

“Participant” Perceptions of Twitter Research Ethics

Casey Fiesler¹ and Nicholas Proferes² 

Social Media + Society
January-March 2018: 1–14
© The Author(s) 2018
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/2056305118763366
journals.sagepub.com/home/sms


Abstract

Social computing systems such as Twitter present new research sites that have provided billions of data points to researchers. However, the availability of public social media data has also presented ethical challenges. As the research community works to create ethical norms, we should be considering users' concerns as well. With this in mind, we report on an exploratory survey of Twitter users' perceptions of the use of tweets in research. Within our survey sample, few users were previously aware that their public tweets could be used by researchers, and the majority felt that researchers should not be able to use tweets without consent. However, we find that these attitudes are highly contextual, depending on factors such as how the research is conducted or disseminated, who is conducting it, and what the study is about. The findings of this study point to potential best practices for researchers conducting observation and analysis of public data.

Keywords

Twitter, Internet research ethics, social media, user studies

Introduction

In recent years, research ethics has become a topic of greater public scrutiny. This is particularly the case for research taking place on or using data from social computing systems (McNeal, 2014; Wood, 2014; Zimmer, 2010a). However, as Vitak, Shilton, and Ashktorab (2016) point out in a study of ethical practices and beliefs among online data researchers, there are not agreed upon norms or best practices in this space. Organizations such as the Association of Internet Researchers (AOIR) have offered guides for researchers (Ess & AOIR Ethics Working Committee, 2002), although these guidelines have had to be revised as new platforms present novel research contexts and as data collection practices evolve alongside them (Markham, Buchanan, & AOIR Ethics Working Committee, 2012). It is also rare that the people whose data are being studied are involved in the development of such guidelines. However, public reaction to research ethics controversies (McNeal, 2014; Wood, 2014; Zimmer, 2016) suggests that these individuals have opinions about how researchers should study and use online systems and data. Brown, Weilenmann, Mcmillan, and Lampinen (2016) suggest that research ethics should be grounded in “the sensitivities of those being studied” and “everyday practice” as opposed to bureaucratic or legal concerns. We therefore suggest users can help inform ethical research practices.

Traditional interventional human subjects research involves informed consent and direct interaction with participants who are aware they are being studied. Under certain regulatory conditions, research that falls outside of this scope may not be strongly regulated. For example, in the United States, the use of publicly available data (e.g., tweets) may not meet the criteria of research involving human subjects as per the Code of Federal Regulations (45 CFR 46.101). However, there is a lack of consensus among individual university institutional review boards (IRBs) on this point (Vitak, Proferes, Shilton, & Ashktorab, 2017). Data obtained from sources such as Twitter most often do not constitute research that requires their oversight or informed consent practices. Most researchers who use data sets of tweets do not gain consent from each Twitter user whose tweet is collected, nor are those users typically given notice by the researcher. Although Twitter's Privacy Policy now mentions that academics may use tweets as part of research, this update was

¹University of Colorado Boulder, USA

²University of Kentucky, USA

Corresponding Author:

Nicholas Proferes, School of Information Science, University of Kentucky, 352 Lucille Little Library, 160 Patterson Drive, Lexington, KY 40506, USA.
Email: nproferes@uky.edu



not included until revisions were made to the policy in the Fall of 2014.¹ Moreover, Internet users rarely read or could fully understand website terms and conditions (Fiesler, Lampe, & Bruckman, 2016; Luger, Moran, & Rodden, 2013; Reidenberg, Breaux, Cranor, & French, 2015).

Although there are a great many open questions around research ethics, including around experimental manipulation on online platforms (Schechter & Bravo-Lillo, 2014), for this study, we focus specifically on the issue of researchers' use of public social media content. Because of the prominent use of Twitter data in research (Zimmer & Proferes, 2014), in this study, we ask: how do Twitter users feel about the use of their tweets in research?

Our goal with this work is not to suggest whether or not use of Twitter data is ethical or not based entirely on user attitudes. Instead, we aim to inform ethical decision-making by researchers and regulatory bodies by reporting on how user expectations align with the actual uses of their data by researchers. Therefore, we designed our survey instrument to probe *contextual factors* that impact whether Twitter users find studying their content acceptable. This study is exploratory, providing initial insights into a complicated space.

In brief, the majority of Twitter users in our study do not realize that researchers make use of tweets, and a majority also believe researchers should not be able to do so without permission. However, these attitudes are highly contextual, differing based on factors such as how the research is conducted or disseminated, who the researchers are, and what the study is about. After describing the study and findings, we conclude with a discussion of what these findings might suggest for best practices for Twitter research and offer potential design interventions that could help mitigate some user concerns.

Background

Research Ethics for Social Computing

Globally, policy and guidelines around human subjects research come from a variety of sources, such as the Declaration of Helsinki (World Medical Association, 2013). In the United States, across all disciplines including computer science and social sciences, a significant source of ethical standards is the Belmont Report, created in 1979 in the aftermath of the Tuskegee syphilis experiment (Bruckman, 2014; Office for Human Research Protections, 1978). The report laid out three guiding principles: (1) respect for persons; (2) beneficence (minimizing harm); and (3) justice (Office for Human Research Protections, 1978). Research ethics governance varies in different contexts (e.g., internationally or for industry), but US universities that accept federal funding are required to maintain IRBs where human subjects research must go through an approval process.

The Internet introduced a host of additional complications to research ethics around issues such as anonymity and confidentiality (Bruckman, 2002; Zimmer, 2010a), informed consent (Barocas & Nissenbaum, 2014; Hudson & Bruckman, 2004), terms of service (TOS) (Fiesler et al., 2016; Vaccaro, Karahalios, Sandvig, Hamilton, & Langbort, 2015), relationships with and dissemination of findings to participants (Beaulieu & Estallea, 2012), and the definition of public spaces and public data (Bromseth, 2002; Hudson & Bruckman, 2004; Zimmer, 2010a). The last is particularly important to this study because one of the major challenges for researchers is even determining when a project constitutes human subjects research. Under US federal regulatory definitions, studying public information on the Internet is typically not considered "human subjects research," which is characterized by an investigator obtaining "data through intervention or interaction with the individual" or "identifiable private information" (Bruckman, 2014). Therefore, many IRBs consider the study of public data to not be under their purview (Bruckman, 2014; Vitak et al., 2016). A study of IRB attitudes toward social computing research revealed that a majority of IRB respondents feel that researchers do not need informed consent to collect public data (Vitak et al., 2017). Similarly, an analysis of several hundred published papers using Twitter data noted that very few of these studies discuss undergoing ethics review (Zimmer & Proferes, 2014).

However, regardless of regulatory standards (or lack thereof), there are still potential harms that can occur in social computing research (Collmann, Fitzgerald, Wu, Kupersmith, & Matei, 2016). For example, Crawford and Finn (2015) point out that Twitter users may publicly share personal, identifying information of themselves or others during crisis events as a way to find assistance and support. Another early controversy surrounding the ethics of using and disseminating public data was around the "Taste, Ties, and Time" paper that released profile information from Facebook users (Lewis, Kaufman, Gonzalez, Wimmer, & Christakis, 2008). Despite good faith efforts to protect the privacy of those users, they were quickly identified, putting their privacy at risk (Zimmer, 2010a). Furthermore, harm from online research can occur not just at the individual level but to classes of people and communities (Hoffmann & Jones, 2016).

