independently coded the regrets to break ties. In cases where all three coders disagreed, two coders reached a consensus. A regret could be coded for multiple, different means of awareness.

Participants became aware of regret using different means on Twitter and in person (Table 3.3). This is partially explained by the different contexts for Twitter and conversational regret. Audience body language is usually immediately available in person but typically absent on Twitter. Thus, 19% of conversational-regret participants described using audience body language to become aware of regret. Participants often realized the regret immediately when they saw their audiences' facial expressions. For example,

Descriptions of Means of Awareness	
Self realization	The individual realizes either by thinking about it or by just feeling bad that they should regret the message
Audience says something	The intended audience says something to imply that the person should regret the message
Audience takes an action	The intended audience does something to imply that the person should regret the message (e.g., stops speaking to the individual)
Audience body language	The individual realizes they should regret the message based on the intended audience's body language (e.g., smile, frown)
Third party says something	A person other than the intended audience says something to imply that the person should regret the message
Third party action	A person other than the intended audience does something to imply that the person should regret the message
Third party body language	A person other than the intended audience uses body language to imply that the person should regret the message

Table 3.2: Codes for means of awareness

Participant-Reported Means of Awareness

	Twitter		Conversation	
Self realization	58%	275	39%	294
Audience said	29%	138	17%	126
Audience action	7%	32	26%	191
Audience body lang	0%	1	19%	143
3rd party said	7%	33	5%	39
3rd party action	1%	5	1%	8
3rd party body lang	0%	1	0%	3
Other	1%	6	0%	3
Total		474		747

Table 3.3: Means of awareness for Twitter and Conversation

one participant reported calling “his cousin an asshole in-front of our entire family" and realized he should regret it “[w]hen everyone glared at me."

Conversational-regret participants were also more likely to report relying on audience actions to become aware of regret (26% for conversation, versus 7% for Twitter), also likely due to the intended audience’s physical presence. Such actions included storming out of a room, laughter, or sitting silently, which are difficult to convey over Twitter. Offline followups to Twitter messages, such as job termination or laughter, led to awareness for Twitter regrets, as did Twitter-specific online actions, such as being unfollowed or ignored.

Comparatively, Twitter-regret participants more frequently became aware of regret on their own (58%, versus 39% for conversational regrets). Participants in both conditions would often realize that the regretted message was something that they should not have said or tweeted, either after thinking about it or because they felt bad. As one participant put it: “Something inside just told me it was wrong." However, on Twitter, messages also remain available over time. Several Twitter-regret participants reported re-reading the message later and realizing that they should regret it, an option that is rarely available in person. For example, one participant tweeted, “Absolutely pointless," about her relationship and realized she should regret it when she “read over [her] tweets the next morning and thought it was dumb."

Twitter-regret participants were also more likely to report that their intended audience said something to imply that they should regret the message (29% of Twitter, versus 17% of conversational). This may partly reflect the wider audiences targeted by Twitter users but also how, on Twitter, people helped participants realize they should regret a message. Often, a friend or co-worker saw the message and contacted the participant to tell them

that they should regret it. For example, one participant tweeted “Having fun on my day off. #callinginsick" and realized he should regret it when “[o]ne of [his] friends told [him] it wasn’t a good idea."

Time until awareness Conversational-regret participants also became aware of regrets more quickly than participants on Twitter. Based on wording used by Meyer [61], we asked participants "how much time passed between" when they tweeted or spoke and when they became aware they shouldn't have tweeted or said the message. We found that the majority of conversational respondents became aware immediately (62%), with many of the remaining participants becoming aware within a few minutes (18%). Of the remaining 20%, the majority became aware the same day or the next day (13%). On Twitter, participants reported taking longer. Only 11% were immediately aware, while 29% realized within a few minutes, 33% at some point the same day, and 16% the next day. The majority of the remaining 11% became aware of the regret within a few days.

For some types of awareness, participants were more or less likely to become aware immediately. On Twitter, participants were significantly less likely ($p = 0.028$) to become aware of the regret immediately (4%), rather than later, when the audience said something to imply that they should regret the tweet. This is consistent with users tweeting and audience members later informing them that they should regret the content, implying a time delay. For conversational regrets, participants were significantly more likely ($p < 0.001$) to learn immediately (84%) from audience body language about a regret. They often reported realizing as soon as they spoke that they should regret the message due to the audience's physical reactions. As one participant reported, "The moment it slipped out, I knew I shouldn't have. The awkward looks and silence that followed confirmed that it was as bad as it sounded." In contrast, conversational-regret respondents were significantly less likely ($p < 0.001$) to become aware immediately (13%) when a third party told them something to imply that they should regret the message. The person about whom they were talking, or who was impacted by the message, often contacted them, delaying awareness. For example, one participant "told a coworker that I intended to leave my job in an open area" and regretted it "[w]hen I went to meet with my boss she told me she had heard rumors."

3.4.4 Repair strategies

After becoming aware of a regretted message, people often employ strategies to repair the impact, or potential impact, of the message. We asked participants about the repair strategies they used after tweeting or saying the messages, as well as the impact of these repair strategies. We found that conversational-regret participants most often chose to apologize, while Twitter-regret participants most often chose to delete regretted tweets. As occurred in regret awareness, Twitter-regret participants also took longer to repair regrets than conversational-regret participants.

Participant-Reported Repair Strategies

	Unsuccessful		Successful	
	Twi.	**Conv.**	**Twi.**	**Conv.**
Delete	111	–	134	–
Apology	53	173	72	218
Act like nothing hppnd.	44	70	38	42
Excuse	36	92	34	55
Justify	38	89	30	64
Say something to offset	17	77	22	67
Deny	10	50	10	31
Non-verbal behavior	–	40	–	30
Other	11	21	5	21
Apology and delete	30	–	38	–
Apology and justify	15	49	16	43
Apology and offset	5	52	12	45
Apology non verbal	–	25	–	19
Total (participants)	191	329	196	302

Table 3.4: Repair strategies for Twitter and Conversation

Frequency of repair strategy We asked each participant to select repair strategies they used from a list taken directly from the conversational-regrets literature [**?**]. Participants in both conditions were provided with the options: "I tried to say something to offset the harm done," "I tried to justify or defend what I said to minimize its offensiveness," "I apologized for saying it," "I just acted like nothing had happened," "I denied or tried to take back what I said," "I offered an excuse for why I said it," "I didn't do anything." Conversational-regret participants were also offered the option "I employed a nonverbal behavior to indicate that I regretted it" (from the regrets literature), while Twitter participants were offered "I deleted the tweet."

Overall, we found that a similar proportion of Twitter- and conversational-regret participants took actions (did not report doing nothing) to repair regrets (82% and 84%, respectively). However, the distribution of repair strategies varied (Table 3.4). Conversational-regret participants most frequently chose to apologize (34% of strategies). Twitter-regret participants most often chose to delete regretted tweets (37%), an option unavailable in person. Both conversational and Twitter participants were relatively likely to try to make an excuse (11% of Twitter and 13% of conversational strategies), justify their messages (10% and 13%), and act like nothing had happened (12% and 10%). However, conversational participants were more likely to try to say something to offset the harm (12%, versus 6%

for Twitter).

Success of repair strategies These different repair strategies also met with varied levels of success (Table 3.4). Participants rated, on a five-point Likert scale, how successful or unsuccessful their repair strategies were. Participants who ranked their strategies as "successful" or "very successful" were categorized as having successfully repaired the regret. Approximately half of each of Twitter- and conversational-regret participants who took repair actions were successful. Controlling for seriousness of regret at the time of the message, several repair strategies emerged as significantly more likely to be successful or unsuccessful.

On both Twitter and in conversation, using an apology significantly increased the probability of success ($p = 0.043$ and $p < 0.001$ respectively). In person, making an excuse significantly decreased the probability of success ($p = 0.002$), while on Twitter, deleting the tweet significantly increased the probability of successful repair ($p = 0.038$).

Participants who apologized on Twitter varied in their use of online and offline apologies. Online, they apologized using a variety of means, including tweets, instant messages, and text messages. Offline, they apologized face-to-face or by calling impacted individuals. This choice of online or offline strategy seemed to depend on level of personalization and context. Several participants chose to apologize offline because they were confronted about a regretted tweet in an offline environment. For example, one participant apologized when his tennis coach confronted him about an insulting tweet and and told the coach that he "would delete the tweet immediately." Other participants reported apologizing in person to make the apology more personal, writing, "It was personal," so "I called them personally."

Twitter is often a relatively public forum, and, as the regretted tweets often reached wide audiences, apologizing online could also allow participants to reach larger audiences. Participants reported using online apologies to add additional information to their original tweets or add corrections. For example, one participant described accidentally posting misinformation about an animal rescue. After realizing her mistake, she tweeted a correction and an apology. Online apologies were also used to reach large groups of people. One participant described how she "tweeted back so everyone could see my apology and called the person" that she had upset.

Apologies after regretted tweets were also often paired with other online actions. Of the regretted tweets participants apologized for, 54% were also deleted. After posting "something passive-aggressive about someone," one participant described how she tried to repair the situation by telling her "friend that I'd acted immaturely and that I was sorry." She also "deleted the tweet because [I] was embarrassed by my actions."

For in-person regrets, apologies tended to be offline and verbal, often face-to-face to a

single person involved with the regret. For instance, one participant jokingly “insulted a friend only to find out his mother had passed away earlier in the week and hadn’t told anyone." Once he found out, the participant “immediately apologized stating that [he] didn’t know and offered [his] condolences." Such apologies were often paired with justifications (23% of conversational apologies) or explanations that tried to offset the harm (25%). One participant described criticizing how her husband had done the household chores. She explained that she “apologized, and I think maybe explained that I hadn’t meant to sound as rude and critical as it sounded. I also thanked my husband for the work he had done and said that I was glad he was so helpful."

Time to repair Varied amounts of time passed before participants addressed the regretted messages. Participants responded in free-text to “When did you take these actions?" Two coders coded responses for all participants who used repair strategies other than acting like nothing had happened (1127 participants), based on the indication of the first repair. The coders reached a consensus on any disagreements. The categories were: Immediately/a few minutes after the regret (15 minutes or less), the same day, the next day, more than a day but less than a week, more than a week but less than a month, and one month or more. For 32 participants (29 for Twitter and 3 for conversation), the time period was unclear.

Conversational-regret participants tended to respond more quickly, as might be expected because they also become aware of the regret more quickly. Of conversational-regret participants who actively tried to repair their regrets, 392 (67%) did so within a few minutes. The majority of the remainder did so the same day (78 participants, 13%) or the next day (49 participants, 8%). Alternatively, only 98 Twitter-regret participants (26%) who actively tried to repair their regrets did so within minutes; 131 (34%) tried to do so the same day, and 74 (19%) did so the next day. The majority of the remaining 10% took less than a week.

3.5 Limitations

There are limitations in our study design. We performed this study using Mechanical Turk. Although this potentially biases our sample, MTurk’s population biases have been documented [72]. Samples and results from MTurk workers have also proven comparable to other online sources [19, 36]. We also took several measures to ensure quality responses. However, such quality control measures may also have biased our participant pool, potentially electing for more diligent or intelligent workers. It is unclear how this impact might differ from quality-control measures used for other survey methodologies. However, previous conversational-regrets work drew from an undergraduate population [61, 62]; using MTurk allowed us to expand to a large, cost-effective sample relative to offline pools

or alternative online sources.

Our survey design had additional, inherent limitations. We used the basic design from the conversational-regrets literature [61, 62] in which each participant recalled a single, regretted message. Thus, we don't have a true analysis of the frequency of different types of conversational or Twitter regrets. Based on the conversational-regrets design, we asked participants for the regret that first came to mind, rather than the most recent or strongest regrets. However, certain regrets may come to mind more easily or may be more or less embarrassing to detail in a survey. Thus, we may have an overrepresentation of memorable regrets and an underrepresentation of deeply shameful regrets.

The survey format was also a limitation. We asked participants for self-reported, recalled data. Participants may attribute more meaning to events occurring in the past when reporting on them in a survey. There was also potential for reverse causality issues. We tried to limit causality questions, but participants may have attributed factors like states of being to the regret, when they were actually caused by the regret. We could offer more conclusive results if we tracked participant behavior over time and noted actions, like repair strategies, as they occurred. For example, a diary-study approach could be used to supplement this work.

3.6 Discussion

We found that Twitter- and conversational-regret participants differed in the types of messaging regrets they reported, how they became aware of the regrets, and how they tried to repair the harm caused by the regrets. Time delays on Twitter, as well as lack of face-to-face communication with audiences, also caused awareness and repair on Twitter to occur more slowly than for conversational regrets. Based on these findings, we offer several early potential design directions for helping users prevent and repair Twitter regret.

Detecting and preventing regret on Twitter Although our participants took measures to repair harm caused by the regretted messages, they often would have liked not to have tweeted the messages. One way to potentially prevent regret on Twitter would be to develop tools to detect potentially regrettable messages and provide users with suggestions for when they might want to reconsider tweeting. Behavioral economics offers a potential direction to help prevent users from sending such tweets by using behavioral "nudges" to help people identify tweets they might not want to post [1, 7]. Such nudges are cues that suggest that users should alter a behavior without forcing them to do so.

We found that several negative emotions, including anger, stress, and frustration, tended to lead to regret on Twitter. A recent study of deleted tweets also found a slightly

higher frequency of negative-sentiment keywords in tweets that were deleted [3], a common strategy for coping with regretted tweets. Prior to a tweet being sent, such negative states could potentially be detected using tools like sentiment analysis or word frequency. Word analyses could potentially also be combined with environmental cues, such as location, especially when users tweeted from mobile devices; 45% of regrets reported by Twitter-regret participants were made from mobile devices. Once a negative mood was detected, it might be possible to provide feedback to the user about the negative emotion, or, in a manner similar to Google Mail Goggles [68], lock them out until they could think more clearly.

We also found that certain types of regret related to broadcasting thoughts to wide audiences were more common on Twitter. Twitter-regret participants tended to report regretting revealing too much, revealing expressive/cathartic thoughts, and sharing with unintended audiences. Such types of regret might be preventable through better audience awareness or management on Twitter. Participants often regretted tweets that revealed too much or that were expressive/cathartic because they were seen by people they didn't want to see them, or because people saw the tweets and were hurt. For these regrets, it might be possible to indicate more clearly who might view a tweet, for example by showing images of a user's followers. Interestingly, several tweets were sent by participants who had protected accounts at the time of the regretted message (25% overall, and 21% for unintended audience). Participants tended not to accidentally tweet to the general public. Rather, their tweets were viewed by people they didn't initially anticipate would view the posts. This is in line with Acquisti and Gross' concept of "imagined communities" [2] and the concept of tweeting to an "imagined" audience [**?**]. One way to visualize the actual audience might be to show images of people who could view the tweet, potentially prioritizing by interaction level. For instance, Lieberman and Miller's Facemail prototype uses this approach for email [48].

Promoting regret awareness To address a regretted message, users must first realize that they should regret the tweet. We saw several methods for becoming aware of regretted messages that were unique to in-person conversation and could potentially be adapted for Twitter, as well as several techniques that were unique to Twitter and could be further emphasized.

In person, participants often quickly became aware of regretted messages, typically through physical cues. For instance, one conversational-regret participant experienced regret after his girlfriend "instantly became upset and started to cry." Other participants saw audiences storm out of the room or laugh. Twitter users, physically separated from their audiences, usually lack instant audience feedback.

One possibility for improving Twitters users' awareness of regret would be to improve their abilities to gauge potential audience reaction absent physical feedback. Work has been performed to visualize sentiment conveyed in electronic communications. For example, Liu et al. prototyped an "EmpathyBuddy" for email that presents a line-drawn face that reacts to the emotion in the text [53]. Similar visualizations showing the sentiment conveyed by tweets might help Twitter users more quickly become aware of potentially regrettable tweets before tweeting them. A visualization that persisted after a user tweeted might also allow awareness to occur more quickly after a tweet.

We also found that Twitter-regret participants often reported being informed by their communities (e.g., friends, family, and co-workers) that they should regret messages, often over electronic means like text messages, or on Twitter itself. Lampinen et al. discussed how users of social networks collaboratively control disclosure [47]. Their participants used collaborative strategies to protect each others' privacy. Similarly, other individuals helped our participants become aware of regretted content. In some cases, these individuals were impacted by the message. In other cases, they were not. Developing easy mechanisms for people to tell someone about potentially regrettable tweets could mitigate potential regret.

Throughout our results, we saw that Twitter had a time delay compared to conversation, both in terms of time to awareness and time to repair. This was somewhat due to the lack of immediate audience feedback; in cases where Twitter regret was informed by others, this response often came hours or days later. On Twitter, users cannot typically see immediate feedback, and audiences sometimes cannot immediately access messages, delaying regret awareness and potential repair. However, unique to Twitter, even when there was no negative reaction, participants regretted tweets because of the record provided by Twitter. Participants re-read their tweets and realized the message was regrettable. Creating tools that better help users review past tweets may also help them become aware of, and purge, possibly regrettable content.

3.7 Conclusion

In this chapter we used a Mechanical Turk survey to examine Twitter users' regrets during in-person conversations and on Twitter. We found that Twitter users tended to regret similar types of messages both on and offline, included a variety of types of critical messages and blunders. However, reflecting Twitter's broader reach, participants reporting regrets on Twitter tended to report regretted messages targeted at broad audiences, including expressive or cathartic messages, messages that revealed too much, or messages that reached unintended audiences.

In general we also observed that participants describing regrets on Twitter tended to

become aware of the regrets more slowly than conversational-regret participants. Absent the physical cues and reactions available from offline audiences, Twitter-regret participants relied on becoming aware of regret through eventual self realization or when others told them they should regret tweeted messages.

These differences between online and offline regret offers some insight into potential mechanisms for improving sharing mechanisms on Twitter, or more broadly, online, with the goal of reducing regret. Because Twitter-regret participants tended to regret messages broadcast to broad audiences, behavioral nudge, or other educational mechanisms that help users become aware of potential audiences for their content may help reduce regret, especially for emotionally-charged content. It might also be possible to add some of the absent offline emotional cues to online messaging to help users realize more quickly when they should regret content.

In the next chapter (Chapter 4) we expand on these insights by examining another potential way selective-sharing mechanisms may fall short, specifically when users may choose to self censor, or not share, because mechanisms do not meet their sharing needs.

4 | Understanding shortfalls: Self censorship on Facebook

In Chapter 3 I described examining regretted messages on Twitter versus regrets in conversation to identify causes and impacts of regret that could potentially be addressed by changing online sharing mechanisms to reflect methods people use to identify or address regret for offline messages.

However, in some cases, when online sharing mechanisms fall short, instead of a user regretting a shared message, the user may be unable to target a desired audience and choose, instead, not to share the message (to *self censor*). Looking at the content users self-censor, and specifically the content users self-censor because selective-sharing mechanisms may not meet their needs, can reveal where selective-sharing mechanisms could be adapted to help users better target desired audiences.

In this chapter I describe a diary- and interview-based study I performed (with colleagues) to examine the types of content participants chose to self censor on Facebook, as well as the portion of that content they may have shared if provided with optimal selective-sharing mechanisms. These results complement the results on online and offline regrets described Chapter 3. They reveal potential shortcomings in current selective-sharing mechanisms that are not apparent when only looking at content that is actually shared and then regretted. Looking at unshared content provides insights into the types of selective-sharing mechanisms that might be necessary to allow users to selectively share currently self-censored content with desired audiences.

The majority of this chapter was previously published for CSCW 2013 [76].

4.1 Introduction

Social Networking Site (SNS) users make decisions about what content to share and with whom. Sharing inappropriately can result in consequences ranging from regret to job loss [96]. SNSs provide tools that allow users to share content with some people and block other people from viewing content. However, sometimes instead of targeting a particular audience, users will self-censor or choose not to share.

Lampinen et al. describe self censorship as one of the techniques SNS users rely on

to manage the co-existence of different social groups on SNSs [46]. Self censorship is an important ability; SNS users choose not to post content for a variety of reasons, including to protect their own and others' privacy and to prevent regret [46, 47, 96, 100]. In this paper we explore users' self-censorship decisions on Facebook, as well as the types of content they choose to self-censor.

While self censorship can be a desirable behavior both on- and offline, users sometimes choose to self-censor on SNSs because available access-control tools don't meet their needs. For a subset of self-censored content, users choose not to share because they would like only specific audiences to see the content, and those audiences are difficult, or impossible, to target given current interface design. We focus on understanding this subset of self-censored content and the potential impact of optimizing selective sharing tools to allow users to share this content with their preferred audiences.

Selective sharing [38] occurs when users can share with only their desired audiences, by selecting people to share with or block. We look specifically at sharing that could potentially have occurred if participants had been able to target exactly their desired audiences (optimal selective sharing). Our intention is to explore the potential ability of tools to allow users to share a subset of currently unshared content.

This chapter has two primary contributions. Self censorship has been established as a means for preserving SNS privacy but has not been thoroughly examined. We seek to expand understanding of types of, and reasons for, self censorship on SNSs by examining self censorship on Facebook. Second, we provide insight into the subset of self-censored content users could potentially share given improved SNS selective sharing mechanisms, as well as the types of tools that would be necessary to allow users to share this content. Previous work tended to focus on shared content; by focusing on unshared content, we provide additional insight for creating selective-sharing tools.

To address these issues we examined the types of Facebook content that users were not sharing, and why. Specifically, we looked at the following research questions:

- **Q1:** What types of content are users currently not sharing?
- **Q2:** Why do users choose not to share different types of content?
- **Q3:** What subset of content that users currently don't share (unshared content) could potentially be shared if they could exactly target their intended audiences (i.e., given optimal selective sharing mechanisms)?
- **Q4:** What attributes typify the groups with whom users would like to selectively share currently unshared content?

We ran a weeklong, 18-participant, diary study during which we asked participants to send us Facebook posts they thought about posting but decided not to share. We used

an in-lab interview to gather additional information about the content. We found that participants chose not to share a variety of types of content, especially entertainment and personal content. Participants would have shared approximately half of the unshared content if they were able to share with or block some combination of specific individuals, groups of individuals, and more ambiguous, attribute-defined groups.

4.2 Methodology

We wanted to determine what users were not sharing, and why (Q1,2), as well as the subset of unshared content that could potentially be shared using optimal selective sharing (Q3). We also wanted to explore attributes of the groups with whom our participants would have wanted to selectively share or block from viewing this subset of unshared content (Q4).

The study had two phases and took place in April and May of 2012. First, participants took part in a weeklong diary study during which they used SMS messaging to report all instances of unshared content on Facebook (i.e., content intentionally self-censored). Participants also filled out nightly surveys to further describe unshared content and any shared content that they decided to post on Facebook. Next, qualified participants took part in in-lab interviews. The interview provided more details about reported, unshared content and a better understanding of participants' decisions on when to share. We asked about participants' reasons for deciding against sharing, as well as the people, if any, participants hoped would see or wanted to block from viewing their content.

