Ownership, Privacy, and Control in the Wake of Cambridge Analytica: the Relationship between Attitudes and Awareness

Frank M. Shipman
for Submission
City, Country
e-mail address

Catherine C. Marshall
for Submission
City, Country
e-mail address

ABSTRACT

Has widespread news of abuse changed the public's perceptions of how user-contributed content from social networking sites like Facebook and LinkedIn can be used? We collected two datasets that reflect participants' attitudes about content ownership, privacy, and control, one in April 2018, while Cambridge Analytica was still in the news, and another in February 2019, after the event had faded from the headlines, and aggregated the data according to participants' awareness of the story, contrasting the attitudes of those who reported the greatest awareness with those who reported the least. Participants with the greatest awareness of the news story's details have more polarized attitudes about reuse, especially the reuse of content as data. They express a heightened desire for data mobility, greater concern about networked privacy rights, increased skepticism of algorithmically targeted advertising and news, and more willingness for social media platforms to demand corrections of inaccurate or deceptive content.

Author Keywords

Social media attitudes, ownership, privacy, data use, data monetization, Facebook, LinkedIn, Cambridge Analytica.

CSS Concepts

• **Human-centered computing~Collaborative and social computing;** *Social media; Empirical studies in collaborative and social computing*

INTRODUCTION

Social networking applications like Facebook, Instagram, LinkedIn, and Snapchat play a central role in today's Internet. Many people rely on these applications to communicate—to keep in touch with friends and family—but also as platforms and infrastructure that fulfill a larger range of business, civic, entertainment, and news dissemination functions. In 2018, 68% of adults in the US reported Facebook use [42]. Even if younger users now prefer other platforms, they may continue to communicate with extended family and maintain lingering weak ties using Facebook [4].

How do Facebook users feel about reuse of their data? In previous studies, participants have demonstrated an ambivalence both toward the personal content they have accumulated online [25] and toward how others—social contacts, strangers, and corporations—may reuse it [31]. The terms of use for most online services present individuals with a perplexing mixed message about whether they own what they put online, whether they should view their personal data as a commodity, and whether once they've shared content via social media, they have lost control of it [17].

Users' responses to the complexities of data ownership have been mixed. According to the technology press, people are now sharing less information over social media, and what they are sharing has become less personal; data privacy has become a nearly ubiquitous concern [3]. But users' reported desire for online privacy may not translate into action: it's difficult to put privacy aspirations into practice, particularly if one must sacrifice convenience or understand security arcana to do so [37]. Although there have been numerous reports of online privacy breaches and data misuse over the life of popular social networking platforms like Facebook, user discontent seems to be on the upswing (e.g. [19]). In this paper, we examine evolving social norms for ownership of social network-based content through the lens of a controversial news story involving UK consulting firm Cambridge Analytica, with the idea that attention to the news may correlate with diverging norms.

Cambridge Analytica acquired Facebook data captured in 2014 by a researcher, who used a voluntary quiz app to collect several hundred thousand Facebook users' personal data. Quiz-takers consented to academic use of this data. However, their participation also gave the app access to information about their Facebook friends, which exposed up to 87 million more people to the data harvesting efforts. Over 70 million of those affected lived in the US [19]. This information enabled Cambridge Analytica to create psychographic profiles of individuals based on relevant details such as location, likes, birthdate, and other profile elements. The profiles were subsequently used by political campaign organizations to target ads.

CHI 2020, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6708-0/20/04...$15.00.
DOI: https://doi.org/10.1145/3313831.3376662

This paper examines whether knowledge of the CA/FB event correlates with attitudes about content reuse. We break the Cambridge Analytica story (referred to as ***CA/FB***) into two constituent parts: (1) data acquisition (***data*** for short) and (2) the use of the harvested data to profile users and target ads (***use*** for short). Although the psychographic profiles were used in several countries, for clarity, we limited our study to effects experienced by US-based participants.

Naturally, not everyone paid the same degree of attention to CA/FB, and not everyone felt the data use was scandalous. Instead, CA/FB may serve as a litmus test for how people perceive the tacit and actual agreements they enter into when they contribute content to social media platforms. Familiarity with CA/FB, either ignoring the news, or following it closely, may correspond to how people react to hypothetical ownership quandaries (e.g. individuals' rights to their account profiles or social networks). Specifically, we examined whether awareness of CA/FB correlates with differences in participants' attitudes about:

- Ownership and control of content in a social network setting;
- Monetizing personal social network content as data;
- Responsibility for content accuracy;
- How social networking platforms can use or provide access to user-contributed data.

We begin by discussing related work. Then we present the method we used to collect and analyze data, and describe the study's participants. The results follow. We conclude by discussing how the results address our four primary research questions, and giving a brief overview of future work.

RELATED WORK

Our work is related to broader discussions of social media ownership, digital identity, and online privacy as well as more specific examinations of user reactions to CA/FB and other episodes of gathering and using data from social networking platforms. We discuss each area of related work.

Social Media Ownership

Our earlier studies, run from 2010 to 2013, explored participants' attitudes toward ownership, control, and reuse of different types of user-contributed online content, and whether they varied by media type [31]. The studies used hypothetical statements about specific situations to identify norms for acceptable and unacceptable reuse of content and data. The work presented here extends our 2013 survey that focused on social networks (such as Facebook and LinkedIn) as a distinct content genre [30]. We found that participants' attitudes toward reuse were influenced by the boundaries introduced by a social network, a phenomenon dubbed networked privacy by Marwick and boyd [33]. Specifically, we found that participants were more conservative about what they could do with content contributed by someone they were socially connected with (even at a remove) than they were with content retrieved from the Internet writ large. Our study also examined the use of social network content as data by corporate entities and as part of the historical record (by memory institutions). The norms we identified can serve as a baseline for our current exploration of how ownership norms may diverge in line with attention to salient news.

Underlying these studies of social media ownership is the basic premise that ordinary behavioral norms can be more effective in governing individuals' actions than complex legal definitions and tests [13]. In the case of reusing online material, norms are especially important because users often misunderstand legal concepts like copyright [17], fair use [15], and the tension between them [34]. Tehranian [45] points out the absurdity in the application of current copyright law in everyday types of online reuse, and suggests that in this context, norms more appropriately express accepted behavior. Consider, e.g., the fluid appropriation of visual memes and their use as elements of online speech [23]. This is especially important in the case we are addressing: the reuse of user-contributed material, rather than conventionally published, DRM-protected content. Zhao et al. [48] identified varying content and participation norms on different social media services; e.g., a particular genre of photo might be considered more appropriate for Instagram than for Facebook. Norms have been explored on services as diverse as Facebook and LinkedIn [30] and Reddit [7].

