Privacy, Anonymity, and Perceived Risk in Open Collaboration: A Study of Service Providers

Nora McDonald
Drexel University
nkm39@drexel.edu

Rachel Greenstadt
New York University
greenstadt@nyu.edu

Benjamin Mako Hill
University of Washington
makohill@uw.edu

Andrea Forte
Drexel University
aforte@drexel.edu

ABSTRACT

Anonymity can enable both healthy online interactions like support-seeking and toxic behaviors like hate speech. How do online service providers balance these threats and opportunities? This two-part qualitative study examines the challenges perceived by open collaboration service providers in allowing anonymous contributions to their projects. We interviewed eleven people familiar with organizational decisions related to privacy and security at five open collaboration projects and followed up with an analysis of public discussions about anonymous contribution to Wikipedia. We contrast our findings with prior work on threats perceived by project volunteers and explore misalignment between policies aiming to serve contributors and the privacy practices of contributors themselves.

CCS CONCEPTS

• **Human-centered computing** → *Collaborative and social computing theory, concepts and paradigms*; • **Social and professional topics** → *Computing / technology policy; Privacy policies*;

KEYWORDS

Tor, Wikipedia, Anonymity, Peer Production

ACM Reference Format:

Nora McDonald, Benjamin Mako Hill, Rachel Greenstadt, and Andrea Forte. 2019. Privacy, Anonymity, and Perceived Risk in Open

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.
CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300901>

Collaboration: A Study of Service Providers. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 12 pages.
<https://doi.org/10.1145/3290605.3300901>

1 INTRODUCTION

If someone wants to add information to a Wikipedia article, contribute a Linux patch, or map their town on OpenStreetMap, do these contributions need to be attributed to a particular individual? To an IP address? To a real world identity? Seeking anonymity is a common privacy management strategy among internet users [21], but the extent to which anonymity is possible depends on the design of technical and social infrastructures.

Participation is a cornerstone of online production; however, the openness that facilitates participation is also a source of threats to internet businesses and online communities. As a result, services like Google, Yelp, and Wikipedia turn to third party blacklists, real-name policies, and banning users of anonymity networks like Tor. Such security measures may punish all privacy-seeking contributors because of a few bad actors. Anonymous blacklisting systems such as Nymble [40] promise cryptographic alternatives to blocking entire anonymity networks [18], but adopting these requires resources and coordination. The ways and conditions in which service providers might support experimentation with novel technological interventions is not well understood.

One reason that service providers might value anonymity is diversity. Prior research suggests not only that open collaboration projects like Wikipedia or open source software struggle with underrepresentation from vulnerable categories of contributors [10, 26, 28], but also that policies limiting anonymity as a privacy strategy may censor individuals whose identities (e.g., gender, race, ethnicity) create vulnerabilities. For example, Forte et al. investigated threats experienced by contributors to online projects and discovered that those who sought anonymity often did so to protect themselves from threats like surveillance, opportunity loss, violence, and harassment; those who did not perceive threats “enjoyed privileges due to their gender, nationality, race, or

the scope of their interests” [11]. They found that these threats can inhibit behavior to the point of self-censorship or withdrawal. Identity features like gender and sexual identity or interest in topics that might be considered controversial can play a role in creating vulnerabilities and decisions to withdrawal from projects [26].

Although service providers determine many parameters of access and privacy protections available to users, service providers and users may have different goals and perceive different risks. Understanding when and why these goals and perceptions conflict means not only understanding the values of each group, but also how those values are translated into design and policy. To better understand decisions about technical and social infrastructures that support anonymous participation, we conducted semi-structured interviews with employees of service provider organizations that often support one or more open collaboration projects.

Open collaboration projects include open source software, open content, citizen science, citizen journalism, collaborative mapping, and similar efforts. For example, Github, although not one of our research sites, is a service provider that hosts infrastructure for many open source software projects. We refer to such organizations collectively as *open collaboration service providers*. We asked interview participants about risks and threats they perceive to the open collaboration projects they support and how they address them. Interview data suggested that perceptions of anonymous contributors’ identities played a role in service provider decisions and policies related to anonymity and privacy. In a second study, we built on this finding by examining discourse around perceptions of anonymous contributor identities in a public discussion forum about English Wikipedia.

We make several contributions with this work. We describe three major categories of threats perceived by service providers. We describe how these perceived threats inform policies related to the use of privacy infrastructures and data collection. We compare our findings about threats perceived by service providers with previous research on threats perceived by contributors. Finally, we highlight how reliance on social norms as an analytical tool can blind HCI and online community researchers to minority experiences.

2 RELATED WORK

Anonymity Online

Prior work has shown that people may seek anonymity online in response to real or potential privacy violations [33] or to talk about personally meaningful controversial or sensitive topics [24, 35]. In the case of open collaboration projects, anonymity may facilitate joining a community of practice [5, 20] or mitigating perceived threats linked to stigmatized

or marginalized identities or to personal positions [11]. Sociologist Gary Marx defined “anonymity” as a state that requires thwarting multiple forms of identity knowledge [25]. When people use different online services, they may choose to reveal some types of identity knowledge while obscuring others. Different types of identity knowledge (such as name, location, or behavior patterns) are associated with forms of data that can be stored in online systems. One straightforward example is when sites impose a “real name policy” that allows them to collect and store legal names. Other examples of data with implications for identity knowledge include: IP addresses or EXIF data that are linked with physical location [13], relational data which group membership can be inferred [29], and user agent strings that serve as a form of persistent pseudonym [9].

Technical tools can help people protect different kinds of identity knowledge. For example, Tor is an anonymity network that allows people to browse the web without revealing their IP address [7], plugins/extensions for web browsers may shield users from third-party tracking [19], and anonymous remailers obscure the origin of a message. Service providers play a critical role in either facilitating anonymous participation using such tools or erecting barriers to doing so; for example, by blocking contributions from the Tor network. In this work, we examine how service providers choose to encode and reveal identity knowledge about contributors in the open collaboration infrastructures they manage.

Infrastructures for Open Collaboration

Open collaboration projects support a variety of activities from open source software to open content to citizen science but share in common “an online environment that (a) supports the collective production of an artifact (b) through a technologically mediated collaboration platform (c) that presents a low barrier to entry and exit and (d) supports the emergence of persistent but malleable social structures” [12].