We iteratively coded each piece of unshared and shared content that we were able to ask participants about in the final interviews (122 piece of unshared and 83 pieces of shared content) for types of content, the types of groups the participant wanted to share with or block from viewing the content, and the participant's reasons for not sharing.

4.2.1 Recruitment and Demographics

We recruited 18 participants from a campus participant pool website, Craigslist, flyers, and a targeted Facebook ad. They were screened online for high English proficiency, a minimum age of 18, at least 6 months of Facebook use, frequent Facebook use (more than once per week), texting regularly (at least once per week), and having frequently held back content on Facebook (at least 3 pieces of unshared content over the past week). Thirty potential, qualified participants were sent online instructions for participating in the diary study, including 8 students. Nineteen of the 30 recruited participants completed at least one nightly survey, and 16 out of 30 completed the full study including the final interview.

Code	*Age*	*Gender*	*Occupation*	*Unshared items*	*Shared items*
P01	20	F	engineering student	7	4
P02	26	M	engineering student	8	4
P03	20	M	bus admin student	1	3
P04	33	F	social science student	24	0
P05	30	M	dental student	4	3
P06	26	M	unemployed	2	2
P07	23	F	non-profit	13	5
P08	29	F	art/writing/journalism	3	11
P09	25	F	non-profit	1	8
P10	28	M	human resources	10	4
P11	26	M	unemployed	6	7
P12	25	F	art/writing/journalism	4	5
P13	51	F	bus/mgt/fin	9	7
P14	24	F	lab mgr	8	1
P15	24	M	art/writing/journalism	2	7
P16	32	F	unemployed	12	8
P17	22	M	architecture student	4	2
P18	21	F	engineering student	4	2

Table 4.1: Participant demographics

Two additional qualified student participants received the link to the instructions from friends and participated in the full study, resulting in 18 participants.

Participants ranged in age from 20 to 51. Ten were female, and seven were students. Table 6.1 summarizes participant demographics. Participants were compensated $20 for the final interview and $2 per nightly survey completed, up to a total of $34. We also reimbursed $6 for parking.

4.2.2 Diary Study

The diary study lasted seven days. Participants had continuous access to a set of online instructions. Participants sent SMS text messages whenever they thought "of things that they would like to post on Facebook but decide[d] not to post." They were asked to describe the potential post and include the type of post it would have been (e.g., wall post, photo, link, etc). This SMS-based approach was based on the technique used by Brandt et al. [**?**].

Every night, each participant was also sent a link to an online survey, which contained questions for each piece of unshared content. Participants could provide more detailed descriptions of unshared content and reasons for not posting. Participants were also prompted with questions about the people with whom they would have liked to share or would have liked to block from viewing each item. These questions were open ended, allowing participants to either name specific individuals or define their own notions of

the people that would have constituted a “group" for sharing. The interface allowed participants to add additional unshared content, so they were not bound by the SMS messaging system. The survey also asked participants to describe content they had shared that day (shared content). If a participant had not shared any content, they were asked to fill out an auxiliary question about why they had not shared. We hoped to ensure a baseline level of effort and minimize incentives not to report.

We used this diary study and survey system because users can think of unshared content throughout the day, and we wanted to capture this as it occurred. This technique also allowed a participant to provide a quick “digest" of unshared content through the SMS system, and, if they were busy, return to the survey at a more convenient time to provide details.

4.2.3 Semi-Structured Interview

Participants who completed at least four surveys qualified for a final, in-lab interview (18 participants). We chose a semi-structured approach, which allowed us to capture similar types of data across all the interviews while maintaining the flexibility to explore the varied content reported. The interviews each lasted approximately one hour and occurred in a lab. One researcher served as the primary interviewer and interacted with the participant. A second researcher served primarily as a note taker. All interviews were audio recorded.

We used participants' shared and unshared content to explore our four research questions. We went through each piece of unshared content, and the participant's nightly survey responses, and probed for details on the content, reasons for wanting to share and not sharing, and, when relevant, details about the groups the participant would potentially have wanted to share with or block. For example, we asked the participant to describe the unshared content in more detail, to further explain why they decided not to post it, and to expand on their relationships with or common characteristics of the people they would have wanted to share content with or block. We also asked participants about their willingness to share each item, given selective sharing mechanisms, as well as for additional details on shared content and a series of questions on SNS usage and privacy habits.

Prior to study launch, we refined our methodology by piloting with 10 additional participants who are not included in analyses.

4.2.4 Data coding and analysis

To analyze the data, we looked at each piece of nightly survey content that we were able to discuss with participants during the final interviews. We removed any content that we

were unable to discuss in the interviews, either because the participant did not complete enough nightly surveys to qualify for the interview or because the participant submitted too many items to allow for discussion of all content (only P04, who submitted 52 pieces of unshared and 32 pieces of shared content). Table 6.1 lists the number of shared and unshared items included in the analyses for each participant. We coded each item for the type of content, the participant's reason for not sharing, and the types of people with whom the participant would have wanted to share and/or block (where relevant). Our coding process was based on that used by Kairam et al. to code content shared on Google+ [38] as well as the technique used by Naaman et al. to code Twitter data [65].

To create codes, two researchers each independently coded a random selection of 50 items, using data from the nightly surveys and notes from the interviews. Based on those codes, the researchers collaboratively created a set of high level codes and independently coded the majority of the remaining data. The two researchers then iteratively coded all the data with updated codes two additional times. Between each iteration, the researchers updated the coding scheme based on shortcomings from the previous round. Using the final codes, the researchers went through their independent codings and discussed and agreed on any codes that differed. This process produced the set of codes used in the analyses.

Analyses presented in this paper are intended to be entirely qualitative. Numbers are intended to illustrate results from the sample but are not meant to indicate statistical significance or quantitative generalizeability. Examples are only intended to illustrate trends seen during the study.

4.3 Results

Participants self-censored a variety of types of content, especially those related to external material (content unrelated to the participant), like entertainment. They most commonly chose not to share because they were trying to control how they presented themselves, and they would have shared about half of unshared content, given optimal selective sharing. The groups participants wanted to use for this optimal selective sharing included specific individuals, specific groups of individuals, as well as more dynamic groups that depended on context.

In this section we outline the types of unshared content, the reasons participants chose not to share, how much of that content would have been amenable to selective sharing, and the characteristics of the groups that would have been necessary to allow participants to have selectively shared that subset of unshared content.

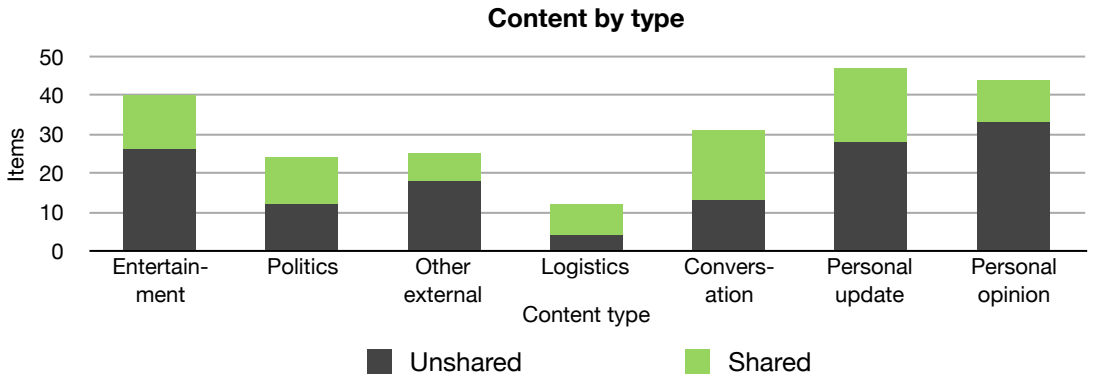

Figure 4.1: Shared and unshared content by type

4.3.1 Types of unshared content

We coded shared and unshared items into one or more of seven categories. We split items based on whether they were external content (e.g., entertainment, politics, or other external content), personal information, or related to planning or conversation. This led to three external categories: **entertainment**, **politics**, and **other**; two personal content categories: **personal update** and **personal opinion**; and categories for **conversational** content and **logistics**. Figure 4.1 shows the number of unshared and shared items in each category. These categories were roughly based on those Kairam et al. used for reasons to share Google+ content [**?**]. Items could be coded in multiple categories.

External content

External content included references unrelated to the participant. It could be intended to entertain, inform others, or allow the individual to express an opinion about the outside world. There were three subcategories: **entertainment**, **politics**, and **other**. We also noted when it included an opinion.

Entertainment: Examples included references to or articles about movies, television, sports, or music. This category contained 21% of unshared and 17% of shared content (26 and 14 items respectively). Unshared entertainment content tended to contain more material that could potentially offend. Several items contained explicit language or drug references. P17, for example, considered sharing a “weird" video that included drug-related content, but decided not to because her “family in Austin is really religious…they would’ve called [her] about it."

Half of unshared entertainment items included opinions, as opposed to three shared items. In contrast, shared entertainment content tended to advertise without a stated

opinion. For example, several participants posted to advertise concerts.

Politics: Content that referenced politics, current events, or activism was coded as politics, which included 10% of unshared and 15% of shared content (12 items each). The majority of the unshared political content was considered potentially controversial. P04, for example, decided not to post a “Link to article about young black republicans" to try to avoid controversy. On the other hand, shared content trended more toward current events. P12, for example, “posted a link to an article about the slow recovery from the BP Oil Spill in Louisiana." She explained “it was one of the few instances when there was something kind of political and I put it up anyway," because it “was the true story and what’s seen."

Other external content: This category included items that referenced content not related to the participant, entertainment, or politics. It included facts, quotes, pictures, and jokes, and included 15% of unshared and 9% of shared content (18 and 7 items respectively). Many sharing decisions depended on context. For example, P18 considered sharing “a recipe for a cake I saw posted by a friend from high school" but decided not to share because “I haven’t spoken to her in a while and it would be awkward."

Personal content

Personal content related to a participant’s life or general opinions and included **personal updates** and **personal opinions**.

Personal updates: These were items that described something that happened in a participant’s life. Examples included content about the participant’s day or about events the participant took part in, including photos. Personal updates made up 23% of the unshared and shared content (28 and 19 items respectively). Participants often decided not to post personal updates because they were too “frivolous" or not “creative" enough. For example, P10 thought about posting “Kicking ass and taking names!!! Happy Monday!!!" but decided not to because it was “very vague very generic, didn’t think it was very creative." Participants also didn’t post because they felt their personal updates were too negative or sounded like they were “whining." P16, for example, thought about posting about a fight between her and her boyfriend but decided not to because it was “grumpy." Shared personal updates tended to be relatively positive or straightforward.

Personal opinions: Opinions unrelated to external content were coded as personal opinions. These included how the participant generally felt about life, such as “having a stressful day," or more general opinions such as “We are way too old to be celebrating 420 day." Personal opinions included 27% of unshared and 13% of shared items (33 and 11, respectively). As with updates, many unshared personal opinions tended to be negative. Participants also worried that some might offend or start an argument. For example, P05

considered posting about how she disapproved of the Pokemon tattoo her brother-in-law was considering but decided against the post "because he wouldn't have liked it and it really wouldn't have made a difference anyway."

Conversation and planning

Conversational: This category included conversational niceties without additional content, such as birthday wishes or replies to posts that did not include additional content. This category included 11% of unshared and 22% of shared content (13 and 18 items respectively). Participants tended to not post conversational content based on potential social awkwardness. For example, P07 thought about wishing a friend happy birthday but decided not to because she hadn't "talked to him in a long time."

Logistics: Logistics included posts related to making plans. More were shared (8 items, 10%) rather than unshared (4 items, 3%). When participants didn't share, it tended to be due to offline, social reasons. For example, P05 decided not to discuss lunch plans because he didn't know one of the people involved in the conversation well enough.

4.3.2 Reasons for not sharing

We were also interested in reasons for self censorship. We asked in the surveys and the final interview, for each unshared item, why participants decided not to share. Responses tended to fall into one or more of five categories:

- **Argument/discussion**: Didn't want to start or participate in an argument or discussion.
- **Offend**: Didn't want to offend or hurt someone.
- **Boring/repetitive**: Felt the content was redundant, boring, or not interesting enough.
- **Presentation of self**: Felt the content went against the way the participant wanted to present him/herself (e.g., "seemed silly" or "don't like to post that kind of thing").
- **Inconvenient**: Prevented from posting due to time or technology (e.g., location made it difficult to post).

Figure 4.2 summarizes the number of items in each category. Presentation-of-self issues were most common (34%, 41 items) by a small margin; however, the remaining reasons each applied to approximately 20% of items. Percentages add up to over 100% because some items were not shared for multiple reasons.

Several reasons emerged more frequently for different types of content. Approximately half of entertainment and personal updates weren't shared because of presentation-of-self concerns, and slightly over half of political items weren't shared because participants

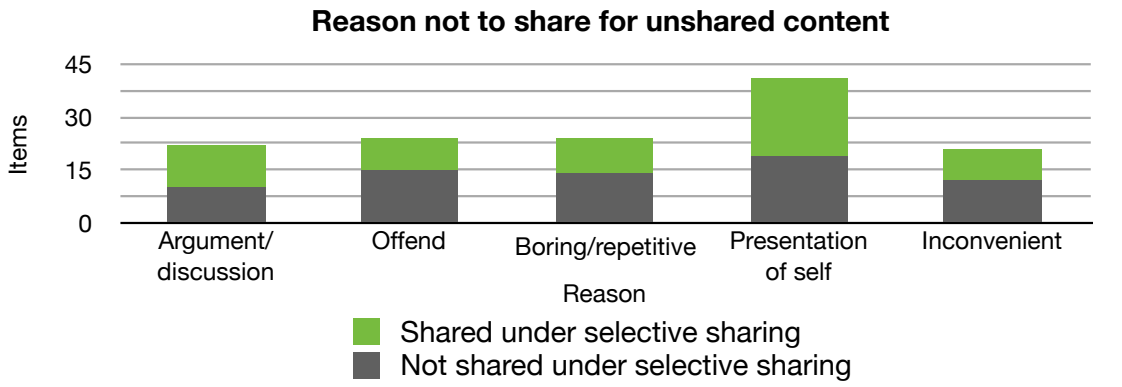

Figure 4.2: Reasons not to share unshared content

didn't want to start or participate in arguments or discussions. Almost half of personal updates also weren't shared because participants were worried that the items would be boring or repetitive.

4.3.3 Potential for selective sharing

Participants would potentially have shared a subset of the unshared content if they could have exactly targeted particular audiences under optimal selective sharing. To isolate this subset, we used two Likert scale questions to judge participants' willingness to share given an optimal ability to selectively share with desired audiences. For each item of unshared content for which a participant provided a potential group that they would have liked to have shared with or blocked, we asked the participant to imagine that they either "could have shared this content only with" the people they wanted to share it with or could have "prevent[ed]" the people they didn't want to see it from viewing the content. To increase generalizability, we did not specify the interface that would be used to share the content, only that it would exactly target desired audiences. Responses were on a five-point Likert scale where a one was "very unlikely" and a five was "very likely." We consider a participant who indicated above a three for either question to have been potentially willing to share given optimal selective-sharing mechanisms. If a participant answered above a three for sharing and/or blocking selected people, we analyzed the people with whom the participant indicated they wanted to share and/or block. Overall, 60 out of 122 unshared items (49%) would have potentially been shared given optimal selective sharing. Of those, 57 would have been shared if the participants could have shared with only a desired set of people, and 25 would have been shared if the participants could have blocked people from viewing content. Figure 4.3 shows each participant's potential willingness to share given optimal selective-sharing tools.

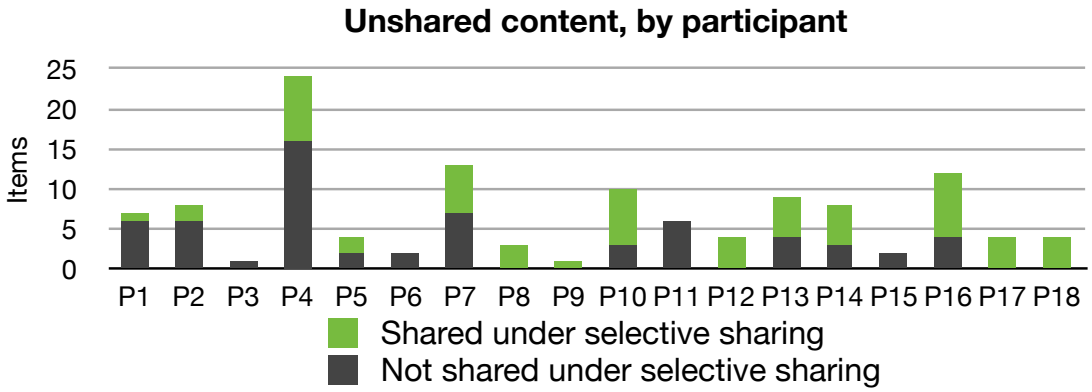

Figure 4.3: Unshared content participants were willing or unwilling to share, given optimal selective sharing

Although our small sample limits the generalizeability of these results, this indicates that participants could have potentially shared a relatively large subset of their self-censored content if they could have exactly targeted desired audiences.

Types of content for selective sharing

Participants would potentially have selectively shared approximately half of each type of content. External content tended to be amenable to selective sharing because participants wanted to share items with people who would have been interested and block people who might have been offended. For example, P08 considered posting “a lot of angry status updates" during a hockey game she was watching but decided not to because others were already doing so. She would have posted if only her hockey friends had been able to see the posts because they would have been interested. This was common for unshared entertainment content, which participants often felt only a subset of people would be interested in and/or other people might be offended by. A similar dynamic occurred with political content, with more of an emphasis on avoiding debate. This might have been expected because participants often decided not to share such political content to avoid argument or discussion.

Approximately half (15 out of 28 items) of personal updates would have potentially been shared with optimal selective sharing. Participants often wanted to share with people who were “close friends" or who they saw regularly who would understand or appreciate the posts. For example, P16, who considered posting about a fight with her boyfriend, only wanted to share with a small group of friends because “they can relate, because they know more about me and we talk about more personal things with each other." Participants

would have potentially shared 42% of personal opinions (14 items) given optimal selective sharing. In several cases, participants wanted to share opinions with people who would understand the context. P07, for example, wanted to post “don’t have pets if you’re not prepared to take care of them!!!" after a bad experience cat sitting but only wanted mutual friends of the person who the post was directed at to see it.

Items for which participants wanted to control presentation of self or didn’t want to start an argument or discussion were most amenable to selective sharing (slightly over half). Participants tended to want to share items with presentation-of-self issues with close friends or people who would be interested in or understand the content. For example, P07 thought about posting about her frustrations at her babysitting job to get advice, but chose to self-censor because she didn’t think babysitting was “cool." She would have preferred to share only with particular people who also babysat. Participants who chose not to share because they didn’t want to get involved in an argument or discussion tended to want to share with people who agreed or thought the same way about potentially controversial content. P04, for example, considered posting a link to an article about “cohabitation and divorce." She decided against posting because she had a lot of Facebook friends who were religious Christians who disapproved of cohabitation, and she wanted to avoid a long discussion.

4.3.4 Types of groups

For participants to selectively share the desired subsets of content, they would need to be able to specify, using the interface, the individuals or groups with whom they wanted to share. We asked participants to specify who they did and did not want to view each unshared item, so we could understand the kinds of groups participants would need to create to express their optimal selective-sharing preferences. We looked at the number of people in, and characteristics of, the groups.

Number of people in group

We coded the people with whom participants did or did not want to share each item into one or more of the following: a **specific person** (e.g., “my sister," “Tim"); **specific people** defined as a countable set of people (e.g., a group of ten close friends); or an **ambiguous group** defined by one or more attributes or relationships (e.g., “hockey friends"). Percentages add up to over 100% because participants sometimes specified multiple sets of people they wanted to share with or block (e.g., a specific person and an ambiguous group).

Participants specified individuals or groups with whom they wanted to share for 92 out of the 122 unshared items (75%). For the remaining items they were willing to share

with everyone or weren't willing to share with anyone. Of the groups associated with the 92 pieces of content, we looked at those with which participants would have shared given optimal selective sharing (53/92 items). Of the groups that would have been useful for selective sharing, 47% (25) were ambiguously defined, 30% (16) were groups of specific people, and 33% (17) were specific individuals. Participants specified individuals or groups to block from viewing content for 57 items (47%). Of these, participants said that blocking 23 of the groups would allow them to share the content items under optimal selective sharing. Of these groups that would have been useful for selective sharing, 74% (17) were ambiguously defined, 13% (3) were specific groups of people, and 26% (6) were specific individuals.

These results partially imply that our participants were not using Facebook's current custom privacy settings. Participants indicated that they wanted to share with single individuals or specific people, which could be done on Facebook. Participants' reasons for not doing so are further addressed in the Discussion. More ambiguous groups also accounted for a relatively large percent of potentially useful groups. They tended to be attribute-based and consisted of both concrete groups (e.g., classmates) and more context-specific groups (e.g., people who would disagree with a post). Such groups would require more extensive user effort or new tools. For example, a user could set up a school-based group ahead of time but might have more difficulty creating a group defined by people's feelings toward a topic.

Group characteristics

We also looked at characteristics associated with the individuals and groups with whom the participants would have liked to have selectively shared. We coded each individual or group into one or more of the following categories:

- **Work/school**: Work or school at any stage of the participant's life (e.g., coworkers, classmates, high school).
- **Demographics**: Age, gender, geography, race (e.g., younger relatives, male/female).
- **Family**: Relatives (e.g., mother, extended family).
- **Close friends**: Close relationships (e.g., close friends, people seen on a regular basis, boyfriend/girlfriend).
- **Not close friends**: Lacking close relationships (e.g., "not close to," someone never met, "frenemies").
- **Relationship to post**: Interested in the post, felt a certain way about the post, personally relevant to the post (e.g., "feel the same way as me," person the post was

Shared	Total	Specific person	Specific group	Ambiguous group
Work/school	17	0	6	11
Demographics	10	2	4	6
Family	6	3	2	3
Close	9	3	6	0
Not close	2	1	1	1
Relationship to post	33	11	11	16
Total items	**53**	**17**	**16**	**25**

Table 4.2: Characteristics of groups participants wanted to share with for optimal selective sharing, by type of group

Blocked	Total	Specific person	Specific group	Ambiguous group
Work/school	8	1	2	6
Demographics	4	0	0	4
Family	2	0	2	1
Close	1	1	0	0
Not close	7	2	0	6
Relationship to post	13	5	1	10
Total items	**23**	**6**	**3**	**17**

Table 4.3: Characteristics of groups participants wanted to block for optimal selective sharing, by type of group

directed at, interested in the content).

A summary of the characteristics of the groups associated with the items participants would have been willing to share if they could have targeted or blocked specified people is in Tables 4.2 and 4.3. These categories are similar to those that emerged in other work on grouping [38, 42, 46, 99].