Privacy

Throughout the evolution of social media, researchers have noted the tensions between participation and privacy. In particular, users' privacy intentions are often not in sync with their behavior [37]. Nissenbaum has clarified that privacy is not so much just about control over personal information; it is about controlling the context in which it is seen [36]. By centering their study on posts that users regretted, Wang et al. confirmed that audience, tone, and topic are all vital elements of determining contextual privacy boundaries [46]. Child et al. [8] studied how users' perceptions of their own social network behavior predict their privacy management behaviors and showed that individual understanding of concepts like the differences between public and private communication can have a significant impact on both. Stutzman and Kramer-Duffield [44] examined how these tensions played out in undergraduates' use of Facebook in its early days, and how the tensions manifested in emerging strategies for maintaining privacy. Kang et al. [24] examined how users' mental models of the technology affected their use of available security options. Their study showed that knowledge of technology was less indicative of user behavior than was users' sense of risk (e.g. belief that they have nothing to hide). Phelan et al. [39] describe how the actions of social media users reveal a gap between their high-level privacy concerns and risk assessments of individual actions. Duffy and Chan [11] explored how college-aged users adapted their use of social media, given an underlying belief that their online activities will be monitored by family, educators and future employers. They found that decisions of what to share with whom were made alongside the use of privacy settings and pseudonyms to reduce their perceived

exposure. In our work, we are sensitive to the distinction between users' expectations of a right to privacy, and their strategies for working around the absence of privacy guarantees.

Digital Identity

Contributing content to social networks is often deeply connected with managing and controlling one's digital identity. Social network profiles present a persona that is developed with certain goals [6]. When assumptions about the social network change, for example when new classes of users join the network, potentially incongruous goals may be added [21]. Users manage their identity by trying to limit the information about themselves available. Madden and Smith [28] found users changed privacy settings, deleted unwanted comments that others attached to their profiles, and deleted images. They also noted that this form of identity curation was more common among younger users.

Farnham and Churchill [14] observed that identity creation is not a unitary whole; people reveal different and possibly conflicting aspects of their identities in different online venues. Users put considerable effort into managing these multi-faceted online identities [47]. As users adopt different social networking services to achieve different goals [48], identity management becomes more difficult by the automatic and often unintentional bleed of information from one social network to others. It is little wonder that study participants express such vehement and negative responses to online aggregators who assemble unwanted identity patchworks [29]. Digital identity research motivates our exploration not only of users' reactions to everyday reuse by other people, but also their reactions to corporate reuse of content as data.

Reactions to Cambridge Analytica

Even before the Cambridge Analytica scandal occurred, one set of concerns expressed by Facebook leavers (about one-fifth surveyed) was that Facebook had a propensity to misuse their personal information [2]. Although to those surveyed by Baumer et al., this was just a suspicion, the roots of some of the reactions we see in our data predate CA/FB. Fiesler and Hallinan [16] examined the "in the wild" responses to two other user data sharing controversies (WhatsApp's sharing arrangement with Facebook and unroll.me's sale of data to Uber) that provide some insight into the how users understand FB/CA. Because their analysis reflects *in situ* reactions to the controversies, we assume our participants were exposed to some of the same themes and ideas (e.g. that trafficking in users' data underpins Facebook's business model), but were variably influenced by these ideas.

Analyses of the Cambridge Analytica affair have led to the recommendation of new corporate data sharing policies [22] and have motivated participatory design exercises for social media platforms that include greater user control over the data they produce [35]. Such interpretations are based on initial media and user reactions to the story. We explore whether the initial reactions resulted in substantive changes to attitudes and behaviors and whether changes have persisted as coverage of the event gradually dissipated.

In a related effort, Edwards [12] surveyed Facebook users about their use of Facebook and their understanding and use of its data privacy settings. The study also monitored Facebook use over four months to identify changes to behavior. Edwards found that the decrease in use was very small and potentially attributed this result to Facebook's aging user population: older users perceive less personal risk in their social media content and have greater investment in creating their social networks, thus making adoption of alternative personal communication methods more costly. Edwards' results are expanded by a survey by Pew Research Center in May/June of 2018 that found 54% of respondents had adjusted their Facebook privacy settings and 26% had deleted the Facebook app from their phone during the last year [38]. The Pew data confirms that younger users had a much higher incidence of both adjusting privacy settings (~64% for those 49 and younger vs ~40% for those 50 and older) and removing the phone app (with 44% of those 18-29 vs. ~20% for those 30 and over). Deleting the phone app and adjusting privacy settings may be consistent with the idea that people may keep their Facebook accounts, but may be compensating for adjusted participation norms.

METHOD

To collect the data we used in this analysis, we fielded a four-part questionnaire. The first three parts of the questionnaire were based on our earlier study to identify norms for reuse and control of social network data [30]. We added a fourth part to gather reactions to statements about the CA/FB story, which would have been fresh in participants' minds during the first data collection interval (April, 2018), and fading during the second data collection interval (February, 2019).

In addition to collecting demographic and background information, the baseline study presented participants with a series of six brief stories (scenarios), each followed by 2-5 hypothetical statements to explore their reactions to various aspects of ownership, control, and reuse of Facebook and LinkedIn data. Figures 4a, 5a, and 6a summarize the three scenarios salient to our interests in this study (collecting social network data, correcting inaccuracies, and monetizing personal data) and their associated hypotheticals. Square brackets in the figures denote summaries. The three scenarios used in this paper were originally documented in [30]. They were developed using the combined results of fieldwork, personal experiences, and news stories; we include them primarily for intelligibility of the results.

Hypotheticals are a conventional technique for exploring social norms in law and legal education. For example, when cases are argued in front of the US Supreme Court, attorneys and justices often propose hypotheticals to advance a slippery-slope argument or to test its boundaries [27]. In essence, hypotheticals vary elements of a case's fact patterns [41]. In our studies, participants judged each hypothetical on

a 7-point Likert scale (from *1-disagree strongly* to *7-agree strongly*).

The final five questions elicited participants' attitudes to CA/FB (also assessed on a 7-point Likert scale), and asked them to self-report on their familiarity with the story:

- *My friends ignored the Facebook/Cambridge Analytica story.*
- *It is okay for Facebook to analyze user data to target news*
- *It is okay for Facebook to analyze user data to target ads*
- *It is okay for Facebook to give apps access to user data*
- *How much have you heard about Cambridge Analytica's access to and use of Facebook data?*

For the final question, participants were offered five possible responses to reflect various levels of awareness and interest: they hadn't heard about it (referred to as ***nothing***); they had heard about it, but didn't follow the details (***heard***); they knew personal data had been gathered (***data***); they knew how the collected data had been used (***use***); or they had followed the story closely and were aware of the data involved and how it had been used (***both***).

The questionnaire was fielded as an Amazon Mechanical Turk (AMT) HIT in April of 2018 and in February of 2019. Participants were required to be US-based, English-speaking Facebook users; they were also required to have a 95% acceptance rate on previous HITs to reduce spam responses, in accordance with reported best practices and our prior experiences with AMT [30]. Participants were paid if they completed the survey, regardless of whether they passed the data quality criteria described below. Research by Bentley, Daskalova, and White [5] shows AMT to be comparable to other methods of soliciting survey participation.

