

What (or Who) Is Public?

Privacy Settings and Social Media Content Sharing

Casey Fiesler¹, Michaelanne Dye², Jessica L. Feuston³, Chaya Hiruncharoenvate², C.J. Hutto², Shannon Morrison⁴, Parisa Khanipour Roshan², Umashanthi Pavalanathan², Amy S. Bruckman², Munmun De Choudhury², Eric Gilbert²

¹University of Colorado Boulder, ²Georgia Institute of Technology,

³Northwestern University, ⁴Syracuse University

casey.fiesler@colorado.edu, {mdye, chaya, cjhutto, khanipour, umashanthi}@gatech.edu,

jes.feuston@u.northwestern.edu, shmorris@syr.edu,

{asb, mchoudhu, gilbert}@cc.gatech.edu

ABSTRACT

When social networking sites give users granular control over their privacy settings, the result is that some content across the site is public and some is not. How might this content—or characteristics of users who post publicly versus to a limited audience—be different? If these differences exist, research studies of public content could potentially be introducing systematic bias. Via Mechanical Turk, we asked 1,815 Facebook users to share recent posts. Using qualitative coding and quantitative measures, we characterize and categorize the nature of the content. Using machine learning techniques, we analyze patterns of choices for privacy settings. Contrary to expectations, we find that content type is *not* a significant predictor of privacy setting; however, some demographics such as gender and age *are* predictive. Additionally, with consent of participants, we provide a dataset of nearly 9,000 public and non-public Facebook posts.

Author Keywords

privacy; content analysis; Facebook; dataset; machine learning; Mechanical Turk; mixed methods; prediction; research methods; social media

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Content privacy control is a central issue for users on social media platforms. However, prior work has shown that privacy setting options are often difficult to understand and

use, and that default privacy settings are the norm [18,32,39]. Nevertheless, selective sharing is one strategy for users searching for ways to better manage their online content and identities [31,46]. As a result, the content of many social networking sites (SNSs) have become a mixture of public content and selectively shared, non-public content. This is especially true for sites like Facebook with multi-level privacy settings, allowing the public (viewable to anyone) and non-public (viewable to only certain people, e.g., friends) to cut across a single user’s content. Though even on Twitter, Instagram, or Tumblr, where users have a single privacy setting, these two types of content may cut across the site or stream as a whole. The “public” view of an SNS therefore provides an incomplete picture of the content on the site. A better understanding of what is public and what is not public on these sites could provide insight into user practices for privacy settings—useful for both site designers and privacy researchers.

Beyond privacy, both HCI researchers and social scientists studying online behavior use social network content to study a vast array of phenomena. Twitter is a particularly common platform for study. Tufekci suggests that it has become the “model organism” of social computing research; like the fruit fly, we use it to answer many types of questions in part because it is the easiest way to do so [43]. Why might this be the case, when Facebook users represent a larger proportion of the world population, with nearly 1.65 billion monthly users?¹ One possible explanation is that on Twitter, data is mostly public and easy to access through an API. Facebook, in contrast, not only has less public data,² but it is more difficult to retrieve. Interesting studies using large-scale quantitative data from Facebook often come from researchers with direct access to this proprietary data (e.g., Facebook employees) [2,3,9].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CSCW '17, February 25–March 01, 2017, Portland, OR, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4335-0/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/2998181.2998223>

¹ As of March 31, 2016. Source: <http://newsroom.fb.com/company-info/>

² As of 2013, private accounts made up just 4.8% of Twitter [29], in contrast to the 75% non-public Facebook content revealed by our data.

For the overwhelming majority of researchers working outside of social media companies themselves, if we want a complete picture of an SNS, either for descriptive purposes or to ensure a representative sample, we should understand the difference between what we *see* and what we do *not* see. Therefore, we began with this research question: Are there systematic content differences between public and non-public content on Facebook? If so, are there certain *types of content* that tend to be posted publicly versus non-publicly, and/or certain *characteristics of users* who tend to post publicly versus non-publicly?

Though privacy behavior is complex, as examined extensively in prior work in this area, these research questions make it appropriate to discuss *public* and *non-public*, as operationalized by Facebook privacy settings, as a binary. Specifically distinguishing *public* content under this operationalization is appropriate not only because of ease of access for scraping, but because “viewable to anyone” is a common ethical metric for data that can be collected without consent [48]. Because this work is motivated in part by practical questions around the study of social media data, our findings and claims apply to these objective measures of privacy settings that users employ for controlling the visibility of their posts.

To collect data to answer our research questions, we deployed a survey via Mechanical Turk and asked participants to provide recent Facebook posts, along with information about the privacy settings for each. This resulted in a dataset of nearly 11,000 posts, a subset of which we hand-coded qualitatively for content. We also applied statistical machine learning techniques (i.e., regression and classification) to determine the link between privacy settings and content.

We found that for most Facebook users in our dataset, their content is either all public or all non-public, and that public posts represent only a quarter of this content. Content type as determined by our coding scheme does not predict privacy setting at the post level, but at the user level some demographics *are* predictive. In other words, in characterizing privacy settings of content across Facebook, it is not necessarily the case that certain *posts* are public or not, but more often that some *people* tend to post publicly or not. The fact that content differences do not differentiate public and non-public posts within a dataset of thousands is somewhat surprising relative to previous work, and suggests that if systematic content differences do exist, they are likely subtle ones.

We urge caution, however, for technologists applying these findings. This paper describes the “what” around public versus non-public data, but little about the “why”: under the current state of Facebook privacy settings (which includes, e.g., how easy these settings are to navigate), post content *may not currently differ* between public vs. non-public content. This is intuitively surprising. But the result says little about how people *intend or want* to control access to

their posts. For example, perhaps the lack of content difference arises because many people do not understand just how public they are on Facebook. Future work should investigate these issues—especially if it aims to build new privacy interfaces or tools.

Additionally, one of the motivations of this work was a recognition of the difficulty in obtaining non-public Facebook data. Therefore, with consent of participants, we openly release a portion of our dataset to the research community. The dataset consists of nearly 9,000 public and non-public Facebook posts, and includes *demographic information* about the content authors, details about the size of their social networks and their general Facebook *usage and experience*, the *type of media* contained in each post (e.g., text, image, video, web-links, and location check-ins). We include statistics about the date of each post, the number of Likes and Comments, the specific privacy setting employed, and the author’s reported reasons for *why* the privacy setting was used for each post. It is our hope that this dataset will benefit social media researchers without access to this kind of data.³

RELATED WORK & MOTIVATIONS

In the field of HCI, a major contextual backdrop to people’s interactions with technology is privacy. A large body of research addresses questions about how equipped individuals are to navigate privacy in the information age, and there is increasing public concern over this issue [1,23]. Meanwhile, social media has provided a unique context for studying privacy behavior due to voluntary self-disclosure in combination with often unintentional data disclosure. Privacy has therefore become a major area of inquiry for social computing researchers, particularly with respect to Facebook, currently the most prevalent SNS. Here, we focus largely on the prior work that is relevant to *privacy settings* as one measurable component of privacy behavior.

Privacy Strategies on Facebook

A number of researchers have examined the way that Facebook users navigate privacy concerns such as worry over keeping up with constantly changing privacy policies [40] or control over self-presentation on Facebook [27]. One potential way to mitigate privacy risk is with targeted disclosure—that is, using different privacy settings for different posts [46]. For example, one study found that participants were most concerned about controlling who could see emotional or self-expressive content [51]. Facebook users might also use ad-hoc strategies to mitigate privacy threats [19,47]. Overall, privacy management on SNSs is complex, encompassing a range of strategies (including, but not limited to, use of different privacy settings) [49].