Based on the 53 items that participants said that they would have been willing to share if they could have shared with selected individuals, the most frequent attribute was the person or group's relationship to the post (62%, 33 items). It was slightly more likely to occur for a specific person or ambiguous groups. Participants tended to want to share only with people at whom the content was directed or people who would be interested in an item. For example, P08 "had tickets to an advanced screening of The Avengers and almost posted about how excited [she] was to see it using a bunch of profanity." She wanted to share it with her friends who liked comic books and video games and was "sure I would have posted it if it was just like the people I know like it would've seen it." But, as she

pointed out, “I don’t have a group for comic book friends, mostly because I don’t know who would like it, there are people who like things I don’t know about." Determining these more complex, ambiguous relationships to posts that rely on time-of-post decisions would be relatively difficult.

The second most common attribute was work/school (32%, 17 items), which only occurred for groups of specific people and ambiguous groups. For specific groups of people, this attribute tended to be associated with a close group of friends that included people from school. Such specific groups would be relatively easy to define using a selective-sharing mechanism, because they are at least partially defined by a concrete common attribute. When participants defined more ambiguous groups using work/school, they tended to be people who would be interested in the content and who either currently went to school or worked with the participant or went to school with the participant in the past. P08, for example, wanted to share content about a hockey game with “hockey friends," who also tended to be college friends. Defining these more ambiguous groups would be more difficult with current tools and might not be encompassed by the work/school attribute.

Relationship to post also occurred most frequently for the 23 items that participants would have been willing to share if they could have blocked a specific group of individuals (56%, 13 items). Again, it was more likely for specific people and ambiguous groups. For specific people, participants tended to want to block the person who originally posted the content they were planning to comment on or people who might be offended. For example, P12 considered posting “some links to articles I read on NPR and WeArePowerShift.org - very political stuff." She didn’t mind the general public seeing the content, but wanted to block her boyfriend’s dad and other conservative friends from viewing it.

Work/school was also the second most common attribute for people participants wanted to block. However, for blocking selected people, but not for sharing with selected people, “not close friends" emerged as the third most common attribute. This attribute characterized specific individuals and ambiguous groups. Participants tended not to want to share more personal content with people who didn’t know them as well. For example, P14 considered posting about a stressful day but didn’t want to share it with people she wasn’t as close to. As she put it “if they’re better friends with you then they don’t necessarily care if you’re venting or complaining." Such groups would be relatively difficult to capture using current tools because they tended to be context-specific. They ranged from friends-of-friends to the “frenemies" P13 considered too “weird" to know about her evening plans.

4.4 Limitations

This study had several primary limitations. First, it was qualitative, limited to a small sample, and did not consider unshared content in a cross-cultural context. Conclusions, therefore, lack broad generalizability. Our sample also skewed young. This age skew partly reflects SNSs; in 2010, approximately three-quarters of SNS users were 35 or younger [31]. However, future work examining differences in self-censorship across age levels would also be interesting.

Using a diary study also introduced bias. Participants were aware of the purpose of the study; as part of a “study on Facebook usage" they were asked to report “everything you think about sharing on Facebook but decide not to post." Texting in content and filling out surveys likely primed them to think about Facebook, unshared content, and audiences. When asked, participants did not feel they had changed their behavior due to the study. However, about a third mentioned being more aware of what they posted and unshared content. This may have pushed them to think more about self-censorship.

The study structure also relied on self-reported data based on hypothetical scenarios. Actual behavior does not always match what participants say they will, or mean, to do. These issues could be partially addressed in future studies by designing studies to focus on actual behavior. One possibility would be to examine the differences in types and levels of sharing that occur under different interface designs or when a user is instructed to share in different manners (e.g., posting only for oneself, for close friends, etc). Focusing on behavior might reduce the limitations of self-reported hypothetical data and could allow for less priming.

Finally, this study was only able to capture a subset of self-censored content. There is a spectrum of how likely a user would be to post an unshared item, which ranges from content they are almost prepared to post (e.g., at the keyboard and have fully composed) to vague ideas that they decide they probably shouldn’t post. Responses to this study mostly included more fully-thought-out ideas, although there were some vague thoughts. It likely missed more of what people self-censor before ideas are fully developed.

Participants may also have been less likely to report sensitive or embarrassing content. To reduce participants’ sensitivity, we avoided face-to-face interaction until after the diary study. We believe this was at least partially successful; participants reported some potentially sensitive items that included profanity, political opinions, and drug references. There is also likely content that is so sensitive that it is self-censored in an ingrained way and was not captured. Future work might accompany an approach like this by using a survey to try to probe more ingrained self-censorship by asking participants if they would consider posting content on a variety of more extreme topics (e.g., sexual content, violence,

etc.).

4.5 Discussion

Participants self censored content, often because they wanted to manage how they presented themselves to various audiences or to avoid argument or discussion. They indicated that they would have potentially shared about half of this self-censored content, across content types, given the ability to optimally target audiences. The people participants wanted to share with, or block, ranged from those captured by current Facebook privacy controls to ambiguous, context-specific groups that would require more sophisticated mechanisms. We discuss why participants seemed not to use Facebook's custom privacy settings, participants' uses of alternatives to self censorship, and some high-level design suggestions for capturing selective sharing preferences.

4.5.1 Reasons for not using Facebook custom privacy settings

In general, participants didn't use Facebook's custom privacy settings to control who could see content. At the time of the study (April-May 2012), Facebook offered the ability to set the visibility of a post to the general public, friends-only, lists of friends defined either automatically by Facebook or manually by the user, or a post-specific list of people. Users could set a default to public, friends, or a custom list. Most participants used a friends-only default setting. Some, like P03, felt that friends should be able to "see everything," while others, like P15, assumed that anything posted on Facebook was available for a general audience. A few participants had set up friend lists at some point but tended to have used them once or set them up and then stopped maintaining them. P18, for example, had used the friend lists feature when it first appeared but hadn't continued to actively use them. This behavior is consistent with the literature. Kelley et al. found that users tended not to want to use previously created groups for sharing [42], Strater and Lipford found that users both had trouble understanding Facebook privacy settings and tended not to revisit them once set up [81], and Karr-Wisniewski et al. found that users did not use provided grouping tools [41].

Several participants found Facebook's grouping and privacy features too confusing or difficult to use. P01, for example, hadn't recently adjusted her settings, even though she realized Facebook had changed their privacy settings, and said that "it kind of worries me that I haven't messed with it." She both found the settings confusing, admitting that "I'm not really sure how lists work" and felt that Facebook was something she quickly logged on and checked, not something she sat down and used long enough to bother with the

settings. P08 pointed out that she frequently posted status updates from the Facebook mobile app on her phone, where “it’s easier to just not post than to go in and mess with the settings."

Other participants didn’t trust Facebook. Some didn’t trust Facebook to maintain their privacy settings. P16 put it “Like maybe one day they’ll just take off all the permissions, like just for fun…so I never know if that’s going to happen, since Facebook seems to have a negative track record in most people’s minds, I just try to censor myself." Other participants didn’t trust themselves to configure the privacy settings and understand how they would propagate.

This suggests that for preferences that could have been captured by Facebook’s current tools, users might require better, built-in, education about Facebook’s privacy controls, a better interface, or an overall increased level of trust in Facebook’s data privacy. One potential direction might be to increase the visibility of available tools and their impact on sharing. Many participants seemed confused about available custom privacy settings and friend lists and how they could be used; increasing transparency could increase their abilities to use such tools and might potentially increase trust in Facebook.

4.5.2 Alternatives to self censorship

Lampinen et al. describe strategies SNS users rely on to mitigate the co-presence of multiple social groups on SNSs. One of these strategies is self censorship. However, users also rely on other strategies, including choosing “channels of communication" and dividing up who can see what content [46] both of which our participants described.

Consistent with “dividing the platform," Facebook includes “group pages" that allow users to post content to particular groups. Unlike the other privacy features, participants used the group pages to post content for particular, self-selected groups. Several participants used groups affiliated with interests, school, or work to post and read content, participate in discussions, and advertise events to interested people.

Stutzman and Hartzog also describe how SNS users can choose to use different social network services to maintain “privacy, identity, utility, and propriety" [83]. Participants mentioned using different channels of communication as ad hoc privacy controls to varying degrees. Some used chat for more private communications. Others used locked and unlocked Twitter accounts to post personal content they felt was unfit for Facebook. Several felt they could better limit who was following their content on Twitter. P08 for example, was friends with her young sisters on Facebook, and said “I have been kind of watching things I post [on Facebook] because they’re on it a lot, so I’m trying not to swear as much or post a whole lot of crazy things." Instead, she would “post it somewhere else like on my

Twitter or on my blog or something" where she felt her sisters couldn't find the content as easily. Such behavior was in line with Stutzman and Hartzog's observation that SNS users relied on a strategy of "practical obscurity" to make it difficult to find certain accounts and maintain privacy [83]. Other participants felt they could better track who viewed Twitter content, even with public accounts. This might indicate a desire for a simpler sharing interface. Participants tended to be wary of how their content would be shared through friends of their friends on Facebook.

4.5.3 Potential improvements to selective sharing

Participant interest in selectively sharing currently self-censored content (approximately half of currently unshared items) suggests a desire for selective sharing tools that would allow them to share with the groups encompassed by their desired audiences. To allow users to share such content would require interface tools that captured the more ambiguous groups participants wanted to target for selective sharing. They would require context-specific information or information often unknown to the user. As outlined in Related Work, machine learning solutions are being developed to help users dynamically create groups. Facebook provides a rich dataset for machine learning, including posts, group pages, likes and a user's own and friends' profile data.

Participants often wanted to share with or block audiences that were relevant to posts. This might require tools that could target groups related to topics people are interested in or might disagree with. Defining these traits could require discovering traits, like "geeky comic book friends" or friends with liberal political views that users might not know themselves but might find useful for sharing. One potential method would be to rely on self-identification. For example, a user could indicate that she wanted to share with "comic book lovers" and wait for people to indicate an interest. Alternatively, a user or algorithm could try to identify characteristics that typified a trait. Such a system might be similar to the Hummingbird system for Twitter. Hummingbird uses Twitter hashtags to allow users to indicate the topics of their tweets and then request and approve others' requests for access by topic [21]. Participants also sometimes wanted to just target individuals involved in a conversation around a post; future tools could make it easier for users to limit a post's audience to people who had previously been involved in a conversational thread. Some combination of these tools could help facilitate users' abilities to target content to people relevant to posts.

4.6 Conclusions

In this chapter I used a diary- and interview-based study to examine the types of content participants self censored on Facebook and found that participants commonly self censored external content, especially entertainment items, closely followed by personal content. We also found that participants most often tended to self censor to better control their presentation of self.

I used the reported self-censored content to explore how much of this self-censorship behavior occurred because of limitations of selective-sharing mechanisms versus a broader desire not to share the content. I found that participants thought they might share approximately half of the currently self-censored content if they could have exactly targeted their desired audiences.

In this context, self censorship, like regret (outlined in Chapter 3), represents a metric for measuring shortfall of current access-control or selective-sharing mechanism. Thus, to improve selective-sharing tools to allow users to share more, it might be possible to focus on creating mechanisms that allowed participants to usably target the audiences they described.

In some cases participants described wanting to target specific individuals or groups of people that they would have been able to target using currently available tools. They tended to choose not share using current mechanisms because they felt tools were not trustworthy or were inadequately usable. Chapter 5 examines how this dynamic can lead to choice between channels based on factors including trust in the service, people involved, and the task-at-hand.

In other cases, however, participants wanted to target audiences that could not easily be captured by current selective-sharing tools. They tended to want to share with, or block, ambiguously-defined, often context-dependent groups. These groups included people who were characterized by their participation in a conversation around a post, or their interest in an item. In Chapter 6 I examine how one of these characteristics, specifically interest in a topic, could potentially be applied to Facebook.

5 | Channel choice in everyday sharing decisions

Chapter 3 and Chapter 4 describe how shortfalls in current selective-sharing mechanisms on Twitter and Facebook may lead to regretted or self-censored content. Examining these suboptimal outcomes provides insight into potential improvements to selective-sharing mechanisms on these services that might allow users to avoid regret or post currently self-censored content.

However, selective-sharing rarely takes place on a single application or service (*channel*). For example, a user may choose not to share content on Facebook because the selective-sharing options do not meet their needs. But, in realistic everyday situations a user would then have the option to share on a different channel, or combination of channels, that better met their needs. If they weren't able to easily share with a group of ten friends on Facebook, the user could instead share with those friends by text message, or could share with a few of those friends by text, a few friends on Facebook, and a few friends by email.

This choice of channel depends on selective-sharing needs, as well as other dynamics, including the audience with whom the user want sto share, the broader task they're performing, and selective-sharing and task-related features available on different services. Selective-sharing tools on individual services should, therefore, both seek to meet selective-sharing needs at the channel level (e.g., prevent regret, allow users to share content they may not currently share on the channel) while simultaneously accounting for how and why users may move between, and combine, channels to meet broader task, audience-driven, and selective-sharing sharing needs.

In this chapter I examine the role of selective-sharing tools at the level of the everyday-sharing-application *ecosystem*. I draw on a diary- and interview-based study (performed with colleagues) to focus on how people decide between different channels to share content. I then discuss how selective-sharing tools should account for the task, feature, and audience-drive dynamics that emerge from the overall everyday online sharing environment

The majority of this chapter was originally published for CHI 2016 [78].

5.1 Introduction

People often want to share personal content (e.g., photos, videos, documents) online with particular audiences. Online services and platforms, referred to in this paper as *channels*, offer mechanisms that allow users to target desired audiences, ranging from the ability to set file access-control rules in Google Drive or Dropbox, to the ability to manually create a list of people to whom to send email, to the ability to share with known friends or followers on social networking sites (SNSs) like Facebook or Twitter.

Prior work examined use patterns for online selective-sharing mechanisms. This work tended to explore dynamics for individual types of communication or systems, focusing, for example, on social communication patterns [34, 35, 73, 85, 87], traditional or cloud-based file systems [11, 16, 20, 93, 94, 98], or SNSs [38, 47, 89, 92, 100].

However, in today's multi-device and multi-application environment, users are typically not limited to a single site's sharing mechanisms [69, 79, 80, 98]. Instead, when sharing options on one channel don't meet a user's needs, the user can move to another channel, or can combine channels to usably share content with their desired audience. In this study we focus on *personal content sharing*, defined as content shared for non-work purposes. In this context, institutional guidance is largely absent, and users can draw on both traditional file-sharing and more socially focused services, such as texting applications or SNSs. This allows *personal-sharing ecosystems*: combinations of channels that together approximate a user's desired features and audiences better than any one channel's sharing options.

For example, a user might take photos on a trip with friends. She might want to share most of the photos privately with those friends, but a few photos more publicly. She might typically share photos using Instagram, but wouldn't be able to share privately using the service. So, instead of not sharing, she might share publicly visible photos on Instagram and the remainder just with the group of friends, using Google Drive.

We document the dynamics that emerge from the ability to choose between and combine different channels, focusing on two research questions: 1) What factors impact channel choice for sharing with particular audiences? and 2) What sharing behavioral patterns emerge from the ability to combine or switch between channels?

To address these questions we performed a three-part, qualitative study (n=17) that consisted of a preliminary interview focused on general sharing practices, a weeklong diary study tracking self-reported shared and accessed content, and a final interview following up on the diary entries.

We found that the task during which sharing took place, for example collaboration or conversation, combined with the type of content being shared, tended to shape the use of different services' features, both specifically related to selective sharing as well as to

other task-driven needs. Audience attributes, such as access to different services or social dynamics, also shaped channel choice. In some cases, participants could meet sharing needs with one service; in other cases they shared across multiple channels to create composite sharing features unavailable on any one service, such as the ability to share at multiple access levels, or to express urgency on a service that lacked notification capacity. Participants also shared across multiple channels to target composite audiences unavailable on any one service. We document these ecosystem-level dynamics and discuss the design implications of these observed behaviors for creating selective-sharing mechanisms that account for broader task and audience dynamics.

5.2 Methodology

We focused on how and why participants use different channels to share and access content with different people, for personal ("non-work-related") purposes. Our goal was to elicit both high-level reasons for channel choice and reasons for channel choices during specific activities. The study took place in three parts. In an initial semi-structured interview participants explained, at a general level, how and why they choose different services to share and access content. We then used a diary study to ask participants to report their actual content-sharing and access behaviors over a week. We finished with a semi-structured interview in which participants explained how and why they used different services during the reported activities.

5.2.1 Recruitment and demographics

Participants were recruited by posting a link to a screening survey on Pittsburgh and Washington DC-area Craigslist sites, as well as on Carnegie Mellon University's experiment recruitment board. We screened for English proficiency and to include a mix of genders, ages, and occupations. We also screened for participants who regularly used the Internet for non-work purposes and had a personal smartphone they accessed regularly, so they could participate in the diary portion (Table 6.1). Participants were compensated with a $50-65 Amazon gift card based on level of participation in the diary-study portion.

5.2.2 Interviews and diary study

Initial interview

The initial interview focused on services participants used to share and access personal content. The structure was based on Voida et al.'s work on file sharing [94], expanded to re-

Code	*Age*	*Gender*	*Occupation*	*Hrs/wk*	*Services used*
P01	31-40	F	Art/writing	10-20	Drive, Dropbox, Email, Text, Google Hangouts, NextDoor, Soundcloud, Physical device
P02	18-25	F	Admin. support	20+	Drive, Email, FB, FB Groups, Text, Physical device
P03	26-30	F	Fitness instructor	5-10	Drive, Dropbox, Email, FB, FB Groups, Instagram, Text, YouTube, Soundcloud, Other sites/blogs/discussion boards
P04	31-40	F	Science/engineering/IT	20+	Drive, Dropbox, Email, Flickr, Text, Pandora, Steam
P05	51-60	M	Unemployed	20+	Drive, Dropbox, Email, FB, FB Messenger, FB Groups, Text
P06	26-30	F	Admin. support	10-20	Drive, Email, FB, FB Groups, Twitter, Text, Physical device
P07	18-25	F	Americorps	20+	Drive, Dropbox, Email, FB, FB Groups, FB Messenger, Instagram, GroupMe, Text
P08	26-30	F	Business/management	10-20	Drive, Dropbox, Email, FB, FB Groups, FB Messenger, Instagram, Text, Wedpics, WhatsApp, Google Hangouts
P09	31-40	M	Student (medicine)	20+	Dropbox, Email, FB, FB Groups, Snapchat, Text, Google Hangouts
P10	31-40	M	Legal	5-10	Drive, Dropbox, Email, FB, Twitter, Google+, Text, Google Hangouts, Soundcloud, Bandcamp, YouTube, GitHub, Other sites/blogs/discussion boards
P11	31-40	M	Service	20+	Drive, Email, FB, FB Groups, FB Messenger, Instagram, Tumblr, Twitter, Text, Bandcamp, YouTube, Physical device, Other sites/blogs/discussion boards
P12	31-40	F	Other professional	5-10	Dropbox, Email, FB, Pinterest, Text
P13	18-25	F	Unemployed	20+	Drive, Dropbox, Email, FB, FB Groups, FB Messenger, Flickr, Instagram, Pinterest, Tumblr, Twitter, Snapchat, Text, YouTube, Physical device
P14	41-50	M	Skilled Labor	1-5	Drive, Email, Text, Physical device, Other sites/blogs/discussion boards
P15	26-30	F	Student (management)	20+	Drive, Email, FB, FB Groups, FB Messenger, Pinterest, Text, WhatsApp, Physical device
P16	18-25	M	Student (management)	20+	Drive, Dropbox, Email, FB, FB Groups, FB Messenger, WhatsApp, Physical device, Other sites/blogs/discussion boards
P17	18-25	F	Student (comp. science)	1-5	Drive, Dropbox, Email, FB, FB Groups, FB Messenger, Instagram, WhatsApp

Table 5.1: Participant demographics: participant code, self-reported age range, gender, occupational category, typical hours per week online for non-work purposes, and services described in the initial interview (most participants also described showing someone content on a device)

flect modern services and devices. Participants were interviewed in the lab (7 participants) or remotely with video chat (10). Interviews were audio-recorded and transcribed.

Participants first described the devices they "used to connect to the Internet." Next, we introduced them to personal content sharing: "any time you create content and share it" or "times when other people you know or interact with create content and then share it with you," limited to "content you typically share for personal use...outside of work."

As an initial probe we used a predetermined list of services developed from pilots and asked participants whether they used each to share or access content for non-work purposes. The list included: Google Drive, Dropbox, email attachments, Instagram, other photo sharing services, text messaging, Facebook Groups, Facebook Messenger, Facebook (general), Twitter, Google+, instant messaging, Snapchat, repository services, physical devices, and showing someone a device. If participants mentioned other services we asked about those as well. We then asked participants if there were any services we hadn't talked about that they used for a variety of types of content (e.g., music, videos, genealogy data). Finally, we asked participants if we'd missed any services. By cueing both services and content types, we tried to prompt a relatively comprehensive set of services. For each service participants used we asked probing questions, including:

- What they typically used the service for; why and how they used the service
- Who they shared content with or accessed content from using the service; whether/how they shared with specific people/groups
- What types of content they shared using the service
- What their typical sharing/access pattern was; whether they typically looked at the content once or multiple times
- What their typical notification activities were (e.g., how they knew content was available/how they told others content was available)
- Why they used that service versus others (for types of content, people, etc.)

Diary study

We next asked participants to fill out a diary of personal content sharing and access activities over a 6- to 7-day period (exact length depended on interview timing). We used the Paco smartphone-based experience-sampling application[1] to send five brief surveys at random intervals each day. Each survey asked the participant if they had shared or accessed content since the previous response. The participant described the service they used to share or access the content, who they shared the content with or received the

[1] pacoapp.com

content from, and the type of content. Each participant's survey contained personalized multiple choice options, created based on the initial interview. They also provided a brief free-response description of activities. Participants had to complete at least ten surveys to participate in the final interview.

Final interview

We reviewed the reported sharing activities with the participant during a final, approximately half-hour-long, semi-structured interview within a few days of the diary study. For each activity we asked the participant for more details about who they shared with or received content from, the type of content, and the dynamics of the activity. We also asked about why they chose specific services for specific activities, and, where applicable, why they chose between different services for handling similar people or content.

5.2.3 Data analysis

The initial interviews resulted in high-level descriptions of how and why participants used services. The diaries and final interviews resulted in 223 content-sharing activities for which participants described how and why they shared or accessed specific items.

We qualitatively coded the initial interviews. Personal sharing activities are embedded in information management [11, 16, 20, 70, 93, 94, 98], SNS [38, 47, 89, 92, 100], and communication-related activities [34, 35, 73, 85, 87], so the two interviewers first drew themes from prior work to affinity diagram their interview notes and create an initial codebook. One coder then coded the initial interview transcripts, iteratively updating the codebook. This resulted in three codebooks, related to tasks, behaviors, and reasons for channel choices, which were used to create Tables 5.3, 5.4, and 5.5. A second coder coded a random set of 50 items from each codebook, resulting in Kappa values >0.65 for each (>0.6 indicates "substantial agreement" [88]). One coder used a subset of the codes to code the diary items. We report themes seen in the initial interviews and reflected in the diaries. Participant counts from the initial interviews are sometimes included for illustrative purposes.