Although open collaboration projects are built on diverse infrastructures, policies, and tools, they often rely on the visibility of individuals’ contribution histories. In general, the designs and policies that support “persistent but malleable social structures” require contributors to reveal information that establishes them as trusted citizens of a community and helps community volunteers manage threats from bad actors.

How do open collaboration projects navigate the need to support collective production and governance while managing the threats of bad behavior and also respecting contributor expectations of privacy? Although a considerable literature has been dedicated to understanding governance mechanisms and undesirable behavior on participatory sites and the architectures designed to mitigate them [6, 22], we don’t know how service providers perceive these types of threats nor how these perceptions shape policy and design.

Conceptual Framing

Threat Modeling. To explore how service providers think about threats to their projects and organizations, we draw on the concept of “threat modeling” from the security literature, which involves identifying assets and attack vectors as well as archetypical attackers, their motivations, goals, and knowledge of the system/organization [27]. Aiming to understand the basis of organizational culture and policy is a departure from classic frameworks like Solove’s “privacy taxonomy” that aims to carefully elucidate and connect types of socially recognized threats and violations oriented toward an individual that might be called privacy violations [37]. Our approach shares Nissenbaum’s attention to local (and potentially idiosyncratic) cultural features of groups and organizations [30]. We aim to characterize the mechanisms of an organization as its representatives perceive threats, interpret them, and contribute to organizational responses, including policies. This entails as a first step understanding how service providers see the people who cause problems, and the types of things that go wrong. For instance, what can service providers tell us about the types of actors that pose a threat to their organization, the motivations of such actors, and their activities? What can they tell us about their organizational responses to these “threats?”

Privacy and Anonymity. As a matter of conceptual clarity, it is important to understand how anonymity and privacy differ. Privacy, as Solove notes, is an “umbrella term” [37]; it includes diverse cultural experiences sometimes described using terms like control or aloneness [2, 41], whereas anonymity is a state in which one’s actions are dissociated from their identity. Anonymity can facilitate privacy by constructing a barrier around one’s activity that renders it visible but unattributable. Theoretical work about online anonymity often focuses on its role in freeing people to behave in counter-normative ways, both negative (e.g., harassment, discrimination) and positive (e.g., self-disclosure, intimacy) [32, 38, 39]. Privacy theories often frame privacy as a tradeoff between individual needs and social values and norms [23, 43].

As we analyzed our data about open collaboration service providers’ privacy-related decision-making and compared it to literature on privacy concerns of open collaboration contributors, we noticed systematic differences that didn’t neatly fit within commonly invoked privacy framings such as contextual integrity [30], boundary regulation [1], or Westin’s categories of privacy strategies [42]. We observed that differences in contributors’ and providers’ threat models often hinged on how contributors were identified. To frame our analysis of threats, then, we used identity as an analytic tool. We examined how open collaboration service providers talk about contributors’ identities in the context of policy discussions and decision-making and how these perceptions

of contributors’ identities shape perceptions of threats and policy making.

Identity and Anonymity. Literature on identity includes fractured and overlapping definitions; however different schools of thought consistently recognize forms of self-identification as well as identities perceived or bestowed by others. For example, Irving Goffman describes expressions given (identity features we wish to present) vs. expressions given off (identity features beyond our control that others perceive) [16]. Sometimes people wish to be viewed in a particular way, but are perceived differently. Goffman’s dramaturgical perspective attributes this to poorly controlled aspects of self-presentation (e.g., looking nervous, not knowing the dress code). Donath extends this dichotomy to online worlds using signal theory: assessment signals are features of self that reliably indicate aspects of identity whereas conventional signals are less reliable but more easily controlled by the signaler [8]. For example, using an MIT.edu email address is a more reliable signal of membership in the MIT community than including this information in an email signature. Donath uses these concepts to explore identity deception online, and specifically notes efforts to remain anonymous by circumventing signals (for example, using anonymous remailers) are not always well received.

James Gee defines four aspects of identity including identities derived from some natural state (I am deaf), from institutions (I am a student), from discourse with others (I am charming), and from an affinity with a particular kind of experience or group (I am a Trekkie) [14]. Identities can be simultaneously held that conflict with one another in interesting ways. Gee uses the example of a child diagnosed with ADHD. An underlying natural state gives rise to a diagnosis (institutional identity) that can enforce certain identities over others, for example by curbing a potential discourse identity of the child as a troublemaker. In this example, we see the potential for anonymity, while often framed as freeing people from identity-related constraints [38], to also result in a reduction of control over one’s own identity narrative.

The above notions of identity feature opportunities for rifts between concept of self and perception of self by others. This rift became increasingly salient as we examined how service providers perceived contributors and compared it to contributors’ own identity-related privacy concerns.

3 RESEARCH DESIGN

We began this research by asking:

- (1) What challenges and threats to their mission are perceived by decision makers at service provider organizations that host online collaboration projects?
- (2) How do decision-makers perceive project contributors and their privacy practices?

- (3) How do these challenges and perceptions influence privacy and security-related decisions about policy and design?

To answer these questions, we conducted two complementary studies. The first, a semi-structured interview study, investigated decision-making about privacy and security at five different service provider organizations. The organizations we recruited from were selected because they host a diverse set of open collaboration projects where volunteers contribute content, data, code, and media. This study was intended to provide a subjective, phenomenologically grounded account of decision-makers' perceptions of threats and the reasoning behind privacy and security-related decisions [36]. We did not define "threats" for participants, but instead elicited the experiences that caused them to recognize threats to their organization. To bolster our ability to reliably interpret interview participants' accounts, we reviewed privacy-related policies and discussions associated with our research sites when available. The second study involved discourse analysis of public discussion forum threads on a mailing list about English Wikipedia. Whereas the interview study yielded abstractions derived from experiences at a variety of open collaboration service providers, the content analysis aimed to serve as a validity check for the findings from interview data, and to provide a concrete context in which to apply and discuss abstract concepts. By including a follow-up study of public data, we are able to overcome some of the limitations associated with holding interview participants' organizational affiliations in confidence.

Interview Study Design

Interview Participant Recruitment. Our recruitment strategy targeted people who work for service providers that support one or more open collaboration projects. We set out to obtain two to three interviews per site to ensure we acquired multiple perspectives, and because no one person has "ownership" of security- and privacy-related decision-making.