Data quality. We used two criteria to discard surveys. The first checked responses to two reading comprehension questions, ensured that completion speed was over a minimum threshold, and identified missing answers to individual questions. Each anomaly was assigned a point. If participants accrued two or more points, their surveys were discarded. Additionally, in 2019, four surveys that contained suspicious response patterns to the Likert scale questions (e.g. all 7s) were also discarded (this phenomenon was not observed in 2018). The authors independently checked the dataset to ensure data quality rules were applied uniformly.

We recruited 500 participants for each of the 2018 and 2019 studies. Application of the data quality process described above eliminated data from 26 participants in 2018 and 62 in 2019, leaving us with 474 participants in 2018 and 438 in 2019. The remaining surveys appeared to have been completed in good faith. As Ahler, Roush, and Sood noted in 2019 [1], recent study data gathered on AMT requires additional quality measures; we believe our data quality tests compensated for this apparent uptick in bots and scammers.

Data analysis. To test correlations with awareness of CA/FB, the 2018 and 2019 data was merged and the respondents with the greatest and least self-reported attention to the event were analyzed using Mann Whitney U tests. The Holm-Bonferroni correction was applied to the 12 hypotheticals exploring privacy-centric activities (found in Figures 4, 6, and 7 below) and a separate Holm-Bonferroni correction was applied to the 3 hypotheticals exploring data veracity (in Figure 5). The correction adjusts the p values individually, limiting the overall p value for the family to .05. The p values presented in the paper are the Holm-Bonferroni adjusted values. The responses were then charted using simple visualizations to spot trends among each scenario's family of hypotheticals. We have included these visualizations.

Open-ended responses were coded using conventional qualitative techniques to reveal themes and patterns [43]. The codings were documented with associated inclusion heuristics (what determined membership in a given category) and a growing set of examples to ensure consistent interpretation of participant responses across study years. In this paper, we use the qualitative responses primarily as triangulation with responses to hypotheticals covering the same topic, and to ensure we were getting responses consistent with those reported in the related work.

PARTICIPANTS

Table 1 summarizes demographic characteristics of the participants. Survey-takers on AMT appear to represent workers in the so-called gig economy who are Internet-savvy, fairly well-educated (more than half report having earned a bachelor's degree), and span a range of ages. They appear demographically comparable to the workers described by Difallah, Filatova, and Ipeirotis [10].

Participants' responses to an open-ended question about what they do online, in tandem with their selections from a nine-item list (email, IM/chat, photo sharing, video conferencing, video sharing, shopping, massively multiplayer online games, social media, and Twitter), reveal that the types of online activities they report have not changed appreciably since those collected in 2013 [30], although as smart phones have become more ubiquitous, "being on the Internet" is more a matter of course; it is easy to check social networking apps throughout the day.

Many participants have experience with multiple social networking platforms. Screening requirements specified that

study group	% female	% born before 1980	% born after 1990	% current students	% BS, BA, or higher
2018	59%	31%	25%	14%	52%
2019	55%	34%	25%	13%	58%

Table 1. Participant characteristics during the study years

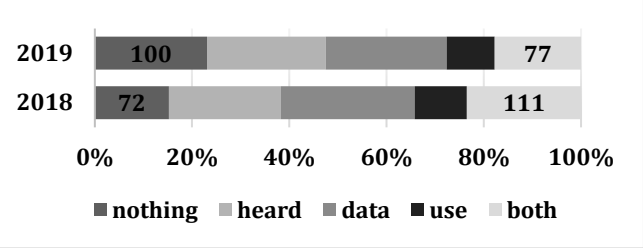

Figure 1. Two subpopulations based on CA/FB awareness.

participants must be Facebook users, so it is no surprise that 99% of participants in both survey years said they still use Facebook. By 2019, the second most popular platform was Instagram; 60% of the participants (264/438) had active Instagram accounts; 45% had LinkedIn accounts; and almost a third (139/438) had Snapchat accounts. Snapchat use is most prevalent in the cohort born in or after 1990: 50% of the participants in this youngest group use Snapchat, while only 34% of those born in the 1980s, and 16% of those born in 1970s or earlier use the app.

Participants listed many other social networking platforms they use in addition to the six choices offered, including YouTube, Pinterest, Reddit, Whatsapp, Slack, NextDoor, and Tumblr; they also mentioned a number of specialty platforms, including Twitch (the app for watching gamers), VKontakte (a Russian Facebook), Gab (a Twitter-like app for the alt-Right), Fetlife (a fetish-based meetup app), and myLot (a "social networking for pay" site). This apparent diversity serves to illustrate the role Facebook is playing; e.g. while one's extended family may not be on Snapchat, they are more likely to have Facebook accounts.

RESULTS

The distribution of participants' responses to the survey's final question about their awareness of CA/FB is shown in Figure 1. For the purpose of this analysis, we are interested in the responses at the ends of the spectrum: those who claimed to be unfamiliar with CA/FB (i.e. members of the

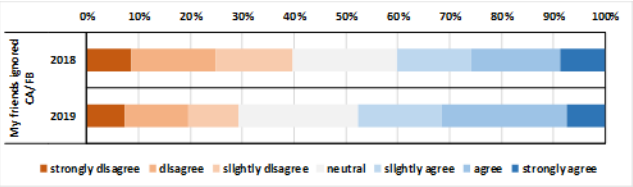

Figure 2. Friends' unfamiliarity with CA/FB story increases from 2018 to 2019

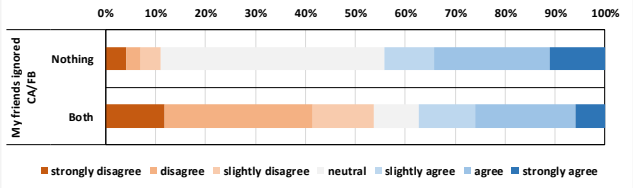

Figure 3. Friends' unfamiliarity with CA/FB story greater for those unaware of the story themselves (*nothing* group)

nothing subpopulation), and those who claimed to have followed the story closely (i.e. ***both***). We aggregated these subpopulations across the 2018 and 2019 surveys, which resulted in a pool of 172 participants in the ***nothing*** group, and 188 in the ***both*** group. The remaining 552 participants are omitted from this analysis.

A year proved to be long enough that participants' reported awareness of CA/FB decreased substantially. Figure 1 shows how the relative size of the ***both*** and ***nothing*** subpopulations essentially flipped between 2018 and 2019.