5.3 Results

Participants shared personal content during a variety of *tasks*. Accomplishing these tasks, with different *types* of content, required *features* supplied by varied services. For example, collaborating on content may require a user to be able to both share with a particular audience and edit collaboratively. To choose channels for sharing, participants, therefore,

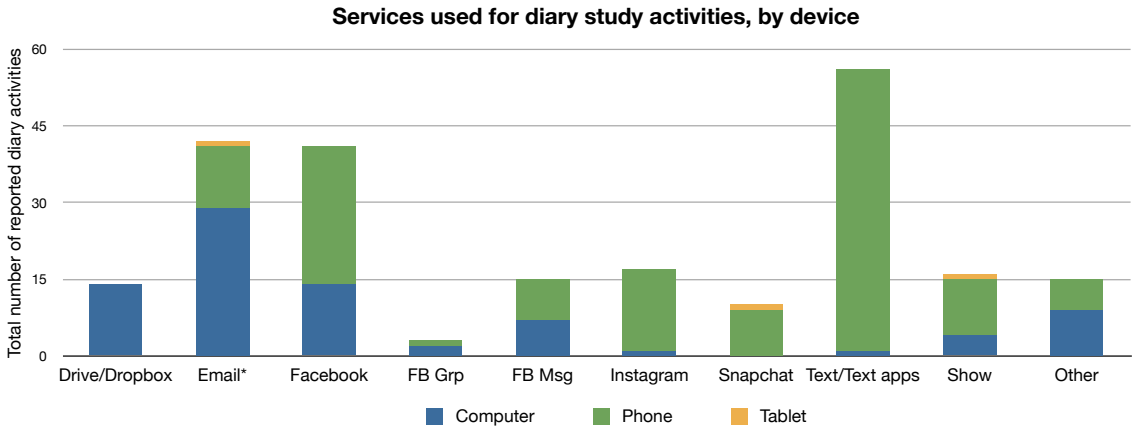

Figure 5.1: During the diary study portion of the study participants reported sharing and accessing content on a number of services, across devices including personal computers, phones, other computers (work, library), and tablets (*one email activity is excluded because the participant didn't remember the device)

sought desired attributes in features of available services. However, participants also wanted their audiences to have access to the content in a timely manner. Thus, choices were also constrained and shaped by *audience attributes*, such as access to services or tech-savviness. One service sometimes provided the features necessary for sharing in the desired manner while performing a task and accessing the desired audience. However, participants also sometimes combined multiple services to achieve these goals.

We describe *tasks* and *types of content* that tended to shape participant consideration of services' *features*, as well as how *audience attributes* tended to constrain choices within the available feature space. We then discuss participants' strategies for combining multiple services to reach *composite audiences* or to create *composite sharing features*.

5.3.1 Personal content sharing in an ecosystem of services

Participants chose between, and combined, varied services for personal content sharing (Table 5.2). Overall, participants described using between five (P12) and fifteen (P13) services (Table 6.1). They typically shared on some combination of a personal computer and phone. Some services were primarily used on a computer (e.g., Google Drive), others primarily on phones (e.g., texting), and some were used across devices (e.g., email) (Figure 5.1). Most participants also described sharing content by showing it to others.

Channel type (Participant count)	*Examples*
Email (17)	Gmail, school email systems
Text/instant messaging (17)	Google Hangouts, Facebook Messenger, texting, GroupMe
Social networking or photo sharing sites (15)	Facebook, Twitter, Instagram, Tumblr, Google+, Flickr, Pinterest
Discussion boards/platforms (15)	Facebook Groups, Slack, NextDoor, Ancestry.com
Music/video sharing (6)	YouTube, Bandcamp, Soundcloud
Physical devices (6)	USB/thumbdrive, personal hard drive
Repository (1)	GitHub
Showing someone (15)	Showing content on a laptop or phone

Table 5.2: Types of services participants described using to share personal content during the initial interviews

5.3.2 Content-sharing decisions are embedded in tasks

Participants chose channels based on how service features met their sharing needs. They focused both on ensuring content reached desired audiences and accomplishing broader activities not directly related to sharing. For example, when planning events some participants sought tools that would allow them to share content with fellow planners and perform collaborative editing. Thus, channel choices often relied on matching services' features to both selective-sharing and broader, activity-based needs. Table 5.3 outlines types of tasks participants described as interacting with personal-content-sharing channel choice. The role the sharing component of the task played in shaping channel choice varied.

Many participants (13) described personal content sharing as interacting with *archival or synchronization* tasks. Here, the ability to share tended to shape channel decisions by defining content location. P07 shared content through Dropbox because "a lot of my like pictures and stuff that are on my Mac go straight to Dropbox."

Other tasks required sharing in combination with other task-related activities, which together shaped desired features. While the task might not primarily be focused on making

Types of tasks	
Archival or synchronization	Use a service for storage, backup, archival
Browsing	Browse posted content (e.g., on a newsfeed)
Collaboration	Ongoing project collaboration, feedback, review
Connect w/friends or family	Connect, stay in touch with friends or family
Conversational sharing	Use content to facilitate or as part of a conversation; brief, conversational-style sharing
Documentary sharing	Document a life event (e.g., with photos or videos)
Resources	Receive or provide information or a resource
Planning or logistics	Plan an event; organize logistics in real time
Publicity	Publicize events or promote oneself

Table 5.3: Tasks that shaped participants' desired channel features for personal content sharing

content viewable by others, sharing tended to play a role in accomplishing the activity. For example, participants (12) integrated sharing into *collaborative* activities. They shared content during ongoing projects or to allow editing or feedback. P01 used email to edit "a book cover" with collaborators, and P10 used Google Drive with "script ideas that would be passed back and forth" with friends.

Similarly, content sharing played a role in event-*planning* for some participants (8). They described using Google Docs to plan events, for example coordinating a party by "getting everyone to fill out a form for their availability" (P08) or using a shared document to "see what other people are bringing" (P06). A few participants also used interest-driven Facebook Groups for planning. For example, P02's college club sports team coordinated "serious business, schedules, practices" through a Facebook Group.

A few participants also shared content, using text messaging, to facilitate real-time logistics. P01 described coordinating with a friend to find each other at a baseball game by sending "a picture of what she's standing next to."

Other tasks that shaped channel choices were more directly focused on content sharing. These tasks, and sharing needs, tended to vary by type of content and audience. A few

participants focused on *documenting* events. They tended to want to share large amounts of content with limited audiences to fully document events, and described using Dropbox, or specialized applications like WedPics that allowed them to upload large amounts of content while controlling the audience. P03 explained, "We were at my brother's graduation and there were like 500 pictures, and my mom really wanted to see them...so we thought she should just start using Dropbox."

Participants (6) also sometimes shared in a more *publicity*-focused manner to promote events or self-promote. They focused on reaching more public audiences, often by broadcasting. P07 explained that she'd publicize content on Facebook: "if there's an event on campus I want everybody to go to."

Many participants (10) also performed more ambient, *conversational*-style, content sharing. They tended to use small amounts of content, like photos or videos, to supplement ongoing conversations, for example, showing a picture on a phone to demonstrate something that's "too hard to explain" (P13). Participants also used services like texting or IM to send brief pieces. P10 explained how he and his friend would "send photographs of things kinda throughout our day, if we're walking around and something seems relevant." They also included content in ongoing conversations on services like email. For example, P08 used email for conversations with her mother-in-law: "that's like our primary mode of communication...have a whole conversation with her on e-mail back and forth and then send a picture of something."

Similarly, participants (13) described *browsing* content friends or people they followed posted on SNSs like Facebook or Instagram. P09 explained, for Facebook, "I just log on, every, every other day or every day just to read through the News Feed."

5.3.3 Matching features of services to content for a given task

Participants sought services with features that met both the sharing and broader, task-based needs required to perform tasks-at-hand, tending to draw on different services for different types of tasks. Features participants wanted were also shaped by the type of content shared in the task (Table 5.4).

Selective-sharing features/affordances Participants often wanted to target desired audience(s) at particular levels of access. Services provide different *selective-sharing* mechanisms that afford varied levels of control over content sharing, including the ability to:

- Share one-on-one with specific individuals (e.g., email, text, instant messenger)
- Create pre-defined groups for ongoing sharing (e.g., mailing lists, Facebook grouping tools, traditional access-control lists, Google+ Circles, texting applications like

Features of services	
Access to content	Content is available on that platform
Ability to edit	Support for editing or collaborative editing
Additional content	Ability to add additional information to the content
Broadcast	Whether the service will make the content public; level of reach
Connectivity requirements	Whether the service is dependent on an Internet connection or requires data
Control over file hierarchy and structure	Ability to control how a file hierarchy and structure are created and maintained
Recipient contact information	What information is needed to reach a recipient
Selective sharing	Available selective-sharing features
Simplicity	Number of steps required to access the service; convenience
Size/number	Ability to manage large files/number of files
Speed/notifications	Speed of receipt; whether the service provides notification
Support for file type	Support for different types of files or content
Trust	General level of security, privacy, trust in the service

Table 5.4: Available features of services were matched to both the desire to target particular audiences and accomplish broader tasks.

WhatsApp or GroupMe)

- Create on-the-spot groups by filling in names or contact information (e.g., email, texting)
- Send a link or share an ID with others to share content, sometimes with the option of additional privacy control (e.g., Google Drive, Dropbox, WedPics)
- Broadcast content to known friends, followers, or contacts (e.g., Facebook, Twitter, Instagram, Snapchat Stories)

Participants used these mechanisms to achieve desired control over their content (Figure 5.2). Most participants wanted *explicit control* for some content and used services that allowed them to set access-control rules or explicitly limit who could view content. Some participants used access-control mechanisms paired with predefined groups or sent access-controlled content to a group using a link. P12 started using Dropbox, for example, "when my son was born...because we could allow who we wanted to, like, view it, and give them like access to everything...we didn't want them publicly shared."

Other participants sought control by sharing with *specific individuals* or small groups, using services like text or Facebook Messenger. P03 used text messaging, for example, to share with an individual when "it's something that I don't really want to be known public." Some participants, on the other hand, felt they had control when they shared with *smaller, known, friend or follower groups* on broadcast platforms: P13 explained that she used Snapchat because "I think it's a lot more private...you very definitely know who's gonna see it and who isn't, based on who's in your contact list."

In parallel to these methods, participants (7) also *trusted* different services to enforce selective sharing or provide general security. P08 moved the password information file she shared with her husband off Google Drive, because she "just stopped feeling like it was safe, and so now I keep it in Dropbox. I guess I perceive that as a little bit safer." Similarly, P03 shared a photo by text, because she wasn't "comfortable with putting pictures of my godson online."

For some tasks, like publicity, participants wanted content to reach broader audiences and tended to share on more public-facing platforms (e.g., Facebook, Instagram) that allowed them to *broadcast*. For example, P13 used Facebook for "stuff I want to have a general outreach rather than a private outreach." They sometimes tried to call specific audience-members' attention to the content, using mechanisms like tagging recipient(s), adding hashtags, or posting on walls.

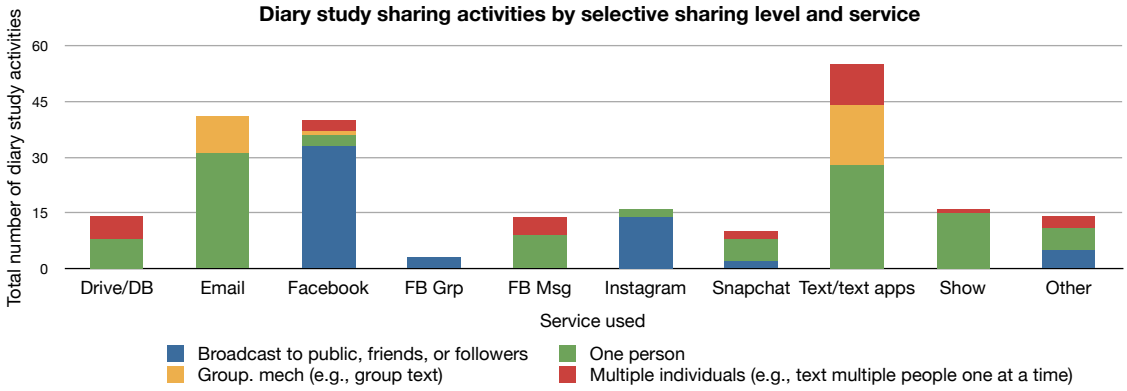

Figure 5.2: During the diary, participants used services to share with audiences at different levels. They tended to use some services to broadcast to friends, followers, or the general public (e.g., Facebook or Instagram), and others primarily to share with individuals, for example sending an email to one person. They also used group sharing mechanisms on some services, such as group texts, or sending email to a group. On other services they shared with multiple people, one at a time, for example texting several people individually.

Content-type-driven feature choice

Some tasks also required sharing or accessing different types of content. Channel decisions were, thus, partly driven by support for *content formats*. P03 explained, for example, switching from Google Drive to email to more easily open a fillable PDF. For many participants (12), channel choice was also driven by services' support for content *size or number* of files. Some participants tended to move from services like email to cloud-based platforms such as Google Drive or Dropbox as the amount of content increased: "If I'm gonna share a lot, I'm gonna use something like Dropbox" (P04).

In line with work on boundary management [28, 84], some participants (5) also tended to associate different services with different *content tone* or levels of formality: P13 used Snapchat, for example, for "slightly goofier stuff than Instagram." This dynamic sometimes interacted with available access-control options: P08 shared "the more intimate photos [from her wedding] like pictures of me getting dressed or with my mom crying in the dressing room" on Dropbox, with family, while "the ones where I'm walking down the aisle and they're just more of the typical wedding pictures I shared on Facebook."

Some participants (6) also drew on services that provided fast *notification*, or that they perceived to deliver content more quickly when they wanted to share content they considered important or urgent, for example when planning or as part of an urgent conversation. Several participants, for example, used text messaging when they wanted content to reach others quickly, especially WhatsApp, because it provided notification of receipt. Or, P04 perceived Dropbox as faster than email and used it when her husband

"went to go get a medical procedure done and he had forgotten to take a copy of his notice from his other doctor to say that this needs to be done. So I took a picture of it and I posted it to Dropbox and then when he was at the doctor's office he pulled it down."

Features to support other task-driven activities Participants also considered needs related to activities required for the task-at-hand, but not directly related to selective sharing or access control. A few (4) considered services' *data or network connectivity* requirements, especially when they expected be in locations where they didn't know whether they would have good connectivity. P01 explained, for example, that she used a USB drive to bring content to an artist's studio for collaboration, because she didn't want to risk not having an Internet connection after traveling.

Participant channel choices were also impacted by how much *additional information* a service allowed users to add to the content (4 participants), and how easily *editable* or *accessible* the service made the content (8 participants). For tasks like collaboration, for example, participants tended to want to add information or edit: P03 used Google Drive for feedback on an essay because it gave recipients "access to edit because I prefer all the comments on, like, one document."

Participants also tended to use services because they were more convenient or required fewer steps to share or access content (the level of *simplicity*). P17 explained choosing Dropbox for sharing because, "I can store [content] in my laptop, I can add and edit right from my laptop."

Creating a usable group organization schema can also be challenging when sharing [71]. When sharing large amounts of content, for tasks like archival or collaboration, some participants (5) described choosing services they felt allowed them to create a file *hierarchy and folder structure,* organize content, and control or maintain the organization. P03 chose to use Dropbox instead of Google Drive for photo sharing, for example, because she felt it gave her more control over file structure: "I sorta like to keep everything organized...because I find that when people share photos with me in Google Drive, they just go anywhere, they don't go to like a specific folder."

5.3.4 Influence of audience characteristics on channel choice

Participants also tended to want to reach their audience(s) in a timely manner. They shared with varied types of individuals and groups, including friends, family, acquaintances, classmates, professional contacts, interest- or activity-based contacts (e.g., friends interested in fitness, people interested in finding coupons or deals, etc.), and the general public. Thus, while participants tended to draw on services for the features they offered for particular

Attributes of individuals/groups	
Access/ availability	Access to the service, availability on the service
Experience, tech-savviness	Experience with a service; general perceived tech-savviness
Proximity	Geographic location/proximity
Level of interaction	How often one interacts or expects to interact with the audience

Table 5.5: Audience attributes also shaped channel choice and dynamics.

tasks, channel choices were constrained by the services desired recipients used, and were shaped by broader social dynamics.

Access to services

Most participants (16) described *audience access* as a factor in their choice of service. P08 used Google Drive, for example, because, "everyone's already on Gmail and most people that I share files with are, like, in one of my Hangout chats or something." Similarly, participants tended to want to use services that they knew audiences checked frequently, especially for urgent content. They tended to take audience experience or comfort with different technologies into account, or considered general levels of tech savviness. For example, P03 used text messaging "for communicating with my mom, that's the only service I can use where she can get it instantly."

Reflection of social dynamics

For more social tasks or services, for example conversational sharing or text messaging, channel considerations also tended to reflect broader social dynamics. Participants' service choices partially reflected geographic *proximity* or typical level of social *interaction*, in line with dynamics seen in work on communication patterns [16]. Many participants (11) used certain services with people who lived nearby or with whom they interacted regularly: P02 explained that she used text messaging to share content with "friends that I see in person, at least once a week."

Grouping tools and one-off or repeat interactions

Setting up a maintained sharing environment can also incur costs. Participants used channels that provided more or less permanent sharing mechanisms that, in some cases,

partially aligned with levels of interaction they expected to have with audiences: one-time, repeated but time-limited (e.g., for an event), or repeated and ongoing.

Participants sometimes used faster, lower-cost, sharing methods during expected one-off interactions. At an extreme, participants sometimes described showing someone content on a device when they didn't think the recipient would want repeat access. Some participants also sent quick, one-off messages on services like email; P10 used email "if it's not expecting a response necessarily, but just kinda like, oh, here's a picture I wanna share with you."

Some participants also described expecting time-limited interactions, for example when planning a game night or trip. They sometimes described using services that allowed them to set up temporary sharing environments with lower coordination costs and that could be easily closed after use, such as group emails or group messages on Facebook Messenger. For example, P11 described using Facebook Messenger to organize game nights: "you just start a new group and then add three or four people...when the event's coming up, and then, then that group will stop" after the game night occurs. Similarly, P08 used WedPics, a specialized photo application, at weddings. She sent out a private ID to everyone attending so they could consolidate photos for the event, with minimal overhead. She explained: "I've done that for every other wedding I've been at. They have some cute catchphrase or something and then that's the ID."

In other cases participants described groups for which they expected ongoing, longer-term interactions, such as project groups or friends and family. For these groups they tended to describe services that allowed them to set up pre-defined or ongoing groups with more archival structure, such as Facebook Groups (e.g., for school clubs or with classmates), mailing lists (e.g., for ongoing conversations with friends), Google Drive or Dropbox folders (e.g., for longterm collaborations), or WhatsApp or GroupMe groups. For example, P16 checked in with a group of office colleagues every three or four weeks, so he maintained an ongoing group chat.

5.3.5 Combining channels to meet needs

Sometimes one service met task-, content-, and audience-driven needs. However, participants also combined multiple services to create composite sharing features that allowed them to share with varying access levels or meet a task's selective-sharing and other activity-driven needs. They also combined audiences available on different services to target all desired recipients.

Composite features: indicating urgency

For some tasks participants wanted to make sure content reached recipients quickly, sometimes with acknowledgement of receipt. Voida et al. describe the use of "out of band" channels, like the phone, to speed up notification [94]. The current environment also provides synchronous tools like text messaging or IM that tended to be perceived as quickly delivering content and can provide explicit receipt notification. However, these services may lack functionality desirable for other task-driven activities, such as support for editing.

Thus, to accomplish tasks (e.g., sending detailed information, collaborating on a document, etc.) and also indicate urgency, participants (11) sometimes described using text messaging or other secondary channels in parallel to more task-driven channels to alert recipients to the presence of the content. For example, P12 texts her husband to let him know "that I sent him an email instead, with more detailed information."

A few participants described similar dynamics for multi-tasking. They wanted to collaborate on projects and also have conversations. Some services, like Dropbox or Github, allow collaboration but not synchronous conversation. Participants, therefore, used one service to edit or share content while using another for discussion. For example, P10 worked on coding projects with friends. He would chat on email or Google Hangouts and share content from Github.

Composite features: redundant selective-sharing groups

Some participants also shared with the same groups of people using multiple services, because both services allowed them to reach the desired audience, but neither services met all other task-driven needs. For example, P12 shared with the same group of friends using email and text messaging. She chose between the services based on how much information she included in the content. Similarly, P13 used Snapchat with her friends, but switched to text message when the content had more substance.

Other participants switched between services based on content tone. For example, P07 shared with, and received content from, college friends and family on Facebook and Google Drive or Dropbox. The Facebook content was "less personal to me than whatever is on Dropbox or Google Drive."

Composite features: chaining services

Some participants (8) moved from one service to another to access desired features. For example, for some tasks, participants wanted to send content to people for whom they didn't have particular contact information (e.g., phone number, email). They lacked

secondary channels for contact information present in organizational settings [94], so they sometimes fell back on channels that required less information to initiate contact. Then, they would move the communication to channels with additional features they wanted for tasks. For example, P16 described starting a conversation on Facebook Messenger and then moving to WhatsApp or email.

Participants also sometimes moved from one channel to another as task-related needs changed. P13 described sharing a document over Facebook Messenger. She then realized the recipient needed an editable version of the document, so she moved to Google Drive. Similarly, participants sometimes started interactions by showing someone a piece of content, for example in the midst of a conversation, when it was the fastest method. Then if the person wanted a copy of the content they moved to another channel to send it.

Targeting a composite audience through multiple channels

Desired audiences were also not always available on a single channel. All desired recipients might not use a single service, or participants might want to share content at varied levels of access for different portions of the audience. Participants sometimes used multiple services, in combination, to create a *composite audience* by:

- Increasing broadcast reach of content by cross-posting it on multiple publicly facing services
- Sending pointers to publicly available content to specific people or groups to make sure it was available to non-overlapping audiences
- Sharing some content on a channel with more access control and a subset more publicly

This dynamic sometimes occurred when participants tried to broadcast content as widely as possible, for example when focused on publicity. Some participants had different followers, friends, or known audiences on different social networking sites or public platforms. For example, P13 shares content on a variety of broadcast services with different audiences on each. On Facebook she has "a lot more friends like I don't really interact with," on Instagram it's "more people that like I generally want to know what's happening in their life," and on Tumblr her audience is mostly people she doesn't know.

To increase the reach of content, some participants (10) described cross-posting on multiple platforms. Services provided features that supported this behavior, including the ability to embed content, and to allow pointers or links to content from or on another site. Participants used these features to create broader composite audiences, sometimes while taking advantage of an initial service's features. For example, P11 shared music-related

content on Bandcamp, which allowed him to upload music, and then also posted a link on Facebook. Some participants also used Instagram for video and photo editing and then cross-posted the content. P03 explained, "usually what I'll do is have the video like go to Instagram and then, you know, having Instagram post it to Facebook." A few participants also described expanding an audience to include specific people they knew were unavailable on certain services. For example, P12 knew that her parents didn't use Facebook, so she shared with them separately after posting on Facebook.

