Privacy, Anonymity, and Perceived Risk in Open Collaboration: A Study of Tor Users and Wikipedians

Andrea Forte, Nazanin Andalibi, Rachel Greenstadt
College of Computing and Informatics, Drexel University
Philadelphia, PA, USA
{aforte, naz, rachel.a.greenstadt}@drexel.edu

ABSTRACT

This qualitative study examines privacy practices and concerns among contributors to open collaboration projects. We collected interview data from people who use the anonymity network Tor who also contribute to online projects and from Wikipedia editors who are concerned about their privacy to better understand how privacy concerns impact participation in open collaboration projects. We found that risks perceived by contributors to open collaboration projects include threats of surveillance, violence, harassment, opportunity loss, reputation loss, and fear for loved ones. We explain participants' operational and technical strategies for mitigating these risks and how these strategies affect their contributions. Finally, we discuss chilling effects associated with privacy loss, the need for open collaboration projects to go beyond attracting and educating participants to consider their privacy, and some of the social and technical approaches that could be explored to mitigate risk at a project or community level.

Author Keywords

Wikipedia; Tor; Risk; Privacy; Identity

ACM Classification Keywords

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

INTRODUCTION

Bassel Khartabil—open source developer, Wikipedia editor, and founder of Creative Commons Syria—was detained by Syrian authorities in March 2012 and jailed for three years before disappearing in October 2015. His contributions to open collaboration projects include founding the #NewPalmyra project in 2008 to digitally preserve the world heritage site at Palmyra by using satellite and high-resolution imagery to create open 3D models of ancient structures. Many of these structures would be razed in 2015 by Daesh. The reasons for Bassel's detainment and disappearance are

not documented, but are thought to hinge on his activism and open collaboration projects. As of this writing, his location is unknown. [12]

CSCW researchers have often investigated aspects of open collaboration projects like open source software or Wikipedia, but seldom frame participation in such projects as a process of negotiating risk. The groundswell of open collaboration projects throughout the 1990s and 2000s became an unprecedented source of data for those interested in understanding computer-mediated cooperative practices. Yochai Benkler theorized commons-based peer production as a new economic and organizational model [5]. Researchers asked questions like, "How do contributors organize their efforts?" [13, 23] and "What do contributors learn through participation?" [18]. Threaded throughout such studies were implicit and often privileged assumptions about the virtues and value of participation and about the safety of online spaces for participants. Wikipedia in particular has been viewed as archetypical of a paradigm shift toward freeing information from traditional economies of production and "democratizing" knowledge by offering widespread opportunities to contribute. Yet, even Wikipedia is not a uniformly safe place for everyone.

Bassel Khartabil's story is a dramatic and public example of the kinds of risks that contributors may face when they participate in online projects, but there exist many more examples of people whose efforts to contribute to open collaboration projects put them at risk, some mundane and some exceptional. Said Wikipedia founder Jimmy Wales when announcing the 2015 Wikipedian of the year, "I learned of a remarkable case of a Wikipedian, but to name this person would put them into danger, and I've been asked not to do it. And so the Wikipedian of the year 2015 is ...in pectore." *In pectore* is a Latin term used by the Catholic Church to refer to those whose recognition by the Pope might put them in danger of harm or discrimination by their local communities.

In this paper, we use a phenomenological approach to examine the threats that people perceive when contributing to open collaboration projects and how they maintain their safety and privacy. Our goal is to help inform policies and technical infrastructures that support open collaboration projects in attracting, sustaining, and protecting a diverse volunteer base. Our findings are anchored in interviews with 23 individuals, each of whom have considered how

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

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

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

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

DOI: <http://dx.doi.org/10.1145/2998181.2998273>

and whether to protect their privacy when participating in online projects. Their experiences range from men and women whose participation in open collaboration projects attracted death and rape threats, to citizens who fear sanctions from their government, to people who are concerned about having facets of their identity “outed” through their participation, to people who perceive few if any threats. We examine the experiences and tactics of these individuals within the context of the policies, norms and technological infrastructures of the projects to which they contribute in order to answer the questions:

RQ1. What kinds of threats do contributors to open collaboration projects perceive?