If fewer participants are aware of the story, we should also see a drop in their assessment of their friends' awareness. Figure 2 shows that this corresponding drop took place. In 2018, participants were split evenly (40% agreed to some extent that their friends had ignored CA/FB, and 40% disagreed to some extent). By 2019, 48% agreed that their friends had ignored CA/FB and only 29% disagreed. Figure 3 shows the difference in this assessment between the ***both*** and ***nothing*** groups: Participants familiar with CA/FB (***both***) report that their friends are familiar with the story too. Since the test statement is framed as a negative (*my friends ignored the story*), those reporting no knowledge (***nothing***) generally either agreed with the statement, or reported a neutral rating (which probably means that they didn't know), while those reporting considerable knowledge (***both***) disagreed with the statement to a much greater extent; they also seemed more aware of whether their friends had ignored the story or not (MW: $p < 0.0001$, $r = .28$).

Data Collection. Participants who reported paying close attention to CA/FB—members of the ***both*** group—are more apt to be aware that social ties figured into the data gathering aspect of the story. Hence they might be more sensitive to the networked privacy aspect of saving profiles, network elements, and reachable content from Facebook.

The scenario in Figure 4a posits a long-term user of Facebook, Susie, who began using the platform in college and has accumulated a great deal of personal content and social connections on the platform. In this scenario, Susie is planning to delete her account as she looks for a job, and the hypotheticals' fact pattern varies what she should be able to take with her as a lightweight litmus test for ownership and competing interests in privacy and control of data belonging to one's social ties.

The hypotheticals distinguish between content Susie has created and contributed such as photos (H3) and content she has created in the context of her account on the platform such as her profile information and social connections (H1). Competing interests come into play when participants assess the other three hypotheticals (H2, H4, and H5). H2 tests a person's ability to download her friends' friends' contact information, recognizing that network connections are one reason people maintain accounts they would otherwise delete. H4 and H5 tests rights to save friends' content, given different privacy conditions; H4 tests saving any content a

Scenario 1 (Saving): [*Susie has been using Facebook since 2005, when she was a college freshman. She has accumulated hundreds of friends, photos, 'likes', posts, and mail. She would like to delete her account prior to seeking a job; she would also like to save her 'friend' connections and the content she has posted*.] Before she deletes her account:

H1: Susie should be able to save her Facebook profile, including who's connected to whom in her network of friends. **(label: User can save profile and friend list)**

H2: Susie should be able to save contact information for her friends' friends. **(label: User can save friends' friends)**

H3: Susie should be able to save the content associated with her account, including the pictures in her galleries and her message box (including email to/from her). **(label: User can save account content)**

H4: Susie should be able to save anything that she can ordinarily reach on Facebook, including pictures from her friends' and friends' friends' galleries. **(label: User can save all encounterable)**

H5: Susie should only be able to save photos from her friends' Facebook accounts if the photos are public. **(label: User can save only public photos)**

Figure 4a. Scenario (summarized) and hypotheticals to test norms for saving social network content and profiles

person can reach through her Facebook friends network, while H5 tests the more restricted situation of saving only that content which is explicitly public (making the download more analogous to saving content on the open Web, which participants have viewed—and continue to view—as an uncontroversial "no harm done" right [32]).

Generally, members of the ***both*** group—people who had paid attention to CA/FB—reacted more negatively than the ***nothing*** group to the two hypotheticals H2 and H4 that involved trading off between networked privacy rights (respecting that friends' friends might find downloading their data to be an overreach) and the ability to save content that might no longer be available to them. For H2, saving contact information at 2-step distance (*User save friends' friends* in Fig. 4b), the difference in responses between the

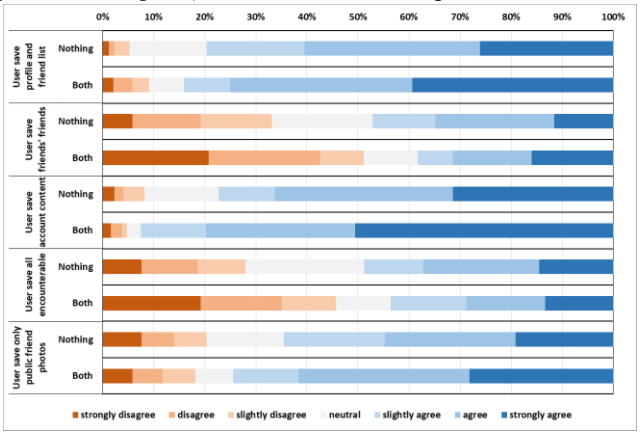

Figure 4b. Discrepancies in views between CA/FB-aware participants (*both*) and those who ignored the story (*nothing*) using saving as a lightweight litmus test for ownership.

two groups is significant (MW: p<.03, r=.14), although for some participants, awareness of CA/FB also seemed to provoke a counter-reaction (some members of the ***both*** group were emphatically positive that Susie had a right to act in her own interest and save this valuable information that would no longer be accessible to her if she deleted her account). For H4, the hypothetical test that is closest to the CA/FB data collection situation (*User save all encounterable* in Fig. 4b), the difference between the two groups is also significant (MW: p=.04, r=.13). H4 elicits more negative reactions from members of the ***both*** group, who seem to assess networked privacy rights as outweighing norms for saving user-contributed content.

Hypothetical H5 (labeled *User save only public friend photos* in Figure 4b) tests the importance of abiding by another user's desires (the hypothetical proposes that the user should **only** be able to save photos shared publicly). The ***both*** group agreed more strongly that a user should abide by this limitation (MW: p<.04, r=.13).

By contrast, the two hypotheticals most connected with personal ownership, H1 (*User save profile and friend list* in Figure 4b) and H3 (*User save account content*), demonstrated that while both groups had generally positive reactions to the premise that users should be able to download their own information, the ***both*** group responded more adamantly than the ***nothing*** group (H1 MW: p<.04, r=.14; H3 MW: p<.001, r=.21). As we might expect, those in the ***both*** group found saving platform-specific information like profiles and lists of friends more controversial than uploaded content like photos.

Note that for all of these hypotheticals, the number of participants expressing neutral feelings about ownership or the tradeoffs between ownership and networked privacy shrinks for the ***both*** group; attitudes at either end of the spectrum (strong agreement and strong disagreement) seem to be more common. The differences between ***both*** and ***nothing*** for this family of hypotheticals are significant.

Data Veracity. Networked privacy rights and content ownership norms weren't the only things at stake. Although the veracity of information online and off- has been a long-term concern, of late, there are fewer universally trusted sources. According to a Pew Internet survey conducted in the summer of 2018, 43% of Americans use Facebook as their main conduit for news, which is often supplied by undifferentiated sources [32]. Moreover, many people rely on their own judgment to identify untrustworthy information, regardless of whether they are capable of doing so [18].

The second family of hypotheticals we examined for potential correlation with the two subpopulations' attitudes concerned responsibility for content veracity. The scenario presented participants with the three hypotheticals (H6, H7, and H8) shown in Figure 5a. The hypotheticals offer three alternative (but not mutually exclusive) solutions to the inequity posed by the scenario: H6 places the authority for

correction in the hands of the whistleblower, and H7, in the hands of his manager. H8 offers a more draconian solution: the platform has the right (or possibly, the responsibility) to moderate disinformation or inaccuracies by shutting down the offending account.