We sent email recruitment messages to people at service provider organizations to invite them to participate in a 30-90 minute interview. We made initial contacts through professional connections, at privacy and security-related events, and snowball sampling [31]. We cast a wide net by asking for people "involved in privacy and security-related decisions." We directed potential participants to an online consent form and scheduled interviews in follow-up emails. We interviewed participants using technologies they were comfortable with, such as Skype and Google Hangouts. They were offered \$25 in cash or Amazon gift card as a thank you, and seven participants declined the offer.

Although each interview participant sent out internal communications to encourage others to participate, they were

not always successful. When we spoke with multiple people at a single site, we found that they shared similar perspectives on perceived threats and goals of the open collaboration projects they support. In all, recruitment efforts yielded 11 interviews with participants from five service provider organizations, which we refer to as five distinct research sites. We spoke with two women and nine men, ages 30 to 47 for an average of 55 minutes. We spoke with four people at Sites 1 and 4, and one person at Sites 2, 3, and 5 respectively.

Because interview participants spoke about security issues, we ensured confidentiality for participants and the organizations they represented. This level of anonymization was a condition of our consent process and so we present aggregate identity characteristics of research sites in Table 1. In reporting results, we refer to interview participants as P1-P11 and have removed identifying details from quotes.

Data Collection and Analysis. Interviews were guided by a list of questions designed to explore challenges that service providers perceive in supporting open collaboration projects and how they address them. In each interview, we first asked interview participants to describe open collaboration projects supported by their organizations and how people participate. Borrowing from the privacy and security literature's reliance on threat models [27], the second part of the interview focused on perceptions of threats by asking participants to describe activities that cause problems for their sites and their responses to these activities. Demographic data were collected at the end of the protocol. The research was approved by IRBs of all authors.

Interviews were audio-recorded and transcribed with participant permission. The methodological basis for our analysis is the constant comparative method in which researchers iteratively collect and code data to identify concepts that are "integrated, consistent, plausible, close to the data" and ideally ready to be at least partially operationalized for further testing [15]. In initial open coding of interview transcripts, the first author flagged text that described threats. Then, in subsequent iterations, grouped the flagged text into themes using the qualitative data analysis software Dedoose. Separate iterations focused on distinct emergent concepts. The first and last author discussed initial codes and emergent themes that required further refinement and development and informally tested agreement on interpretations. All authors reviewed the transcripts to discuss the validity of themes and resolve any questions or differences of interpretation or emphasis.

We refer throughout to the people who contribute to open collaboration projects as "contributors," though our interview participants do not consistently refer to them that way. This distinguishes them from "users," who may encompass a broader group of consumers or readership to which a different set of privacy requirements and policies may apply.

Table 1: Characteristics of research sites

What's produced	Security Software, Other Software, Scientific Data (Citizen Science), Educational Content
Number of employees	<50: 1 site, 51-200: 2 sites, >200: 2 sites
Age of organization (in years)	19, 14, 14, 11, 7
Roles of interview participants	Researcher: 4, Engineer: 3, Security manager/director: 2, Developer: 1, Director: 1
Requires account to contribute	None of the sites
Allows contributions from Tor	Yes: 4 sites, No: 1 site

Content Analysis Study Design

Our first study yielded interpretations about how perceived contributor identity shapes privacy and security decision-making. We learned that the perceptions of contributors' identities—for example as “newbies” in need of anonymity protections—had influenced service providers, and we wondered if we had a full picture of how anonymous contributors were viewed. For our second study, we gathered posts from the WikiEn-l public mailing list archives where both volunteers and Wikimedia Foundation employees discussed various issues related to English Wikipedia between September 2001 and July 2017. Because Wikipedia has other venues where discussions occur, these data do not capture all communications among Wikipedia employees and volunteers.

We used an iterative approach to identify a sample of threads related to anonymous contributions. First, we identified all posts that included the term “anonymous” and familiarized ourselves with the discussions. We quickly learned that the term “anon” is commonly used to refer to anonymous contributors. We then identified all thread titles that contained the term (or prefix) “anon” and also selected a random sample of 50 posts containing the term “anon.” After coding a sample of 50 messages from each set, we concluded that only messages in threads with titles that contained the term were consistently relevant and had sufficient depth to contribute to our understanding of decisions. We identified 35 threads in which anonymous contributions were discussed between 2001 and 2010. The volume of correspondence diminishes in 2010, after which there appeared no thread titles with references to “anon.” We determined that early discussions held particular value because they happened during formative years for the project when norms and policy were being established [17]. Ultimately, we read a total of 655 messages, 605 of which were included in the final analysis of the 35 threads.

We analyzed these data in two parts: first, we developed a limited codebook based on the themes that were generated in study 1 to test for the presence of *contributor-perspective* and *organization-perspective* and applied those codes. Second, the first author iteratively open coded messages and produced memos in Microsoft Excel. Each message in each thread

was coded for concepts that related to value of anonymous participants, perceived problems, and solutions. The first author also wrote memos about features of discourse including rhetorical strategies, conceptualizations of anonymity, and perspective taking. The first and last author discussed interpretations and memos during this coding process. Although formal measures of interrater reliability are not required for the inductive analyses, the first and last author regularly tested their agreement on application of codes to specific data, and to verify conceptual integrity of emergent themes.

4 STUDY 1: SERVICE PROVIDER INTERVIEWS

The three themes that emerged from analysis of service provider interview transcripts included three major categories of threats:

- (1) *community norm threats*: violations of community norms for interaction and engagement such as harassment,
- (2) *volunteer threats*: loss of failure to attract volunteers to the project, and
- (3) *low quality contribution threats*: contributions that drain community resources.

We report participants' explanations for each of these types of threat, including examples of data that were coded as a threat description as well as interview participants' accounts of how they typically deal with various threats. We emphasize privacy-related decisions as well as how these three types of threats intersect with participants' thinking about privacy on the sites they support. Second, we discuss how perceived threats and conceptualizations of contributor identity inform service providers' policies related to use of privacy enhancing technologies and data collection.

Findings: Threats Perceived by Open Collaboration Service Providers

Community Norm Threats. The most common and often first threat identified by interview participants was harassment targeting gender, race, or other perceived features of contributors' identities. This excerpt reflects a common theme:

We sometimes see [contributors] being generally mean to each other, in a way that people can be mean in

open-source projects. Make attacks on people based on their perceived identity. If they think that someone is a black person, or a woman, and they think that those people should not be participating in software projects in any sort of way. They might be rude, so that's a problem. (P6)

The assessment that contributors “might be rude” stands in understated contrast to reports in the literature by contributors to projects that harassment can result in serious life changes like job loss and psychological distress [11].