Part of this complexity is due to the fact that many users have difficulty in understanding privacy settings on

³ This dataset is available at <https://github.com/compsocial/WhatWhoCSCW2017>

Facebook. For example, one study found that despite a desire from Facebook users to selectively share, they rarely used custom privacy settings because they found them confusing, resulting instead in self-censorship [39]. A longitudinal study of Facebook users examined changes in the privacy settings on their profiles and found that over time, users disclosed less, but that this trend reversed itself after Facebook made changes in default settings [41]. This led the authors to conclude that much of the difficulty users have in managing their privacy is due to the power that the providers have over the interface and system defaults. Based on this prior work, in the current study we investigate the prevalence with which Facebook users intentionally change (or not) the *default* privacy settings initially provided by the system.

Some prior work has shown that, given the choice, people prefer strong privacy settings as a default [18]. However, they may also make mistakes in setting privacy options. In one study, every single Facebook user-participant confirmed at least one inconsistency between their sharing *intentions* and their *actual* privacy settings [32]. Therefore, it is unclear to what extent privacy intentions align with the actual settings applied. Though in the current study we include some basic information about users' sharing intentions, our findings and claims are limited to *actual privacy settings* observed.

Context collapse and self-disclosure heuristics also provide insight into privacy setting choice in social media. Hogan argues that self-representation on social media should be thought of in terms of an exhibition of many performances in different contexts, and proposes the lowest common denominator theory, which states that people “need not consider everyone when submitting content but only two groups: those for whom we seek to present an idealized front and those who may find this front problematic” [17]. In support of this claim, Vitak and Ellison found that two strategies many Facebook users employ to avoid unintended sharing with their entire network is to either abstain from posting at all, or to only share content they deem appropriate for everyone in their network [44]. Vitak et al. also found that granular privacy settings and multiple friends lists was a strategy for controlling disclosure and audience [45]. Other studies have revealed complex strategies around avoiding context collapse on Facebook, though these do not always involve privacy settings—for example, culling Facebook friends, self-censoring, or maintaining multiple accounts are strategies that Facebook users commonly employ [19,47].

Patterns of Privacy Settings

In light of the complexities of sharing strategies, other researchers have considered whether there might be patterns to content sharing behavior. For example, one study revealed that social network size correlated with disclosure: Facebook users with more friends tended to reveal more information [50]. The study also reported both

topic-based and *gender*-based differences in disclosure behavior: notably, that women reported disclosing political views and their current address less frequently than men. Another study found evidence of women displaying less open sharing behavior (more with friends, less with strangers) than men [32], and a 2012 Pew survey confirmed that women are more likely to choose private (i.e., friends only) settings for profile access than men [31].

Also with respect to topic-based differences in sharing behaviors, the study found that certain sensitive topics (such as sex, drugs, and alcohol) were less likely to be shared widely [32]. Others uncovered that photos were more likely to be shared friends-only, compared to video, text, and links [28]. These studies are both informative and insightful; they motivate our own investigation of differences in privacy settings from both *people*-focused and *content*-focused perspectives. However, in contrast to our work, these previous studies are largely founded on self-reported topic-based sharing restraints rather than detailed content analyses of actually observed privacy settings.

A series of studies by Bazarova and colleagues examined self presentation and self disclosure for messages with different privacy intentions, both on Facebook and Twitter [4,5,8]. Though there was some content-based coding specifically around the presence or absence of disclosure, the measures in these studies were self reports of goals and intimacy [4,8], and in one study, positive or negative emotion as expressed by language [5]. They revealed some differences based on intended audience or network (e.g., less self disclosure and less negative emotion on status updates over one-to-one private messages) [5,8]. However, because their analyses bucketed all types of Facebook status updates together, they note the need for future work considering message-level privacy settings [8], which is a motivation for the work described in this paper.

As noted by Leary and Kowalski, “public” versus “private” is not a binary, but rather publicness is “the probability that one’s behavior will be observed by others and the number of others who might see or learn about it” [26]. As a result, how comparatively “private” a friends-only Facebook post is depends on the size of one’s friends list and how many of those people the user truly knows. As Vitak et al. and others have pointed out, the conception of one’s audience can have an impact on privacy attitudes and behavior [45]. For the purposes of this paper, we distinguish *public* posts (viewable to anyone) from *non-public* (accessible to certain or all Facebook friends), but suggest the potential for future work in considering network size in more detail.

Technological Intervention and Prediction

Other researchers have looked at potential technological interventions in this space, such as developing recommendations for privacy settings. Fang and LeFevre built a “privacy wizard” that uses a machine learning model based on a user’s past privacy preferences to configure

future settings [11], and Ghazinour et al. built a tool that would suggest privacy settings based on those of similar other users [13]. As part of our discussion, we will consider how our findings could (or rather, could *not*) be used as part of tools such as these.

Others have considered privacy setting *prediction* as well. Stutzman and Kramer-Duffield surveyed 444 undergraduate Facebook users, then used regression analysis based on a number of factors including demographics and Facebook use to model the odds that an individual had a friends-only (versus public) Facebook *profile* [42]. Their only significant demographic factor was male users being 59% more likely than female users to have a friends-only profile (in contrast to the findings in [31,32]). They also found that more friends increased the odds of a friends-only profile (which contradicts the finding in [50]). Contradictions between study findings could be based on a number of factors, notably population differences, suggesting that considering data from a larger set of users might reveal more information. We hope that the current study contributes to this body of knowledge about user privacy control by considering a host of qualitative content measures as well as objective attributes at both the post and user level.

DATA COLLECTION

Other researchers have used multiple methods for gathering non-public Facebook data, including bringing Facebook users into a laboratory setting [14] or creating an app that will gather participants' data [20,22,36]. Another option is to ask Facebook users to provide their own posts, remotely and manually. For example, in Morris' 2010 study, in addition to collecting survey data about question-asking behavior, they asked each of their 612 Facebook user participants to provide an example status update from their own Facebook feed [34]. Similarly, Bazarova et al. asked 79 university students to each provide 6 Facebook status updates [5], and in another study Choi and Bazarova asked 164 students in a university class to each provide 5 updates [8] in order to construct datasets.

For this study, we used a similar method to Morris and Bazarova, but on a larger scale, via Mechanical Turk. We conducted a survey of Facebook users (approved by a university IRB) in which we asked them to provide us with recent posts, as well as information about privacy settings. We first conducted a pilot study of 28 (uncompensated) participants recruited through our personal social networks in order to clarify any confusion in the survey format, questions, and process. We also asked them an open-ended question for each post: "Why did you choose this privacy setting?" We used the patterns of responses to formulate multiple-choice answers for our final survey.

In the final survey comprising the main data in this paper, we asked each participant to provide us with their six most recent posts. We described a post as content *they* wrote (as opposed to for example other users posting on their wall, or automated Facebook messages). We asked for the text of

the update, a description of any included pictures or videos, the URLs of any links, descriptions of check-in locations, number of likes, number of comments, and date of the post. We also asked them for the privacy setting of that post, as well as *why* they chose that privacy setting. Based on our pilot testing, the number of posts (six) was chosen to give us a reasonable amount of data per person without making the task overly time consuming or tedious. This number is also similar to those used to create datasets in Bazarova and colleagues' previous work [5,8]. We also instructed participants to provide most recent posts in order to make post choice simple and objective.