RQ2. How do people who contribute to open collaboration projects manage risk?

We will present risks associated with participation in open collaboration projects and strategies that participants use to mitigate these risks. We conclude by discussing how sites might develop infrastructures to successfully control for abuse and promote high-quality contributions without identifying volunteers who wish to remain anonymous.

RELATED WORK

Participation in Open Collaboration Projects

Much of the vitality and innovation credited to the Internet is a product of volunteerism and participation on a massive scale. From product and service reviews, to open source software, encyclopedias, citizen science and journalism, open collaboration projects involve volunteers generating valuable content and products. The prototypical open collaboration project has been defined as “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” [19]. Commons-based peer production is often invoked to understand such collaboration as an organizational phenomenon that challenges traditional corporate structures and gives rise to new economic opportunities [6].

The perceived value of open collaboration projects has prompted research on how they succeed and how to sustain them including attracting volunteers [28], enculturating new members into the norms of the project [11], and sustaining their participation over time [16]. One particularly vexing problem for Wikipedia has been to encourage participation from a diverse set of contributors to ensure representation of diverse content. Research has demonstrated that inequities along demographic lines such as gender can lead to inequities in the content [24]. Because Wikipedia’s content reflects its authors’ interests and expertise, maintaining editor diversity is an issue of encyclopedic quality and coverage.

Understanding barriers to diverse participation is difficult because the sampling frame can be impossibly large, but researchers have studied some reasons people choose not to participate in online venues. In a scholarly investigation of lurking practices, Preece et al. [31] identified reasons for non-participation in online communities: not needing to post, needing to know more about the group before posting, the belief that not posting is more helpful than posting, poor usability, and not liking the dynamics of the community. Ardichvili [2] recently reviewed research on participation in online communities and provided four categories of reasons for not participating: interpersonal (e.g., fear of criticism), procedural (e.g., not knowing best practices), technological (e.g., lack of technological skill), and cultural (e.g., in-group orientation, saving face).

In the literature on non-participation and lurking, concerns about privacy, confidentiality, and security are sometimes identified explicitly, for example in Baumer et. al’s examination of why people leave social media [4], but often privacy concerns have been implicitly connected to factors such as fears of criticism or discomfort with power relationships. However, there exists complementary literature that focuses on online privacy and perceptions of risk and can help illuminate these concerns.

Online Privacy and Perceptions of Risk

Palen and Dourish have observed that “active participation in the networked world requires disclosure of information simply to be a part of it” [29]. They suggested that managing privacy involves the continuous negotiation of boundaries along three dimensions: disclosure (i.e. need for publicity in tension with need for privacy), identity (i.e., presentation of self in tension with disparate audiences), and temporality (i.e., past/present/future are in tension). Negotiation of these boundaries is context dependent. Because it entails disclosure, all online participation raises issues of privacy and risk, but in some contexts privacy negotiations are trivial or routinized to the point of being unnoticeable, whereas in other contexts they require extensive attention and effort. Nissenbaum explains this variability in terms of “contextual integrity” [27]. She argues that although norms and expectations about privacy and constraints on the flow of information can vary dramatically, they are systematically related to the context in which information is disclosed.

We view participating in open collaboration projects like Wikipedia or open source as a set of related contexts in which people establish privacy expectations and concerns that may be dissimilar to concerns about sharing photos, chatting with friends, taking online classes, or other forms of online participation. Participation in projects must be analyzed as contexts that are systematically distinct from other contexts in which people share information online. Moreover, projects themselves are different contexts within which people establish expectations.

Perceived characteristics of technical communications infrastructures and the organizations that deploy them are important features of contexts in which privacy concerns are embedded. An individual's digital trail can be captured at different resolutions depending on how technologies are configured: clicks, keystrokes, and even cursor position can be captured and analyzed, but participants in open collaboration projects may have discrepant understandings of the technologies they use. Horsman observed that Edward Snowden's revelations about government surveillance mechanisms have sensitized people to vulnerabilities when they use communications technologies and called this the "Snowden effect" [20]. Yet, it's not only people's perceptions of technologies that influence privacy concerns, but also their perceptions of the organizations that deploy the technologies. Smith and colleagues identified four facets of privacy concerns about organizational information privacy practices: collection and storage of large amounts of personal data, unauthorized secondary use of personal data, errors in collected data, and inappropriate access to personal data [33]. For example, Wikipedia editors may not know that when they contribute to the site, their actions on the site are publicly logged and also made available by the Wikimedia Foundation in public datasets.