Sites nevertheless take toxic interactions among members seriously not only because it is socially undesirable but because it can raise barriers for newcomers and harm retention of existing members. It is notable that interview participants frequently described this behavior as being perpetrated by established project contributors, rather than anonymous users or newcomers. A related toxic community behavior identified in interviews was the tendency for existing members to be untrusting and dismissive towards newcomers.

Some interview participants suggested that many of the undesired behaviors they observed were typical of the ways that people act online. For some, “kids” (P1, P4) and new contributors with no obvious malicious intent were scapegoats for deviant or unwelcome behaviors. Misogynist (P1, P6) or racist (P1, P6, P9) worldviews were seen as more insidious. As one interview participant pointed out, affecting change would require a more fundamental shift in community outlook and makeup:

We can't start kicking people out for being assholes or being too white or too male or too American or whatever... That, to some extent, can be [changed] directly if we can recruit more people who have... different cultural, gender, whatever backgrounds. But, there's no one intervention there that fixes a problem. We can't replace one community with another. (P9)

In general, service providers explained that they try to leave it to the members of project communities to decide what is abusive, and moderate it themselves. While some have formal policies about what constitutes abusive behavior, others said that they typically leave it to the project community to decide:

I don't know if we have sort of full written policies. For the most part, this is a sort of 'you know it when you see it' kind of thing. For a lot of our projects... we try to let each project sort of each community-manage itself... So it's really for each project community to sort of decide for itself what they consider to be abusive behavior. (P1)

This can result in problems where some contributors don't get reported because there is social pressure not to do so, particularly if abusive contributors are established community members. One interview participant mentioned that the

lack of tools for reporting certain types of abuse is a problem because it can mean that only the worst cases rise to the level of harassment:

Harassment causes problems... The fact that you have to be a community member for a while to figure out the ways to report your problems is a problem. (P8)

For some sites, moderating abuse requires access to contact information. For others, it requires giving moderators access to information about problem contributors, like IP addresses, that can be used to investigate and make a determination about whether to take action, like banning. As we discuss in later sections, some interview participants are agnostic on the subject of storing IP identifiers, while others feel it is the only way to address certain types of threats.

Volunteer Threats. Another major threat perceived by service provider sites was failure to attract new contributors or failure to retain existing ones. We found that these concerns led service providers to offer both more and less privacy protection in different contexts. Each of our research sites allowed various degrees of anonymity and offered a way to contribute pseudonymously; they rarely erected barriers to anonymity-seeking strategies. In all cases, pseudonymous accounts were easy to create. Although some sites required validated email addresses, none of the sites blacklisted email providers at the time of the interviews. All research sites provided a way to contribute without registering. However, as we will describe in detail in this section, account creation was required to use certain features of projects.

Offering contributors ways to protect their privacy was often seen as a way of lowering barriers to participation. To make it as easy as possible, sites create, “a whole bunch of different ways” (P6) to contribute, including ways that do not require account creation: “Yeah, so we just want to make it as simple as possible. We wouldn't want individuals put off contributing to a project by having to log in” (P2).

Some interview participants recognized that connecting contributors' identities with their contributions could be seen as a barrier and that concealing contributions helps newcomers avoid worry about how their contributions reflect on them. One participant opined that allowing pseudonyms helps because “people are intimidated and wouldn't want to ask questions or offer opinions if it was attached to [their real name]” (P3). That said, one interview participant was not sure that allowing people to contribute anonymously was worth the trade-off of potentially “dehumanizing” (P7) contributors and worth the risk that their acculturation might be hindered without establishing an identity on the site. This participant characterized the first problem as “being treated like a number” and further explained that other contributors will believe “there's a good chance you're a troll, there's a good chance that you're not someone who we'll have to deal

with again” and consequently are “more rough” with anonymous contributors (P7). This narrative suggests that those who can’t be identified are treated as second-class citizens of the community—taken less seriously at the outset and unable to move through the steps people typically go through as they become more central members of the community.

On some of the projects we spoke with, contributors who do not register cannot engage with other members using project-supported communication channels or get credit for their work. Some interview participants reported a preference for allowing only registered volunteers to get involved because of a perception by project leaders that “people will be better if they are logged in” (P3). Another described the value of registering (even with a pseudonym) to establish consistent identity:

I think the value of registration, generally, is to establish consistent identity. Whether that identity is pseudonymous or not, or anonymous or not... It just helps to establish identity when you are communicating with someone on a regular basis. (P6)

In general, although allowing anonymous contributions appeared to lower barriers, anonymous participation was often perceived to have negative effects on efforts to serve more established contributors and help people become part of a project community. For example, interview participants at multiple sites reported that supporting a better contributor experience and engagement for regular contributors equated with less privacy. Logging in with a persistent identity was seen as necessary for tailored feedback to improve quality of contribution, engagement, and efficiency:

I might see that you’re getting bored and decide to send you something more interesting to work on... we’re just assuming that if you’re not logged in, you’re getting random work. And none of the optimization happens to you. (P3)

Data about contributors can also be used to assess the quality of contributions. Interview participants suggested that having access to contributors’ histories is useful for understanding their commitment to the project both to reward them and to customize tasks based on their “abilities and their talents”:

At the moment, there are a couple of projects that will give feedback to the [contributors] dependent on their type of contributions, and we’re hoping to be able to provide, in the future, interventions for the individuals. So for example... things like badges for particular contributions. Obviously, if you weren’t logged in, you wouldn’t be able to build a profile of the contributions that an individual is making and be able to reward them. Also, if you were looking to do something more interesting in a project, such as passing tasks to an individual based on their abilities

and their talents. And you wouldn’t be able to do that if they weren’t logged in. (P2)

Low Quality Contribution Threats. Because the sites we studied exist to support collaborative production, low quality contributions are a threat to their success. People sometimes make low quality contributions such as buggy code or inaccurate data (which are rooted out by community and project leader oversight), but interview participants generally did not attribute these contributions to anonymous contributors. Rather, as one participant put it: “I can’t imagine people deliberately not logging in to be able to give contributions that are of poorer quality. That’s not really something that’s come up” (P2). An interview participant from a different field site said that low quality contributions were often by “people who either don’t know the rules, and then eventually follow them or choose to ignore them, but don’t care enough and keep coming back” (P7) regardless of identity status. Another participant reported that they don’t trust people who contribute to projects anonymously and suspect other people don’t either. They explained that some people:

explicitly don’t create an online identity within the community... I think that a lot of people distrust them. I certainly can say I distrust them. That’s not to say that I assume that every single person who does that is doing it with bad intentions. I assume that many of them have perfectly good reasons to do that. But, they’re an unknown quantity... trust is developed over time and through a history of interactions. Demonstrations of good intentions and proficiency and expertise and buy-in. By not having a persistent identity, there’s no way to establish that. (P9)