We likewise collected information about participants' Facebook use—how many Facebook friends they have, how much time they spend on Facebook, whether they joined before or after 2014 (when the default privacy settings changed), and answers to Ellison's "Facebook Intensity Scale" [10]. We collected basic demographic information, and used Hargittai's web use scale [16] and Litt's Facebook use scale [27] to measure how web- and Facebook-savvy our participants were. We also included as a check a question about whether they skipped any posts in reporting them to us (assuring them that this answer would not affect their compensation). Finally, we asked for optional permission (again, noting that their answer would not affect compensation) to include their posts in a public dataset available for other researchers to use; our analysis for this paper, however, is based on the entire dataset.

To facilitate recruitment of a large number of participants, we used Amazon's Mechanical Turk (mturk) crowdsourcing service. This allowed us to cast a wider demographic net than localized recruitment such as through a university. Studies have shown that mturk users may be more demographically diverse than other Internet samples [6] and perform comparably to laboratory subjects in traditional experiments [37]. Though a known limitation of mturk is that participants may be less likely to pay attention to experimental materials [15], many techniques exist to ensure high quality data is collected from micro-labor markets like mturk (c.f., [33,38]). Such steps were unnecessary for the current study because our intensive manual inspection and in-depth qualitative analysis of the data allowed us to check for poor quality responses in situ.

Our pilot study showed that the survey usually took 10 and no more than 15 minutes to complete, so we paid participants \$1.50 per survey, to ensure a rate of greater than \$0.10/minute, considered the baseline for fair mturk pay. Due to the geographically binary nature of the mturk population (with the majority of workers located in either the United States or India), we chose to limit our survey to U.S. turkers. A sample of half U.S. and half Indian turkers would not generalize to a global population; our limitation at least allows some generalization to U.S. Facebook users.

We deployed the survey on mturk first to a set of 55 workers on April 3, 2015, which we used to create our

initial codebook (described in more detail later in this paper). We threw out the initial data once our codebook was finalized, and then deployed the survey in 3 batches of 605 mturk users. These batches occurred between April 27 and May 13, 2015. This resulted in 1,815 participants, for a total of 10,890 posts in our final dataset. We collected this large sample for the purposes of rigorous quantitative analysis; due to the labor-intensive nature of hand-annotating this data, we used a subset for our detailed qualitative analysis. Specifically, we chose 500 participants at random, for a total of 3,000 posts, and conducted a meticulous qualitative content analysis (our “coded sample”, described in detail later). We used statistical machine learning techniques to fully assess the entirety of the content in our 3,000 sample dataset (i.e., regression and classification; seeded first by our hand-coded data, and then by computational language models). In addition to this post-wise content analysis, we calculated additional quantitative measures on the full dataset using authors (rather than posts) as the unit of analysis.

DATA ANALYSIS AND FINDINGS

Describing the Dataset

The population of our dataset is drawn from the intersection of U.S.-based turkers and Facebook users. Though we had 1,815 total participants, only 1,706 completed the survey in its entirety, so we used this sample for quantitative measures.

Demographics

Table 1 includes descriptive statistics, including age, number of Facebook friends, and established scales for web and Facebook use. Both Litt’s Facebook Skill scale [27] and Hargittai’s Internet Skill scale [16] are on a 5-point scale, with 5 being the highest. Our population shows a high Facebook skills average (4.57), though this is a similar score to Litt’s non-turk population (4.4). Our population also displays a higher Internet skills score (4.01) than both Hargittai’s 2010 sample (3.24) and Litt’s 2014 sample of Facebook users (3.4). Unsurprisingly, turkers have higher Internet use skills than the general population, but their skill with Facebook is not markedly higher than other users.

	μ	σ	min	max
Age	30.74	9.15	18	75
# of Friends	387.36	424.45	0	5K
Facebook Skill	4.57	0.65	1.38	5
Internet Skill	4.01	0.84	1.17	5

Table 1. Descriptive statistics of our population

With respect to other demographics, though well gender-balanced, our limitation to U.S.-based turkers resulted in a non-ethnically diverse sample:

- **Gender:** 993 female, 701 male, 12 undisclosed

- **Ethnicity:** 1,218 White, 235 Asian, 93 Latino, 87 Black, 46 Multiracial, 8 Native American, 8 Other, 9 undisclosed
- **Education:** 6 less than high school, 139 high school graduate, 71 technical or vocational school, 601 some college, 631 college graduate, 254 post-graduate, and 4 undisclosed

Limitations and Data Quality

As noted in the description of our data collection, employing mturk as a solution to the difficulties in obtaining non-public Facebook data has its own trade-offs. This population may be non-representative in the following ways: (1) turkers may be unusually tech-savvy (as evidenced by the web use scale described above) or privacy-savvy (as implied by Kang et al. in a study of turkers [21]); and (2) our population is limited to turkers in the United States and is not ethnically diverse.

Additionally, as with any survey work, this dataset potentially contains some response bias. We attempted to measure how much of a problem this might be by including two checks for the accuracy of the data provided to us. First, we asked participants at the end of the survey if they had skipped any posts in providing them to us (noting that their answer would not affect their compensation); only 128 participants (about 7%) answered yes to this question. Additionally, in our qualitative coding of posts, we marked any posts as “junk” if we thought that they were gibberish or not real, or was not actually a status update (e.g., an automated Facebook message). Out of 3,000 posts, only 54 were so marked, and for only 5 participants (1% of our sample) did these represent all of their provided posts. Prior work with mturk surveys that incorporated attention checks showed a similarly low number of participants submitting in bad faith [12]. Based on our sample and evidence from prior work, we can posit that our larger dataset (though not extensively analyzed “eyes-on”) is mostly truthful.

Of course, it is still possible that participants were not truthful about skipping posts, or that turkers self-selected into the task based on their privacy preferences. However, given the constraints of this problem (collecting information that is not available by any method other than asking people for it), we felt that the potential for self-selection was less for this method than, for example, asking Facebook users to download an app that would scrape all of their data. However, we acknowledge that response bias is inherent in any dataset obtained with consent, and our results should be interpreted with this in mind.

Public Dataset

The demographics described above (along with the rest of this paper) represent the entirety of our data. With respect to our public dataset, 1,393 participants (or 82%) agreed that their data could be anonymized and included. Anonymization is being done by two human readers: one to do an initial anonymization, and a second to verify that it was done correctly.

We compared descriptive statistics for our full dataset versus the public dataset, and found very little difference. Gender and age breakdowns were nearly identical, and the only notable discrepancy was a total of 27% public posts in the public dataset versus 25% public posts in the full dataset. Though this dataset does not include network data for our participants, we attempted to collect most other information relevant to content. We see the potential for many uses beyond studies of privacy, both qualitative (e.g., topic-based content analysis) or quantitative (e.g., natural language processing).

Privacy and Sharing Behavior

Our large dataset provides us with a great deal of objective information about privacy setting choice on Facebook, as we have both the privacy setting for each of our nearly 11,000 posts and basic information about *why* the user chose that setting.

It is worth noting that in 2014, Facebook changed their default privacy setting for a new account from “public” to “friends only.” Less than 3% of our participants (51 of 1,815) joined Facebook after 2014. Because this is such a small portion of our dataset, any differences based on account creation date would not be meaningful. Also, only a small number of participants (< 300, about 16%) stated that they had *never* changed their privacy settings.