Another point of disagreement regarding privacy is the potential for content quoted verbatim to be tied back to the person who created it. In some early proposed ethical guidelines for Internet research, Bruckman (2002) put forth levels of "disguise" for using online content. For example, "light disguise" includes quoting content but does not include usernames or pseudonyms; she points out, however, that although it would take some effort, someone could unmask the content creator's identity (e.g., through a search engine). This re-identification could happen with a public tweet as well. In at least one case, a journalist in reporting on a research study

searched for a tweet quoted in the paper and contacted the Twitter user, who was unaware of the study (Singer, 2015).

In recent years, there have also been a number of research ethics controversies that gained public attention, such as the “emotional contagion” experiment that studied Facebook users without their consent (Kramer, Guillory, & Hancock, 2014; McNeal, 2014) and the public release of non-anonymized OkCupid data scraped by a researcher against the site’s TOS (Zimmer, 2016). Although the Facebook study involved direct experimentation rather than collection of existing public data, the public outcry highlights the need for user voices in consideration of ethics (McNeal, 2014; Schechter & Bravo-Lillo, 2014).

Although participatory research with respect to these kinds of decisions is more common with respect to privacy (De Cristofaro & Soriente, 2013; Martin & Shilton, 2015), there has been some prior work that considers user attitudes toward research ethics. For example, M. L. Williams, Burnap, and Sloan (2017) conducted a survey of British Twitter users about how they feel about tweets being published verbatim by researchers along with other actor such as the government; they found that users are most comfortable when asked for consent and when the tweets are anonymized. There has also been domain-specific research, for example, in medical ethics. In a study of user attitudes toward using Twitter data to monitor depression, users were far more comfortable with aggregate level health monitoring than assessment of individuals (Mikal, Hurst, & Conway, 2016). These are the kinds of contextual factors that our study examines with respect to Twitter research more broadly.

In addition, more public scrutiny toward ethics has inspired increasing discussions within the online research community, which in turn has highlighted the lack of consistent norms. Vitak et al. (2016) found that much of this disagreement is around how data subjects should be treated and represented, and similarly, Brown et al. (2016) emphasize the need for understanding the definition of *harm* in ethics. In addition to research that considers these norms and harms, there are also broader initiatives dedicated to helping researchers navigate increasingly complex ethical situations. For example, the Connected and Open Research Ethics (CORE) initiative in part connects researchers and regulatory bodies to help create guidelines, particularly, in the context of mobile health (Nebeker et al., 2017).

Research, Audience, and Privacy on Twitter

Twitter is an appropriate venue for considering use of public data due to its prominence in social computing research. This popularity is partly due to availability. For example, although Facebook has many more users than Twitter, Facebook has considerably less public data and has stricter limits on its application programming interface (API; Fiesler et al., 2017). Tufekci calls Twitter the “model organism” of big data. Researchers study Twitter partly because (as with the fruit

fly) it is easy to study; the platform makes data easily available (Tufekci, 2014).

Twitter in turn has proven itself a highly successful platform for research, resulting in studies and knowledge gained on topics such as disease tracking (Lamb, Paul, & Dredze, 2013; Signorini, Segre, & Polgreen, 2011), disaster events (Kogan, Palen, & Anderson, 2016; Sakaki, Okazaki, & Matsuo, 2010; Starbird & Palen, 2011), and stock market and election prediction (Bollen, Mao, & Zeng, 2011; Tumasjan, Sprenger, Sandner, & Welpe, 2010). This “data gold rush” has resulted in researchers using Twitter content (both tweets and profile information) to examine all sorts of aspects of human interaction (Felt, 2016; S. Williams, Tarras, & Warwick, 2013; Zimmer & Proferes, 2014).

However, prior work suggests that many Twitter users have limited understandings of how their public content might be used. Proferes (2017) found that users are broadly unaware of the fact that APIs exist, that Twitter sells access to tweets via the “firehose,” and that the Library of Congress archives tweets. As with many social media platforms, Twitter users also often have an “imagined audience” that may not match to reality (Bernstein, Bakshy, Burke, & Karrer, 2013; Litt, 2012; Marwick & Boyd, 2011). Misalignment of imagined and actual audience can be dangerous when there are more eyes to catch a potential faux pas (Litt, 2012). Moreover, some users may not realize that their tweets are publicly viewable at all (Proferes, 2017), and even users who tweet under protected accounts (which limit Tweet visibility to approved followers only) may still endure privacy violations through re-tweets (Meeder, Tam, Kelley, & Cranor, 2010). Privacy is also an issue on Twitter due to accidental information leaks by users, such as divulging vacation plans, tweeting under the influence, or revealing medical conditions (Mao, Shuai, & Kapadia, 2011).

As we will explore in more detail with our survey findings, issues of privacy are heavily entwined with research ethics. As Vitak et al. (2016) note, weighing privacy risks is part of the guiding research value of beneficence from the Belmont Report. However, attitudes toward privacy, especially on social media, are highly context-dependent. In proposing “contextual integrity” as a benchmark for privacy, Nissenbaum (2004) emphasizes that information gathering and dissemination should be appropriate to context. For example, despite the dominant norm among researchers that Twitter is public, the consent of a user for the public to view their tweets may not imply consent for them to be collected and analyzed by researchers (Zimmer, 2010b). In addition, Twitter users’ expectations for privacy could be very different based on factors such as how long they have been engaged with the platform or how they compare it to other platforms or contexts (Martin & Shilton, 2015).

These open questions, points of controversy, and subjects of debate within the social computing research community led to our conclusion that to move forward with norm-setting, we need more data—which include an empirical understanding

of how the “participants” in this type of research understand it and feel that it may impact them.

Methods

This study elicited Twitter users’ responses to hypothetical situations involving the use of tweets for academic research. As this is an area that lacks extensive prior work, we used an exploratory approach. In exploratory surveys, the research question remains open-ended without a specific testable hypothesis driving the study (Adams, 1989). Instead, research proceeds through inductive analysis using descriptive statistical analysis and inferential statistical analysis techniques to explore response trends.

Population of Interest

This study’s primary interest is in Twitter users. A challenge in studying this population is that true random sampling is difficult, and furthermore, recruitment presents bias toward those interested in participating in research. In fact, any method for recruiting participants to study how people feel about research would be necessarily and systematically biased. Logically, those willing to participate in a research study may have different attitudes about research than those who are not.

We therefore turned to a recruitment method for which we had a clearer idea for the type of potential bias we could encounter: Amazon Mechanical Turk (MTurk).² MTurk is a crowdwork system where workers (or “turkers”) complete small tasks for micro-payments (Hitlin, 2016). Because it is a platform for general crowdwork, turkers may complete research tasks as part of their work but are likely not there for the express goal of participating in research. They have the additional extrinsic motivation of completing more tasks and being paid. We also speculated that if Twitter users familiar with academic research are hesitant about the idea of being studied, this will tell us something interesting about general trends that would give us important directions for further research. Although MTurk has its drawbacks, turkers have performed comparably to laboratory subjects in traditional experiments (Paolacci, Chandler, & Ipeirotis, 2010). A known limitation is that participants may be less likely to pay attention to experimental materials (Goodman, Cryder, & Cheema, 2013), although this can be somewhat mitigated by the use of “attention checks” such as the one we will describe in our survey design (Hauser & Schwarz, 2016).