Findings: Implications for Contributor Data and Privacy Enhancing Technologies

The types of threats described above have implications for how service providers approach the question of collecting and storing contributor data and their perceptions of privacy enhancing technologies that allow contributors to achieve different degrees of anonymity. During interviews, we prompted participants to reflect on the kinds of contributor data they collect and store, and on anonymous contributions involving the use of Tor in particular, to elicit their thoughts on different types of anonymous participation.

Because service providers largely reported the value of anonymity as lowering barriers to contributing (rather than safeguarding identity) they tended not to equate anonymous participation with privacy enhancing tools like anonymous proxies. When prompted, one interview participant explained that “[i]f we can’t connect a thing that was added to an identity, then we have very little way to identify the likely motivation behind the contribution, which ends up being a really useful shorthand for vetting the quality of the

contribution” (P9). An interview participant with experience in developing privacy technologies summed up a critical problem for sites in general that want to allow contributions from anonymous proxies: “the biggest challenge that any corporation faces with privacy tools is that they’re not able to tell the difference between malicious [and non-malicious activity] this is what they’re trying to solve, that they’re not able to identify malicious activity versus real user interaction” (P11). Despite the fact that anonymous contributors were not described as common sources of low quality contributions, they were viewed by one site as a potential threat because service providers often relied on IP data as a tool to help identify and eliminate sources of bad contributions. In one case, an interview participant explained that:

We don’t have a lot of high profile abuse coming in over Tor. We don’t have any high profile abuse effectively at any time, from any user. So we don’t have to make a lot of hard decisions. I guess it’s easy for us to say we love our anonymous and Tor [contributors], but I’m pretty sure we would fight pretty hard for people’s ability to access using any system if it came to it. (P5)

Despite the apparent conflict with some of their pragmatic concerns, interview participants unanimously spoke of a commitment to privacy and touted the value of privacy enhancing technologies. Some participants expressed the belief that contributing to their projects was not something that would occasion the need for anonymity. As one participant put it, “we’ve always thought of anonymous users as the crowd of people who haven’t bothered to log in as opposed to a group of people who have chosen for whatever reason to be more anonymous on the internet and then thus using Tor or whatever” (P3). This perception of the contributor as not getting around to logging in fits neatly within the policy of not requiring users to log in to make it “as simple as possible” (P2) to contribute. It underscores service providers’ own concerns, and their difficulty imagining a broader range of experiences that might prompt their contributors to seek out anonymity. The idea that anonymity might be required elsewhere “on the internet” but not on their sites is mirrored in another interview participant’s belief that their site is “maybe different from other services” and that contributors are “probably not trying to hide their identity from us, they are just trying to hide their identity from other people on the internet” (P5). However, this same participant explained that their site had been approached by a government request for contributor records—one they successfully fought. This experience did not elicit reflection on reasons that people might wish to remain anonymous on their site.

The sites we spoke with had different policies for IP collection and storage. At least one site made the decision to hash IPs to protect user data from public exposure and government

surveillance. An interview participant from a site that doesn’t hash IPs acknowledged that it had caused some controversy, but also felt that IP hashing would prevent them from effectively addressing abuse. Another participant opined that “I think sort of our internal approach, was if we are not comfortable publishing this data in some form, then we should be very reluctant to collect it in the first place” (P7).

Several interview participants also said storing IPs is ineffective for addressing persistent abuse and limits contribution tracking. When someone is abusive toward other contributors, service providers may ban the offending IP address, but sometimes find that committed offenders find a way to return. Multiple sites acknowledged that anonymous contributions make it difficult to track contributor behavior or keep accurate contributor counts if they use a different IP for each session.

A Contradiction of Perspectives

We found that one of the biggest threats service providers perceive is when contributors make remarks that target the identity of another individual. These threats can trigger changes to policy at two levels. First, they may change what information the community decides to share (e.g., IP addresses) as a consequence of these threats. Second, service providers may decide to require or collect contributor data like IP address or real names to deal with harassment, although concealing their identity is precisely how contributors often avoid or respond to harassment by others.

The contradiction implicit in service providers requiring identifying data to address harassment and contributors concealing these same data to avoid harassment is an important tension identified in our analysis. Prior research has demonstrated that contributors to open collaboration projects who face risks include people whose identities create vulnerabilities; for instance, being female, being from an ethnic minority, or being transgender [11]. Other research has also identified cases in which open collaboration projects may exclude contributors who require privacy because it conflicts with norms of transparency. For example, in their study of a citizen science project, Bowser et al. found that the norm of “openness” is defined by those who feel that people who are “extremely privacy-conscious” simply cannot contribute and prompts the authors to label the community as a “self-selecting” group. [4].

5 STUDY 2: DISCOURSE ANALYSIS

Our initial study raised questions about the role that perceived contributor identity plays in shaping policy. To test and provide context for our interpretations of interview data and further explore (mis)alignment between contributor and service provider perspectives, we collected and analyzed mailing list discussions related to anonymous contributions

Table 2: Counts of threads and messages that included contributor or organizational perspectives

Total Number of Threads	35
threads w. contributor perspective	6 (17.1%)
Total Number of Messages	605
messages w. contributor perspective	35 (5.8%)
messages w. organization perspective	541 (89.4%)
messages w. neither perspective	29 (4.8%)

to English language Wikipedia. Critical discourse analysis is useful for looking at language relative to social, political, and cultural formations [34]. We consider how language used to construct the contributor-identity mediates relations of power and privilege in policy decisions.

Our prior analysis of interviews sensitized us to the divergent perspectives of service providers and their contributors. “Perspective-taking” became a sensitizing concept [3] that informed our analysis of posts in that we considered cases where organization or contributor perspectives informed discussions about anonymous contributions. These two categories of perspective taking constitute a concise codebook:

- *contributor perspective-taking*: consideration for the motives, knowledge, and needs of the contributor. E.g., if a message considers that a contributor might want to remain anonymous to avoid harassment.
- *organization perspective-taking*: consideration for the motives, knowledge, and needs of the organization. E.g., if a message talks about the threat of vandalism.