Could it be that users simply do not change the default setting? 342 of our participants used all public status updates, and of those 326 joined Facebook before 2014. Among those 326, 122 (37.4% of sub-sample, 7.2% of all participants) report that they never changed their privacy setting. In future work, it would be interesting to collect a dataset of users who joined after the change in account default for new users and investigate what percent leave the new (friends only) default unchanged for all posts. However, for the purposes of this paper, we are solely concerned with *what* they do, though intentionality continues to be an intriguing question for future work.

Additionally, less than half of the participants in our dataset (864) show *multiple* privacy settings across their 6 provided posts. In other words, it is slightly more common to maintain a single privacy setting than to switch back and forth between posts; posts are usually either all public or all friends-only, at least in the span of six updates. This becomes important in determining whether there are content differences based on privacy settings, discussed later.

Table 2 shows the observed frequency for each privacy setting in our dataset: 25% public and 75% non-public (i.e., posts that are restricted to viewing by friends only, by friends and tagged users, by friends of friends, or by a custom filter). We discuss how this breaks down per person in more detail later, but it should be noted that overall, our sample of Facebook posts contains far more non-public updates than public. However, even in this Internet-savvy

sample, using custom Facebook lists (which allow a much greater level of control over selective sharing) is rare.

Privacy Setting	N	%
Public	2,711	25%
Friends only	6,680	62%
Friends and tagged users	651	6%
Friends of friends	399	4%
Custom	353	3%

Table 2. Total posts by privacy setting

We also asked our participants *why* they chose the privacy setting for each post. Their options were multiple choice, based on patterns of response from our pilot study. Table 3 shows the observed frequency of each response. Notably, it is common for users to simply have their own default privacy setting, rather than thinking about the appropriateness for an individual post, which could in part explain why switching is rare. Additionally, despite prior research revealing that privacy settings are difficult to understand, very few posts had privacy settings that were reported as unintentional: only about 1% of all posts. Of these, 75% were unintentionally public. In other words, when users do make mistakes on privacy settings, it is likely that they make it more *public* than they intend.

Reason	N	%
I only wanted to share with my FB friends (not the public)	4,647	43%
It is my default setting	4,030	37%
I wanted to share this with as many people as possible (public)	1,221	11%
It was a personal preference based on my level of comfort for this content	562	5%
Relevancy/appropriateness based on people interested in this content	173	2%
I didn't mean to use this setting	117	1%
Other	80	< 1%
Someone else suggested I do it this way	22	< 1%

Table 3: Frequency of rationale for selecting privacy settings

We provided a free-answer “other” option for privacy setting rationale. The most common of these were variations on “saving for my own reference.” We did not provide a specific option for the “me only” custom privacy setting, so these represent some small number of custom-marked posts in our dataset. Having posts that are entirely private to the user appears to be a rare but not unheard of behavior. We also offered a free response to describe their custom Facebook list, if applicable. The most common were “close friends”, “family”, or “all but X” (e.g., “everyone I am friends with except my employees” or “friends minus family members”). There were also a few

specific groups mentioned, such as “gaming friends” or “dance class.” However, the use of custom privacy settings is relatively rare (about 3% of total posts), which follows Sleeper et al.’s finding of a lack of use of this feature despite users’ desires to selectively share [39].

In sum, our sample of nearly 11,000 posts reveals that public posts only make up about a quarter of Facebook content, and of non-public posts, most are set to “friends only” rather than making use of options that provide more granularity and selective sharing, with most users sticking to their owned defined defaults. Additionally, at least in this relatively tech-savvy population, mistakes in privacy settings are rare.

Characterizing Facebook Content

Codebook Creation

As noted in our data collection section above, we used our initial data (165 posts) to create our codebook. Three researchers worked together in an inductive, iterative coding process [7,30], coding for emergent phenomena as well as considering coding schemes from previous related work. They worked cooperatively on a subset of the data to create an initial set of codes. They then independently coded before coming back together to adjudicate differences and iterate on the codes. They repeated this process, considering the codebook finalized when there were no new properties discovered (meaning that the code categories were saturated) and no further substantial adjudication needed [7]. This resulted in a final set of 35 coded categories, including 11 high-level major categories (every post was coded in at least one of these), 8 lower-level type subcategories (6 of which were “children” of higher-level categories), and 16 topic categories. We also incorporated helpful meta codes in our dataset, including media type in the post, whether the post was targeted towards an individual or group, and whether the post should be anonymized if included in our public dataset. Using guidelines for codebook development in teams [30], our codebook included these elements for each code: a mnemonic, a brief definition, rules for inclusion and exclusion, inclusion examples, and exclusion examples.

In creating the codebook, we also drew from previous work categorizing social media *content type*. The initial high-level categories were drawn from Naaman et al.’s study in which they categorized Twitter content [35]. The categories they identified were: Information Sharing, Self-Promotion, Opinions/Complaints, Random Thoughts, Me Now, Questions to Followers, Presence Maintenance, Anecdote (Me), and Anecdote (Others). We identified one new, emergent category: Well Wishes. We added this to the set, and also clarified Questions to Followers as Mobilization Requests and conflated Thoughts and Opinions/Complaints due to conceptual similarities. These major content type categories were not always mutually exclusive (though they were more often than not), and every post in our dataset was coded with at least one.

We also considered previous work from Lampe [24], Ellison [9], and Morris [34] on mobilization requests and question asking, and included their categories of mobilization requests: Recommendation Request, Factual Knowledge, Social Coordination, Favor/Request, and Opinion/Poll. We added one emergent category that was unique to our Facebook data: Share Request.

We also coded for the *content topic* of posts. Independent of the *content type* category, each post included between 0 and 16 topics that we identified as being the most common and consistent in our dataset. We found our emergent topics to be similar to those identified in previous work categorizing Facebook content [32,34,39]. In finalizing topics, we made sure to include those identified in previous work as being “sensitive” social media topics that might impact sharing behavior: drugs and alcohol, sex and relationships, politics, and religion [32,39]. Our codes along with descriptions and examples are detailed in Table 4.

For the qualitative coding of the posts, there were a total of six coders, three more in addition to the three who created the codebook. Our subset of the full dataset for qualitative coding was 3,000 posts, from 500 participants chosen randomly. The initial coders coded a sample of 10% of this dataset (300 posts), coming to a consensus on the correct codes. The three new coders were trained on the same set of initial 165 posts, and then were tested against the 300 coded posts for inter-coder reliability to ensure consistency of coding heuristics. We required a Cohen’s Kappa threshold of at least .61 (considered “substantial” agreement [25]) before they could continue; after training sessions, all coders exceeded this threshold on their first try. New coders’ codes were also reviewed to ensure that discrepancies were due to reasonable subjective judgments rather than systematic misunderstandings. The rest of the 3,000-post dataset was split up among the six coders and coded independently.

What Type of Content is on Facebook?

Though no classification scheme could capture all of the nuances of Facebook content, our method of analysis made us confident that our codebook covers a great deal of it. Table 4 includes a description and observed frequency of both the *types* and the *topics* of content from the 3,000 sample posts that we hand-coded. A detailed description of the boundaries of each category would take more space than available for this paper, but our codebook will be available in its entirety along with our public dataset.