Greater awareness of CA/FB (i.e. membership in ***both***) is correlated with a stronger sense that the platform (in this case, LinkedIn) can force the user to correct his resume information (H8 MW: $p<.01$, $r=.17$). Figure 5b illustrates these differences.

At the same time, those in the ***both*** group expressed less support for either another user (H6 $p<.06$, $r=.10$) or a manager (H7 $p<.06$, $r=.09$) being able to force such a change. Comparing responses to the three hypotheticals, participants seem to agree that the platform is responsible for preventing the spread of inaccurate information. As in the other family of hypotheticals, participants who didn't follow the CA/FB story are more likely to respond neutrally; H6 and H7 tended to polarize the ***both*** group, with larger portions of the group expressing strong views at either end of the spectrum.

Data Monetization. Data monetization tests competing interests in content ownership, control of identity, and personal privacy rights. From the responses to the first family of hypotheticals (H1-H5), we might expect that there will be a split between the two groups' attitudes about data monetization, manifested by a stronger assertion of personal ownership and increased polarization of attitudes expressed by those in the ***both*** group. The scenario and hypotheticals H9-H12 are shown in Figure 6a.

The most negative of the monetization hypotheticals—Facebook's right to sell user data to Amazon (In Figure 6b, *FB can sell user info to Amazon*)—does not elicit a significantly different reaction from the two groups. This hypothetical also elicited the most negative response of all of the hypotheticals we tested in our 2013 study; at that point in time, over 80% of the responses were negative. Negative responses still stand just below 80% for both groups. Note,

Scenario 2 (Data Veracity): [*Greg used his extensive LinkedIn profile the last time he was looking for a job. Eventually he went to work for a company, Xiblix. Greg works with a peer, Homer, at Xiblix.*] As Greg looks around at other peoples' LinkedIn profiles at Xiblix, he discovers that his colleague Homer has lied about his education and his job title. Greg feels that these fraudulent details will make Homer look more qualified for jobs they will compete for in the future.

H6: Greg should be able to force Homer to change his erroneous profile. **(label: User can force edit to fix error)**

H7: Homer's manager should be able to force Homer to change his erroneous profile. **(label: Boss can force edit to fix error)**

H8: LinkedIn should be able to shut down Homer's account for misrepresenting his credentials. **(label: LinkedIn can force edit to fix error)**

Figure 5a. Testing correction responsibility

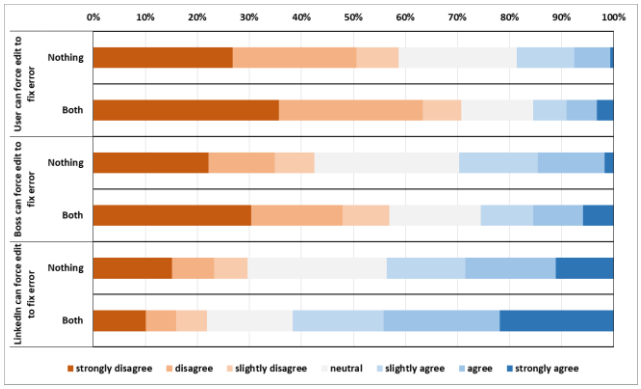

Figure 5b. Discrepancy in data veracity views between CA/FB-aware participants (*both*) and those who ignored CA/FB (*nothing*).

Scenario 3 (Monetization): Facebook's business model begins to falter. The company needs to develop additional ways to make money. Unlike Susie, her friend Greg has not deleted his Facebook account; he has accumulated a lot of Facebook data over the years.

H9: Facebook should be able to sell the information in Greg's user profile to Amazon so Amazon can create a better profile of Greg's interests. **(label: FB can sell user info to Amazon)**

H10: Facebook should need Greg's permission to sell his profile. **(label: FB needs permission to sell info)**

H11: Greg should be able to sell the information in his Facebook profile to Amazon to get a cash rebate on his Amazon purchases. **(label: User can sell own FB info to Amazon)**

H12: Facebook should be able to analyze the content of Greg's Facebook-internal communication so it can create a better profile of Greg's interests for its own use in selling targeted advertising. **(label: FB can use internal messaging to improve ads)**

Figure 6a. Testing ownership and monetization

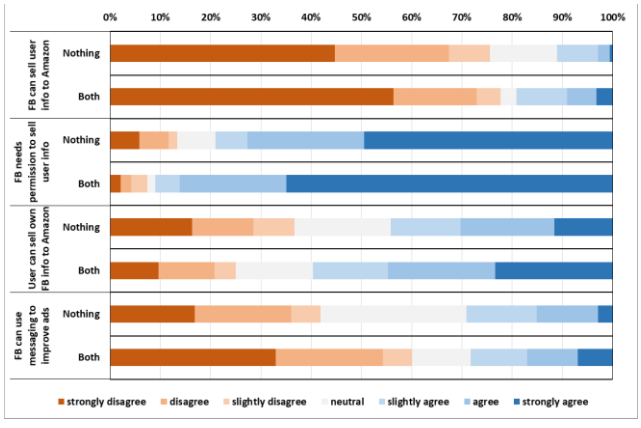

Figure 6b. Discrepancy in data monetization views between CA/FB-aware participants (*both*) and those who ignored the scandal (*nothing*).

too, the relatively small number of neutral responses. It is possible that participants' interest in privacy and identity management overwhelms any apparent difference between attitudes of the ***both*** and ***nothing*** groups.

The other three monetization hypotheticals were viewed differently by the two groups. Facebook's right to analyze messages to generate targeted advertising (H12) is viewed more skeptically by those in the ***both*** group than those in the ***nothing*** group (MW: p<.04, r=.13). Also, again we see polarized responses from the ***both*** group. In Figure 6b, this comparison is labeled *FB can use messaging to improve ads*.

When asked whether a user should be able to give Amazon access to his Facebook data in exchange for reduced prices on Amazon's goods (H11), which amounts to a test of selling personal data, the ***both*** group was significantly more positive than the ***nothing*** group (MW: p<.01, r=.16). In Figure 6b, this comparison is labeled *User can sell own FB info to Amazon*.

Finally, the most popular panacea in situations of contested rights, permission, tests more positively with the ***both*** group than it does with the ***nothing*** group; in Figure 6b, see *FB needs permission to sell user info* (H10 MW: p<.01, r=.17). As a point of comparison, a hypothetical requiring that the Library of Congress to solicit permission before archiving public user information elicits an even bigger difference between ***nothing*** and ***both*** (MW: p<.001, r=.22). Members of the ***both*** group, interestingly, view archiving with permission significantly more positively than those in the ***nothing*** group. Institutional archiving presents a more complicated tradeoff between personal privacy, public good, content ownership, control of identity, and potential misuse of data at access time, which makes it a more difficult scenario for participants to evaluate.