With respect to generalizability, research comparing MTurk and representative panels found that MTurk responses for users aged 18–49 years with at least some college were largely representative of the general US population (Redmiles, Kross, Pradhan, & Mazurek, 2017). Turkers do skew younger, less ethnically diverse, and less educated than all American working adults (Hitlin, 2016), although a study comparing demographics of Twitter users to the general

adult population showed similar patterns, with the exception of education levels (Blank, 2016). Twitter users overall tend to be more educated than the general population, although as reported below, 67.9% of our participants reported some college education compared to 69% of American Twitter users reported in prior work (Blank, 2016). In sum, although there are some demographic similarities between MTurk and Twitter, we cannot say how representative our sample may be of all of Twitter (particularly, beyond English-speaking Twitter users), and our results should be interpreted with this in mind.

Survey Instrument

The survey instrument begins by asking participants about demographic characteristics and Twitter use. Participants provided information about how often they tweet, when they first signed up for Twitter, and an open-ended question about how they perceive their Twitter audience. Next, the survey provides some information about how researchers use tweets as part of research, using the following language:

Sometimes scientists and other academic researchers conduct studies using Twitter. Sometimes they study tweets in order to understand things like how people communicate within social groups, what people are saying about a given subject, or are trying to predict things like changes in the stock market. Most often, the people whose tweets are used are never told.

The survey then asked respondents to consider different contextual factors and rate their level of comfort with the scenario on a Likert-type scale. Factors included, for example, how many tweets were in the dataset, whether or not users whose tweets were used were told about the study, whether the tweets would be quoted, whether the tweets used had since been deleted by the users, and so on (see Table 3 for a full list). The survey then provided respondents with a number of qualitative open-text responses to allow them to expand on their answers.

The survey questions, and, in particular, the choice of these contextual factors, were based on the tension points around use of public data that have surfaced in prior work (Shilton & Sayles, 2016; Vitak et al., 2016) and workshop discussions (Fiesler et al., 2015). Next, the survey asked respondents about whether they think researchers must receive different kinds of permission to use data from Twitter. Finally, it questioned whether respondents had been aware before they took the survey that researchers use tweets in academic research and whether, now that they had taken the survey and knew for sure, they would opt out of having their tweets used for research if given the option.

Participants were allowed to skip questions, with the exception of the initial informed consent. To improve reliability and validity, we used an attention check question in the latter third of the survey. Attention checks are a common

Table 1. Twitter Profile Statistics of Sample.

I.I.I.I.I.I.I.I	M	SD	High	Low
Tweets	2,107.2	7,319.6	79,500	0
Following	372.1	1,039.4	11,500	0
Followers	782.4	4,956.5	58,800	0
Influence score (following/followers)	1.6	6.2	98.6	0
Likes	869.3	3,418.7	46,000	0
Photos and videos	134.3	915.1	13,500	0

M: mean; SD: standard deviation.

mechanism in MTurk surveys to ensure that respondents are reading the questions and instructions rather than simply checking random boxes (Goodman et al., 2013; Hauser & Schwarz, 2016).

Procedures

In recruiting participants through MTurk, we specified only that they had to be current or former Twitter users. We deployed the survey in July 2016. To increase the validity of the survey, we piloted it with roughly a dozen undergraduate and graduate students at the authors' institutions. These students (not involved with the project) provided feedback on wording and survey flow, which we incorporated into the final survey design. Student responses were not included in the dataset and were used for piloting purposes only. We deployed the survey on MTurk and paid participants US\$1.50 for completing the survey. Our piloting revealed that it took between 5 and 15 min to complete, ensuring a payment rate that corresponded with MTurk best practices. In total, 394 respondents completed the survey; 26 participants either failed or did not answer the attentiveness question or provided "gibberish" responses. The resulting dataset, therefore, included 368 respondents.

Data Analysis

We used descriptive statistical analysis to summarize numerical responses. Since the majority of questions contained in the survey elicited responses on an ordinal scale of measurement, we used chi-square tests for independence to explore whether response choices are independent, or if there is a relationship between the two. To not violate the expected cell counts requirement of chi-square tests, we recoded several questions to combine semantically similar response categories (e.g., collapsing some Likert-type scale responses). For correlational analysis with categorical and scale responses, we used Spearman's rho to test for relationships.

In addition to descriptive statistics, we conducted inductive, open coding on open-ended survey responses (Charmaz, 2006). Researchers deliberated on conceptual themes and used these to supplement and organize statistical findings. Quotes included here are representative examples of broader themes.

Findings

Demographic Characteristics of Sample

Of the 368 valid survey responses, 268 participants stated that they use a "public" account as their most frequently used Twitter account, and 100 of the respondents indicated they used a "protected" Twitter account as their main account. This proportion of protected accounts is far higher than the general Twitter population, which, as estimated in 2013, is closer to 5% (Fiesler et al., 2017). To limit the number of potentially confounding factors, we have chosen to focus our attention on the responses of users who mainly use public accounts. For the purposes of space and focus, we do not report the findings from protected account respondents but will note that overall, their attitudes about researchers using tweets were less favorable across the board.

Of the 268 public users, the sample's average age is 32.25 years ($SD=8.80$), and (with gender reported using an open-text answer) men (60.4%) outnumber women (39.2%). A majority of participants (67.9%) reported either some undergraduate education or a completed undergraduate degree as their highest level of education. The majority (92.9%) of participants reported being in the United States.

Twitter Use

Generally, our participants are fairly active Twitter users who have been using the service for some time. Most respondents (80.2%) indicated they currently have only one Twitter account. When asked how often they access Twitter, 7.8% indicated they access it "Almost Never," 29.3% indicated "Occasionally," 35.0% indicated "Semi-regularly," and 27.8% indicated they access it "All the time." In total, 99.3% have sent a Tweet since becoming Twitter users, 86.2% have sent one in the last year, 64.6% in the last month, and 48.9% in the last week. Almost half of the sample reported using Twitter for 4 or more years.

We also asked participants to give us the number of tweets, following, followers, likes, and photos and videos they have on their account. Although as Table 1 shows, there was a high degree of variance, the average participant had sent roughly 2,000 tweets, follows about 350 accounts and has twice as many followers, had "liked" almost a 1,000 tweets, and had uploaded about 100 photos and videos. We also calculated an "influence score" determined by dividing the number of followers by the number of accounts following. A mean influencer score of 1.6 indicates that the average participant had 1.6 times as many accounts that follow them than accounts they follow.

General Awareness of Research Using Tweets

At the end of the survey, we asked participants whether, prior to this study, they knew that researchers sometimes use tweets for research purposes. Almost two-thirds (61.2%) of respondents

Table 2. Comfort Around Tweets Being Used in Research.

Question	Very uncomfortable	Somewhat uncomfortable	Neither uncomfortable nor comfortable	Somewhat comfortable	Very comfortable
How do you feel about the idea of tweets being used in research? (<i>n</i> = 268)	3.0%	17.5%	29.1%	35.1%	15.3%
How would you feel if a tweet of yours was used in one of these research studies? (<i>n</i> = 267)	4.5%	22.5%	23.6%	33.3%	16.1%
How would you feel if your entire Twitter history was used in one of these research studies? (<i>n</i> = 268)	21.3%	27.2%	18.3%	21.6%	11.6%

Note. The shading was used to provide a visual cue about higher percentages.

(*n* = 268) indicated they did not. If this pattern is representative of Twitter users more broadly, it suggests that users are often unaware of how the content they produce is collected and used by those beyond their followers. This finding complements prior work that suggests users are broadly unaware of how the content they create flows within a larger informational ecosystem that Twitter supports (Proferes, 2017).