Independent of privacy settings, this provides us with a broad view of the type of content that is common on Facebook (at least, as provided by our Spring 2015 dataset). For example, in contrast to Naaman’s findings regarding categories on Twitter [35], “Me Now” posts (current state or activity) are less common proportionally on Facebook than what they observed on Twitter. Additionally, personal content taken together (“Me Now,” “Anecdote Me,” “Thought/Opinion,”) is far more common than information

Code	Type	Description	Example	N
Thoughts / Opinions	Major	thought, observation, or opinion without anecdotal or informational context	If you try to concentrate on a squirrel you'll be more distracted by important things	803
Personal Thought	Secondary	thought or observation about oneself	My life would be easier if I were in New York	174
Fortune Cookie	Secondary	proverb, inspirational thought, or quote	Friendship is like a mirror. Even when broken it still reflects the truth.	117
Complaint / Rant	Secondary	complaint about something specific or general ranting	Learn to drive jerkwards	91
Anecdote Me	Major	anecdotes or photographs of poster detailing past or future events	Red Robin tonighttttt woooo	738
Information / Content Sharing	Major	informing others or sharing information or content that is not personal	15 things to do for more responsive website design	705
Clickbait / Memes	Secondary	clickbait, otherwise "silly" information, or meme sharing	I'm Jon Snow! Take this quiz to find out which Game of Thrones character you are.	178
Anecdote Other	Major	anecdotes about or photographs of other people besides poster	My poor puppy is resting after hurting his leg	649
Me Now	Major	poster's current activity	watching Daredevil on Netflix	288
Well Wishes	Major	celebratory or well wishes targeted at someone or everyone	Happy Pi Day!	251
Mobilization Request	Major	questions, or requests for help or action	check out kickstarter from my friend	167
Request / Favor	Secondary	request for help or action from one's network	Please pray for A***** and her family this morning.	80
Social Coordination / Invitation	Secondary	search for others with similar agendas, invitation, or coordination of meeting/goal	Anyone want to go to YC's tonight?	35
Share Request	Secondary	specific request to share or spread content	LIKE and SHARE to wish Israel a peaceful and blessed 67th year.	22
Factual Knowledge	Secondary	question posed that assumes and expects a correct, objective answer	Does anyone know if toddlers can visit suplex city?	15
Opinion / Poll	Secondary	request for opinion, vote, or general solicitation	So I guess Lawrence has Uber now? Has anybody tried it?	10
Recommendation Request	Secondary	subjective, open-ended request for suggestions, referrals, etc.	Movie suggestions for my intermediate English class?	10
Self-Promotion	Major	promotion of something for poster's benefit	Come see me play at the Vinyl tonight	64
Presence Maintenance	Major	meta information about poster's presence on Facebook	I broke my Facebook hiatus to let everyone know they are loved this Valentine's Day	2
Entertainment	Topic		Hillary Swank is a wonderful actress	472
Family	Topic		Happy Birthday to my wife!!	386
Food	Topic		Easter isn't complete without Jelly Beans!	217
Humor	Topic		My white and gold laptop [image of a blue-and-black laptop]	199
Friends	Topic		Spring Break 2015 with my besties	160
Pets	Topic		I LOVE PETS	151
Current Events	Topic		sad about the ppl in Nepal ☹	130
Work / Academic	Topic		Just another day in the office	130
Religion	Topic		Just going to be grateful to God for another year with my family and friends!	115
Science / Technology	Topic		"How Egyptians moved pyramid stones [article URL]"	110
Sex / Relationships	Topic		"How did I marry someone more dramatic than me? How is that even possible???"	96
Weather	Topic		Rainy days suck with little 2 year old boys who need to burn off energy	91
Shopping / Consumerism	Topic		Christmas shopping, dinner date, more Christmas shopping!	83
Politics	Topic		Oklahoma Libertarian Party needs to gather 41,242 signatures to gain ballot access. Can you help?	51
Drugs / Alcohol	Topic		Last night was a blur... too many drinks no food	48

Table 4. Frequency of content codes, as expressed 3,000 qualitatively coded posts. Examples are from our dataset or shortened approximations. Codes listed by type and ordered by frequency. Secondary categories are children of the major categories above them in the table.

sharing, also in contrast to Naaman’s findings for Twitter. The largely personal content of Facebook posts, therefore, could explain why public posts are less frequent in that medium as compared to Twitter.

Quantitative Analysis towards Privacy Prediction

Now knowing what kinds of *people* and what kinds of *content* are in our dataset, we turned to our initial research questions regarding the relationship between these factors and privacy settings. To this end, we used regression techniques to model this relationship, asking what variables could reasonably predict privacy settings at either the post or user level. We considered three primary axes: (1) post content as instantiated by the qualitative codes described in Table 4 (with the *post* as a unit of analysis), (2) computational linguistic features; and (3) person-level demographic features (with a *user* as the unit of analysis). We call out the major results in subsections below.

Result: Content Codes Have Little Bearing on Privacy Settings

First, we asked how much information the codebook provides in modeling participant-selected privacy settings. In other words, with a post as the unit of analysis, do certain types of content (information sharing, family, etc.) typically co-vary with non-public posts more often than public ones? To answer our research questions about the objective measure of publicly viewable content, for the purposes of this analysis, we binarized privacy setting as *is public* (visible to everyone) and *is not public* (any other level of restricted visibility). This also serves to smooth potential biases in the data. Because we are examining content codes, we began with only the 3,000-post qualitative dataset.

We performed two logistic regressions, with the first used as a control for the second:

$$M_1: \text{privacySettings} \sim \text{WorkerID}$$
$$M_2: \text{privacySettings} \sim \text{WorkerID} + \text{coding}$$

For this analysis, our initial finding that most users in our dataset do *not* change privacy settings, but instead are public or non-public for all of their posts, is important. If their setting does not change at all, then it necessarily is not changing based on content. For the purposes of modeling, most of the time knowing *who* wrote the post (modeled as the fixed effect of knowing a turker’s *WorkerID*) is all the information that is needed. Therefore, in order to examine the impact of content codes, we looked to only the subset of posts from people whose privacy settings vary at least once over the course of our data collection ($N_{posts} = 864$, $N_{participants} = 145$). Within this subset, we find that simply knowing the identity of the author of the post still confers an impressive amount of information. The regression containing only *WorkerID* reduces deviance to 498.5 from 897, a significant reduction: $X^2(N=864, df=144) = 389.3, p < 10^{-10}$. In other words, even for users who switch privacy settings, they are still *mostly* post either publicly or non-publicly, aiding in prediction.

Next, we add to the fixed effect *WorkerID* all the content codes assigned by human raters (Table 4). The omnibus test determines whether the addition of the codes provides significant information *beyond* what the author information confers. We find that the addition of the content codes does improve on the first model, albeit only marginally: $X^2(N=864, df=40) = 64.9, p = 0.008$. (The *df* here represents only the added content codes.) While statistically significant, this is only a modest gain in real terms for a dataset of this size. Further, none of the content codes described in Table 4 possess significant coefficients. In sum, there is not enough of a relationship between content *type* or *topic* and privacy setting for us to use one to estimate the other.

Given this finding, we also considered whether we might learn something from *atypical* posts. Prior work suggests that privacy behaviors may be punctuated by life events [51]. Therefore, if privacy settings do not generally vary systematically with topic, then perhaps there is a pattern to posts where a privacy setting appears to be purposefully changed. To test this theory, we examined users who have a “typical” privacy setting (more than half public or more than half non-public) and then considered those posts that diverge from this norm. Out of our set of 3,000 topic-coded posts, fewer than 100 fell into this category. Of these, it is more common for a user to switch from non-public to public, than vice versa. However, there was no significant difference in content *topic* or *type* distribution among these groups.