As the first author examined policy-related discourse and decisions to identify rhetorical strategies and outcomes, the codebook was used to determine the prevalence and rhetorical roles for each type of perspective-taking.

Findings

Our analysis of the mailing list threads supports our findings that service provider perspective-taking tends to support policies that overlook the identity-related vulnerabilities that contributors report [11]. Most debates about anonymous participation invoked the organization perspective (See Table 2), often centered on lowering barriers to participation and ensuring that contributors have a clear path from peripheral to core participation. The invocation of “vandal” to describe problematic anonymous contributors frequently swayed policy discussions whereas the conceptualization of anonymous contributors as vulnerable individuals coping with identity-related threats was rarely evoked and garnered little sympathy when it was.

When the perspective of contributors was taken into consideration, it was mainly to discuss anonymity as a privacy

strategy, with many comments pointing out that it is safer to login than to publicly expose one’s IP address. There was little discussion of threats that might prompt privacy seeking and when it was discussed, those who invoked the organizational perspective often did not feel that the benefit of providing more stringent protections for contributors who felt vulnerable outweighed the costs associated with such measures.

In the next sections, we describe these perspectives in more depth by analyzing discourse about specific policies. In doing so, we look at the ways in which contributor-identity is conceptualized to support certain groups over others. We adopt the Wikipedian nomenclature of referring to contributors who edit while not logged into an account as “anons.”

Banning Anonymous Edits. A proposal to ban anonymous edits resulted in articulations of why anons are valuable. When doubts were raised that anonymous contributions should be allowed at all, others responded that anons are valuable because they may eventually create accounts while “vandals” wouldn’t—an allusion to a popular stance that IP/anonymous edits make vandals easier to spot. Other arguments elaborate on how low barriers to participation make it easy for “newbies” to contribute by making anonymous edits before deciding to register.

Appreciation of anons has clear limits. They are valued insofar as they eventually legitimize themselves by creating a persistent identity that differentiates them from vandals. Additionally, some posters expressed views that anons should not have the same status as registered contributors and that anonymity cannot support the kind of trust necessary to build a reliable encyclopedia. The conceptualization of anons as having potential but not legitimate membership echoes service provider interview descriptions of anonymous contributors as second-class citizens.

New Article Creation. There were several instances of Wikipedians discussing whether to bar anons from specific tasks like creating articles and whether to hide “red links” (an affordance for new article creation) from them to limit vandalism and shift their focus to editing existing articles. While hiding red links is framed as a soft method of discouraging page creation, the argument was raised that it appears to go against wiki values and that Wikipedia’s success owes a lot to contributions from anons. Others point out that anons have the right to edit, but that they should prioritize improvement of existing articles. In this argument, the fact that anons have demonstrated value in building Wikipedia doesn’t translate into a right to choose the nature of future contributions. Someone proposed that a kinder and more effective approach than excluding anons from creating articles might be to display the reason for deletion if a new article disappears. If an anonymous “newbie” sees that their

newly created article was deleted for suspected vandalism, they might understand what happened and modify their behavior. In this case, the conceptualization of the anon as a valuable “newbie” who does not know how to participate effectively overpowers the “vandal” narrative resulting in softer, preemptive policies.

Blocking Anonymous Proxies. In one discussion, Wikipedia’s blocking policy is challenged. Specifically, it is pointed out that some sysops routinely block contributors who use anonymous IPs and an argument is raised equating editing anonymously with the right to free speech. The community is mostly outspoken in their rejection of this assertion, arguing that the law does not require Wikipedia to do anything. Although there ensues general agreement that concealing one’s identity promotes expression of diverse political ideas, and that Wikipedia should encourage that, there is a strong push to separate Wikipedia from any legal obligation to protect free speech. Although posters believe that Wikipedia should attempt to protect anonymous free speech, it should only do so insofar as does not harm Wikipedia itself.

Several other threads arise that discuss blocking anonymous proxies. One raises the question whether there are valid uses for anonymous proxies and posits that, because they are only used by people who are interested in doing nefarious things, a permanent block on all anonymous proxies would not be amiss. A further observation is made that concerns about contributor safety do not make valid use cases for determining Wikipedia policy or design—a view that goes largely uncontested.

These reflections on specific policy discussions provide insight into and concrete examples of the ways that organization perspective-taking result in policies that don’t allow for certain types of anonymous contribution with implications for potential contributors who perceive threats.

6 DISCUSSION AND IMPLICATIONS

We have described three major types of threats perceived by service providers who support open collaboration projects and explained how these threats affect service provider decisions about collecting and storing contributor data and how to handle privacy enhancing technologies. We also show how threats are linked to policies and technical decisions that limit anonymous participation and suggest that limitations stem from the way that contributors are conceptualized by service providers.

From these findings, we draw three key insights. The first is that with few exceptions, anonymity seekers—for example unregistered contributors or Tor users—are not perceived to pose the greatest threats to sites with whose representatives we spoke. Anonymous users were not discussed as frequently as registered users when it comes to violations

of community norms, the most commonly described type of threat. Second, concerns about maintaining sustainable participation levels by diverse contributors led open collaboration service providers to accommodate some forms of anonymous contributions. That is, more permissive policies about anonymous contributions were seen as advancing service providers’ goal of lowering barriers to contribution, but less permissive policies allowed service providers to improve contributors’ experiences and protect community norms. Third, the service providers with whom we spoke tend to emphasize the value of anonymity as it affects the process of entry to open collaboration projects and only secondarily as a protection once members have joined.

We triangulated these findings in our analysis of mailing list posts: contributors to policy discussions almost always consider the perspective of the organization and not the contributor, resulting in support for policies that don’t support certain types of anonymous participation. These findings lend further weight to our interpretation that how contributors are conceptualized and the identities bestowed upon them by service providers influences what protections they are entitled to and plays a role in policy decision-making.

Identities and Perspectives in Future Work

We framed identity as a pluralistic concept and perspective-taking as a window to understanding how identities are constructed and used by others. This proved to be a productive analytic lens for understanding how open collaboration projects can reproduce and reify systems of inequity through the development of privacy-related norms and policies that privilege the experiences of some contributors over others.