Some participants also described using multiple services to selectively share content at mixed access-control levels. They sometimes wanted to share more content with a subset of a larger audience. For example, participants might share photos with a smaller group of people using a service that allowed more access control, such as texting, email, Drive, or Dropbox. Then they would share one or two photos more publicly, for example on a broadcast service like Facebook. P07 described sending her mother ten photos of her brother's graduation by text and then posting one to Instagram that she "wanted her friends to see."

5.4 Limitations

We focused on ecosystem-level behaviors, which limited deep insights into specific themes (e.g., specific channels or strategies, participant background knowledge, etc.). Our choice of interviews and diary studies also relied on self-reported behaviors and motivations. This allowed us to probe participants' reasons for their behaviors; however, the results may also reflect participant biases regarding their motivations, as well as potential unwillingness to discuss sensitive topics. Our use of a smartphone-based diary also means that participants were relatively tech-savvy, which may be reflected in the somewhat young sample. Additional work could include a broader sample, more focused scope, or more observational insights to expand on themes that emerged in this work. It could also explore how personal-content sharing varies for participants with different backgrounds (e.g., demographics, technical or privacy knowledge, etc.).

5.5 Discussion

Participants used one or more channels to match features to personal-content-sharing needs in the context of task and audience dynamics. Designers should account for these task, feature, and audience dynamics, as well as the potential for multi-channel behaviors, to create selective-sharing mechanisms that account for realistic ecosystem-level behaviors.

5.5.1 Designing embedded selective-sharing mechanisms

When developing selective-sharing mechanisms for personal content, designers should account for the broader task and audience context in which mechanisms will be used. We observed three interacting factors that should be considered.

Designers should consider the *level of access control* users may want when sharing content; for example, a desire to broadcast to the public, share with known followers, or limit access to specific groups of people or individuals. They should also account for *costs users may be willing to incur* to set up and maintain the access control. These access-control needs also interact with *broader, task-related features* users may need, such as the ability to share files of different sizes or content types, the ability to collaborate, or the ability to share with limited network bandwidth.

The need to account for these factors, in context, is driven by task and audience dynamics we observed in this paper. Different tasks may require different levels of control over access and may necessitate activities that may not be tied directly to access control. Similarly, users may be more willing to incur costs for tasks or audiences for which they expect to have ongoing interactions and less willing to incur costs when they expect shorter interactions. Based on the degree to which designers expect these dynamics to be present in a given context, they can draw on design dimensions, such as those we observed, that address each factor (Table 5.6). Including design features that consider each of these factors may help designers create selective-sharing mechanisms that account for access control in the context of broader task and audience dynamics.

5.5.2 Facilitating multi-channel or single-channel strategies

Participants drew on multiple services when one channel was not sufficient to meet their needs; however, participants also relied on single channel strategies to maintain boundaries or control over content, or for simplicity. Designers of sharing mechanisms should, therefore, also consider when to add features on a single service to support tasks or audiences and when to instead facilitate multi-channel strategies to meet task and audience needs. We observed several dimensions designers can consider.

Including features to facilitate tasks

Participants drew on multiple services when one service was insufficient to complete a task or provided insufficient audience reach. Designers should, therefore, consider whether one service will provide the desired features and audience for tasks users may wish to perform, or whether users may reach out to additional services. In some cases it may

Design features

Factor	Sample design dimensions/affordances
Level of access control	One-on-one sharing, pre-defined or on-the-spot grouping tools, links to access-controlled content; ability to broadcast to a defined group of friends, a known group of followers, or the general public
Willingness to incur costs	Length of time access control will be supported (e.g., one-off link vs ongoing interaction environment), speed of setup/takedown of access-controlled environment, level of support for archival/organization, ability to control or delegate control of organizational schema or access
Task-related features	Support for editing, synchronous notification, bandwidth availability, types and sizes of content

Table 5.6: We observed a number of design features and affordances that could help designers account for task and audience dynamics when creating selective-sharing mechanisms.

reduce complexity for a service to facilitate multi-channel strategies rather than trying provide all features within one service. For example, a text messaging service might choose to facilitate content export to a service that provided broader editing features if its users wanted to use it during collaboration, rather than adding an editing tool. In these cases designers can draw on a number of design features that facilitate multi-channel strategies, including cross-posting, allowing content to be embedded across sites, and facilitating content export or download. However, in other cases, integrating features into one service may be simpler for users.

Clarifying boundaries

We also observed that, in line with prior work [28, 30, 84], participants used services to maintain boundaries between audiences or types of content. While users may draw on multi-channel strategies to achieve task- or audience-based needs, moving content across services may risk lowering boundaries, through potentially unintended audience access to content, or unclear or unintended data flow across platforms.

Designers should, therefore, consider when multi-channel strategies may lead to undesirable or unexpectedly lowered boundaries or leaked data. When the consequences

of performing activities across sites may be unexpected, designers should seek to clarify, and limit the potential negative impacts of, the user's actions. For example, when data, or content, is cross-posted or linked across sites, the new audience and access-control policies should be clear. A number of design mechanisms could potentially facilitate this process, including notifications that clearly communicate when and how data leaves a platform, defaults that maintain access-control policies from the initiating service or the service with stricter policies (e.g., cross-posting a link to access-controlled content instead of automatically embedding content), and allowing users to limit metadata shared with content (e.g., not automatically including user identifiers specific to one service when sharing on another).

5.5.3 Understanding channel-based mental models of trust

We also observed that participants tended to choose channels partially based on trust in individual services or in their percieved abilities to provide control over content. Consistent with prior work [40], participants had variable mental models for deciding when to trust different platforms, ranging from considering paid platforms to be more secure to trusting services that provided clearer selective-sharing tools. These models may or may not be accurate in context, but, to encourage use of services' sharing features, designers should also seek to understand users' mental models for trust and control for individual channels and across channels. Further work is needed to explore factors that may drive these mental models of trust by channel, as well as how these factors may interact with task, audience, and multi-channel dynamics.

5.5.4 Evaluating sharing mechanisms

Researchers should also evaluate sharing mechanisms at the ecosystem-level when considering their ability to meet selective-sharing needs. Prior work has tended to assume the decision to share content occurs based on how well individual channels' selective-sharing mechanisms meet users' needs (e.g., [76]). However, we observed that participants move between and combine aspects of tools that are substitutable at the task- or audience-level to meet personal-content-sharing needs. Researchers should, therefore, consider that people may use tools that do not have all the features they need as long as the tools can be incorporated into multi-channel strategies. Or, users may exclude a tool that doesn't work well in their overall ecosystem. Thus, channels' selective-sharing tools should be evaluated based on their role in broader ecosystem context.

5.6 Conclusion

Personal-content-sharing decisions take place at the ecosystem level. Users rarely only consider selective-sharing mechanisms on one service when deciding how to share. Instead they move between channels and combine channels to meet sharing needs.

In this chapter I describe a study that examined how participants chose between and combined channels to meet their personal-content sharing needs. I found that selective-sharing needs were one of a variety of factors, including task-at-hand, audience attributes, social dynamics, and available task-driven features (e.g., support for collaboration. support for different file sizes, etc.) that drove participants to choose different channels to share personal content. Participants also combined channels when one channel didn't meet their sharing needs, for example when all their desired audience members weren't available on a single service.

In this context, when designing, or trying to improve selective-sharing mechanisms, it is necessary to consider a tool's role in the broader sharing ecosystem. Selective-sharing is embedded in audience context as well as the context of other task-driven features provided by services. For example, a user may want to be able to explicitly target a particular audience, but they may also want to use a service that allows them to edit content. This dynamic will shape which service, and which selective-sharing tools, people find usable.

We draw on these insights in the next chapter (Chapter 6) when prototyping and evaluating potential topic-based sharing mechanisms for Facebook. We evaluate the mechanisms both in terms of their ability to supplement current Facebook sharing tools, but also their ability to supplement and substitute for tools used on other services in participants' sharing ecosystems.

6 | Exploring topic-based sharing mechanisms for Facebook

6.1 Introduction

Selective-sharing preferences can vary around different dimensions, including time of sharing, a user's relationship with their audience, or the user's location [57, 86, 99]. Access-control mechanisms on different services seek to capture these dimensions to allow users to share content with their desired audiences. As outlined in Chapters 3 and 4, when users aren't able to use existing mechanisms to selectively shared in their desired manner, they may turn to compensatory behaviors, such as self censorship, or may share in a manner they come to later regret.

As observed in Chapter 4 users may want to share based on the content's topic, or with other people who have a shared interest in a topic (referred to in this chapter as *topic-based sharing*). For example, a user might want to share content related to dogs with other dog lovers or content related to sports with other fans of the same sports team.

Some services provide mechanisms that facilitate topic-based sharing. For example, Twitter allows users to add hashtags; photo-sharing and blogging sites also often allow users to tag content. Similarly, users can join mailing lists on specific topics or can participate in discussions on topic-specific forums using services like Reddit or Slack. Alternatively, on social-networking sites like Google+, users can manually create groups of friends centered around known interests to selectively share content.

In this chapter I explore the potential impact of adding topic-based sharing to Facebook. Current Facebook mechanisms do not provide a way to directly identify others with shared interests, absent a shared group identification or trait. For example, a user who wants to share content about hockey on Facebook might want to share it with their friends interested in hockey. However, they might not know which of their friends are interested in hockey, and might not want to join a group related to hockey, making it difficult to share this content with friends who share their interest on Facebook. In cases like this, when people aren't able to easily identify common interests using available mechanisms, they may

compensate by not sharing, or self censoring [76].

We explore adding topic-based sharing to Facebook, focusing on the general concept of being able to selectively share with people interested in a topic as well as the potential ability to opt out or opt in to viewing content by topic. We also look at the potential ability to share on a topic without identifying oneself.

We use a semi-structured interview centered around a retrospective diary to probe how participants ($n = 16$) currently share about topics on Facebook and other services, as well as how they feel they would share for each topic using a general, hypothetical topic-based sharing mechanism. We use walkthroughs of three mockups of potential topic-based sharing mechanisms to examine how participants think they might use topic-based sharing that included: 1) the ability to opt out of topics; 2) the ability to opt in to topics; and 3) the ability to share on topics without identifying oneself.

We specifically focus on:

- Exploring how participants currently seek interest-driven audiences for topics and target people interested in different topics, on Facebook and other services
- Understanding use cases participants perceive for topic-based sharing mechanisms on Facebook
- Understanding participants' perceived strategies of engagement with topic-based sharing

We find that participants currently intend content for interest-based audiences on a variety of topics, and use a range of techniques to try to target the content. Participants tended to feel that topic-based sharing mechanisms on Facebook might allow them to avoid the risk of oversharing or offending others for some topics, and might allow them to target narrower or better audiences to share improved content, more content, or facilitate better discussions. However, topic-based sharing wouldn't meet all of participants' sharing needs, for example they sometimes wanted privacy guarantees or other features they felt were not provided by Facebook.

6.2 Background

Interest in a topic (referred to in this paper as *topic-based sharing*) is one dimension around which people may want to selectively share. For example, a person might only want to share pictures of food with other people interested in food. Services offer different mechanisms that facilitate the ability to share content with other people who have a common interest.

6.2.1 Properties of topic-based sharing mechanisms

Topic-based sharing mechanisms tend to have four properties that allow content to be selectively shared along topic-based dimensions: the ability to tag content to facilitate search by topic, the ability to allow people to opt in or opt out of content by topic, whether topics are associated with content or with an individual, and the ability to be identified or not when sharing on a topic.

Tagging with topics for search or filtering

Adding topics to content, often in the form of *tags* can allow a user, or a user's audience, to find content on a topic [5, 54]. A number of services provide tagging mechanisms that allow users to tag content by topic, and that facilitate filtering or search by topic. For example, Twitter allows users to add hashtags, which can be used for search. Blogging platforms, like Tumblr, or photo services like Flickr also provide the ability to tag content by topic.

Ability to opt in or opt out of content by topic

Some services also allow audiences to opt in or out of viewing content by topic. At a high level, users can opt in or out of viewing certain topics on discussion boards by opting to view the forum or not. Similarly, users can decide whether or not to join mailing lists on particular topics.

Identifying topics with users or content

Topics may also either be associated with a trait of the individual sharing, for example someone who identifies as being interested in dogs, or with the piece of content they share, for example a piece of content someone shares that they identified as being about dogs.

Some topic-based sharing mechanisms rely on a user associating themselves with a topic, for example, when a user becomes a member of a group to share content on a topic. This is the modality used by Facebook Groups, or by some discussion boards. Other services rely on a user assigning others to a group based on an assumed interest or trait, for example using Facebook's user-assigned lists or Google+ Circles.

On the other hand, some mechanisms allow users to tie topics to content, rather than associating the topics with users. Tags may be added to content without explicitly associating traits with the user sharing the content. For example, a user who identifies as a Republican could share content that they tagged as being about Democrats.

In this study we focus on exploring topics identified with content, rather than associated

with individuals. Identifying topics with content provides the freedom to share or express interest in content without explicitly creating or assuming ties to one's own, or one's audience members', identities.

Real name or pseudonymity

Users may also consider some topics sensitive, and may wish to share in a de-identified manner [14, 39]. Services may also allow users different degrees of online anonymity or pseudonymity when sharing on topics. For example, some discussion forums, like Reddit are primarily pseudonymous. Other forums, like Twitter, have mixed real name and pseudonymous use, while other forums, like Facebook require real names [82].

6.2.2 Selective and topic-based sharing on Facebook

Facebook currently provides a number of mechanisms for selectively sharing content, including the ability to share only within a friend network, to share with a limited list of people, or to share within public or private groups. However, Facebook currently provides limited topic-based sharing mechanisms that allow users to associate individual pieces of content with topics (e.g., tag posts with topics).

From the perspective of a content viewer, Facebook also provides limited direct user control over opting in or out of content by topic. Users can choose to follow, unfollow, or hide friends, or can request to see more or less of certain content. However, Facebook Newsfeeds tend to be controlled algorithmically, and past work has found that users have varied perceptions of the impact of their actions [26].

We focus on the perceived impact of topic-based sharing by focusing on adding content-level topic-based sharing on Facebook. We chose this mechanism because prior work has found that Facebook users might be able to better target desired audiences if they could share content with people interested in individual topics [76].

6.3 Methods

We used an interview ($n = 16$) grounded around a pre-work retrospective diary to probe the topics around which people share on Facebook and other online services. We then explored the potential impact of adding topic-based sharing to Facebook by first asking about the hypothetical impact of a *general topic-based sharing mechanism* for each topic. Then, we used three design mockups, presented to participants as workflows, to explore the potential impact of three possible design mechanisms: a mechanism that would allow users to tag content with a topic and then allow their audiences to *opt out* of viewing it,

Code	*Gender*	*Age*	*Occupation*	*Num. FB Friends*
P01	F	53	Art/writing	101-500
P02	F	23	Student (History, Political Science)	501-1000
P03	F	20	Student (Cognitive Science)	501-1000
P04	F	26	Scientist/engineer	101-500
P05	F	26	Unemployed	101-500
P06	M	23	Unemployed	501-1000
P07	M	27	Business/Mgt./Fin.	101-500
P08	F	26	PhD Student	501-1000
P09	F	50	Admin. support	501-1000
P10	F	25	Education	501-1000
P11	M	30	Unemployed	101-500
P12	F	25	Student (User experience)	
P13	F	33	Unemployed	501-1000
P14	F	65	Retired	51-100
P15	M	23	Student (Information Science)	101-500
P16	F	63	Service	101-500

Table 6.1: Participant demographics: participant code, self-reported gender, age occupational category, range of Facebook friends

a mechanism that would allow users to require their audiences to *opt in* before they saw content on a topic, and a mechanism that would allow users to share on a topic without identifying themselves as the one sharing (*de-identified sharing*).

6.3.1 Recruitment

We recruited participants from the Pittsburgh area using a Craigslist ad, flyers, and an ad on the Carnegie Mellon participant recruitment board. Each recruitment ad provided a link to a recruitment survey. Based on responses to the recruitment survey, we sampled participants for a variety of ages, and occupations (e.g., not all students). We also required that participants be over 18 years of age, self-report as highly proficient English speakers and active Facebook posters, and have access to a computer on which they could complete the pre-work retrospective diary. Table 6.1 provides a list of the participants.

Participants were compensated with a $40 Amazon gift card for completing the pre-work retrospective diary and the in-lab interview.

6.3.2 Pre-work retrospective diary

To prompt discussion of the topics around which participants shared during the interview, approximately two days before the interview, we asked participants to complete a retrospective diary of their online sharing activities for the previous week.

They were sent a link to a survey that asked them to report all the online services they'd used in the previous week to share content online for personal (non-work-related) reasons. For each service, we asked them to open the service and use it to list all the content they had shared in the past week. They briefly described each item shared as well as details such as the type of content, how they shared (e.g., share, reshare, etc.), who they shared with, and why they chose that service to share the content.

6.3.3 Interview

Next, we asked participants to complete an approximately hourlong, semi-structured interview in a lab on the Carnegie Mellon campus. One interviewer performed all the interviews. The interviews were audio recorded and transcribed.

The interview was intended to: 1) examine the topics around which the participant shared on Facebook and other services; 2) explore hypothetical impacts of general topic-based sharing; and 3) examine hypothetical use cases for three, specific mocked-up topic-based sharing mechanisms (tag-based opt-in, tag-based opt-out, and de-identified topic-based sharing).

Topic elicitation

The initial portion of the interview explored the topics around which the participant shared online, and the benefits and shortcomings participants currently associated with sharing on different topics using different services. In this portion of the interview, the interviewer asked the participant to describe the topics they typically shared about on Facebook and the other services they reported in their retrospective diary.

The interviewer used two probes to facilitate this process. The participant was provided with a copy of their retrospective diary to prompt recall of recently shared topics. The interviewer also kept track of topics the participant described on a *topic discussion sheet* that was visible to the participant throughout the interview, and around which the interviewer based questions.

The interviewer began by asking the participant to describe ten topics they typically shared about on Facebook. Prompting for ten topics was drawn from piloting. If the participant was not able to describe ten topics, or wanted to describe more, they were

encouraged to share the number they felt was appropriate. Participants were told to use the provided diaries as reference when describing the topics; however, they were also told that they could include topics that they had not shared about in the previous week.

After the participant described the topics they typically shared on Facebook, the interviewer went through each topic and asked the participant to:

- Describe the content in more detail (e.g., "Tell me more about that?" "What kinds of things do you share related to that?")
- Describe who could see the content (e.g., "Who do you share that with?" "Is that friends only or public?")
- Describe the intended/active audience (e.g., "Is there anyone you particularly want to view it?" "Is there anyone you think would be particularly interested in that?" "Is there anyone you don't want to view that?" "Who do you think views that?")
- Rate how happy or unhappy they were with who currently viewed the content on a five-point scale from very unhappy to very happy, and explain why

The interviewer repeated this topic elicitation with each of the services the participant had mentioned using in the previous week. For each additional service the interviewer began with the topics the participant had previously described sharing, using the topic discussion sheet as a prompt. The interviewer then asked about any additional topics for the service.

Potentially sensitive topics

After the topic elicitations, the interviewer probed whether the participant considered any of the topics they mentioned sensitive, specifically whether there were any that they "would prefer not everyone knew you shared about."

The interviewer finished up this section by asking about any topics the participant self censored, for example any topics the participant tended not to share about online or had considered posting online but decided not to share. The interviewer added these topics to the topic discussion sheet.

General topic-based sharing

The next portion of the interview explored how the participant viewed general, hypothetical topic-based sharing on Facebook for topics they currently shared on Facebook, on other services, or chose not to share online. The interviewer asked the participant to imagine that they could "post on Facebook for just [their] friends interested in the topic." For this portion of the interview no specific mechanism was give.

For each topic the participant had previously described sharing about on Facebook, the interviewer asked the participant to rate, on a five-point scale, whether they would be happier, unhappier or feel the same if they were able to post content on the topic on Facebook just for their friends interested in the topic.

For each topic the participant had described sharing about on services other than Facebook, or had described choosing not to share about online, the interviewer then asked them to rate whether they would be less likely, the same, or more likely to post it on Facebook, instead, if the topic's content would only be shared with their friends interested in the topic.

For each topic the interviewer followed up with probing questions about the participant's feelings and who the participant thought would or would not be interested in the topic. If the participant indicated that they might want to use topic-based sharing, the interviewer also asked about whether they would change anything about how they posted on the topic.

Design mockups: opt-in, opt-out, de-identified topics

In the next portion of the interview the interviewer probed perceived use cases for more concrete versions of topic-based sharing with specific aspects of topic-based sharing mechanisms.

The interviewer presented participants with walkthroughs created from design mockups of three topic-based sharing mechanisms. The mockups were created in Balsamiq and were presented to the user as hypothetical scenarios.

The design presented in all three mockups allowed users to "tag" a status update with a topic, in a manner similar to adding a location or emotion tag in the current Facebook interface. The three mockups presented slightly different mechanisms for interacting with the tagged content. In the sample workflow scenario the interviewer asked participants to imagine that they wanted to "share content related to food" with their friends interested in food:

- **Opt-out topic-based sharing:** In the first workflow the interviewer asked the participant to imagine that after tagging the post with "food" they could share it with all their friends, and a hypothetical friend would have the option to go to a "topics page" (modeled after Facebook's Groups editing interface), see that they were viewing food content from their friends, and decide to stop viewing food content for either specific friends, or all their friends (Figure 6.1)
- **Opt-in topic-based sharing:** In the second workflow the interviewer asked the participant to imagine that they didn't want to share with anyone unless the audience

Figure 6.1: Participants were presented with mocked-up versions of potential topic-based sharing mechanisms created in Balsamiq. The mockups were presented as a series of screens in a scenario. Three screens from the opt-out scenario are presented above. Participants were told they would tag a post with the topic, which, in the example is "food," and then, from their friend's perspective the content would appear tagged as being related to food. Their friend would then be able to view their "topics page" and choose if they wanted to continue viewing this food-related content, or not.