Thus, it is more informative to focus on the author of the post, rather than on the content itself, when trying to model which privacy settings will be selected. To eliminate the possibility that our qualitatively derived codebook might be insufficient to capture nuanced linguistic distinctions for content topics or types, we used machine learning and classification techniques to inspect the dataset one word at a time (unigrams model), as well as one phrase at a time (using two-word phrases, or bigrams model). We describe this work in the next section.

Result: N-gram Models only Marginally Improve upon the Previously Presented Codebook Models

We next investigated whether there might be other potential content effects. Could nuanced linguistic constructs of posts, beyond the content codes, improve our predictions of privacy choices? In some sense, by doing this we ask whether there exists some codebook (not necessarily ours) that would better capture privacy setting correlations. Focusing on general linguistic constructs also allows us to examine whether sentiment or similar attributes of content could be predictive of an individual’s privacy settings. We built two different language models utilizing the *n*-grams technique of posts ($n=2$; unigrams and bigrams), the demographic attributes of the post authors, and the length of the posts. The first model focused on post-level predictions (i.e., whether the *n*-grams of a post could predict its privacy

setting). The second model involved user-level predictions—whether the n -grams of all six posts from each participant could predict the chosen privacy setting (for participants who used modal privacy settings, all public or all non-public). Note that to control for user-centric dependencies on post authorship, as above, we included a categorical variable for the *WorkerID*.

Since both of these predictions involve binary classification, we used regularized logistic regression as in the previous section, with regularization controlling for collinearity and sparsity. Because of an unequal number of posts and users in the public and private categories for the two prediction tasks (fewer public than private posts/users), we randomly sampled from the private post/user pool to obtain an equal sized set as the public class. On these balanced samples, we performed k -fold cross validation ($k=10$) for model fitting and evaluation.

Post-level prediction				
	Precision	Recall	F1-score	Accuracy
non-public (0)	0.60	0.73	0.66	73.30%
public (1)	0.66	0.52	0.58	51.60%
average	0.63	0.62	0.62	62.50%
User-level prediction				
private (0)	0.60	0.86	0.71	85.70%
public (1)	0.75	0.43	0.55	57.10%
average	0.68	0.64	0.63	64.20%

Table 5. Performance of two logistic regression models in predicting post-level and user-level privacy choices. Numbers reported are aggregated across all 10 folds on heldout data.

Model fits for both the prediction tasks using n -grams results in significant reduction of deviance values compared to the null model (or the intercept only model) for all ten folds. On average, the differences are statistically significant: $X^2(N=1,072, df=216) = 1572.3 - 861.3 = 711, p < 10^{-8}$ for the post-level logit model and $X^2(N=620, df=200) = 5.13e+17 - 1.05e+05 = 5.13e+17, p < 10^{-10}$ for the user-level regression. However, for the heldout set (10% of the balanced samples for the two prediction tasks), we did not see noticeable improvement in performance over the baseline change model. We report the results of our two models on heldout data in Table 5, combined across the 10 folds. Performance is reported in terms of a number of metrics including accuracy, precision, recall, F1-score, and AUC (area under curve). The mean accuracy for the post-level prediction model is 62.5%, which improves on a chance model by a small margin: 12.5%, whereas that of the user-level model is 64.3%, an improvement of 14.3% over the chance model (since our class sizes are equal, the accuracy of the chance model is 50%). In other words, these gains are small in real terms, as observed above. We also find that prediction performance for the private class (posts/users) is consistently relatively better than those of the public class. This might indicate an inherent bias in our participant pool toward sharing content tailored more towards less public sharing. It is also worth noting that the second set of predictions (user-level) are marginally better

in terms of accuracy. This again suggests that privacy choices are driven more by the attributes of the *person*, rather than by the *content* of posts (confirming our earlier result). We further explore this idea in the next section.

Result: Some User Characteristics Vary Strongly with Modal Privacy Settings

Having established that codes provide little information in modeling privacy setting, and that keying in on the author is very informative for that purpose, we next set out to determine whether reported demographic characteristics of (rather than posts) have predictable privacy posting behavior: that is, we now shift to users rather than posts as the unit of analysis. With content codes not relevant for this analysis, we moved to our full dataset over nearly 11,000 posts. For each user, we computed the *modal privacy setting*: the majority label they gave to their posts. For example, if a person gave us 5 public and 1 non-public posts, we will call them *mostly public* for the purposes of the analyses that follow. Though six posts may represent only a “snapshot” of a user’s privacy behavior, it provides sufficient enough context to characterize the interaction between demographics and typical privacy settings.

Table 6 presents a user-level, penalized logistic regression modeling the dependent variable *mostly public* vs. *mostly private* from the same demographic independent variables discussed in our demographics section above. Note that penalized models offset intra-correlations by including only those variables that add information above and beyond other variables; however, our focus here is on the omnibus predictive power of demographic variables relative to content ones. We see that this user-focused model, unlike the purely content-focused models above, significantly informs modal privacy setting: $X^2(N=1,706, df=37) = 1,275, p < 10^{-10}$. In particular, female participants are much more likely to usually post non-public ($\beta = -0.84, p = 0.001$); on the other hand, older users of Facebook (those 65 and above) are much more likely to usually post public ($\beta = 2.97, p = 0.019$). We see that this user-focused model, unlike the purely content-focused models above, significantly informs modal privacy setting: $X^2(N=1,706, df=36) = 1,273, p < 10^{-10}$. In particular, female participants are much more likely to usually post non-public ($\beta = -0.81, p = 0.001$); on the other hand, older users of Facebook (those 65 and above) are much more likely to usually post public ($\beta = 3.04, p = 0.016$).

The additional “default: changed” variable is based on their answer to the question of whether they have changed their privacy settings from the default. The significance for that variable suggests that users who change the default privacy settings are more likely move towards being *less* public rather than *more* public. Interestingly, and perhaps surprisingly, we find that increased Facebook skill is related to a higher likelihood of posting publicly, albeit marginally ($\beta = 0.37, p = 0.039$).

	β	std err	z	p
(intercept)	-11.9	623	-0.02	0.985
default: changed	-0.65	0.31	-2.08	0.037
gender: female	-0.81	0.24	-3.31	0.001
gender: omitted	-0.34	1.65	-0.2	0.838
ethnicity: asian	10.4	623	0.02	0.987
ethnicity: black	8.37	623	0.01	0.989
ethnicity: latino	9.1	623	0.01	0.988
ethnicity: multi	9.21	623	0.01	0.988
ethnicity: native	10.65	623	0.02	0.986
ethnicity: white	9.05	623	0.01	0.988
ethnicity: other	9.15	623	0.01	0.988
ethnicity: omitted	9.03	623	0.01	0.988
edu: < high school	2.46	1.48	1.66	0.098
edu: high school	-0.12	0.42	-0.29	0.775
edu: some college	0.46	0.27	1.68	0.094
edu: post grad	-0.27	0.35	-0.77	0.441
edu: vocational	0.18	0.55	0.33	0.745
edu: omitted	2.22	1.52	1.46	0.145
age: 25-34	0.04	0.28	0.13	0.895
age: 35-44	0.66	0.37	1.8	0.073
age: 45-54	0.71	0.49	1.46	0.145
age: 55-64	-0.31	0.74	-0.42	0.676
age: 65+	3.04	1.27	2.4	0.016
facebook skill scale	0.35	0.18	1.96	0.049
internet skill scale	-0.02	0.15	-0.12	0.907
fb intensity scale	0.17	0.12	1.45	0.147
number of friends	0.0003	0	-1.11	0.267
N	1,706			
null deviance	1,882	df	1,705	
residual deviance	608.2	df	1,669	
$X^2(N=1,709, df=36)$	1,273	p	$< 10^{-10}$	

Table 6. Logistic regression predicting participants’ modal post-level privacy setting as a function of demographic and Internet use questions.