We also asked whether participants think researchers are permitted to use a tweet without permission from the user (*n* = 267). Slightly over half (57.3%) indicated they believe researchers are permitted, while 42.7% indicated they believe researchers are not. Those who indicated *not* (*n* = 114) were given a follow-up question as to why they believed this. In total, 23.2% stated that Twitter's TOS forbid it, 10.1% believe researchers would be breaking copyright law, 60.9% believe that researchers would be breaking ethical rules for researchers, and 5.8% gave an "other" response. This finding suggests that many Twitter users incorrectly believe that researchers cannot use tweets at all or must ask permission from the users, and a majority of that group believe that this is an ethical rule for researchers.

Respondents who thought that Twitter's TOS prohibits researchers from using tweets were also incorrect; Twitter's policies actually specifically state that researchers do have access to public tweets. Twitter's privacy policy³ (current as of a June 2017 update) states,

Twitter broadly and instantly disseminates your public information to a wide range of users, customers, and services, including search engines, developers, and publishers that integrate Twitter content into their services, and organizations such as universities, public health agencies, and market research firms that analyze the information for trends and insights. When you share information or content like photos, videos, and links via the Services, you should think carefully about what you are making public.

Of course, it is not surprising that our respondents were unfamiliar with this warning. A great deal of prior work shows that people do not read TOS or other website terms and conditions (Reidenberg et al., 2015), even turkers (Fiesler et al., 2016).

Attitudes About Research Using Tweets

After the demographic questions, participants were asked how they feel about the idea of tweets being used in research. As shown in Table 2, the majority of respondents are somewhat comfortable or are ambivalent about the idea of tweets being used in research. However, participant responses shifted to much higher levels of discomfort when "your entire Twitter history" became the subject of study.

At the end of the survey, we asked respondents if they were given the possibility to opt-out of having their tweets used in all academic research, would they? A plurality (46.3%) indicated that they would not, 29.1% indicated they would, and 24.6% indicated that it would depend (*n* = 268). This suggests that contextual factors are important in users' decisions about wanting to be part of research.

We also asked respondents, "Regardless of whether you would want them to use your tweets specifically, do you think that researchers should be able to use tweets in research without user permission?" A majority (64.9%) indicated they should not, and 35.1% indicated that they should (*n* = 268). This result suggests that users feel strongly about the desire to have researchers seek consent or permission. Several respondents left qualitative feedback about this subject. For example, one wrote,

I would not want my tweets to be used in a study without being informed prior to such use.

Others highlighted the contextual nature of such a decision:

If my tweets were being used in a large scale study, I really wouldn't care. If anything was being personally picked out about me in a small study, I would care.

I would want to know how it was to be used, who would see it, whether my information would be kept anonymous and how long the tweet would be kept.

When asked if a university researcher contacted them and asked for permission to use a tweet of theirs as part of a

Table 3. Percentage of Respondents Checking Each Contextual Factor, Ordered Highest Percentage to Lowest Percentage. ($n=268$).

“Would any of the following conditions change how you feel about a tweet of yours being used in a research study?”	Total checked
Whether or not you are asked permission	67.4%
Whether or not you are informed before the research took place	59.6%
What the study is about	56.6%
The size of the dataset (i.e., is your tweet one among millions or are only a small number of tweets being analyzed)	45.7%
Whether the researchers are also analyzing other information about you from Twitter, such as your profile information or geo-location information	48.3%
Who is doing the research	44.9%
The type of analysis (i.e., whether a human is reading your tweet or it is only being analyzed by a computer program)	39.7%
Whether your tweet is quoted verbatim in a research paper	38.6%

research study, a majority (53.4%) indicated they would give permission, only 13.8% indicated outright they would not, and 32.8% indicated it would depend on some contextual factor ($n=268$).

Contextual Factors

The majority of survey questions focused on different specific situational factors that would impact a respondent's level of comfort with the idea of their tweets being used in a research study. We asked contextual questions in two ways. First, we asked participants to place a check mark next to specific situations that would change how they feel about a tweet of theirs being used. Table 3 provides a breakdown of these responses. We also asked participants their level of comfort (on a Likert-type scale) if their tweet was used in a research study and a specific situational factor was in place. Table 4 reflects the results of those questions.

Table 3 shows that the subject of the study is also important context. Table 4 further reflects high levels of discomfort around the idea of Tweets being collected from “protected” accounts and the idea of previously deleted Tweets being used for research, as has been done in studies such as Almuhammedi, Wilson, Liu, Sadeh, & Acquisti (2013) and Zhou, Wang, & Chen (2016). In qualitative feedback about this particular facet, one respondent wrote,

I think people at times tweet heatedly, and sometimes regret speaking so candidly, so I'd be a little concerned that faulty conclusions about one's nature or intent were being extrapolated by those tweets, especially if they had been later deleted.

Respondents were uncomfortable with the idea of profile information also being analyzed along with tweets. These tables also suggest that generally, users do not want to be directly attributed if quoted in a research article but are amenable to the idea of being quoted anonymously.

Other responses suggest that some users, despite sharing their content publicly, do not necessarily think that this means they are giving it away to the public. They may view their tweets as property that they are in the position to derive exclusive (or semi-exclusive) benefit from. In fact, several respondents mentioned that a reason why they do not want their tweets used is that they own them or might expect some benefit or compensation for their use. For example,

They are taking property and using it without permission, just don't think that is right.

[o]nly if they were offering me some kind of compensation. It is not fair to profit from MY ideas and offer me nothing.

Finally, we also asked respondents about two additional contextual factors: the content of the tweet being used and what the study is about. With respect to content, respondents were primarily concerned with private, personal, or offensive tweets, for example,

If it's personal, has identifying information, or embarrassing/offensive/private I don't want my tweets used.

Some of my tweets are personal and some are jokes. Using the jokes would usually be fine.

The purpose of the research study was also a dependency for many respondents. A number of respondents simply wanted to know whether it would help people. Others had more specific concerns, such as whether it were controversial, reflected a specific ideology, or would paint the user in a bad light.

If it was for a conservative cause I would be more forgiving. But that is unlikely with academic researchers, who are inherently biased, i.e., liberal slanted.

If I was made out to look poorly this would be against my wishes.

Attribution and Dissemination

As there has been some discussion within the research community about whether or not to attribute quotes when reporting the results of research, we asked participants whether they would want to be quoted in a published research paper with their username attributed or not (instead remaining anonymous). As seen in Table 4, a majority (58.2%) indicated they would not want their usernames attributed, 18.3% indicated they would want attribution, and 23.5% indicated

Table 4. “How Would You Feel If a Tweet of Yours Was Used in a Research Study and . . .” (n=268).