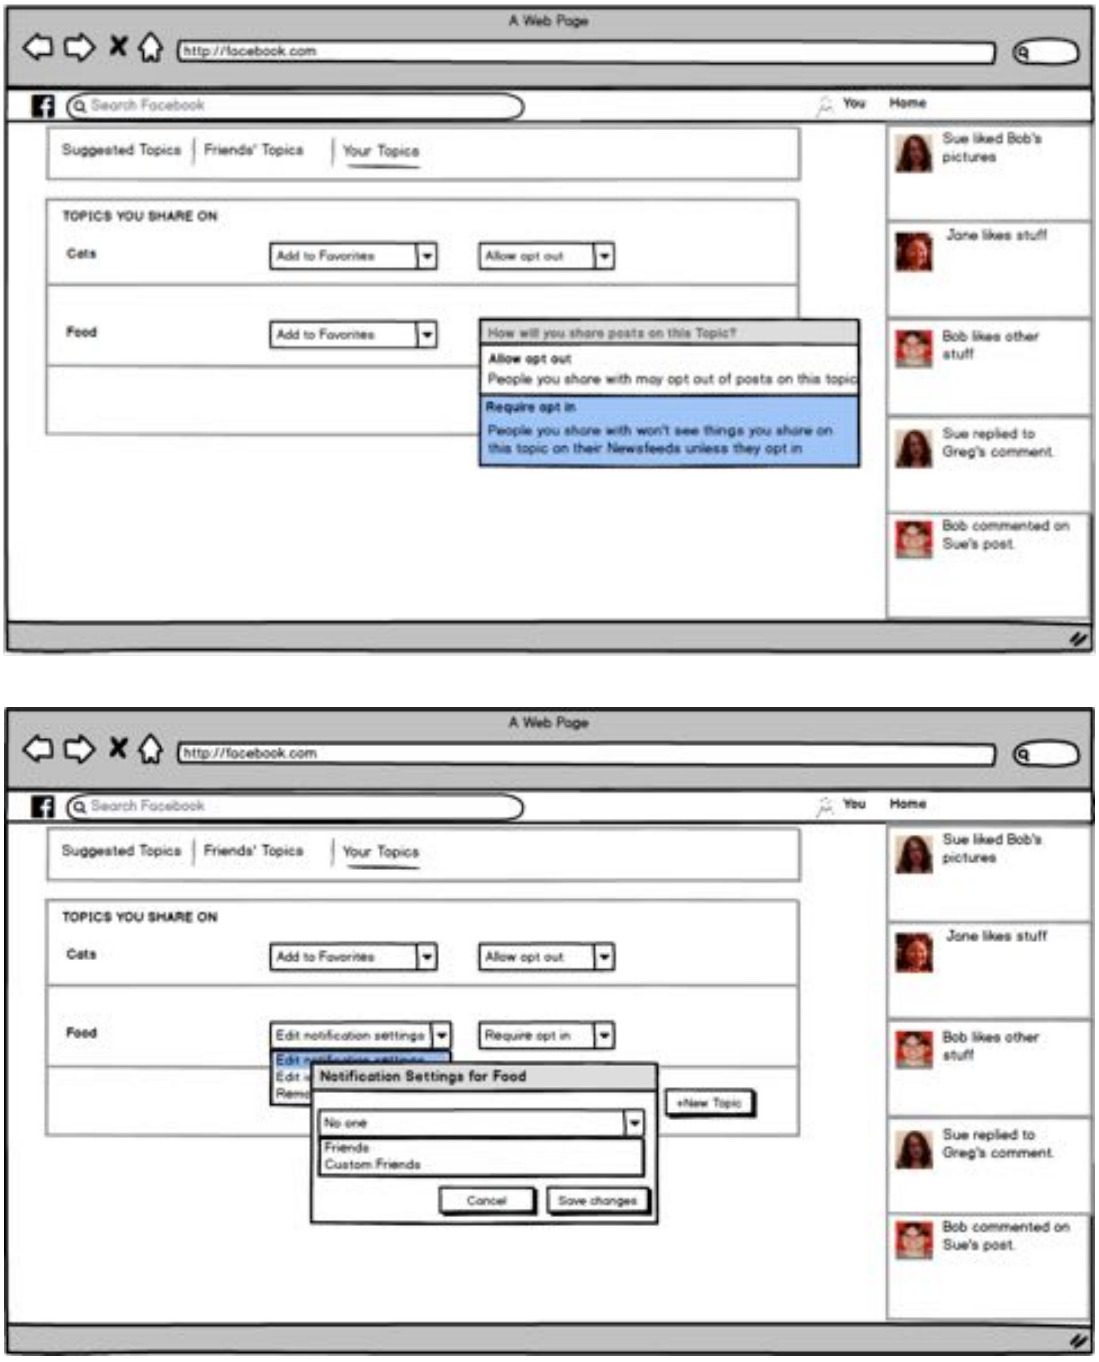

Figure 6.2: In the opt-in workflow, participants were told to imagine that they wanted to share content with their audience members but didn't want the audience members to able to view it unless they opted in to viewing it. When they tagged the content with the topic they could go to their "topics page" and for the topic, set it so that it required people to opt in to view it. They were also shown that they would have the option to notify others about sharing on that topic, or not.

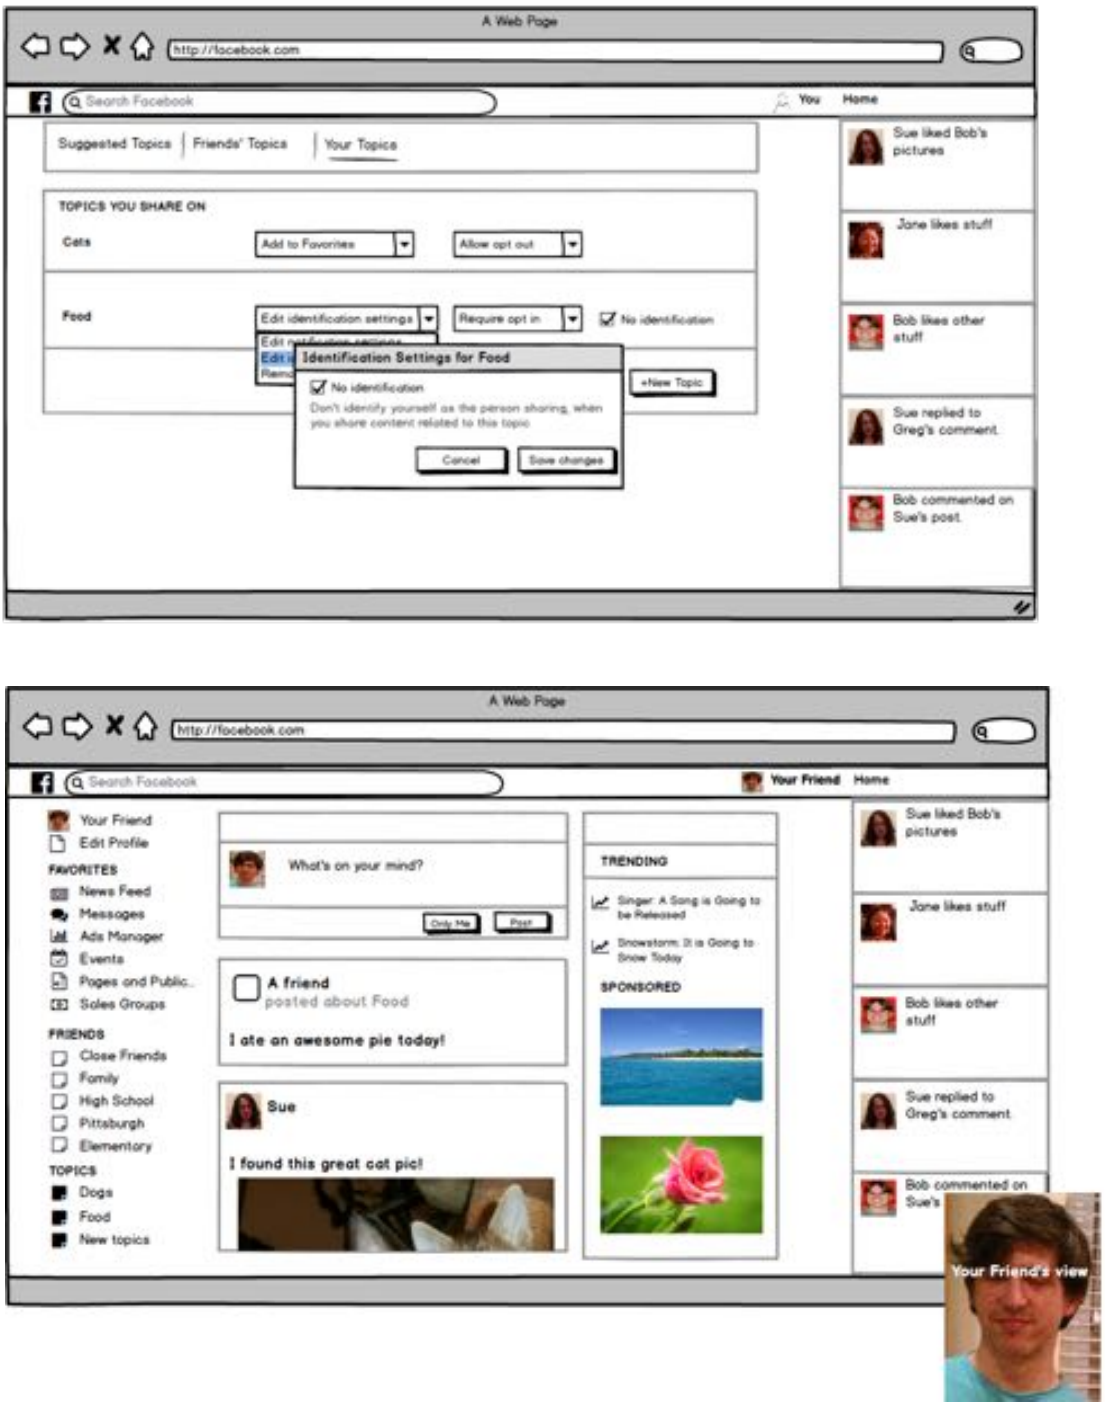

Figure 6.3: In the de-identified workflow, participants were told to imagine that they wanted to share content related to food, but didn't want others to know that they were the ones sharing about "food." They were shown that, for a given topic, they were provided with the option to check a box to share without identifying themselves. Then, the content shared on the topic would appear on a friend's NewsFeed as coming from "a friend" rather than as coming from the participant.

members specifically said they wanted to view content about the topic. In this workflow the participant went to their "topics page," on which they could set a topic to require people to opt in to the content on that topic (in the example the default was opt out). They could also decide whether or not they want to notify people that they were sharing on the topic (Figure 6.2). The workflow user decided to make the topic opt in and to not notify anyone. They then posted the food content, and the workflow demonstrated how a hypothetical friend would need to go to their topics page, and opt in to viewing "food" before it would appear on their NewsFeed.

- **De-identified topic-based sharing:** In the final workflow the interviewer asked the participant to imagine that they could perform opt-in, topic-based sharing but had the additional option to check a box that would de-identify their posts on the topic. Posts on the topic would then appear on their friend's NewsFeed as being posted by "a friend" rather than coming from the participant and would appear with a blank box instead of the participant's photo. In the sample workflow the user checked this box for content related to food, their friend opted in to see food, and the content appeared on their NewsFeed from "a friend" (Figure 6.3).

After completing each mockup walkthrough, the interviewer asked the participant about when the participant any times they thought the mechanism might be useful for for sharing and/or viewing content.

6.3.4 Analysis

We qualitatively coded the interviews. The potential benefits, downsides, strategies, and audiences associated with topic-based sharing relate to broader Facebook, and general online, benefits, risks, and sharing strategies. Thus, we created an initial codebook that drew from related work on broader motivation for use of online services [37, 74, 78, 91], social-networking-site privacy and self-presentation risks [18, 24, 33, 45, 52, 59, 75], sharing and audience-targeting strategies [51, 90], types of online audiences [42, 76, 78], and the benefits and risks of online anonymity [39, 67, 82].

Using this initial codebook, two researchers affinity diagrammed three interviews to create an updated codebook. One researcher then coded the remaining interviews, iteratively updating the codebook. This final codebook is described in the results.

6.4 Results

Participants currently share content with audiences interested in topics, on Facebook and other services, for a variety of reasons including a desire to prompt discussion, inform

others, identify or reach others with a shared interest or avoid bothering people who might not be interested. They currently do so by targeting the content using a mix of access-control tools, choice of channel, sharing content in a generalized format, or choosing to self censor.

When prompted with general hypothetical topic-based sharing on Facebook participants identified a variety of topics for which they thought it might be useful. They tended to feel that sharing with others interested in a topic might allow them to improve their audience, share more, improve discussion, or not bother uninterested parties. However, they also offered a number of cases for which they felt that topic-based sharing on Facebook would not be sufficient or might be detrimental to their sharing goals.

Participants ranged in the types of strategies they thought they might apply to topic-based sharing tools. They varied in the extent to which they thought they would want to filter content based on the people with whom they shared versus the content that they wanted to view. For sharing, participants also tended to feel that some topics had target audiences, so they might want to let others opt out of viewing; whereas, other topics lacked specific audiences, and would benefit less from topic-based targeting.

6.4.1 Current topic-based sharing

A topic-based sharing tool on Facebook would either augment or substitute for participants' current methods for sharing content on different topics with people interested in the topics on Facebook or other services.

Current topic-driven audiences

In line with prior work [38, 76], participants described wanting to share with audiences interested in a variety of topics. The types of interests tended to fall into four broad categories, outlined in Table 6.2 and described in this section.

When sharing, participants described currently intending some topics for audiences they felt would broadly *find topics interesting*. These topics included content shared about animals, funny photos, articles about medicine, feminism, or recipes. For example, P02 described sharing TV-related content, and intending it for "the people that like the fandom."

Participants also intended content on some topics for audiences with particular *general attitudes or opinions*. This was common for topics related to politics, the US presidential election, or religion. Most of these participants tended to explain that they wanted to share on topics with people who held views that were in line with their own. For example, P09 shared memes that made fun of Donald Trump and intended them for her more liberal friends. She explained "I would prefer, actually, that some of the more conservative people

Type of interest-based audience (Audience members who...)	*Sample topics*	*Sample intended audiences*
...would find the topic interesting	cats, television shows, articles related to innovation	people who like cats, fans of TV shows, daughter who is interested in innovation
...have certain general attitudes or opinions	politics, US political election, atheism	conservatives or liberals, Hillary Clinton or Bernie Sanders supporters
...shared traits or life experience with the participant	university topics, museum work, male to female transition support, events	people transitioning male to female, museum professionals, alumni of the university, retirees (interested in the events)
...are involved in an event, activity, organization	organization events, parties, nights out, assignments	members of the organization, people who are attending/attended the party, other people present at the night out, collaborators on the assignment

Table 6.2: Participants described currently intending content on a variety of topics with audiences who had a variety of types of interests in the topics.

weren't looking at them because I'm not trying to make fun of them personally."

Participants also sometimes intended content for audience members who *shared traits or life experiences* with the participant that might make them interested in the topic. Some participants, for example, intended animal-related content for fellow pet owners. P09 explained that she shared a lot of content about her dog with: "other Airedale terrier owners. A lot of my friends are not personal friends. They're not people that I actually quote know. They know about me because my dog has chutzpa. He has a blog. He's not as active anymore, but he has a blog, and that's where a lot of people found me. So I don't actually know a lot of my friends, but they're almost of them are Airedale terrier owners."

Similarly, some participants wanted to share content related to a university or occupation with people who attended the same university or with people with a similar professional background, such as "people I did teaching with in college" (P10). These participants tended to assume those audiences were interested in related topics.

Finally, some participants intended content for others *involved in an event, activity, or organization* with them who would share an interest in related content. They shared content related to party planning, sharing photos after going out at night, collaborating or coordinating when planning an event, or discussing content related to an organization.

Techniques to target topic-driven content

Prior work found that Facebook users generally draw on both audience-reaching and audience-limiting strategies to target desired audiences [51, 90]. Using current mechanisms, participants used both explicit access-control tools and less-explicit audience reaching and limiting techniques to try to target these interest-based audiences for different topics.

Most participants (15) currently tried to share with particular segments of their audiences for some topics. They often used *"platform-based" privacy tools* [90], such as the privacy settings available on Facebook or other platforms, to explicitly share with the limited groups they felt were interested in topics.

For example, participants sometimes shared in private Facebook groups focused on the specific topic: P11 explained that he sometimes shared content related to the Australian Rules football team he played on, and he wanted to share it with "people who played with us or people who used to play with us who want to keep-yeah, up to date with how we're going." He didn't want it to be visible to teams they played against, so he shared in a private, invite-only Facebook group.

Participants also sometimes shared using small group text messaging groups focused on a topic, or sent Snapchat messages to contacts they felt would find the content interesting. At a broader level they used services where they knew that only certain audiences would be present, for example in specific communities on Reddit.

Participants (12) also sometimes tried to avoid the need to limit an audience by perceived interest by only posting content on a service when they considered it acceptable for everyone who might view it (*sharing content appropriate for the "lowest common denominator"*) [51, 90]. P16 explained that she considered everything she posted on Facebook equivalent to "what I would put on the telephone pole."

Within individual topics, participants also sometimes limited the types of content they shared to make the topic acceptable to general audiences. A few participants, for example, tried to only post objective or uncontroversial political content. P12 explained: "I don't really post my personal opinion 'cause my friends could be, like, I know, like they - they're not all like me...some might like Trump, some might not, so I just share some articles I've seen on the Internet."

Participants also drew on "audience-reaching" strategies, similar to those described by Litt and Hargittai [51]. They tended to post content broadly and then try to make sure it content reached audience members who might be interested. In some cases participants (10) assumed that their audience members who were interested in the topic would "self select" and view the content, and the remainder of the audience would ignore it. For example, P13 shared content related to feminist topics. She knows she has some friends

that "are feminists and they like those all the time, and like the stuff that we post all the time, so I think, yeah, it's not like a private interaction, but not everybody, like, read those"

To provide more targeting, participants (12) sometimes posted content and then *"interacted" with the audience* by tagging particular people, highlighting the content with hashtags, posting on peoples' walls, or telling people about the content online or offline. Similarly, a few participants (3) posted broadly but *changed the "script" of the post* so that its meaning would only be clear to certain members of their audience [51]. For example, P06 sometimes posted content on Facebook related to a friend who had passed away. He worded the posts vaguely so that "the people viewing it are either people who know that I'm having trouble with her loss, or just like people who think it's just like some vague poetic thing."

Alongside these content-driven, audience-reaching and limiting methods, participants (8) also targeted their sharing for topics by *controlling overall access to their networks* [90]. Several participants described strictly controlling who could or could not join their Facebook network, so that they would have freedom to assume that their audience members were willing to view topics. For example, P08 "curated [her] Facebook page. I don't have anybody that I wouldn't want to read the stuff. So they read it, okay. If they don't, it's okay." Similarly, on Twitter, P10 explained that "people that I let follow me are generally people my age range so we'd have similar interests."

6.4.2 Use cases for topic-based sharing on Facebook

Almost all participants (13/16) felt that they would benefit from topic-based sharing on Facebook for some topics, as a supplement or substitute for current targeting techniques. Participants tended to feel that they could use topic-based mechanisms to reduce the risk of boring or annoying their audiences as well as potentially reduce the risk of offending audience members by sharing on topics that might be controversial or misperceived. Participants perceived a number of potential benefits for topic-based sharing on Facebook, including the ability to target better audiences, find people interested in topics, and post more or better quality content. Table 6.3 outlines these potential benefits and potentially reduced risks.

Reducing the risk of potential over-posting or offense

As described, participants sometimes posted broadly and assumed audiences could self-select if interested, or self censored to avoid sharing too much or potentially offensive content.

Thus, participants (7) felt that adding topic-based sharing on Facebook could potentially *reduce the risk of oversharing* and of boring or annoying people by posting too much on a topic.

This risk has emerged in prior work as a reason for self-censorship, un-following or un-friending behaviors, and perceptions of violated social norms on Facebook [45, 59, 75, 76].

Participants tended to be concerned about this dynamic for topics that they felt might not be interesting to their general audience. These topics ranged from cats or pictures of the participant's grandchildren, to posts related to the participant's university. As P02 explained, she tries not to post too much about her university because: "I have family members that didn't go to [anonymized university] and...they're not close family members so they probably don't care that much. I try not to like - part of the reason I don't post very often on Facebook is that I try not to be annoying."

Some participants also tended to feel that being able to target interested audiences might help them avoid *the risk of offending others*, a dynamic also seen as emerging as a reason for self censorship or un-following [45, 76]. This tended to occur for political topics, but also for topics participants felt some audience members might misperceive or take out of context.

P09 described, for example, how she wanted to be able to share content related to "subculture" topics on Facebook, such as the occult. She assumed, however, that these types of posts would make some people uncomfortable. She explained that she would be happier sharing on this topic if she could target it at people interested in countercultural content because it would allow her "To be able to say, 'These are - this is my tribe. These are my people. You will understand everything that I'm posting and you won't be offended by it."'

Potential impacts of topic-based sharing

In combination with reducing these risks, participants tended to feel that adding topic-based sharing on Facebook could provide them a number of benefits.

Participants (8) tended to feel that topic-driven sharing mechanisms on Facebook would allow them to specifically post to *narrower, or better-quality, audiences* rather than posting broadly and relying on audiences to self-select. Their perception of the makeup of these improved audience varied. In some cases participants wanted audiences filtered for people who wouldn't be offended by posts on a topic, or who wouldn't be annoyed when viewing the content. P01 envisioned such a mechanism allowing her to share pictures of her grandchildren with people who were interested in them: "You could have a grandma list. A grandma and close friends list that I could send those to, and not bother the rest of the group with."

Participants (9) also sometimes felt a topic-based mechanism would facilitate *locating people interested in topics* to view things they posted on that topic. Some participants felt

that targeting content on a topic at people interested in that topic might generally increase the amount of attention paid to posts.

Participants also sometimes felt that targeting people interested in a topic might lead to more effective assistance or social support. In one case, P12 described using Facebook to ask technical questions. She felt that narrowing the audience to people interested in technical topics would increase her chances of getting useful help: "I think people would be more interested in helping me because...they're closer to me."

Similarly, P06 used Facebook to express grief. As described previously, he currently relies on vaguely-worded posts to reach his target audience. However, he felt that if he could target an audience interested in the topic, more narrowly, he might "write more specifically about her" and would "get more of that support of like the people who are also grieving her loss."

Two participants also described potentially being able to use this type of targeting to share an experience. P10 explained that she liked to share content at sports games and she'd be happy to be able to share with other fans: "if I'm at a hockey game, if people, it was all like [anonymized team] fans and then people like to see that, so I would think that would be kinda cool." Similarly, P04 explained that she might like to talk to her Facebook friends about television shows, and sharing with people interested in the shows might let her find other fans and share without worrying about giving away spoilers.

Participants (5) also sometimes felt that targeting a more focused audience might allow for improved discussions on topics. This dynamic occurred for political content for which participants thought limiting the audience to people interested might lead to a more constructive or informed conversations.

However, participants also felt that limiting content to people interested in a topic might prompt better discussion in other areas. For example, P10, a teacher, liked the idea of being able to post the classroom design ideas she currently posted on Pinterest on Facebook for people interested in them. She explained: "it would be almost like a little teacher forum if you want to go online and find that...'cause it's always nice to bounce ideas off each other."

Similarly, a few participants (4) felt that if they knew they were sharing for an audience who was interested in the topic they might post more, or better quality content. P11 posted basketball-related content, and explained that if he were posting just for people interested in the content he could "sort of cater it towards the people that are interested in it, who want to see it."

Potential use cases for topic-based viewing on Facebook

When presented with the mocked-up opt-out and opt-in mechanisms, participants also sometimes felt that the mechanisms might provide them benefits as Facebook audience members viewing others' content.