It is also interesting to note that although we did find a significant correlation between a user’s number of friends and how many public posts they have ($\rho = 0.094$, $p = 10^{-3}$), number of friends is not a significant factor in the regression model. As we previously noted, the concept of how “private” something is does relate to anticipated audience size [26]. A more in-depth examination of how privacy settings might change with the size one’s network is an intriguing topic for future research.

DISCUSSION

Our motivating research questions for this work center on whether there are systematic differences between public and non-public content on Facebook. Could there be certain *content types* or *topics* that tend to be posted publicly versus non-publicly and/or certain *characteristics of users* who tend to post publicly versus non-publicly? To explore these questions, we first completed a rigorous qualitative analysis of Facebook content on a subset of our data, and then used quantitative and machine learning techniques to explore these questions on a larger dataset.

Our findings show that in this dataset content type or topic has little to no predictive power for privacy setting. This finding is strengthened by the n -gram models that do not improve upon the information provided by our codebook. In other words, even if our codebook is an imperfect or a specific type of representation of Facebook content, *there is likely not some other codebook scheme* that would have provided us with more information about privacy settings.

This finding does seem to contradict some prior work that suggested differences in self-disclosure based on content. For example, Madejski et al.’s survey of privacy attitudes showed that Facebook users wanted to be less private when discussing sensitive topics such as sex and alcohol [32]. However, it is reasonable that findings based on actual posts might differ from those based on self-disclosure of *intention*, since intention could differ from behavior. In other words, people *say* they are concerned about disclosing such topics, but the *observed* data suggests otherwise. Alternately, certain types of posts may be rare enough that even if people have *intention* for their level of privacy for that type of post, it would not appear in our data set. Regardless of the cause, our findings more closely model “big data” approaches to studying public Facebook content.

However, beyond content, our models do show us that at a *user* level we can make some predictions about whose content might be public based on demographic characteristics. This suggests that though a dataset of public Facebook posts may include similar content as a set of all Facebook posts, the user base represented may be different. Supporting prior work based on self-reports [31,32], female users are less likely to use public Facebook settings, and older users more likely.

Another partial explanation for our findings could be that, following prior work regarding privacy strategies around context collapse on social media [19,44,47], users are more often controlling privacy with self-censorship and friends-list management rather than privacy settings. Because we examined objective privacy settings rather than privacy attitudes, we cannot make any claims about privacy behavior more generally, or privacy strategies *outside of* privacy settings that users could be employing. However, our findings do show that users largely stick to their own default settings rather than frequently exercising more granular control over privacy settings for individual posts.

Only about a quarter of the content in our dataset is public, and so though most of their Facebook content is not public, 82% of our participants agreed to have their data included in a public dataset for the purposes of further research. This public data is not substantially different than our full dataset, in either demographics or privacy settings. This suggests that at the very least for our population (which is non-representative in the ways discussed in our demographics section) users are reasonably comfortable sharing anonymized data.

Our content analysis also allowed us to characterize Facebook content, since we have a set consisting of both public and non-public posts. Rather than finding systematic content differences based on privacy setting as we expected, we found that content is similar across public and non-public posts. However, there are some demographic variables predict privacy settings. Therefore, we suggest that researchers studying public Facebook data should attend to the demographics of their users, since a set of public posts could be systematically leaving out some of the user base, particularly around gender and age.

Limitations and Future Research

One implication of this work is that building a predictor for privacy settings appears to be a difficult task, since privacy setting choices may not be driven by the type of content they share. Though our findings also suggest some differences in user base for public versus non-public content, introducing demographics for prediction has the danger of creating bias as well. Therefore, we suggest that privacy-supportive technologies should be particularly sensitive to context.

Some caution is also warranted in interpreting the results of this study. The dataset we collected and analyzed in this paper is large; however, it may not be representative of the general Facebook population or may contain some bias in the ways discussed in our limitations and data quality section. An interesting direction for future work would be to see how our findings generalize to a dataset with a larger set of contributed posts of participants and spreading over a longer timeframe, to examine if privacy settings change due to content shifts over time. For example, it may be that switching privacy settings is a relatively rare event, and would require a much larger window of posts than the six we collected here (i.e., that some privacy setting changes may have slipped away), or that there are effects of posts being close together in time. Furthermore, though there is a danger of response bias in any dataset obtained with consent, our checks for truthfulness suggest that our contributed data reflects our participants' actual Facebook use, at least in the limited context of the six posts they shared with us.

Prior work in this area has also considered other factors that might influence privacy choices, beyond content topic or types—for example, self reports of goals and intimacy [8] or language sentiment [5]. Considering these or other

factors for message-level privacy settings could yield differences or insights that our measure for this study did not.

Finally, our results lend preliminary surface-level insights into *why* participants chose one privacy setting over the other. Because for this particular study we were concerned with observable content measures rather than intentionality, there is a lack of context around other privacy considerations such as perceived audience or privacy from other parties such as advertisers or surveillance. In order to develop deeper insights into user motivations, future work might include qualitative techniques such as interviewing. However, we hope the current study provides insight for other researchers exploring this area.

CONCLUSION

This research provides three main contributions: first, findings about the *nature* of Facebook content. Our content analysis as information about privacy settings (both self-reported and from observed data) adds to the current body of literature around social media content and privacy behavior. Second, our data suggests that the content of public Facebook posts may not be significantly different than content of the entirety (public and non-public) of Facebook posts. However, the user base in a set of public posts could be demographically different. Our findings here are important for researchers studying public Facebook data, as it provides information about how their data set may be similar or different from the whole of public-and-private content.

The third contribution of this work is our dataset. We have checked that all posts are truly anonymous (with human readers reviewing the dataset), removing proper nouns and other identifying details. The dataset will be available on our website with supporting documentation for use by other researchers, and we hope this will make new research possible for others.

ACKNOWLEDGMENTS

Our thanks to the reviewers and our colleagues whose feedback helped to make this a considerably stronger paper. We are also grateful to our participants who allowed us to include their data in our public dataset, which we hope will be beneficial for future research. It is available at <https://github.com/compsocial/WhatWhoCSCW2017>.

REFERENCES

1. Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221: 509–514.
2. Lars Backstrom and Jon Kleinberg. 2014. Romantic partnerships and the dispersion of social ties: A network analysis of relationship status on facebook. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
3. Lars Backstrom, Jon Kleinberg, Lillian Lee, and Cristian Danescu-Niculescu-Mizil. 2013. Characterizing and