	Very uncomfortable	Somewhat uncomfortable	Neither uncomfortable nor comfortable	Somewhat comfortable	Very comfortable
. . . you were not informed at all?	35.1%	31.7%	16.4%	13.4%	3.4%
. . . you were informed about the use after the fact?	21.3%	29.1%	20.5%	22.0%	7.1%
. . . it was analyzed along with millions of other tweets?	2.6%	18.7%	25.5%	30.0%	23.2%
. . . it was analyzed along with only a few dozen tweets?	16.5%	30.3%	24.0%	20.2%	9.0%
. . . it was from your “protected” account?	54.9%	20.5%	13.8%	6.0%	4.9%
. . . it was a public tweet you had later deleted?	31.3%	32.5%	20.5%	10.4%	5.2%
. . . no human researchers read it, but it was analyzed by a computer program?	2.6%	14.3%	30.5%	32.3%	20.3%
. . . the human researchers read your tweet to analyze it?	9.7%	27.6%	25.0%	25.4%	12.3%
. . . the researchers also analyzed your public profile information, such as location and username?	32.2%	23.2%	21.0%	13.9%	9.7%
. . . the researchers did not have any of your additional profile information?	4.9%	15.4%	25.1%	34.1%	20.6%
. . . your tweet was quoted in a published research paper, attributed to your Twitter handle?	34.3%	21.6%	21.6%	13.1%	9.3%
. . . your tweet was quoted in a published research paper, attributed anonymously?	9.0%	16.8%	26.5%	28.4%	19.4%

Note. The shading was used to provide a visual cue about higher percentages.

they wouldn’t care either way. Several respondents left qualitative responses, explaining concerns about whether or not tweets could be traced back to a “real identity.” For example, one wrote,

The vast majority of my tweets are jokes and my username is my real name so I’d like the opportunity so provide explanation or context to the researcher to ensure it was understood if I thought my tweet would be widely quoted.

Another stated,

As long as my name wasn’t tied to it I wouldn’t care.

Another participant remarked on the potential for tweets to be used in a context they had no control over, stating,

It’s a shitty thing to do; you’ll never give the proper context to a tweet if it’s quoted, that’s for sure. There will be a level of judgment for or against the content and you can’t act as if you’re being scholarly; scholars have their biases, too, and just because they have a title or a degree, it doesn’t place them in a place to manufacture an accurate or objective meaning to it.

It is also worth noting that under Twitter’s Developer Agreement⁴ (which applies to anyone who uses the API),

Twitter’s “display requirements” require tweets to be presented verbatim and with usernames.⁵ Although this does not actually seem to be a practice that researchers follow, if Twitter decided to enforce the rule, it would force identification of Twitter users in published research.

Some respondents worried about how a study could bring attention from unwanted publics. For example, one respondent wrote that he or she would be uncomfortable with

how public the information regarding my tweet could become after the research, i.e. in media outlets.

This response is interesting because of the fact that “public” here is not a binary. Rather, it suggests some users conceptualize differing levels of “public,” supporting Nissenbaum’s (2004) idea of the importance of context for privacy norms. At the same time, other users had more binary views of what it means for content to be public:

. . . if I posted something on Twitter I know full well that this is the internet and anyone can come up and read it when i post it publicly like that. If I didn’t want people to know, I wouldn’t post it.

When asked if a tweet of theirs was used by a university researcher in a study, would they want to be informed, 79.5%

indicated they would, 5.6% indicated they would *not* want to know, and 14.9% indicated that they would not care either way ($n=268$). When asked if they would want to *read* the academic article the researcher produced, 83.4% indicated they would, 3.4% indicated they would not want to know about it, and 13.2% indicated they would not care about reading it ($n=268$). Perhaps summarizing the feelings of many of the participants, one respondent wrote,

I honestly wouldn't mind [if researchers used my tweets] as long as I was told up front and I had the option to read the findings.

We hypothesized that there may be an intervening relationship between demographic characteristics about participants and their level of comfort with the use of their data in research. This assumption was based on the stereotype that “Young people don't care about privacy” and the perception that perhaps those who tweet more are less likely to care about how their tweets are used. To examine this, we used chi-square tests for independence to explore whether response choice to the question “How would you feel if a tweet of yours was used in one of these research studies?” and demographic and Twitter use characteristics were independent, or if there is a relationship between the two. For correlational analysis with categorical and scale responses, we used Spearman's rho to test for relationships. We found no statistically significant relationships between questions response and age, $r = .047$, $p = .443$; gender, $\chi^2(4, N=266) = 2.841$, $p = .585$; level of education, $\chi^2(20, N=267) = 26.661$, $p = .145$; how often they access Twitter, $\chi^2(12, N=265) = 9.040$, $p = .699$; the last time they posted a Tweet, $\chi^2(20, N=267) = 15.252$, $p = .762$; the number of tweets they had sent, $r = .034$, $p = .583$; the number of followers they had, $r = .036$, $p = .557$; or the “influence score,” $r = .070$, $p = .257$. We therefore speculate that attitudes about Twitter research are more context-dependent.

Discussion

Our goal for this research was not to answer the binary question of “should researchers be using public tweets or not?” but rather to explore users' perceptions about research on Twitter and how they view contextual factors involved in this practice. Below, we first reflect on the major themes of our findings, then lay out some potential implications for both practice and design, and finally propose important future work implicated by this study.

As previously noted, the decades-old Belmont Report presents guiding principles for human subjects research and is relied upon by many researchers (Vitak et al., 2016). In analyzing our findings, we uncovered themes that tracked well to the Belmont Report (despite, most likely, our participants having little or no knowledge of them), suggesting that these guiding principles are still highly relevant.

Respect for Persons

Three of the strongest and most interesting themes from our data are tied to the Belmont principle of *respect for persons* or the recognition that people should have the right to exercise autonomy. Respect for persons commonly manifests through practices such as informed consent. The other two themes in our data we saw related to beneficence were the idea of *choice* and dissemination of findings.

As noted previously, many online data researchers do not perceive informed consent as relevant for collecting public data such as tweets (Bruckman, 2014; Vitak et al., 2016). However, based on open-text responses to our survey, we saw that the idea of consent or permission came from the underlying importance of *respect* for our respondents. Many respondents' attitudes relied heavily on whether or not permission was sought. For example, one respondent wrote,

The person who tweeted should be respected and asked about it being used first.

Although in a general sense, “asked permission” is not the same as “informed consent,” which Brown et al. (2016) point out is highly ritualized and may not be protecting participants in a meaningful way. This is partially because consent forms, like TOS (Fiesler et al., 2016; Luger et al., 2013), are often difficult to read and understand (Cassileth, Zupkis, Sutton-Smith, & March, 1980). However, absent regulatory requirements, researchers could obtain permission any way they like.

Many open-text responses we received emphasized respondents' desire to understand contextual factors, often remarking that levels of comfort and whether or not they would be willing to grant permission to researchers “depends.” Factors included what and how many tweets are used, what the study is about, who is conducting it, or what methods researchers use. Many participants do not necessarily want to give a blanket answer about how they feel about Twitter research but want the control to consider research case-by-case.

Finally, a number of respondents specifically framed their desire to both *know about* the research and to *see* it when it's finished as an issue of respect. Informed consent can be seen as both *informing* and *consenting*, and for many respondents, the former would be sufficient.

Beneficence

The common ethical principle of “do no harm” is part of the Belmont Report as well in *beneficence*. Minimizing risk and maximizing benefit to participants is a large part of the ethical calculus that researchers often use. Our findings about dissemination could also be seen as part of this theme; given how much our respondents cared not just about being *informed* but about the opportunity to read

published papers suggests that they may want the *benefit* of learning about the study, either for the sake of knowledge or for curiosity.

Our qualitative analysis also revealed a great many respondents suggesting ways to minimize potential harm, particularly with respect to privacy. Primarily, they mentioned things like being careful about anonymization, never using real names, and making certain that nothing could link published data back to a Twitter account.

They also wrote about forms of harm such as being embarrassed by something published about them. One comment that came up frequently in reasons to be wary of research was that single tweets lack context or that quoting and further disseminating tweets makes them more public than they were intended to be.

Finally, the idea of minimizing risk and maximizing benefit is the argument for *not* suggesting that the solution is to stop doing Twitter-based research. After all, many respondents were positive about research, with comments such as “if it is for science why not” and “well research is a noble pursuit.” This suggests that they too are doing ethical calculus about whether a potential invasion of their privacy is worth it, for the benefit of research and science.