The purpose of online participation also influences perceptions of risk. Kang et al. [21] conducted interviews and found that people used anonymity for diverse kinds of online participation and suggest five personal threat models: online predators, organizations, known others, other users on the site or in the community, and unknown others. They found that regardless of technical literacy, half of their participants obtained anonymity by simply not participating, limiting the information they shared, or sharing incorrect information about themselves.

Anonymous Online Participation

When people are concerned about privacy, they take steps to protect themselves from observation. As mentioned above, one approach is to abstain from participation. Another is to seek anonymity. For sociologist Gary Marx, achieving full anonymity means subverting seven dimensions of identity knowledge including legal name, location, behavior patterns, social group membership, identifying personal characteristics, pseudonyms that can be linked with other forms of identity knowledge, or pseudonyms that cannot be linked and serve as alternate identities [25]. We join Marx and other researchers in grounding our work on the premise that anonymity is not a binary construct but instead can be achieved to different degrees. For example, people may use a fake name, temporary technical identities like throwaway accounts [1] or sockpuppets [35] that allow them to participate online in ways that are disconnected from their legal identities, although still may be traceable by service providers.

In this study, we recruited users of the anonymity network Tor. The Tor Project, named after the early project moniker

"the onion router," produces free, open-source software that conceals users' location and Internet use from potential observers [15]. Even tools that subvert most forms of tracking, such as the Tor Internet browser, do not guarantee "true" anonymity as people can still disclose forms of identity knowledge either intentionally or unintentionally through behavioral or linguistic patterns [9].

As noted above, there are differences among open collaboration projects in terms of expectations of privacy and how participants define concepts like anonymity and risk. For example, Wikipedia editors commonly refer to edits that are tagged with a user's IP address as "anonymous edits," whereas IP is considered to be highly identifying by contributors to the Tor Project.

We set out to understand what kinds of perceived risks lead to tensions between anonymity seeking and participation in open collaboration projects and how people manage these tensions.

STUDY DESIGN

To answer our research questions, we conducted semi-structured interviews with 23 participants. Our approach is informed by social phenomenology, which begins with the understanding that scientists study and interpret social worlds comprised of humans who themselves interpret, respond to, and actively construct their social environments [32]. An interview-based phenomenological approach is appropriate for studying social constructs like privacy and related practices because they are linked with people's subjective experience and understandings of the sociotechnical systems in which they participate. In phenomenological research, investigators strive to "bracket" their own understandings of phenomena and privilege the reported experiences of participants [7]. This approach is particularly important to remember when we discuss perceptions of threats or potentially "sensitive" encyclopedia topics: in this study, we bracket our own understandings of these concepts and report on participants' own accounts of what constitutes a threat or a sensitive topic. We used semi-structured interviews to allow participants the freedom to explain their experiences and then used thematic analysis [8] to identify pervasive and salient concepts as explained below.

Recruitment

Our recruitment strategy targeted two groups: Tor users and Wikipedia editors. First, we sought interviews with Tor users who have also contributed to "online projects." By virtue of their Tor use, these participants are known to take steps to protect their privacy while using the Internet, but would likely have diverse experiences with open collaboration. Second, we sought interviews with Wikipedia editors who have considered their privacy while editing. By virtue of their Wikipedia editing, these participants are known to contribute to open collaboration projects, but would likely have diverse definitions and experiences of privacy. By independently targeting both

privacy concerns and participation in open collaboration as primary characteristics for recruitment, we aimed to capture experiences that ranged from heavy to light participation and extensive to minimal privacy concerns.