- curating conversation threads: expansion, focus, volume, re-entry. *Proceedings of the ACM International Conference on Web Search and Data Mining*.
4. Natalya N. Bazarova and Yoon Hyung Choi. 2014. Self-Disclosure in Social Media: Extending the Functional Approach to Disclosure Motivations and Characteristics on Social Network Sites. *Journal of Communication* 64, 635–657.
 5. Natalya N. Bazarova, Jessie G. Taft, Yoon Hyung Choi, and Dan Cosley. 2012. Managing Impressions and Relationships on Facebook: Self-Presentational and Relational Concerns Revealed Through the Analysis of Language Style. *Journal of Language and Social Psychology* 32, 2, 121–141.
 6. Michael Buhrmester, Tracey Kwang, and Samuel D. Gosling. 2011. Amazon’s Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science* 6, 1: 3–5.
 7. Kathy Charmaz. 2006. *Constructing Grounded Theory*. Sage Publications, Thousand Oaks, CA.
 8. Yoon Hyung Choi and Natalya N. Bazarova. 2014. Self-disclosure characteristics and motivations in social media: Extending the functional model to multiple social network sites. *Human Communication Research* 41, 480–500.
 9. Nicole B. Ellison, Rebecca Gray, Cliff Lampe, and Andrew T. Fiore. 2014. Social capital and resource requests on Facebook. *New Media & Society* 16, 7: 1104–1121.
 10. Nicole B. Ellison, Charles Steinfield, and Cliff Lampe. 2007. The benefits of Facebook “friends:” Social capital and college students’ use of online social network sites. *Journal of Computer-Mediated Communication* 12, 4.
 11. Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. *Proceedings of the International World Wide Web Conference*, ACM Press, 351.
 12. Casey Fiesler, Cliff Lampe, and Amy S. Bruckman. 2016. Reality and Perception of Copyright Terms of Service for Online Content Creation. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
 13. Kambiz Ghazinour, Stan Matwin, and Marina Sokolova. 2013. Monitoring and recommending privacy settings in social networks. *Proceedings of the Joint EDBT/ICDT 2013 Workshops*.
 14. Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
 15. Joseph K. Goodman, Cynthia E. Cryder, and Amar Cheema. 2013. Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples: Data Collection in a Flat World. *Journal of Behavioral Decision Making* 26, 3: 213–224.
 16. Eszter Hargittai and Yuli Patrick Hsieh. 2012. Succinct Survey Measures of Web-Use Skills. *Social Science Computer Review* 30, 1: 95–107.
 17. Bernie Hogan. 2010. The Presentation of Self in the Age of Social Media: Distinguishing Performances and Exhibitions Online. *Bulletin of Science, Technology & Society* 30, 6: 377–386.
 18. Chris Jay Hoofnagle and Jennifer M. Urban. 2014. Alan Westin’s Privacy Homo Economicus. *Wake Forest Law Review* 49, 1: 261–317.
 19. Maritza Johnson, Serge Egelman, and Steven M. Bellovin. 2012. Facebook and privacy: it’s complicated. *Symposium on Usable Privacy and Security (SOUPS)*.
 20. Bumsoo Kang, Sujin Lee, Alice Oh, Seungwoo Kang, Inseok Hwang, and Junehwa Song. 2015. Towards Understanding Relational Orientation: Attachment Theory and Facebook Activities. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
 21. Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. *Proceedings of the Tenth Symposium On Usable Privacy and Security (SOUPS)*.
 22. Michael Kosinski, David Stillwell, and Thore Graepel. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences* 110, 15: 5802–5805.
 23. Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. *Privacy Indexes: A Survey of Westin’s Studies*. Carnegie Mellon University, Institute for Software Research International (ISRI).
 24. Cliff Lampe, Rebecca Gray, Andrew T. Fiore, and Nicole Ellison. 2014. Help is on the way: patterns of responses to resource requests on facebook. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
 25. J. Richard Landis and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics* 33, 1: 159.
 26. Mark R. Leary and Robin M. Kowalski. 1990. Impression management: A literature review and two-component model. *Psychological Bulletin* 107, 1, 34–47.
 27. Eden Litt, Erin Spottswood, Jeremy Birnholtz, Jeff T. Hancock, Madeline E. Smith, and Lindsay Reynolds. 2014. Awkward encounters of an “other” kind: collective self-presentation and face threat on facebook. *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMS)*.
 28. Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing facebook privacy settings: user expectations vs. reality. *Proceedings of the AAAI International Conference on Weblogs and Social Media (ICWSM)*.

29. Yabing Liu, C Kliman-Silver, and Alan Mislove. 2014. The Tweets They are a-Changin’: Evolution of Twitter Users and Behavior. *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media (ICWSM)*.
30. K. M. MacQueen, E. McLellan, K. Kay, and B. Milstein. 1998. Codebook Development for Team-Based Qualitative Analysis. *Field Methods* 10, 2: 31–36.
31. Mary Madden. 2012. *Privacy management on social media sites*. Pew Research Center’s Internet & American Life Project, Washington D.C.
32. Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. 2012. A study of privacy settings errors in an online social network. *International Conference on Pervasive Computing and Communications*, IEEE, 340–345.
33. Tanushree Mitra, C. J. Hutto, and Eric Gilbert. 2015. Comparing Person- and Process-centric Strategies for Obtaining Quality Data on Amazon Mechanical Turk. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*.
34. Meredith Ringel Morris, Jaime Teevan, and Katrina Panovich. 2010. What do people ask their social networks, and why? A survey study of status message Q&A behavior. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
35. Mor Naaman, Jeffrey Boase, and Chih-Hui Lai. 2010. Is it really about me? Message content in social awareness streams. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
36. Melissa Niiya, Stephanie M. Reich, Yiran Wang, Gloria Mark, and Mark Warschauer. 2015. Strictly by the Facebook: Unobtrusive Method for Differentiating Users. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
37. Gabriele Paolacci, Jesse Chandler, and Panagiotis G. Ipeirotis. 2010. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5, 5: 411–419.
38. Aaron D. Shaw, John J. Horton, and Daniel L. Chen. 2011. Designing incentives for inexpert human raters. *Proceedings of the ACM 2011 conference on Computer supported cooperative work (CSCW)*.
39. Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013. The post that wasn’t: exploring self-censorship on facebook. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
40. Jaakko Stenros, Janne Paavilainen, and Jani Kinnunen. 2011. Giving good “face”: playful performances of self in Facebook. *Proceedings of the International Academic MindTrek Conference on Envisioning Future Media Environments*, 153–160.
41. Fred Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. *Journal of Privacy and Confidentiality* 4, 2: 7–41.
42. Fred Stutzman and Jacob Kramer-Duffield. 2010. Friends only: examining a privacy-enhancing behavior in facebook. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
43. Zeynep Tufekci. 2014. Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls. *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)*.
44. Jessica Vitak and Nicole B. Ellison. 2013. “There’s a network out there you might as well tap’: Exploring the benefits of and barriers to exchanging informational and support-based resources on Facebook. *New Media & Society* 15, 2: 243–259.
45. Jessica Vitak, Stacy Blasiola, Sameer Patil, and Eden Litt. 2015. Balancing Audience and Privacy Tensions on Social Network Sites. *International Journal of Communication* 9, 1485–1504.
46. Jessica Vitak and Jinyoung Kim. 2014. “You can’t block people offline”: examining how facebook’s affordances shape the disclosure process. *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
47. Jessica Vitak, Cliff Lampe, Rebecca Gray, and Nicole B. Ellison. 2012. “Why won’t you be my Facebook friend?”: strategies for managing context collapse in the workplace. *Proceedings of the International iConference*.
48. Jessica Vitak, Katie Shilton, and Z. Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*.
49. Pamela Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2014. Profiling Facebook Users’ Privacy Behaviors. *Workshop on Privacy Personas and Segmentation*.
50. Alyson L. Young and Anabel Quan-Haase. 2009. Information revelation and internet privacy concerns on social network sites: a case study of facebook. *Proceedings of the International Conference on Communities and Technologies*.
51. Xuan Zhao, Niloufar Salehi, Sasha Naranjit, Sara Alwaalan, Stephen Volda, and Dan Cosley. 2013. The many faces of facebook: experiencing social media as performance, exhibition, and personal archive. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.