Attitudes toward content used as data. Responses to the three statements shown in Figure 7 show how the two participant subpopulations react to elements of the CA/FB story: using user data to target news, using user data to target advertising, and collecting user data via third party apps. The first two responses give us some insight into how immediate knowledge of CA/FB correlates with participants' attitudes

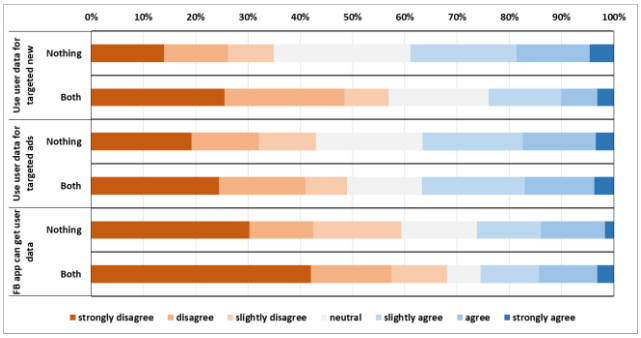

Figure 7. Subpopulation reactions to using personal data to target news, to target advertising, and to collecting user data via third-party applications.

toward related types of data use. The largest effect is on the attitude toward Facebook generating targeted news based on user profiles. Here, those with no familiarity with CA/FB (the ***nothing*** group) and those aware of CA/FB (***both***) are significantly more negative on the idea of targeted news (MW: p<.001, r=.22).

There was no significant effect on how the two subpopulations perceive Facebook's targeted advertising. However, members of the ***nothing*** group are more skeptical of targeted advertising than they are of targeted news (Wilcoxon Signed Rank: p<.06, r=.09), while those in the ***both*** group are more skeptical of targeted news than they are of targeted advertising (Wilcoxon Signed Rank: p<.01, r=.14). Because one aspect of the CA/FB story involves targeting political advertisements, it is interesting that these differences are pulling in opposite directions, with a negative view of targeting news correlating with an awareness of the story.

Unsurprisingly, familiarity with CA/FB is reflected in participants' attitudes toward giving apps access to users' data. While the largest contingent of the ***nothing*** group disagrees strongly, there are other attitudes distributed fairly evenly across the spectrum. Moderate to strong disagreement is more the norm for those in the ***both*** subpopulation.

DISCUSSION

In this paper, we have focused on the extremes of the study participants, those who report paying the most and least attention to the CA/FB story. Results show that these two groups have different attitudes toward the content they have contributed to social networking platforms. Participants the most aware of the story—how personal content was gathered via social network connections and how it was used to target political advertising—generally seem to be more vehement in their attitudes, and more tuned in to some of the complexities introduced by networked privacy; they are more polarized and less likely to be neutral in their views. Awareness of how data is collected and used seems to be vital in helping shape ownership attitudes and norms.

Yet, at the same time, these norms may be evolving. How do they compare, as a whole, with the norms represented in the 2013 dataset, as reported in [30]? How do the differences compare with other changes afoot in the data more broadly?

Evolving norms. Norms are important in governing many aspects of how people who use social networking platforms reuse other users' content, and how they expect their own content to be reused, especially by their peers. As May notes, there are tradeoffs between norms and regulations [34]; these tradeoffs work differently for conventionally published work and content that people contribute to social networking platforms. Conventionally published material may go so far as to incorporate Digital Rights Management (DRM) to codify and flatten nuanced fair use rules. User-contributed content, on the other hand, generally relies on norms to govern reuse. Although participants show some awareness

that copyright applies to the material they have shared online, they also have little sense of the complexities of fair use, nor awareness of the platform's terms of service. Thus norms may drift away from fair use. For example, if something is funny or meme-like, study participants are more willing to reuse it without further thought.

In what direction have norms evolved? Analysis of the 2018 and 2019 data revealed that many of the norms identified using the equivalent 2013 data were stable across this five-to-six year period. In particular, norms associated with peer reuse and removal have not changed significantly. This stability is notable, and stands in contrast to norms associated with corporate and institutional reuse, which have seen significant changes. These changes signal that participants overall have experienced a reduced sense of ownership and control of the data they have contributed to Facebook once they perceive their own data as being in corporate hands. This trend is consistent with the result that people seem to expect less ownership over other types of online content they have contributed to different types of platforms [32].

Privacy rights and social boundaries. The CA/FB story provides news followers with a specific example of how content they have contributed to their social media accounts may be collected and used as data by an organization other than the platform hosting the account. The story may have accentuated new issues or confirmed participants' existing suspicions about how social media data might be used in a way that compromises privacy, control of one's identity, and other entailments of content ownership. Thus it's not surprising that participants in the ***both*** group would have attitudes that diverge from participants in the ***nothing*** group and might begin to expect different ownership norms.

Of course, some participants across the board have lost trust in Facebook regardless of whether they've heeded this particular story; they may have noticed their peers' commentary on this topic, or have gotten wind (in one form or another) of the meme, "When something online is free, you're not the customer, you're the product" [49]. Yet we expected people who were the most aware of CA/FB to respond to specific hypotheticals differently because of the features the hypotheticals were testing.

In three of the five hypotheticals corresponding to Scenario 1 (Saving), an individual is posited as saving content that belongs to a friend or social contact. Part of the CA/FB story is based on the premise that a relatively small number of Facebook users' accounts served as an exploitable conduit for the collection of their friends' data without the friends' explicit permission; social boundaries were more permeable to an outsider than content owners (or their social network) intended. Thus we expected that participants who had been paying attention to the story would be more sensitive to the privacy and control norms suggested by these boundaries, regardless of whether the proposed violation was committed by an individual or an organization. Responses showed this to be true. Our prior studies suggested that participants were more cautious about content available through social connections [31]. This remains true for those who paid attention to the CA/FB story, but it is less so for those who didn't.

The new results are consistent with participants' reduced sense of ownership over content they have contributed to various platforms. Some say they are simply contributing less personal content to the services in acknowledgment of reported privacy violations [3] or uncertainty about perceived audience [9]. *However, this shift is less true for those who paid attention to CA/FB, rather than more*. It is likely that the story attracted the attention of those users who are more concerned with data privacy and greater proponents of privacy rights (rather than just privacy strategies).

Mobility of social media content. Although in many ways, social media accounts are intrinsically tied to an identity on a platform—people present themselves differently in different contexts according to the platform's norms [7] and their goals in using it [14]—the ability to reclaim data (to download it to local storage) is another proxy for ownership. Both of the hypotheticals that test downloading one's own content include platform-specific social network data (e.g. profiles and links to friends). We tested this idea because past field studies have shown that an important part of the value of an account (and the reason some users may keep an account they no longer want) is that they don't want to lose these connections to friends and business contacts, especially if it would be awkward to reestablish them [29]. Do participants feel they own the data the platform helped them create and maintain? Participants familiar with the CA/FB story feel a greater sense of ownership of their accounts than those who ignored the story.