While attempting to lower barriers to participation for their imagined contributor base, service providers also unwittingly erect barriers for others, particularly individuals with more stringent privacy requirements. We argue that this is the result of a narrative around anons as a type of contributor who may require lower barriers to participation in order to get comfortable contributing (e.g., make “newbie” mistakes under the cloak of anonymity), but do not require anonymity as a prerequisite to full participation (e.g., conceal their location for reasons of safety). From prior work, we know that identity facets like race, gender, or sexual orientation can create vulnerabilities that cause people to seek anonymity or curtail their online contributions to protect themselves. For example, in a study of contributors’ privacy strategies, an open collaboration contributor reported using Tor to avoid being outed to his/her/their employer, and another took death and rape threats seriously enough to reduce their editing activity on Wikipedia [11]. These contributors’ experiences suggest that seeking anonymity is an important tool for some would-be contributors; yet, we found that such identity-based vulnerabilities are not often understood

by service providers or may not be perceived as legitimate. Service providers' decisions and interpretations of anonymous users' motivations are largely grounded in a narrative that reflects the experiences of more visible and arguably privileged contributors.

Anonymous users by definition have limited ability to control perceptions of who they are and what their goals are. Our next steps include better characterizing not only contributions of anonymous contributors as a class, but also the ways that individual anons might signal good faith, interest and goals in the absence of a persistent identity. In the first case, we aim to test interpretations of why people seek anonymity with a larger population through natural experiments that compare the type and quality of anonymous or pseudonymous contributions to open collaboration projects in different conditions. In the second case, design experimentation is useful for exploring how "anonymous" contributors to projects might be supported in making more informed choices about revealing forms of identity knowledge and how they might signal their intentions.

As long as little is known about anonymous contributors, their motivations, and the value of their work in open collaboration projects, it is unsurprising that service providers reason from their own experiences and those of central project contributors when considering anonymity seekers. The relative invisibility of anonymous contributors also raises the more general question of how well they are represented in the construction of social norms and technical requirements.

Norms as Analytical Tools

Social norms are a powerful concept in the HCI literature, underpinning important insights about a range of online phenomena. Although they are important features of social systems, when norms become a dominant analytical yardstick by which to assess the "fit" of computing systems with social phenomena like standards of privacy, our conceptual tools may become complicit in erecting selective barriers to participation.

Our studies rendered visible the ways that service providers perceive contributors and threats differently than contributors view themselves, as reported in the literature. These divergent narratives highlighted the limitations of shared norms and expectations as analytical tools. We suggest that, for open collaboration projects, shared social norms can be most useful in understanding the experiences of central community members whose ability to participate with minimal risk helps stabilize the systems in the first place. Conversely, shared norms as analytical tools leave out perspectives of those who may have been alienated from the norm articulation process, calling into question the value of frames that are grounded in community articulation of norms. As a final note, we observe that without using analytical frames that

explicitly consider perceptions of risk and threats as a feature of participation, researchers of sociotechnical systems—like service providers—are likely to overlook certain experiences.

7 CONCLUSIONS AND LIMITATIONS

Our findings are grounded in the experiences of a small sample of decision makers and online discussions of central community members; however, the experiences and values interview participants reported and posted about were mutually supportive of the interpretations presented in this paper. Overall, the service providers with whom we spoke value anonymous contributions and do not see them as a threat to their sites, but neither do they prioritize anonymity in the same ways that people seeking to contribute anonymously do. Our findings illuminated perceived threats and the conceptualizations of contributors that inform open collaboration service providers' decisions about privacy-related technologies and policies. We conclude by raising a critique of social norms as a tool for understanding privacy concerns.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (awards CNS-1703736 and CNS-1703049).

REFERENCES

- [1] Irwin Altman. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding* (Clean & Tight Contents edition ed.). Brooks/Cole, Monterey, California.
- [2] Irwin Altman. 1977. Privacy Regulation: Culturally Universal or Culturally Specific? *Journal of Social Issues* 33, 3 (1977), 66–84.
- [3] Herbert Blumer. 1954. What is wrong with social theory. *American Sociological Review* 19, 1 (1954), 3–10.
- [4] Anne Bowser, Katie Shilton, Jenny Preece, and Elizabeth Warrick. 2017. Accounting for Privacy in Citizen Science: Ethical Research in a Context of Openness. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 2124–2136. <https://doi.org/10.1145/2998181.2998305>
- [5] Susan L. Bryant, Andrea Forte, and Amy Bruckman. 2005. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 ACM International Conference on Supporting Groupwork (Group) (GROUP '05)*. ACM, New York, NY, USA, 1–10. <https://doi.org/10.1145/1099203.1099205>
- [6] Nicholas Diakopoulos and Mor Naaman. 2011. Towards Quality Discourse in Online News Comments. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. ACM, New York, New York, 133–142. <https://doi.org/10.1145/1958824.1958844>
- [7] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The Second-generation Onion Router. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13 (SSYM '04)*. USENIX Association, Berkeley, CA, USA, 21–21. <http://dl.acm.org/citation.cfm?id=1251375.1251396>
- [8] Judith S. Donath. 1998. *Identity and deception in the virtual community* (Peter Kollock and Marc Smith (eds.) ed.). Routledge, London, UK, 29–59.