Justice

The Belmont principle of *justice* involves the assurance of reasonable, non-exploitative research methods that are administered fairly and equally to participants (and potential participants). Part of this, as explained in the Belmont Report, has to do with participant *selection*. However, it also involves fair (or at least equal) compensation to research participants. We are unaware of any studies of public Twitter data where the Twitter users have been monetarily compensated. Some participants stated their willingness to give permission would depend on compensation, or that *commercialization* of the research is a problem. They essentially saw this as a sort of exploitation. However, tying back to beneficence as well, it could be that providing a benefit to users—even if only what benefit is conferred by knowledge of the study and findings—would make these research practices seem less exploitative.

Implications for Practice and Design

The themes above suggest potential best practices or factors that researchers should consider in making ethical determinations about their research design. First, *consider* asking for permission if there is any reasonable way to do so. Since the study of public data may not fall under IRB or other regulatory purview, obtaining consent would not have to be a formal consent form. This would result in, as Brown et al. (2016) suggest, separating the legal from the ethical. Even simply the opportunity to “opt out” would be good practice—for example, tweeting at those whose tweet is included in the research and offering to remove their content from the

dataset if they would prefer. Alternately, if a researcher does not seek permission beforehand, they could still consider informing the Twitter users after.

With respect to privacy, our findings point to some clear best practices: first, consider anonymizing identifying information when quoting tweets. Only a minority of our respondents stated that they would prefer for tweets to be attributed to them. Moreover, although from these questions we did not determine whether participants understand that verbatim tweets can be re-identified through Twitter search mechanisms even if their usernames are not disclosed, prior work suggests that many Twitter users are not aware of how widely available Twitter data are (Proferes, 2017). This suggests that participants who are comfortable with anonymous quotes but not for quotes attributed to them might be uncomfortable with their tweets being re-identified outside the context of a publication as well. We therefore recommend not quoting tweets verbatim without reason and generally to consider Bruckman’s (2002) levels of disguise when directly using content in a published work. This suggestion tracks to conclusions from M. L. Williams et al. (2017) on publishing tweets verbatim, that particularly if there is any personal information involved, researchers should consider obtaining consent.

However, attribution is also a decision where the subject matter of the study and population of participants should be considerations, as there may be contexts when attribution is appropriate and perhaps the most ethical course (Brown et al., 2016; Bruckman, Luther, & Fiesler, 2015). However, those respondents concerned about their privacy and the potential harm of data being traced back to them were most vocal in our data, and in a harm/benefit analysis, it would make sense to focus on minimizing this harm. Therefore, we suggest that publication of user identity should only occur when the benefits of doing so clearly outweigh the potential harms, or with user permission.

Also regarding privacy, respondents were highly uncomfortable with their profile information being analyzed in tandem with tweets. We also recommend *not* using deleted content (which is also prohibited by Twitter’s Developer Agreement [see Note 3]), due to the high level of discomfort expressed by our respondents, unless the benefit of doing so would justify violating users expectations.

In a general sense, our findings suggest users believe strongly in privacy by obscurity (Hartzong & Stutzman, 2013) and that research has the potential to disrupt this. Users often felt more comfortable with research using larger data sets (with the exception of more of a single users’ tweets, as respondents did not want researchers using their entire Twitter history). Furthermore, they felt more comfortable with the idea of tweets being analyzed by a computer rather than read by humans. However, both of these beliefs may be misguided, as re-identification is possible in large data sets (Zimmer, 2010a), and computer algorithms can be biased in both design and application (Friedman & Nissenbaum,

1996). Although we are certainly not advising against using qualitative methods in Twitter research, it is a situation where research design should be considered carefully, particularly, with respect to, for example, the subject matter of the study and content of the tweets.

Our primary proposal for best practices is for researchers to understand and reflect carefully on these contextual factors during study design. We suggest that researchers should most carefully think through ethics when research involves the more problematic factors listed above. For example, a study about sensitive topics such as medical conditions or drug use could be less appropriate for quoting tweets than a study about television habits. Within human-computer interaction (HCI) research, it is already common practice to take special precautions when working with vulnerable populations (Brown et al., 2016), and we suggest that this should extend to Twitter users as well. In sum, this work suggests that in making decisions about ethical use of tweets, researchers should pay close attention to the content of the tweets, the level of analysis with respect to making the content more public, and reasonable expectations of privacy (e.g., deleted or protected content). They should also consider taking steps toward informing users about the research and providing them with opt-out options, if it would not compromise the research or researchers.

In addition to these suggestions for best practices, in line with work done by Bravo-Lillo, Egelman, Herley, Schechter, & Tsai (2013), our findings point to ways in which the development of automated tools could contribute to ethical practices. First, for Twitter itself or anyone designing a new social computing system: consider providing a way for users to opt-in or opt-out of particular forms of research. This could be, for example, a flag set in the user profile or a black/white list included as part of the API. Another potential design would be to build a system that could provide public notices when data collection begins from a specific hashtag, informs users when their tweets are included in a dataset, and/or links those who have had their tweets used back to a published paper based on the results. Both of these designs would be a way to benefit Twitter (or other social media) users as well as to support best ethical practices. Similar to a system like Turkopticon (Irani & Silberman, 2013), we encourage others to think about symbiotic systems that would help empower users and research participants.

It is important to note, however, that we do *not* recommend that platforms solve this problem by making it impossible for researchers to collect public data. As expressed by many of our respondents, science and research is important. Over half indicated that if asked for permission, they would allow their tweets to be used without any dependencies and even more would give permission if they knew it was for scientific research. Therefore, we posit that disallowing use of public data in research altogether would be as poor an outcome as using it indiscriminately without any consideration for ethics.

Conclusion

We consider this exploratory study to be an important step in motivating future empirical studies of research ethics. These findings with respect to Twitter may or may not be generalizable for other platforms or contexts. One next step would be to look beyond Twitter. For example, do people feel different about their data being collected from Reddit, Tumblr, Instagram, or Facebook as opposed to Twitter? What factors affect these potential differences? How much do perceptions of the “publicness” of data impact comfort with the idea of being a research subject? In addition, how does the context of the use change perception? Outside of research, there are similar questions around, for example, journalism. If researchers are disrupting typical imagined audiences (Litt, 2012), then so are marketers and journalists. Is finding out your tweet appears in a research paper more or less concerning than your tweet appearing in a news article? How do attitudes about research connect to other instances of unintended audience? Understanding what factors might make users more comfortable can help inform future best practices.

As noted previously, there is inherent selection bias in any data collected with consent. Therefore, the only way to study people who don't want to be studied is to do so without their consent. For example, in a study of early chat rooms, researchers told chat room participants that they were recording their conversations for a study of language online, when actually they were studying how those participants reacted to the idea of being studied (Hudson & Bruckman, 2004). Consider a study design where the researcher replies to tweets with “your tweet has just been collected by a researcher!” to study the responses. Such a study would have to be carefully and ethically designed since it would require a waiver of consent for deception. We hope that this preliminary work will help in the design of such future work by providing data to help motivate and direct. We also encourage other researchers to explore research ethics from a variety of different methodological approaches. As researchers, we have the tools to tackle the problems that arise in our own practices in creative and rigorous ways.