Yet there is still uncertainty about what Facebook owns (particularly content a user has created in the service like statuses, profile elements, messages, and connections), and what a user can legitimately take away. When we asked participants what they would want to move from Facebook to another service, even a few members of the ***both*** group expressed uncertainty about what they actually owned (in addition to the expected considerations of value). SN2018-203 admitted that she would take, "*My photographs and videos. I would like to take a snapshot of my statuses, but wouldn't expect that I automatically own those*."

Participants also grappled with the expectation of future use by the platform's users, and by the platform itself. SN2018-181, a member of the ***both*** group, responded, "*I would expect to be able to import my contacts but I wouldn't expect to take personal information, pictures, statuses, etc over. I think that it makes the most sense to build a new profile on a new social network site than to just "transfer" over all your existing information. That way, you control which site gets what information. Maybe there are things on FB that you don't want on LinkedIN so by controlling what you publish and submit gives you a bit more control on how your data is used*." Members of the ***nothing*** group generally did not

reflect on the ownership of their Facebook content in their responses; instead, they just addressed its value and its transferability to a new platform. For example, SN2019-200 said, "*I would primarily like to take my identifying information (where I work, my education, past colleges) and my photos. Everything else is trivial.*" This pragmatism seems to align with the idea that content ownership is interpreted less as a right by members of the ***nothing*** group.

Content Accuracy. Although nothing about the CA/FB story explicitly touched on content accuracy, the story raised the specter of potentially unreliable information being deliberately introduced to a social network, and directed to those vulnerable to its influence. The differing reactions to the content accuracy scenario lead us to believe that those who are paying attention to this type of news, and those who aren't, are creating groups with diverging norms. One group (***nothing***) seems less apt to be uncomfortable with individual responsibility (the person who detects the error bears the burden of reporting it), and the other (***both***, those who are reading the news) is in greater agreement with how things have been changing (that the platform itself should be assuming responsibility).

The original interpretation of mid-1990s US legislation (CDA, Section 230) exempted platforms from responsibility for the content they host (making platform owners more like communications infrastructure providers than, for example, publishers), but more recently there's been a perception that platforms should take some responsibility for fallacious or harmful content such as hate speech [20]. It makes sense that those most aware of the CA/FB story would also be aware of reports that social media platforms like Facebook have taken small (but conspicuously reported) steps to remove content and accounts that violate this aspect of their terms of service.

Content Used as Data. One of the most controversial areas highlighted by the Internet era is the application of analytic tools to user-contributed content for various ends: e.g. to target advertising or news. Although people were profiled and classified in a variety of ways well before the Internet was introduced (e.g. some states and the federal government implemented prison classification programs in the 1930s and 1940s to predict recidivism based on a faceted analysis of prisoners' psychological and demographic characteristics [26] and mass-media advertising has long endeavored to connect a desired target audience with a specific product [40]), this type of analysis has become more conspicuous and fine-grained with the ready availability of user-contributed data that can be combined with other data resources.

We anticipated that familiarity with the CA/FB story would signal greater concern for how user-contributed content was collected, shared, analyzed, and used by corporations. We also expected that norms would have changed within the ***both*** group from the collective norms we charted using the same hypotheticals [30]; the public seems to be more aware that their data is used as a means of profiling individuals and their interests. Nonetheless, groups at both ends of the awareness spectrum continue to dislike the status quo (that Facebook can monetize information about people) without soliciting user permission (permission is a popular panacea for a variety of boundary-testing situations, including institutional archiving of user-contributed content). Interestingly, members of the ***both*** group seem to be significantly more sanguine about monetizing this information themselves than those in the ***nothing*** group; the ***both*** group matches 2013 norms. The reduced interest in monetizing personal information seems to reflect the slackening feeling of ownership among the nothing group.

CONCLUSION

The subpopulation of people who are paying attention to news like CA/FB exhibit markedly diverging attitudes toward the ownership of online content, regardless if the content is being saved by individuals (e.g., as personal archives, a start for building a social network on a new platform, or fodder for memes) or corporations (e.g., as big data). Their responses show them to be the activists who notice and protest the misuse of user contributed content at all levels. But as Norberg et al. observed about privacy [37], familiarity with an event and its outcome may change attitudes without changing actual behavior.

Why is it important, then, to understand these social norms and how they are changing? Norms eventually do guide actual behavior [13] (what is acceptable, and what is not, especially in areas like content ownership, where laws lag behind practice in important ways [45] and tensions are introduced by attempts to represent hard and fast rules like DRM [34]). Targeted news and advertising are examples of capabilities that people have become aware of over time; we have shown how experiences and familiarity with these new capabilities change users' attitudes. Looking forward, technologies such as face recognition, and its use to auto-tag people in photographs, have a similar potential to reshape norms.

In future work, we plan to not only chart the evolution of ownership norms over time, but also to explore new types of personal data, some of which is recorded behind the scenes (such as interactions with IoT devices or search engine queries) and reused in an increasingly broad spectrum of applications. We also believe it is important to use other research methods (e.g. fieldwork) to explore how these norms translate into action, and to better understand the reasoning behind the attitudes we continue to observe.

ACKNOWLEDGMENTS

We thank the volunteers for their participation and the CHI reviewers for providing feedback that made this paper better. This material is based upon work supported by the National Science Foundation under Grant No. 1816923.

REFERENCES

[1] Ahler, D.J., Roush, C.E., and Sood, G. (2019) The Micro-Task Market for Lemons: Data Quality on

Amazon's Mechanical Turk. *Proc. 2019 Meeting of Midwest Political Science Assn*, ret. July 22, 2019.

[2] Baumer, E., Adams, P., Khovanskaya, V., Liao, T., Smith, M.E., Sosik, V., and Williams, K. (2013) Limiting, leaving, and (re)lapsing: an exploration of facebook non-use practices and experiences. *Proc. CHI '13*, ACM Press, NY, 3257-3266.

[3] Beck, J. (2018) People Are Changing the Way They Use Social Media. *The Atlantic*, June 7, 2018.

[4] Beck, J. (2019) Facebook: Where Friendships Go to Never Quite Die. *The Atlantic*, February 4, 2019.

[5] Bentley, F., Daskalova, N., and White, B. (2017) Comparing the Reliability of Amazon Mechanical Turk and Survey Monkey to Traditional Market Research Surveys. *Proc. CHI EA '17*. ACM, New York, NY, USA, 1092-1099.

[6] boyd, d. and Heer, J. (2006). Profiles as conversation: networked identity performance on Friendster. In *Proceedings of the Hawaii International Conference on System Sciences*.

[7] Chandrasekharan, E., Samory, M. Jhaver, S., Charvat, H. Bruckman, A., Lampe, C., Eisenstein, J., and Gilbert, E. (2018), The Internet's Hidden Rules: An Empirical Study of Reddit Norm Violations at Micro, Meso, and Macro Scales. *Proceedings of the ACM on Human-Computer Interaction*, Vol. 2, Article No. 32.