Filtering uninteresting or offensive content Prior work found that people sometimes objected to others oversharing, or sharing uninteresting content [45, 59, 75]. Participants, similarly, tended to feel that being able to opt out of friends' posts, or choose when to opt in to select posts might allow them to avoid content they weren't interested in, especially for friend they felt overshared. For example, participants described wanting to opt out of content for friends who posted about content such as babies, exercise, retail, and details of everyday life. Participants felt they could clean up their feeds: "like if it's my Facebook and a ton of people are posting about, like, um, I'll use this example, a bunch of cat pictures, like, I don't really care about cats, so I could go in and get rid of it just to kinda clear up the junk on Facebook" (P10).

A few participants also wanted to use opt in or opt out mechanisms to avoid viewing things they felt were potentially upsetting or that they disagreed with. For example, P16 is a vegan and would prefer to opt out of pictures of meat-related foods: "it would be real nice to not have people's pictures of their nice juicy barbecue, because some of the people that I work with will post pictures of stuff that I would really rather not see."

Preference for opt out For viewing content by topic, participants tended to prefer the idea of allowing opt out to requiring people to opt in. Several participants felt that an opt-in-based mechanism would require too much work. Some participants also felt that an opt-in mechanism might remove the ability to discover unexpected content, which they felt was one of the benefits of Facebook: "I feel like [the opt in] mechanic is just kind of not really what Facebook is supposed to do" (P02).

However, a few participants described topics for which they'd want to know that their audiences had explicitly expressed interest, and for which opt in might be useful. P04 explained, for example, that she might want to use an opt in for sharing about television shows "because that's like you can opt in if you know, or if you know you are not going to be able to watch a specific episode you would be left out so it would be, you would assume that people wouldn't be angry about Game of Throne spoilers if they were opted in to like read posts about it."

Potential uses of de-identified topics on Facebook

We also probed for participants' perceptions of potential uses for sharing on topics on Facebook in a de-identified manner.

Negative perceptions of de-identification Most participants (14/16) generally perceived de-identified sharing on Facebook negatively. Some felt, in line with prior work [39, 67] that sharing without a tie to one's identity on Facebook would be creepy, suspicious, or inappropriate for the platform. Similarly, several (8) felt that removing Facebook's ties to real-life identities would lead to negative behaviors [39]. P14 explained "I think it's important that Facebook retain a level of transparency, because I could imagine that very quickly people would be posting things that are really offensive, and if people can do that anonymously, um, I think we would lose a level of civility."

Some participants (7) also explained that they saw Facebook as a service for which people specifically benefited form having their identities tied to the material they posted: "The people I know, the people that post that kind of stuff want people to find out about it. That's why they use Facebook" (P08).

Use cases for sensitive topics While increased anonymity can cause negative disinhibition, as perceived by most participants, it can also allow for increased freedom to discuss sensitive topics or provide honest feedback [14, 39]. A few participants, pointed out use cases that illustrate how de-identified sharing might be useful for some potentially sensitive topics.

In one case, P06, who is transgender, and was using Facebook for support while transitioning from female to male explained that de-identified sharing might be useful for allowing people to get support when they had sensitive or embarassing questions. For example, he thought he might want to use it: "Like for, like hormone updates, like one of the things, one of the like side effects of hormones is acne. And so like I don't necessarily want people to know that I'm posting about it, but it would be nice to like write about it, and then get some responses like, oh, this is what I tried."

Similarly, P11 had previously worked as a police officer. He explained that it would be useful to be able to use de-identified posts on Facebook to respond to negative statements about police officers: "some people post some pretty nasty stuff, um, so I guess being able to put my point of view forward, um, and it's sort of frowned upon by the police for posting on Facebook, and I guess- 'cause you can get in trouble for certain things."

6.4.3 Needs unmet by topic-based sharing

Most participants described topics for which adding topic-based sharing to Facebook might reduce the risk of oversharing or offense, or provide benefits. However, topic-based sharing on Facebook would not meet all online-selective-sharing needs. All participants also described topics for which they currently shared on Facebook and other services for which topic-based sharing mechanisms on Facebook would not improve their ability to share, or, if used, could be detrimental to their abilities to reach desired audiences.

Privacy needs

Many participants (9) described topics they wouldn't want to share on Facebook because they wanted explicit privacy. This tended to occur for topics like pictures that included other people, personal topics, or funny things that they felt they might find embarrassing to have broadcast more publicly. These privacy concerns included both consideration of their own privacy, as well as consideration of the privacy of other people included or involved in the content, for example other people in photos [52]. In a few cases participants were also concerned about protecting the content from Facebook itself.

Participants tended to want to share these topics on services that allowed them to explicitly access more limited audiences, or, in some cases didn't want to share on these topics online using either current or topic-based mechanisms. For example, P10 used text messaging for "personal things" when she only wanted "things that are really just for that one person that, you know, or that group even, that you don't need to tell other people."

In these cases, participants didn't want to share on Facebook using topic-based sharing mechanisms, because they valued the ability to explicitly limit the audience rather than target an interested audience. For example, P02 wouldn't want to post "personal things" to Facebook using topic-driven sharing because "there are things that I don't want more than a couple of specific people to know."

Peripheral audiences and stratification

A few participants more specifically felt that topic-based mechanisms on Facebook could be counter-productive for certain topics. Prior work found that Facebook users sometimes broadcast on Facebook to reach "peripheral audiences" beyond those the user may specifically consider [51]. Similarly, participants sometimes wanted to reach audiences who might be prompted to become interested in a topic, rather than using topic-based sharing tools to limit the audience to people who were already interested in the topic. For example, P07 described using Facebook for a fundraiser: "we marketed it just to the people we knew in like our community, but we ended up making $700.00 in sales just from this other

organization that, you know, one person saw it, shared it to their group or something. And - I don't know - but we got so many like random people coming in. And a lot of the people that we...we reached out to on Facebook, they just kind of heard about somehow and they were interested. And that's kind of a situation where we didn't even know they were interested."

Similarly, a few participants felt that topic-based sharing might limit the diversity of viewpoints they were able to expose their friends to, increasing undesirable dynamics such as political stratification [64, 74]. P15, for example, wouldn't want to use topic-based sharing for content he posted about atheism, because then, "the other people don't get to see alternative viewpoints, and I think that's important."

Desire for other features and audiences

Some participants (6) also felt that topic-based mechanisms on Facebook wouldn't meet their needs because they wanted to reach audience members who didn't use Facebook, or because they wanted to use features only available on other platforms. For example, several participants described wanting the ephemerality of Snapchat for particular topics. P04 wouldn't want to share content she shared on Snapchat on Facebook, even with the addition of topic-based sharing because she liked: "the disposal aspect [of Snapchat]...I can make a weird face in a picture and unless they specifically save it...it won't be out there forever, probably."

6.4.4 Strategies for topic-based sharing

When discussing topics they felt might benefit from topic-based sharing participants also tended to describe high-level strategies they felt would shape their potential uses of topic-based sharing and viewing mechanisms.

For sharing content, participants tended to view topics as more or less *intended for an audience* or *intended for general sharing*. When discussing potential use of opt-out or opt-in mechanisms for viewing content by topic, some participants tended to base their strategies around a goal of prioritizing their *relationships with the people* whose content they might be viewing rather than wanting to focus on their personal interest in specific topics. Other participants, however, wanted to base opt-out or opt-in decisions primarily around their interest in the *the topics being shared*

General content versus content for an audience

Participants tended to feel they *shared on some topics generally* that would benefit less from topic-based sharing mechanisms. For a few of these topics, participants described sharing

	Description	Examples
Risks potentially reduced	Reduce the risk of oversharing	
	Reduce the risk of annoying or offending others	
Potential benefits	Share with narrower/better quality audiences	Share pictures of the grandchildren with people who wouldn't be bothered by it
	Locate people interested in a topic	Get assistance or social support (e.g., ask a technical question, express grief and get support)
	Share an experience	Share with other fans at a sports event
	Improve discussions on a topic	Have a more constructive political conversation, create a "little teacher forum" for classroom ideas
	Post more or better quality content	Cater basketball-related content toward people who are interested in basketball

Table 6.3: Participants felt that topic-based sharing on Facebook could potentially both reduce the risk of oversharing or annoying/offending others, as well as provide several potential benefits.

as expressing important content, regardless of audience interest. For example, P09 shared about street art and design. She wouldn't want to use topic-based sharing for this topic, because: "it's pretty amazing stuff, and, and it means a lot to me, so I think that a, a lot of what means to me I would want people to see on Facebook."

Participants also felt that there was a lack of topic-driven audiences for topics they tended to assume their audiences were indifferent to, or that they shared about rarely. For example, P07 didn't feel that topic-based sharing would impact her sharing on school events because, "I think anybody really cares either way about school events." Similarly, P05 shared travel photos explaining: "I tend to, if I go someplace, I post pictures. I don't say, oh, well, I don't know if these people want to see it so I'm not going to post."

Participants also tended to share on other topics, however, that they felt were *targeted at particular audiences*, for example because of shared hobbies or passions or to share information or resources. They tended to feel these topics were more likely to benefit from topic-based sharing mechanisms. They felt, for example, that these topics would make specific audiences happy or would get a desired response from specific audiences. P15 explained, for example, that sharing "cute animal pictures" would make friends who like animals happy. Topics that tended to fall in this category included hobbies, like

ultimate frisbee or recipes, passions and interests, like animals, TV shows or movies, and collaborative activities or events, like project work or party planning.

Relationship-driven decisions versus content-driven decisions

When presented with the opt-out and opt-in design mockups participants also tended to describe strategies for deciding whether or not they would choose to view their friends' content. Participants' strategies for these decisions tended to either be more relationship-driven or more content-driven.

Relationship-driven viewing Several participants felt that they would base their decision on whether or not to view content on the person sharing, or their relationship with that person, rather than the topic of the content. For example, P15 explained, "I don't think I filter content based on what it's about, but more related to who's posting it...I attribute more credibility to people rather than the things that are being posted."

Some of these participants didn't like the idea of choosing topics to view using opt in or opt out mechanisms on Facebook. As P14 described: "I want to see what my friends are interested in...And if that means that I'm seeing things about them that I may not like, well, so-so be it. They are who they are." Similarly, P03 explained, "if someone that I know is supporting Donald Trump, I do want to see it 'cause I want to know who's doing it."

Some of these participants preferred to unfriend or unfollow someone if they met a threshold of posting objectionable content, rather than opting out of individual topics. Other participants who prioritized viewing peoples' content on the individual level felt that they might want to opt in or out of some topics but would want to choose to do so on a person-dependent basis. For example, they might know that one friend shared on one or two topics they found annoying or objectionable, so they would opt out of those topics, but only for that friend. P14 mentioned that she'd like to opt out of retail posts from one friend, but "it's very specific."

Content-driven viewing Other participants, however, liked the idea of being able to adjust the content they viewed by turning topics on and off, rather than focusing on viewing posts on a friend-by-friend basis. Some of these participants explained that they found sharing on some topics overwhelming. For example, P09 explained: "you want people to remain your friends per se, but you don't necessarily want to see all the, oh, she's posting about her running a marathon again."

In some cases, these participants also felt that filtering topics for certain people might let them remain friends with someone who shared material they found offensive or annoying: "Some people post things that I find very inflammatory and uninformed and I tend to want

to argue with them. And if I don't see [those posts] then I don't want to argue with them and then I can still be friends with them" (P05).

6.5 Limitations

This study provides high-level insights into when, and what types, of topic-based sharing mechanisms might be useful on Facebook. However, we asked participants to evaluate their hypothetical use of general and specific topic-based-sharing mechanisms. This allowed us to elicit feedback without implementing a full system; however, actual behaviors may not match hypothetical behaviors.

To partially ameliorate this effect we grounded participant responses in topics they actually shared online through retrospective diaries and by beginning the study with a topic elicitation. Actual use of topic-based sharing mechanisms, however, would likely be highly impacted by factors such as the level of effort involved in their use, friends' particiation, and user trust in the mechanisms' effectiveness. Although this work provides a basis for considering use cases, future work grounded in fuller implementations, potentially in a field-study environment, would be necessary to understand the usability of specific mechanisms and the potential impact of topic-based sharing in more realistic scenarios.

6.6 Discussion of design implications

These findings offer a number of design implications for topic-based sharing mechanisms on Facebook, and, more broadly on other services that support sharing through topic-driven mechanisms.

6.6.1 Support for different types of interest-based audiences

Participants described intending current content for several types of interest-driven audiences for different types of topics (outlined in Table 6.2). Services provide mechanisms that support users' abilities to explicitly target these different types of interest-driven audiences to different degrees.

On Facebook, for example, there is support for finding people who are involved in events/activities/organizations, through Facebook Groups, or Events. Where it is easy to create groups associated with a trait or opinion, or where the trait aligns with an identifiable life stage (e.g., high school friends), it may also be possible to create interest-based groups or lists for targeting others interested in topics. However, it is currently relatively difficult, on Facebook to find people who generally find topics interesting, or who have general

attitudes or opinions that they don't identify with on a group level. This dynamic emerged when participants currently tended to use broader targeting mechanisms, on Facebook, to try to reach these audiences, for example depending on interested audience members to self select.

Thus, when there are gaps in a service's ability to allow users to target a specific type of interest-driven audience, users may fall back on more general targeting mechanisms. Adapting selective sharing mechanisms to target gaps around these different types of interest-based audiences may allow users to better target desired audiences for topics. Improving this type of sharing may lead users to be able to share currently self-censored or broadly targeted content with less risk of oversharing, misinterpretation, or causing offense.

6.6.2 Need to combine topic-based and other sharing features

Prior work demonstrated that selective sharing takes place in the context of other tasks and online sharing needs [78]. Similarly, we observed that participants tended not to want to use topic-based sharing for some topics because a desire for audience access, more explicit privacy guarantees, or other services' features overrode the desire to target an interested audience on Facebook.

It's, therefore, important that topic-based sharing mechanisms incorporate the ability to perform other tasks that the user may way to perform in combination with targeting the audience interested in the topic. As seen among our participants, for selective sharing, privacy may be often be important. Users may want to target a particular, topic-driven, audience but may not want to give up their ability to simultaneously explicitly limit the overall audience.

Thus, topic-based sharing should incorporated into mechanisms that provide the ability to transparently, and explicitly, limit audiences. Future work could explore this dynamic by examining the potential for topic-based sharing in more privacy-bounded environments, for example in private Facebook groups or on email lists, or could focus on exploring topic-based sharing and filtering in combination with more explicit privacy controls.

6.6.3 Relationships versus content

We observed that when deciding whether they might want to opt out of viewing content some participants felt that they would want to base viewing decisions on their relationships with others, while other participants would want to base the decision more on the topics they were interested in.

To capture both these types of preferences, topic-based sharing tools could allow users to have the granularity to make opt-out decisions both at the topic level and at the individual level. For example a tool could include both the ability to opt out of food-related content for all friends or opt out of food-related content for just your friend 'Bob'. This increased granularity might allow users to base sharing and viewing decisions on both topic-interest and relationship.

6.6.4 Transparency into viewership

We found that some participants felt that if they knew that people were interested in content posted on a topic they might post more or post better content. However, some participants wanted transparency into whether anyone was interested in a topic, who was interested, or how many people were interested, in order to achieve this benefit.

Topic-based sharing design should consider whether and how to inform people about the number of people in the active audience. Prior work has found that impressions of viewership can depend on cues provided by the service as well as user agency [49]. Future work could examine the impact on potential posting behaviors of informing people when or if people opted in or out of viewing their content for different topics.

Topic-based sharing mechanisms could also consider how best to incorporate notifications for sharing on particular topics. In the mockup workflows we showed participants the user was able to notify their audience about topics that the audience was able to opt in to. Participants had mixed feelings about these notifications. Some felt that they notifications could be bothersome, while others felt that they would want to notify their audiences and not depend on the audience to find out about the topic on their own. A topic-based mechanism could seek to support both types of topic-based sharing - topics that users share and then notify their audience about to give them the option to opt in to, and topics the user might want to share about and then make available for search, without explicit notification.

6.6.5 Tag-based topic management

We presented participants with prototypes in which topic-based sharing was performed by manually tagging posts with a topic, and then viewing one's own, and friends' topics on a "topics page" (modeled off the page used to manage Facebook Groups in the current Facebook interface).

Several participants were concerned that it would take too much time to add topics to posts. As with other tag-based systems, a topic-driven sharing design that depended on tagging posts with topics would need to consider how make this activity fast and usable.

One method for doing so might be to use metadata-related tags as topics, some of which are already included on Facebook. For example, some participants wanted to share with people nearby, or topics such as "photos." Location and or file-type tags could potentially be drawn from system-based sources.

For other topics, it might be possible to infer topics from content. Adding automated tags could increase speed and consistency. However, inferring topics might be also seen as privacy invasive, depending on where the data was take from to infer the user's potential interests. Future work could examine the tradeoffs between inferred tags, having users create tags, or a combined system, focusing on factors including speed, privacy concerns, accuracy and consistency, and usefulness for preventing the risks and achieving the benefits participants described in this work.

Participants also expressed concerns about organizing and finding topics for perform tasks such as opting out or allowing friends to opt in to topics. Thus, a system would also need to consider how to organize the topics in a central location. Options might include organizing by general interest, number of friends sharing on a topic, or user interest. Future work would be needed to address usable topic tagging and organization.

6.7 Conclusions

As described in Chapter 4, adding topic-based sharing on Facebook has the potential to allow people to better share content with their desired audiences. However, as Chapter 5 describes, this dynamic must be evaluated in the context of ecosystem-level sharing behaviors.

In this chapter we used a lab-based study, centered around a retrospective diary, to explore topics that participants shared on Facebook and other services. We also looked at the potential impact of Facebook-based topic-based sharing mechanisms for content shared on Facebook, other services, and currently unshared. We found that there were a number of use cases for Facebook-based topic-based sharing including reducing the risk of oversharing or causing offense, sharing with a narrower or better audience, improving discussions, allowing participants to locate interested audiences, and potentially allowing participants to share better content with a more targeted audience.

To design topic-based sharing mechanisms, however, designers would need to focus on the mechanisms' role within the overall sharing ecosystem as well as the broader task- and audience-driven context (as seen in Chapter 5). Participants also described topics for which they would not use topic-based sharing on Facebook, often for task or audience-driven reasons (e.g., audiences available on particular services, the ability to use a particular feature to accomplish a task, explicit privacy needs). Mechanisms would need to consider

how to incorporate these needs into the ability to target audiences interested in particular content.

7 | Conclusion

People increasingly share content online in everyday contexts across a variety of services and devices. In this space, sharing needs can be complex, multi-dimensional, and challenging. For example, when sharing, people may want content to reach particular audiences, may want to accomplish certain tasks, and may, separately, or simultaneously, want to self present in certain ways.

Online services can provide sharing tools to help users share in a desired manner. For example, tools can help users share with their desired audiences, by allowing them to target specific individuals or groups, or behavioral-nudge-based or other educational tools can be used to help increase audience awareness of their audiences or how they might be perceived by others. Other types of tools, like delete or undo mechanisms can be used to help users revoke potentially regrettable decisions. However, when these tools are insufficient, and users are unable to meet their sharing needs, regret or other suboptimal outcomes, such as self-censorship can occur.

In this thesis I provide several perspectives into online sharing that can provide insights into creating sharing mechanisms to help meet user needs in everyday online sharing contexts. In Chapters 3 and 4 I look at how sharing mechanisms may fall short, leading to regret or self censorship. In Chapter 5 I examine sharing decisions across channels, including the role of selective-sharing-based features, among other task, audience, and feature-driven dynamics. Finally, in Chapter 6 I draw on these previous studies to explore the potential for adding topic-based sharing mechanisms to Facebook. These studies offer several high-level insights around designing sharing mechanisms for online tools.

7.0.1 Drawing on offline metaphor

Everyday online sharing reflects some offline social dynamics and behaviors; however, removing the limits of physical space also removes some of the aspects of offline sharing and communication that people depend on to manage sharing offline. Helping users share online in a way that better reflects offline sharing may provide some opportunities for improving online sharing mechanisms.

Chapter 3 describes how some aspects of online and offline regretted messaging are similar. However, other aspects of regretted messaging emerged in the online environment. For example, online users tended to misunderstand or not fully realize their audience when posting, regretted broadcasting content, or didn't immediately realize they should regret context because of a lack of a physical audience.

Similarly, when we looked at participants' chooses for channels for sharing personal content in Chapter 5 we found that some choices reflected offline communication dynamics, such as channels participants used with people they felt close to or saw regularly. However, other channels reflected online affordances, including the ability to broadcast content, or set up privacy rules. This leads to increasing complexity that wouldn't exist in an offline context.

Differences between online and offline communications, especially those that lead to added complexity or regret, reveal opportunities to help users share by adding missing metaphors or prompts online. These types of prompts may help users avoid misunderstandings, or may allow users to more easily rely on dynamics they would otherwise be missing online.

We observed in Chapter 5, for example, that participants tended to use highly controlled sharing mechanisms, like one-on-one chat for personal or private content, reflecting an offline desire for a private conversation. Providing the ability to share in this manner, or a scaled-up ability to have small, closely-controlled, group sharing, allows users to share in a method that reflects offline dynamics.

Similarly, tools that focus on adding offline metaphors where they are missing online may help users bridge this gap. One example of this is graphic representations of potential audience emotions of online communications [53] Another example is the use of behavioral nudges to remind users of the scope of their audience in broadcast contexts [95].

Future work should continue to focus on the differences between offline and online communications, especially where these differences may result in suboptimal impacts. Where these differences occur there may be opportunities to use mental-model-driven, or nudge-based, techniques to help users better understand, or be aware of, the online dynamics.

This dynamic may become increasingly important as sharing devices becomes more ubiquitous, and move into spaces that have traditionally been offline. Understanding how people expect to be able to react to communications, or their expected audiences could help inform the design of usable notice and choice mechanisms in more ubiquitous environments.

7.0.2 Opportunities for online contexts

Although online sharing may reflect some offline dynamics, it also affords the opportunity to target audiences that might be difficult to reach offline. For example, a user can easily broadcast a message on a social networking site for a large audience, or can send an email to a large number of people. A user can also limit audiences in ways that would be impossible or difficult offline, for example one can set up access control rules for a folder online or limit a broadcast message to a widely dispersed audience. Online sharing mechanisms can, therefore, also move beyond offline metaphors, to also help users target desired audiences by drawing on aspects of sharing that would be difficult to mirror in offline contexts.

In Chapter 4 we found that participants tended to self censor content that they would potentially want to share on Facebook with a variety of audiences, including individuals, defined groups of people, and ambiguous groups of people (e.g., "hockey friends"), often defined by traits or interests. Targeting these groups would require adding sharing mechanisms allowed users to find groups of people with these traits. We explored the idea of this type of targeting and filtering in Chapter 6 and found that most participants tended to like the idea of this type of sharing for some topics.

Thus, sharing tools can potentially offer targeted sharing that may not be available offline. Online sharing affords the ability to broadcast widely, which can lead to privacy challenges. However, this ability also affords the opportunity to provide sharing mechanisms that allow users to filter their audience along a number of dimensions that are not available offline. In Chapter 6 we focused on user-provided topic-driven tags as one method for driving improved targeting on Facebook. However, online contexts provide a variety of other modalities around which sharing could be based and that could be pursued in future work.