Empirical work could also help us determine the reaction of researchers to these ideas to find the real tension points that could keep these practices from being taken up. We also think that it is important that researchers reflect on ethical choices as part of writing up research results. Ethics are not determined by majority opinion, and norms are just one part of deciding what constitutes ethical action. Having open conversations within the social computing research community, bolstered by outside voices such as the people we study, is critical.

Acknowledgements

We would like to thank Dr Katie Shilton at the University of Maryland for feedback on an early version of this work.

Declaration of Conflicting Interests

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

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Ongoing research in this space is funded in part by NSF award IIS-1704369 as part of the PERVADE (Pervasive Data Ethics for Computational Research) project.

Notes

1. <https://twitter.com/privacy>
2. <http://www.mturk.com>
3. <https://twitter.com/privacy>
4. <https://dev.twitter.com/overview/terms/agreement-and-policy>
5. <https://about.twitter.com/en-gb/company/display-requirements>

ORCID iD

Nicholas Proferes  <https://orcid.org/0000-0002-0295-9616>

References

- Adams, R. C. (1989). *Social survey methods for mass media research*. Hillsdale, NJ: Lawrence Erlbaum.
- Almuhimedi, H., Wilson, S., Liu, B., Sadeh, N., & Acquisti, A. (2013). Tweets are forever: A large-scale quantitative analysis of deleted tweets. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)* (pp. 897–908). ACM.
- Barocas, S., & Nissenbaum, H. (2014). Big Data's end run around anonymity and consent. In J. Lane, V. Stodden, S. Bender, & H. Nissenbaum (Eds.), *Privacy, big data, and the public good: Frameworks for engagement* (pp. 44–75). Cambridge, UK: Cambridge University Press.
- Beaulieu, A., & Estallea, A. (2012). Rethinking research ethics for mediated settings. *Information, Communication & Society, 15*(1), 23–42.
- Bernstein, M. S., Bakshy, E., Burke, M., & Karrer, B. (2013). Quantifying the invisible audience in social networks. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (pp. 21–30). ACM.
- Blank, G. (2016). The digital divide among Twitter users and its implications for social research. *Social Science Computer Review, 35*, 679–697.
- Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science, 2*(1), 1–8.
- Bravo-Lillo, C., Egelman, S., Herley, C., Schechter, S., & Tsai, J. Y. (2013). You needn't build that: Reusable ethics-compliance infrastructure for human subjects research. In *Cyber-Security Research Ethics Dialog & Strategy Workshop*.
- Bromseth, J. C. H. (2002). Public places—Public activities? Methodological approaches and ethical dilemmas in research on computer-mediated communication contexts. In A. Morrison (Ed.), *Researching ICTs in context* (pp. 33–61). Oslo, Norway: University of Oslo.
- Brown, B., Weilenmann, A., Mcmillan, D., & Lampinen, A. (2016). Five provocations for ethical HCI research. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (pp. 852–863). ACM.
- Bruckman, A. (2002). Studying the amateur artist: A perspective on disguising data collected in human subjects research on the Internet. *Ethics and Information Technology, 4*, 217–231.
- Bruckman, A. (2014). Research ethics and HCI. In W. A. Kellogg & J. S. Olson (Eds.), *Ways of knowing in HCI* (pp. 449–468). New York, NY: Springer.
- Bruckman, A., Luther, K., & Fiesler, C. (2015). When should we use real names in published accounts of Internet research? In E. Hargittai & C. Sandvig (Eds.), *Digital research confidential* (pp. 243–258). Cambridge, MA: MIT Press.
- Cassileth, B., Zupkis, R., Sutton-Smith, K., & March, V. (1980). Informed consent—Why are its goals imperfectly realized? *The New England Journal of Medicine, 302*, 896–900.
- Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. London, England: SAGE.
- Collmann, J., Fitzgerald, K. T., Wu, S., Kupersmith, J., & Matei, S. A. (2016). Data management plans, Institutional Review Boards, and the ethical management of Big Data about human subjects. In J. Collmann & S. A. Matei (Eds.), *Ethical reasoning in Big Data* (pp. 141–184). Cham, Switzerland: Springer.
- Crawford, K., & Finn, M. (2015). The limits of crisis data: Analytical and ethical challenges of using social and mobile data to understand disasters. *GeoJournal, 80*, 491–502.
- De Cristofaro, E., & Soriente, C. (2013). Participatory privacy: Enabling privacy in participatory sensing. *IEEE Network, 27*(1), 32–36. doi:10.1109/MNET.2013.6423189
- Ess, C., & AOIR Ethics Working Committee. (2002). *Ethical decision-making and Internet research* (Version 1). Retrieved from <http://aoir.org/reports/ethics.pdf>
- Felt, M. (2016). Social media and the social sciences: How researchers employ Big Data analytics. *Big Data & Society, 3*(1), 1–15.
- Fiesler, C., Dye, M., Feuston, J., Hiruncharoenvate, C., Hutto, C. J., Morrison, S., . . . Gilbert, E. (2017). What (or who) is public? Privacy settings and social media content sharing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)* (pp. 567–580). ACM.
- Fiesler, C., Lampe, C., & Bruckman, A. S. (2016). Reality and perception of copyright terms of service for online content creation. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)* (pp. 1450–1461). ACM.
- Fiesler, C., Young, A., Peyton, T., Bruckman, A. S., Gray, M., Hancock, J., & Lutters, W. (2015). Ethics for studying online sociotechnical systems in a Big Data World. In *Proceedings of the ACM Conference Companion on Computer Supported Cooperative Work & Social Computing (CSCW Companion)* (pp. 289–292). ACM.
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems, 14*, 330–347.
- Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making, 26*, 213–224. doi:10.1002/bdm.1753
- Hartzong, W. N., & Stutzman, F. (2013). The case for online obscurity. *California Law Review, 101*(1), 1–49.