[8] Child, J. T., Pearson, J. C., & Petronio, S. (2009). Blogging, communication, and privacy management: Development of the blogging privacy management measure. *Journal of the American Society for Information Science and Technology*, 60 (10), 2079–2094.

[9] Das, S. and Kramer, A. (2013) Self-censorship on Facebook. Proc. ICWSM 2013, 120-127.

[10] Difallah, D., Filatova, E., and Ipeirotis, P. (2018) Demographics and Dynamics of Mechanical Turk Workers. *Proc. WSDM'18*, https://doi.org/10.1145/3159652.3159661.

[11] Duffy, B. E., and Chan, N. K. (2019) "You never really know who's looking": Imagined surveillance across social media platforms. *New Media & Society*, 21(1), 119–138.

[12] Edwards, J.L. (2019) *An Examination of Consumers' Social Media Trust In the Wake of the Facebook and Cambridge Analytica Scandal*. Unpublished Thesis, Kalamazoo College.

[13] Ellickson, R. (1994) *Order Without Law*. Harvard University Press.

[14] Farnham S. & Churchill, E. (2011) Faceted identity, faceted lives: social and technical issues with being yourself online. Proc CSCW '11, 359-368.

[15] Fiesler, C. and Bruckman, A.S. (2014) Remixers' Understandings of Fair Use Online. *Proc. CSCW '14*.

[16] Fiesler, C. and Hallinan, B. (2018) "We Are the Product": Reactions to Online Data Sharing Controversies in the Media. *Proc CHI 2018*.

[17] Fiesler, C., Lampe, C., and Bruckman, A. (2016) Reality and Perception of Copyright Terms of Service for Online Content Creation. *Proc. CSCW'16*.

[18] Flintham, M., Karner, C., Bachour, K., Creswich, H., Gupta, N. & Moran, S. (2018) Falling for Fake News: Investigating the Consumption of News via Social Media. *Proc. CHI '18*, ACM Press, NY.

[19] Granville, K. (2018) Facebook and Cambridge Analytica: What You Need to Know as Fallout Widens. *The New York Times*, March 19, 2018.

[20] Grimmelmann, J. (2019) Continuity and Change in Internet Law. *CACM*, vol. 62, no. 5, 24-26.

[21] Hewitt, A., & Forte, A. (2006). Crossing boundaries: Identity management and student/faculty relationships on the Facebook. *ACM Special Interest Group on Computer-Supported Cooperative Work*, Banff, Canada.

[22] Isaak, J. and Hanna, M.J. (2018) User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 8 (August 2018), 56-59.

[23] Jiang, J.A., Fiesler, C., and Brubaker, J.R. (2018) "The Perfect One": Understanding Communication Practices and Challenges with Animated GIFs. *Proc. CSCW '18*, ACM Press, NY.

[24] Kang, R., Dabbish, L., Fruchter, N., and Kiesler, S. (2015) "My Data Just Goes Everywhere": User Mental Models of the Internet and Implications for Privacy and Security. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)* 2015.

[25] Lindley, S., Marshall, C.C., Banks, R., Sellen, A. & Regan, T. (2013) Rethinking the web as a personal archive. In *Proc. WWW 2013*, 749-760.

[26] Loveland, F. (1960) The Classification Program in the Federal Prison System: 1934-1960. *Journal of Federal Probation*, 24 (June 1960), 7-12.

[27] MacCormick and Summers, (eds.) *Interpreting Precedents*, Ashgate/Dartmouth, 1997, 528-529.

[28] Madden, M., & Smith, A. (2010) Reputation management and social media: How people monitor their identity and search for others online. Washington, D.C.: Pew Internet & American Life Project.

[29] Marshall, C.C. & Lindley, S. (2014) Motivations and Strategies for Self-Search. *Proc. CHI'14*, 3675-3684.

[30] Marshall, C.C. and Shipman, F.M. (2015) Exploring the Ownership and Persistent Value of Facebook Content. *Proc. CSCW'15*, ACM Press, NY.

[31] Marshall, C.C. and Shipman, F.M. (2017) Who Owns the Social Web? *Communications of the ACM 60*, 5, 52-61.

[32] Marshall, C.C. and Shipman, F.M. (2019) The Ownership and Control of Online Photos and Game Data. *Proc. JCDL'19*, IEEE Press.

[33] Marwick, A. E. and boyd, d. (2014). Networked privacy: How teenagers negotiate context in social media. *New Media & Society*, 16(7), 1051-1067.

[34] May, C. (2003). Digital rights management and the breakdown of social norms. *First Monday*, 8(11).

[35] Micallef, N. and Misra, G. (2018) Towards designing a mobile app that creates avatars for privacy protection. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '18)*. ACM, New York, NY, USA, 79-86.

[36] Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

[37] Norberg, P.A., Horne, D.R., and Horne, D.A. (2007) The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126.

[38] Perrin, A. (2018) Americans are changing their relationship with Facebook. Washington: Pew Research Center.

[39] Phelan. C., Lampe, C., and Resnick, P. (2016) It's Creepy, But It Doesn't Bother Me, *Proceedings of CHI 2016*, 5240-5251.

[40] Pierce, J., Lee, L., Gilpin, E. (1994) Smoking Initiation by Adolescent Girls, 1944 through 1988: An Association with Targeted Advertising. *JAMA* 271(8), 608-611.

[41] Rissland, E.L., (1989) Dimension-based Analysis of Hypotheticals from Supreme Court Oral Argument. *Proc. ICAIL-89*, ACM Press, pp. 111-120.

[42] Smith, A. & Anderson, M. (2018) Social Media Use in 2018. Pew Research Center (Internet & Technology), March 1, 2018.

[43] Strauss, A. and Corbin, J. (1998) *Basics of Qualitative Research*, Sage Publications, 1998.

[44] Stutzman F. & Kramer-Duffield, J. (2010) Friends only: examining a privacy-enhancing behavior in facebook. *Proc. CHI '10*, 1553-1562.

[45] Tehranian, J. (2007) Infringement Nation: Copyright Reform and the Law/Norm Gap. *Utah Law Review*, vol. 3, 537-550.

[46] Wang, Y., Komanduri, S., Leon, P. G., Norcie, G., Acquisti, A., and Cranor, L. (2011) "I regretted the minute I pressed share": A Qualitative Study of Regrets on Facebook. *SOUPS'11*,

[47] Woodruff, A. (2014) Necessary, Unpleasant, and Disempowering: Reputation Management in the Internet Age. *Proc. CHI '14*, 149-158.

[48] Zhao, X, Lampe, C, Ellison, NB (2016) The social media ecology: user perceptions, strategies and challenges. *Proc. CHI'16,* 89–100.

[49] Zittrain, J. Meme patrol. Future of the Internet blog post dated March 21, 2012. Retrieved September 13, 2019.