7.0.3 Understanding shortfalls provide insights

The thesis focused on providing insight into sharing mechanisms by understanding when current sharing mechanisms fall short. Specifically, we looked at regretted content (Chapter 3), self-censored content (Chapter 4), and instances in which users switched between or combined channels to meet sharing needs (Chapter 5).

Creating usable sharing tools, especially for everyday contexts, requires examining the full ecosystem of sharing. This includes all the sharing channels people currently use, as well as content they may not share or may share and then regret using current tools. As observed in Chapter 5, selective sharing is one component of a sharing decision. Thus, sharing tools should account for their position in the overall task, audience, and feature-driven ecosystem.

Future work should account for this dynamic by focusing on including selective-sharing tools that incorporate both selective-sharing and broader, task-driven needs. For example, a tool could allow a user to target a particular audience and edit a document.

Sharing tools should also be evaluated at the ecosystem level. Because everyday users can typically move between and combine services to meet their need, a selective-sharing tool should be evaluated assuming that if it does not meet a user's needs in one area, they may switch to a different tool. For example, if a tool allows a user to target a particular audience, but doesn't allow the user to edit a document, the user may switch to different tool that allows for editing. Evaluating tools in this full ecosystem context allows the designer to determine what services they intend to supplement or substitute for, as well as what services they might be able to complement.

Bibliography

[1] A. Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy*, 7(6):82–85, December 2009. Cited on page 30.

[2] Alessandro Acquisti and Ralph Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies*, pages 36–58. Springer, 2006. Cited on pages 6 and 31.

[3] Hazim Almuhimedi, Shomir Wilson, Bin Liu, Norman Sadeh, and Alessandro Acquisti. Tweets are forever: A large-scale quantitative analysis of deleted tweets. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pages 897–908. ACM, 2013. Cited on page 31.

[4] Saleema Amershi, James Fogarty, and Daniel Weld. Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM, 2012. Cited on page 11.

[5] Morgan Ames and Mor Naaman. Why we tag: Motivations for annotation in mobile and online media. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 971–980. ACM, 2007. Cited on page 81.

[6] Oshrat Ayalon and Eran Toch. Retrospective privacy: Managing longitudinal privacy in online social networks. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, pages 4:1–4:13, New York, NY, USA, 2013. ACM. Cited on pages 1 and 12.

[7] Rebecca Balebako, Pedro Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Nudging users towards privacy on mobile devices. CHIPINC 2011, 2011. Cited on page 30.

[8] Lujo Bauer, Lorrie Faith Cranor, Saranga Komanduri, Michelle L. Mazurek, Michael K. Reiter, Manya Sleeper, and Blase Ur. The post anachronism: The temporal dimension of Facebook privacy. In *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society*, WPES '13, pages 1–12, New York, NY, USA, 2013. ACM.

Cited on pages 1 and 12.

[9] Lujo Bauer, Lorrie Faith Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. A user study of policy creation in a flexible access-control system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 543–552, New York, NY, USA, 2008. ACM. Cited on pages 5 and 12.

[10] Lujo Bauer, Lorrie Faith Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. Real life challenges in access-control management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 899–908. ACM, 2009. Cited on page 5.

[11] Ofer Bergman, Steve Whittaker, and Noa Falk. Shared files: The retrieval perspective. *Journal of the Association for Information Science and Technology*, 65(10):1949–1963, 2014. Cited on pages 7, 56, and 60.

[12] Michael S Bernstein, Eytan Bakshy, Moira Burke, and Brian Karrer. Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM, 2013. Cited on page 6.

[13] Andrew Besmer and Heather Lipford. Tagged photos. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems - CHI EA '09*, page 4585, Boston, MA, USA, 2009. Cited on page 6.

[14] Jeremy Birnholtz, Nicholas Aaron Ross Merola, and Arindam Paul. Is it weird to still be a virgin?:Anonymous, locally targeted questions on Facebook Confession Boards. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2613–2622. ACM, 2015. Cited on pages 12, 82, and 98.

[15] Tristan Blanc-Brude and Dominique L. Scapin. What do people recall about their documents?: Implications for desktop search tools. In *Proceedings of the 12th International Conference on Intelligent User Interfaces*, pages 102–111. ACM, 2007. Cited on page 10.

[16] Jeffrey Boase. Personal networks and the personal communication system: Using multiple media to connect. *Information, Communication & Society*, 11(4):490–508, 2008. Cited on pages 7, 56, 60, and 69.

[17] Bianca Bosker. The Twitter typo that exposed Anthony Weiner, 2011. http://www.huffingtonpost.com/2011/06/07/anthony-weiner-twitter-dm_n_872590.html. Cited on page 14.

[18] Danah Boyd. Friends, friendsters, and myspace top 8: Writing community into being on social network sites. 2006. Cited on page 90.

[19] Michael Buhrmester, Tracy Kwang, and Samuel D. Gosling. Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011. Cited on pages 15, 16, and 29.

[20] Robert Capra, Emily Vardell, and Kathy Brennan. File Synchronization and Sharing: User Practices and Challenges. *ASIS&T Annual Meeting*, 2014. Cited on pages 7, 56, and 60.

[21] Emiliano De Cristofaro, Claudio Soriente, Gene Tsudik, and Andrew Williams. Hummingbird: privacy at the time of Twitter. Oakland '12. IEEE, 2012. Cited on page 53.

[22] Sauvik Das and Adam Kramer. Self-censorship on Facebook. In *ICWSM*, 2013. Cited on pages 2 and 8.

[23] Ralf De Wolf and Jo Pierson. Who's my audience again? Understanding audience management strategies for designing privacy management technologies. *Telematics and Informatics*, 31(4):607–616, November 2014. Cited on page 11.

[24] Nicole B. Ellison, Jessica Vitak, Rebecca Gray, and Cliff Lampe. Cultivating social resources on social network sites: Facebook relationship maintenance behaviors and their role in social capital processes. *Journal of Computer-Mediated Communication*, 19(4):855–870, 2014. Cited on page 90.

[25] Motahhare Eslami, Amirhossein Aleyasen, Roshanak Zilouchian Moghaddam, and Karrie Karahalios. Friend grouping algorithms for online social networks: Preference, bias, and implications. In *Social Informatics*, pages 34–49. Springer, 2014. Cited on page 11.

[26] Motahhare Eslami, Karrie Karahalios, Christian Sandvigt, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. First I "like" it, then I hide it: Folk theories of social feeds. 2016. Cited on page 82.

[27] Lujun Fang and Kristen LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide web*, pages 351–360. ACM, 2010. Cited on page 11.

[28] Shelly D. Farnham and Elizabeth F. Churchill. Faceted identity, faceted lives: Social and technical issues with being yourself online. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 359–368. ACM, 2011. Cited on pages 5, 6, 8, 67, and 75.

[29] Bruce Ferwerda, Markus Schedl, and Marko Tkalcic. To post or not to post: The effects of persuasive cues and group targeting mechanisms on posting behavior.

Cited on page 2.

[30] Benjamin M Gross and Elizabeth F. Churchill. Addressing constraints: Multiple usernames, task spillage and notions of identity. In *CHI'07 Extended Abstracts on Human Factors in Computing Systems*, pages 2393–2398. ACM, 2007. Cited on pages 8 and 75.

[31] Keith N. Hampton, Lauren Sessions Goulet, Lee Rainie, and Kristen Purcell. Social networking sites and our lives. *Pew Internet*, June 2011. Cited on page 50.

[32] M. Hart, C. Castille, R. Johnson, and A. Stent. Usable privacy controls for blogs. In *International Conference on Computational Science and Engineering, 2009. CSE '09*, volume 4, pages 401–408, August 2009. Cited on page 10.

[33] Bernie Hogan. The presentation of self in the age of social media: Distinguishing performances and exhibitions online. *Bulletin of Science, Technology & Society*, page 0270467610385893, 2010. Cited on pages 7 and 90.

[34] Yuli Patrick Hsieh. Online social networking skills: The social affordances approach to digital inequality. *First Monday*, 17(4), 2012. Cited on pages 7, 56, and 60.

[35] Ellen Isaacs, Alan Walendowski, Steve Whittaker, Diane J. Schiano, and Candace Kamm. The character, functions, and styles of instant messaging in the workplace. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, pages 11–20. ACM, 2002. Cited on pages 7, 56, and 60.

[36] M. Jakobsson. Experimenting on Mechanical Turk: 5 how tos. http://blogs.parc.com/blog/2009/07/experimenting-on-mechanical-turk-5-how-tos/, July 2009. Cited on page 29.

[37] Adam N. Joinson. Looking at, looking up or keeping up with people?: Motives and use of Facebook. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 1027–1036. ACM, 2008. Cited on pages 6 and 90.

[38] Sanjay Kairam, Mike Brzozowski, David Huffaker, and Ed Chi. Talking in circles: Selective sharing in Google+. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074. ACM, 2012. Cited on pages 6, 8, 10, 36, 40, 48, 56, 60, and 91.

[39] Ruogu Kang, Stephanie Brown, and Sara Kiesler. Why Do People Seek Anonymity on the Internet? Informing Policy and Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2013. Cited on pages 12, 82, 90, and 98.

[40] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. "My data just goes everywhere:" User mental models of the internet and implications for privacy

and security. In *Eleventh Symposium On Usable Privacy and Security*, pages 39–52, 2015. Cited on pages 8 and 76.

[41] Pamela Karr-Wisniewski, David Wilson, and Heather Richter-Lipford. A new social order: Mechanisms for social network site boundary regulation. In *Americas Conference on Information Systems. Detroit, MI*, 2011. Cited on pages 8 and 51.

[42] Patrick Gage Kelley, Robin Brewer, Yael Mayer, Lorrie Faith Cranor, and Norman Sadeh. An investigation into Facebook friend grouping. In *Human-Computer Interaction, INTERACT, 2011*, pages 216–233. Springer, 2011. Cited on pages 11, 48, 51, and 90.

[43] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. Tag, you can see it!: Using tags for access control in photo sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 377–386. ACM, 2012. Cited on pages 1 and 9.

[44] M. L Knapp, L. Stafford, and J. A Daly. Regrettable messages: Things people wish they hadn't said. *Journal of communication*, 36(4):40–58, 1986. Cited on pages 9, 14, 16, and 19.

[45] Haewoon Kwak, Hyunwoo Chun, and Sue Moon. Fragile online relationship: A first look at unfollow dynamics in Twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1091–1100. ACM, 2011. Cited on pages 6, 90, 95, and 97.

[46] A. Lampinen, S. Tamminen, and A. Oulasvirta. All my people right here, right now: Management of group co-presence on a social networking site. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work*, pages 281–290, 2009. Cited on pages 2, 6, 7, 8, 36, 48, and 52.

[47] Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3217–3226. ACM, 2011. Cited on pages 6, 7, 8, 32, 36, 56, and 60.

[48] Eric Lieberman and Robert C. Miller. Facemail: Showing faces of recipients to prevent misdirected email. In *Symposium on Usable Privacy and Security (SOUPS)*, pages 122–131. ACM, 2007. Cited on page 31.

[49] Eden Litt. Knock, knock. Who's there? The imagined audience. *Journal of Broadcasting & Electronic Media*, 56(3):330–345, 2012. Cited on pages 6 and 105.

[50] Eden Litt and Eszter Hargittai. Smile, snap, and share? A nuanced approach to privacy and online photo-sharing. *Poetics*, 42:1–21, February 2014. Cited on pages 1 and 7.

[51] Eden Litt and Eszter Hargittai. "Just Cast the Net, and Hopefully the Right Fish Swim into It": Audience Management on Social Network Sites. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, pages 1488–1500, New York, NY, USA, 2016. ACM. Cited on pages 8, 90, 93, 94, and 99.

[52] Eden Litt, Erin Spottswood, Jeremy Birnholtz, Jeff T. Hancock, Madeline E. Smith, and Lindsay Reynolds. Awkward encounters of an other kind: Collective self-presentation and face threat on Facebook. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 449–460. ACM, 2014. Cited on pages 6, 90, and 99.

[53] H. Liu, H. Lieberman, and T. Selker. Automatic affective feedback in an email browser. In *MIT Media Lab Software Agents Group*, 2002. Cited on pages 32 and 110.

[54] Cameron Marlow, Mor Naaman, Danah Boyd, and Marc Davis. Ht06, tagging paper, taxonomy, Flickr, academic article, to read. In *Proceedings of the Seventeenth Conference on Hypertext and Hypermedia*, pages 31–40. ACM, 2006. Cited on page 81.

[55] Alice E. Marwick and others. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society*, 13(1):114–133, 2011. Cited on pages 6, 7, and 8.

[56] Michelle L. Mazurek, J. P. Arsenault, Joanna Bresee, Nitin Gupta, Iulia Ion, Christina Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, and others. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 645–654. ACM, 2010. Cited on page 1.

[57] Michelle L. Mazurek, Peter F. Klemperer, Richard Shay, Hassan Takabi, Lujo Bauer, and Lorrie Faith Cranor. Exploring reactive access control. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2085–2094. ACM, 2011. Cited on pages 1, 5, 7, 12, and 79.

[58] Michelle L. Mazurek, Yuan Liang, William Melicher, Manya Sleeper, Lujo Bauer, Gregory R. Ganger, Nitin Gupta, and Michael K. Reiter. Toward strong, usable access control for shared distributed data. In *FAST*, pages 89–103, 2014. Cited on page 10.

[59] Caitlin McLaughlin and Jessica Vitak. Norm evolution and violation on Facebook.

New Media Society, 14(2):299–315, March 2012. Cited on pages 6, 90, 95, and 97.

[60] M.L. McLaughun, M.J. Cody, and H. O'Hair. The management of failure events: Some contextual determinants of accounting behavior. *Human Communication Research*, 9(3):208–224, 1983. Cited on page 9.

[61] Janet R. Meyer. Regretted messages: Cognitive antecedents and post hoc reflection. *Journal of Language and Social Psychology*, 2011. Cited on pages 9, 14, 15, 16, 22, 26, 29, and 30.

[62] J.R. Meyer and K. Rothenberg. Repairing regretted messages: Effects of emotional state, relationship type, and seriousness of offense. *Communication Research Reports*, 21(4):348–356, 2004. Cited on pages 9, 14, 15, 16, 18, 29, and 30.

[63] Mainack Mondal, Yabing Liu, Bimal Viswanath, Krishna P. Gummadi, and Alan Mislove. Understanding and specifying social access control lists. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014. Cited on page 11.

[64] Sean A. Munson and Paul Resnick. Presenting Diverse Political Opinions: How and How Much. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1457–1466, New York, NY, USA, 2010. ACM. Cited on page 100.

[65] Mor Naaman, Jeffrey Boase, and Chih-Hui Lai. Is it really about me?: Message content in social awareness streams. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*, pages 189–192. ACM, 2010. Cited on page 40.

[66] Jarno Ojala and Sanna Malinen. Photo sharing in small groups: Identifying design drivers for desired user experiences. In *Proceeding of the 16th International Academic MindTrek Conference*, MindTrek '12, pages 69–76, New York, NY, USA, 2012. ACM. Cited on page 5.

[67] Sai Teja Peddinti, Keith W. Ross, and Justin Cappos. "On the Internet, Nobody Knows You'Re a Dog": A Twitter Case Study of Anonymity in Social Networks. In *Proceedings of the Second ACM Conference on Online Social Networks*, COSN '14, pages 83–94, New York, NY, USA, 2014. ACM. Cited on pages 12, 90, and 98.

[68] Jon Perlow. New in labs: Stop sending mail you later regret, 2008. Official Gmail Blog. http://gmailblog.blogspot.com/2008/10/new-in-labs-stop-sending-mail-you-later.html. Cited on page 31.

[69] Andrew Perrin. Social media usage: 2005-2015. *Pew Internet & American Life Project*, 2015. Cited on page 56.

[70] Emilee Rader. Just email it to me!: Why things get lost in shared file repositories. In

GROUP'07 Doctoral Consortium papers, page 9. ACM, 2007. Cited on pages 7 and 60.

[71] Emilee Rader. The effect of audience design on labeling, organizing, and finding shared files. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 777–786. ACM, 2010. Cited on page 68.

[72] Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. Who are the crowdworkers?: Shifting demographics in Mechanical Turk. In *Ext. Abstracts CHI 2010*, pages 2863–2872. ACM, 2010. Cited on page 29.

[73] Lauren E. Scissors and Darren Gergle. Back and forth, back and forth: Channel switching in romantic couple conflict. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pages 237–248. ACM, 2013. Cited on pages 56 and 60.

[74] Bryan Semaan, Heather Faucett, Scott Robertson, Misa Maruyama, and Sara Douglas. Navigating Imagined Audiences: Motivations for Participating in the Online Public Sphere. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pages 1158–1169, New York, NY, USA, 2015. ACM. Cited on pages 6, 90, and 100.

[75] Christopher Sibona and Steven Walczak. Unfriending on Facebook: Friend request and online/offline behavior analysis. In *System Sciences (HICSS), 2011 44th Hawaii International Conference on*, pages 1–10. IEEE, 2011. Cited on pages 6, 90, 95, and 97.

[76] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pages 793–802. ACM, 2013. Cited on pages 1, 2, 8, 35, 76, 80, 82, 90, 91, and 95.

[77] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. I read my Twitter the next morning and was astonished: A conversational perspective on Twitter regrets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3277–3286. ACM, 2013. Cited on pages 1, 2, 7, 9, and 13.

[78] Manya Sleeper, William Melicher, Hana Habib, Lujo Bauer, Lorrie Faith Cranor, and Michelle L Mazurek. Sharing personal content online: Exploring channel choice and multi-channel behaviors. 2016. Cited on pages 1, 6, 55, 90, and 104.

[79] Aaron Smith. Americans and their gadgets. *Pew Internet Research*, 2010. Cited on page 56.

[80] Aaron Smith. US smartphone use in 2015. *Pew Research Center*, 2015. Cited on page 56.

[81] Katherine Strater and Heather Richter Lipford. Strategies and struggles with privacy in an online social networking community. BCS-HCI '08, pages 111–119, Swinton, UK, UK, 2008. British Computer Society. Cited on page 51.

[82] H. Colleen Stuart, Laura Dabbish, Sara Kiesler, Peter Kinnaird, and Ruogu Kang. Social transparency in networked information exchange: A theoretical framework. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pages 451–460. ACM, 2012. Cited on pages 12, 82, and 90.

[83] Frederic Stutzman and Woodrow Hartzog. Boundary regulation in social media. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pages 769–778. ACM, 2012. Cited on pages 7, 52, and 53.

[84] Frederic Stutzman and Woodrow Hartzog. Boundary regulation in social media. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, pages 769–778. ACM, 2012. Cited on pages 8, 67, and 75.

[85] Norman Makoto Su and Gloria Mark. Communication chains and multitasking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 83–92. ACM, 2008. Cited on pages 56 and 60.

[86] Eran Toch, Justin Cranshaw, Paul Hankes Drielsma, Janice Y Tsai, Patrick Gage Kelley, James Springfield, Lorrie Cranor, Jason Hong, and Norman Sadeh. Empirical models of privacy in location sharing. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pages 129–138. ACM, 2010. Cited on page 79.

[87] Pauline EW Van den Berg, Theo A. Arentze, and Harry JP Timmermans. New ICTs and social interaction: Modelling communication frequency and communication mode choice. *New Media & Society*, 14(6):987–1003, 2012. Cited on pages 7, 56, and 60.

[88] Anthony J Viera, Joanne M Garrett, et al. Understanding interobserver agreement: The Kappa statistic. *Fam Med*, 37(5):360–363, 2005. Cited on page 60.

[89] Jessica Vitak. The impact of context collapse and privacy on social network site disclosures. *Journal of Broadcasting & Electronic Media*, 56(4):451–470, 2012. Cited on pages 10, 56, and 60.

[90] Jessica Vitak, Stacy Blasiola, Sameer Patil, and Eden Litt. Balancing audience and privacy tensions on social network sites: Strategies of highly engaged users. *International Journal of Communication*, 9:20, 2015. Cited on pages 8, 90, 93, and 94.

[91] Jessica Vitak and Nicole B. Ellison. "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*, page 1461444812451566, 2012. Cited on pages 6 and 90.

[92] Jessica Vitak and Jinyoung Kim. "You can't block people offline": Examining how Facebook's affordances shape the disclosure process. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '14, pages 461–474, New York, NY, USA, 2014. ACM. Cited on pages 1, 2, 6, 10, 56, and 60.

[93] Amy Voida, Judith S Olson, and Gary M Olson. Turbulence in the clouds: Challenges of cloud-based information work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2273–2282. ACM, 2013. Cited on pages 8, 56, and 60.

[94] Stephen Voida, W. Keith Edwards, Mark W. Newman, Rebecca E. Grinter, and Nicolas Ducheneaut. Share and share alike: Exploring the user interface affordances of file sharing. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 221–230. ACM, 2006. Cited on pages 1, 2, 56, 57, 60, 71, and 72.

[95] Yang Wang, Pedro Giovanni Leon, Kevin Scott, Xiaoxuan Chen, Alessandro Acquisti, and Lorrie Faith Cranor. Privacy nudges for social media: An exploratory Facebook study. In *Proceedings of the 22nd International Conference on World Wide Web Companion*, pages 763–770. International World Wide Web Conferences Steering Committee, 2013. Cited on page 110.

[96] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. I regretted the minute I pressed share: A qualitative study of regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 10. ACM, 2011. Cited on pages 1, 2, 7, 9, 18, 35, and 36.

[97] Jason Watson, Andrew Besmer, and Heather Richter Lipford. +your circles: Sharing behavior on Google+. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 12:1–12:9, New York, NY, USA, 2012. ACM. Cited on pages 8 and 10.

[98] Tara Whalen, Elaine Toms, and James Blustein. File sharing and group information management. *Personal Information Management: PIM*, 2008, 2008. Cited on pages 7, 56, and 60.

[99] Jason Wiese, Patrick Gage Kelley, Lorrie Faith Cranor, Laura Dabbish, Jason I. Hong, and John Zimmerman. Are you close with me? Are you nearby?: Investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th International*

Conference on Ubiquitous Computing, UbiComp '11, pages 197–206, New York, NY, USA, 2011. ACM. Cited on pages 6, 11, 48, and 79.

[100] Pamela Wisniewski, Heather Lipford, and David Wilson. Fighting for my space: Coping mechanisms for SNS boundary regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 609–618. ACM, 2012. Cited on pages 7, 36, 56, and 60.

[101] Ching-man Au Yeung, Lalana Kagal, Nicholas Gibbins, and Nigel Shadbolt. Providing access control to online photo albums based on tags and linked data. In *AAAI Spring Symposium: Social Semantic Web: Where Web 2.0 Meets Web 3.0*, pages 9–14, 2009. Cited on page 10.

[102] X. Zhang, Q. Gao, C. S. G. Khoo, and A. Wu. Categories of friends on social networking. 2013. Cited on page 11.