- Hauser, D. J., & Schwarz, N. (2016). Attentive turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, *48*, 400–407.
- Hitlin, P. (2016). *Researching in the crowdsourcing age, a case study*. Retrieved from <http://www.pewinternet.org/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/>
- Hoffmann, A. L., & Jones, A. (2016). Recasting justice for Internet and online industry research ethics. In M. Zimmer & K. E. Kinder-Kurlanda (Eds.), *Internet research ethics for the social age: New cases and challenges* (pp. 3–18). Bern, Switzerland: Peter Lang.
- Hudson, J. M., & Bruckman, A. (2004). “Go Away”: Participant objections to being studied and the ethics of chat-room research. *The Information Society*, *20*, 127–139. doi:10.1080/01972240490423030
- Irani, L. C., & Silberman, M. (2013). Turkopticon: Interrupting worker invisibility in Amazon Mechanical Turk. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (pp. 611–620). ACM.
- Kogan, M., Palen, L., & Anderson, K. M. (2016). Think local, retweet global: Retweeting by the geographically-vulnerable during Hurricane Sandy. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)* (pp. 981–993). ACM.
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*, 8788–8790.
- Lamb, A., Paul, M. J., & Dredze, M. (2013). Separating fact from fear: Tracking flu infections on Twitter. In *Proceedings of the HLT-NAACL* (pp. 789–795). ACL.
- Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., & Christakis, N. (2008). Tastes, ties, and time: A new social network dataset using Facebook.com. *Social Networks*, *30*, 330–342. doi:10.1016/j.socnet.2008.07.002
- Litt, E. (2012). Knock, knock. Who’s there? The imagined audience. *Journal of Broadcasting & Electronic Media*, *56*, 330–345.
- Luger, E., Moran, S., & Rodden, T. (2013). Consent for all: Revealing the hidden complexity of terms and conditions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (pp. 2687–2696). Paris, France. ACM.
- Mao, H., Shuai, X., & Kapadia, A. (2011). Loose tweets: An analysis of privacy leaks on Twitter. In *Proceedings on the ACM Workshop on Privacy in the Electronic Society* (pp. 1–12). ACM.
- Markham, A., Buchanan, E., & AOIR Ethics Working Committee. (2012). *Ethical decision-making and internet research* (Version 2). Association of Internet Researchers.
- Martin, K., & Shilton, K. (2015). Why experience matters to privacy: How context-based experience moderates consumer privacy expectations for mobile applications. *Journal of the Association for Information Science and Technology*, *67*, 1871–1882.
- Marwick, A. E., & Boyd, D. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, *13*, 114–133.
- McNeal, G. (2014, June 28). Facebook manipulated user news feeds to create emotional responses. *Forbes*. Retrieved from <https://www.forbes.com/sites/gregorymcneal/2014/06/28/facebook-manipulated-user-news-feeds-to-create-emotional-contagion/#184494f539dc>
- Meeder, B., Tam, J., Kelley, P. G., & Cranor, L. F. (2010). RT @IWantPrivacy: Widespread violation of privacy settings in the Twitter social network. In *Web 2.0 Security and Privacy* (pp. 28–48). IEEE.
- Mikal, J., Hurst, S., & Conway, M. (2016). Ethical issues in using Twitter for population-level depression monitoring : A qualitative study. *BMC Medical Ethics*, *17*, 22.
- Nebeker, C., Harlow, J., Giacinto-Espinoza, R., Linares-Orozco, R., Bloss, C., & Weibel, N. (2017). Ethical and regulatory challenges of research using pervasive sensing and other emerging technologies: IRB perspectives. *American Journal of Bioethics: Empirical Bioethics*, *8*, 266–276.
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, *79*, 101–139.
- Nuremberg Code (1949). In *Trials of War Criminals before Nuremberg Military Tribunals*, Washington, DC: US Government Printing Office.
- Office for Human Research Protections. (1978). *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. Bethesda, MD: Author.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*(5), 411–419.
- Proferes, N. (2017). Information flow solipsism in an exploratory study of beliefs about Twitter. *Social Media + Society*, *3*(1), 2056305117698493.
- Redmiles, E. M., Kross, S., Pradhan, A., & Mazurek, M. L. (2017). *How well do my results generalize? Comparing security and privacy survey results from MTurk and Web Panels to the U.S.* (Technical Reports of the Computer Science Department). College Park: University of Maryland.
- Reidenberg, J. R., Breaux, T., Cranor, L. F., & French, B. (2015). Disagreeable privacy policies: Mismatches between meaning and users’ understanding. *Berkeley Technology Law Journal*, *30*(1), 39–68.
- Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). Earthquake shakes Twitter users: Real-time event detection by social sensors. In *Proceedings of the International Conference on World Wide Web (WWW)* (pp. 851–860). ACM.
- Schechter, S., & Bravo-Lillo, C. (2014). *Using ethical-response surveys to identify sources of disapproval and concern with Facebook’s emotional contagion experiment and other controversial studies* (Microsoft Research Working Paper). Retrieved from: <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/CURRENT20DRAFT20-20Ethical-Response20Survey.pdf>
- Shilton, K., & Sayles, S. (2016). “We aren’t all going to be on the same page about ethics”: Ethical practices and challenges in research on digital and social media. In *Proceedings of the Hawaii International Conference on Systems Sciences (HICSS)* (pp. 1530–1605) IEEE.
- Signorini, A., Segre, A. M., & Polgreen, P. M. (2011). The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. *PLoS ONE*, *6*(5), e19467.
- Singer, N. (2015, Feb 13). Love in the time of Twitter. *The New York Times*. Retrieved from: <https://bits.blogs.nytimes.com/2015/02/13/love-in-the-times-of-twitter/>

- Starbird, K., & Palen, L. (2011). "Voluntweeters": Self-organizing by digital volunteers in times of crisis. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)* (pp. 1071–180). ACM.
- Tufekci, Z. (2014). Big Data: Pitfalls, methods and concepts for an emergent field. In *Proceedings of the AAAI International AAAI Conference on Weblogs and Social Media (ICWSM)*. AAAI.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010). Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *Proceedings of the AAAI International Conference on Web and Social Media (ICWSM)* (pp. 178–185). AAAI.
- Vaccaro, K., Karahalios, K., Sandvig, C., Hamilton, K., & Langbort, C. (2015). Agree or cancel? Research and terms of service compliance. In *2015 CSCW Workshop on Ethics for Studying Sociotechnical Systems in a Big Data World*. ACM.
- Vitak, J., Proferes, N., Shilton, K., & Ashktorab, Z. (2017). Ethics regulation in social computing research: Examining the role of Institutional Review Boards. *Journal of Empirical Research on Human Research Ethics*, *12*, 372–382.
- Vitak, J., Shilton, K., & Ashktorab, Z. (2016). Beyond the Belmont principles: Ethical challenges, practices, and beliefs in the online data research community. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)* (pp. 941–953). ACM.
- Williams, M. L., Burnap, P., & Sloan, L. (2017). Towards an ethical framework for publishing Twitter data in social research: Taking into account users' views, online context and algorithmic estimation. *Sociology*, *51*, 1149–1168.
- Williams, S., Tarras, M., & Warwick, C. (2013). What do people study when they study Twitter? Classifying Twitter related academic papers. *Journal of Documentation*, *69*, 384–410.
- Wood, M. (2014, Jul 28). OKCupid plays with love in user experiments. *The New York Times*. Retrieved from <https://www.nytimes.com/2014/07/29/technology/okcupid-publishes-findings-of-user-experiments.html>
- World Medical Association. (2013). *WMA Declaration of Helsinki: Ethical principles for medical research involving human subjects*. Ferney-Voltaire, France: Author.
- Zhou, L., Wang, W., & Chen, K. (2016). Tweet properly: Analyzing deleted tweets to understand and identify regrettable ones. In *Proceedings of the International Conference on World Wide Web (WWW)* (pp. 603–612). International World Wide Web Conferences Steering Committee.
- Zimmer, M. (2010a). "But the data is already public": On the ethics of research in Facebook. *Ethics and Information Technology*, *12*, 313–325. doi:10.1007/s10676-010-9227-5
- Zimmer, M. (2010b). *Is it ethical to harvest public Twitter accounts without consent?* Retrieved from <http://www.michaelzimmer.org/2010/02/12/is-it-ethical-to-harvest-public-twitter-accounts-without-consent/>
- Zimmer, M. (2016, May 14). OKCupid study reveals the perils of Big Data science. *WIRED*. Retrieved from <https://www.wired.com/2016/05/okcupid-study-reveals-perils-big-data-science/>
- Zimmer, M., & Proferes, N. J. (2014). A topology of Twitter research: Disciplines, methods, and ethics. *Aslib Journal of Information Management*, *66*, 250–261.

Author Biographies

Casey Fiesler, PhD (Georgia Institute of Technology), JD (Vanderbilt University Law School), is an assistant professor in the Department of Information Science at University of Colorado Boulder. In addition to ethics, her research involves forms of governance in online communities, including law, platform policies, and social norms.

Nicholas Proferes, PhD (University of Wisconsin-Milwaukee), is an assistant professor in the School of Information Science at the University of Kentucky. His research interests include users' understandings of information flow on social media, tech discourse, and issues of power in the digital landscape.