- [9] Peter Eckersley. 2010. How Unique Is Your Web Browser?. In *Privacy Enhancing Technologies (Lecture Notes in Computer Science)*. Springer, Berlin, Germany, 1–18. https://doi.org/10.1007/978-3-642-14527-8_1
- [10] Heather Ford and Judy Wajcman. 2017. 'Anyone can edit', not everyone does: Wikipedia's infrastructure and the gender gap. *Social Studies of Science* 47, 4 (Aug 2017), 511–527.
- [11] Andrea Forte, Nazanin Andalibi, and Rachel Greenstadt. 2017. Privacy, Anonymity, and Perceived Risk in Open Collaboration: A Study of Tor Users and Wikipedians. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 1800–1811. <https://doi.org/10.1145/2998181.2998273>
- [12] Andrea Forte and Cliff Lampe. 2013. Defining, Understanding, and Supporting Open Collaboration: Lessons From the Literature. *American Behavioral Scientist* 57, 5 (May 2013), 535–547.
- [13] Gerald Friedland and Robin Sommer. 2010. Cybercasing the Joint: On the Privacy Implications of Geo-tagging. In *Proceedings of the 5th USENIX Conference on Hot Topics in Security (HotSec '10)*. USENIX Association, Berkeley, CA, USA, 1–8.
- [14] James Paul Gee. 2000. Identity as an Analytic Lens for Research in Education. *Review of Research in Education* 25 (2000), 99–125. <https://doi.org/10.2307/1167322>
- [15] Barney Glaser and Anselm Strauss. 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine Transaction, New Brunswick, NJ, USA.
- [16] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Doubleday Anchor, Garden City, NY, USA.
- [17] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2013. The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (2013), 664–688.
- [18] R. Henry and I. Goldberg. 2011. Formalizing Anonymous Blacklisting Systems. In *2011 IEEE Symposium on Security and Privacy*. IEEE, New York, NY, USA, 81–95. <https://doi.org/10.1109/SP.2011.13>
- [19] Daniel C. Howe and Helen Nissenbaum. 2009. *TrackMeNot: Resisting surveillance in web search* (ian kerr, carole lucock, and valier m. steeves ed.). Oxford, New York, NY, USA, 417–436.
- [20] Corey Brian Jackson, Kevin Crowston, and Carsten Østerlund. 2018. Did they login?: Patterns of anonymous contributions in online communities. *Proc. ACM Hum.-Comput. Interact.* 2 (Nov. 2018), 77:1–77:16. <https://doi.org/10.1145/3274346>
- [21] Ruogu Kang, Stephanie Brown, and Sara Kiesler. 2013. Why Do People Seek Anonymity on the Internet?: Informing Policy and Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2657–2666. <https://doi.org/10.1145/2470654.2481368>
- [22] Cliff Lampe and Paul Resnick. 2004. Slash(Dot) and Burn: Distributed Moderation in a Large Online Conversation Space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 543–550. <https://doi.org/10.1145/985692.985761>
- [23] Robert S. Lauder and Maxine Wolfe. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues* 33, 3 (Jul 1977), 22–42. <https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>
- [24] Xiao Ma, Jeff Hancock, and Mor Naaman. 2016. Anonymity, Intimacy and Self-Disclosure in Social Media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3857–3869. <https://doi.org/10.1145/2858036.2858414>
- [25] Gary T. Marx. 1999. What's in a Name? Some Reflections on the Sociology of Anonymity. *The Information Society* 15, 2 (1999), 99–112.
- [26] Amanda Menking and Ingrid Erickson. 2015. The Heart Work of Wikipedia: Gendered, Emotional Labor in the World's Largest Online Encyclopedia. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 207–210. <https://doi.org/10.1145/2702123.2702514>
- [27] Suvda Myagmar, Adam J. Lee, and William Yurcik. 2005. Threat modeling as a basis for security requirements. In *StorageSS '05: Proceedings of the 2005 ACM workshop on Storage security and survivability*. ACM, New York, NY, USA. <https://doi.org/10.1.1.703.8462>
- [28] Dawn Nafus. 2012. 'Patches don't have gender': What is not open in open source software. *New Media & Society* 14, 4 (Jun 2012), 669–683.
- [29] Arvind Narayanan and Vitaly Shmatikov. 2009. De-anonymizing Social Networks. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy (SP '09)*. IEEE Computer Society, Washington, DC, USA, 173–187. <https://doi.org/10.1109/SP.2009.22>
- [30] Helen Nissenbaum. 2010. *Privacy in context: technology, policy, and the integrity of social life*. Stanford Law Books, Stanford, Calif.
- [31] Michael Patton. 2001. *Qualitative Research & Evaluation Methods* (3rd edition ed.). SAGE Publications, Inc, Thousand Oaks, CA, USA.
- [32] Tom Postmes, Russel Spears, and Martin Lea. 1998. Breaching or Building Social Boundaries?: SIDE-Effects of Computer-Mediated Communication. *Communication Research* 25, 6 (Dec 1998), 689–715. <https://doi.org/10.1177/009365098025006006>
- [33] Lee Rainie, Sara Kiesler, Ruogu Kang, and Mary Madden. 2013. *Anonymity, Privacy, and Security Online*. Pew Research Center, Washington, DC, USA. <http://www.pewinternet.org/2013/09/05/anonymity-privacy-and-security-online/>
- [34] Rebecca Rogers, Elizabeth Malanchruvil-Berkes, Melissa Mosley, Diane Hui, and Glynis O'Garro Joseph. 2005. Critical Discourse Analysis in Education: A Review of the Literature. *Review of Educational Research* 75, 3 (2005), 365–416.
- [35] Sarita Yardi Schoenebeck. 2013. The Secret Life of Online Moms: Anonymity and Disinhibition on Youbemom.Com. In *Proceedings of the 7th International Conference on Weblogs and Social Media, ICWSM 2013 (ICWSM '13)*. AAAI, Palo Alto, CA, USA, 555–562. <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/5973>
- [36] Alfred Schutz. 1967. *The Phenomenology of the Social World*. Northwestern University Press, Evanston, IL, USA.
- [37] Daniel J. Solove. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154, 3 (2006), 477–564.
- [38] John Suler. 2004. The online disinhibition effect. *Cyberpsychology & Behavior: The Impact of the Internet, Multimedia and Virtual Reality on Behavior and Society* 7, 3 (Jun 2004), 321–326.
- [39] John R. Suler and Wende L. Phillips. 1998. The Bad Boys of Cyberspace: Deviant Behavior in a Multimedia Chat Community. *CyberPsychology & Behavior* 1, 3 (Jan 1998), 275–294.
- [40] P. P. Tsang, A. Kapadia, C. Cornelius, and S. W. Smith. 2011. Nymble: Blocking Misbehaving Users in Anonymizing Networks. *IEEE Transactions on Dependable and Secure Computing* 8, 2 (Mar 2011), 256–269.
- [41] Samuel D Warren and Louis D Brandeis. 1890. The Right to Privacy. *Harvard Law Review* 4, 5 (15 Dec 1890), 193–220.
- [42] Alan F. Westin. 2003. Social and Political Dimensions of Privacy. *Journal of Social Issues* 59, 2 (Jul 2003), 431–453. <https://doi.org/10.1111/1540-4560.00072>
- [43] Pamela Wisniewski, A.K.M. Najmul Islam, Bart P. Knijnenburg, and Sameer Patil. 2015. Give Social Network Users the Privacy They Want. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 1427–1441. <https://doi.org/10.1145/2675133.2675256>