To reach these groups, we used similar approaches. We sent a message to the general discussion mailing list for the Tor project and posted links on social media to our recruitment materials. In addition, the recruitment materials were reposted and blogged by official Tor Project accounts. To reach Wikipedia users, we posted messages on Wikipedia discussion lists and on social media. Both efforts also yielded participants via word-of-mouth.

Participants

In all, recruitment efforts yielded 23 interviews—12 with Tor users and 11 with Wikipedia editors, all of which were conducted in spring and summer of 2015. (see Table 1).

12 Tor Participants	Gender	Male: 8 Female: 3 1 reported fluid gender
	Age	Min: 18 Max: 41 Avg: 30
	Location	Austria, Belgium, Canada, South Africa, Sweden, United States (7 from northwest, central, midwest, east coast regions)
	Education	High School/Secondary: 3 Undergraduate: 4 Masters: 3 PhD: 2
11 Wikipedia Participants	Gender	Male: 8 Female: 3
	Age	Min: 20 Max: 53 Avg: 30
	Location	Australia, France, Ghana, Israel, Sweden, United Kingdom, United States (5 from west coast, midwest, east coast regions)
	Education	Undergraduate: 8 Masters: 1 PhD: 1 Unreported: 1

Table 1: Participants

Because of some interviewees’ concerns about reporting characteristics like age, gender, education, and location in conjunction with details about their privacy concerns, we report on aggregate features of participants’ backgrounds. Some of our interviewees had concerns about the potential for de-anonymization, and out of respect for these concerns we likewise do not identify participants using consistent pseudonyms. People who responded to Tor recruitment materials are categorized as Tor users and people who responded to Wikipedia materials as Wikipedia editors, although several Tor users had also edited Wikipedia and a

few Wikipedia editors had used Tor. In no cases did participants speak with equal reflection and detail about both Tor and Wikipedia; as described in our findings, the two populations we sampled yielded participants with different experiences and areas of expertise.

Interview Protocol

Potential participants were directed to online consent materials and given the option of either filling out an online form or mailing us a paper consent form if they were uncomfortable with the online form. They were also given the option of receiving \$20 compensation for their time in the form of gift cards or cash, which required further paperwork. About one third of participants declined payment. In most cases, interviews were conducted via phone, Skype, or an encrypted audio channel. In all of these cases, participants agreed to allow the interviewer to make a local audio recording. In one case, the participant met with the interviewer face-to-face and requested that the interviewer only collect written notes. The first author conducted all interviews.

Interviews were guided by a list of topics. Standard questions for Tor users included explaining Tor, what it is for, current examples of use, retrospective examples of use, the story of how and why they first started using it, examples of when they do not use Tor, and questions about their participation in online projects. Tor users who had edited Wikipedia were also asked questions from the protocol for Wikipedia editors: how and why they started editing, examples of privacy concerns associated with editing, steps they have taken to protect their privacy while editing, and examples of interactions with other editors. The interviewer encouraged participants to tell in-depth stories and prompted them for as much detail as possible. Most of our data pertains to open collaboration projects such as editing Wikipedia, contributing to other online information resources, or working on open source software but at times participants also delved into their experiences in discussion groups, political activism or other online groups and activities. Demographic data reported in the previous section were collected during interviews.

Data Analysis

All but one interview were audio recorded and transcribed. The interview that was not audio recorded was captured via written notes. We used thematic analysis [8] to analyze transcripts and notes. The first author used the software Atlas.ti to identify themes in interview transcripts using participants’ own language. After themes were identified, all authors reviewed them and discussed connections among themes. The first author then collapsed themes into affinity groups. For example, themes related to threats were organized into threat.source and threat.type. In some cases, similar categories were merged, for example, threat types of harassment, bullying, and intimidation were then collapsed into one recurrent theme: threat.type.intimidation. The transcripts were read by all team members and emergent

themes were discussed at multiple points; all stages of analyses were done by the first author and reviewed by coauthors. The analysis took approximately 6 months to complete. We report salient themes in the next sections as interconnected features of participants' practice. When participant practices diverge, we use internally consistent features of their experiences and explanations to account for these differences.

Limitations

The need to use communication technologies to interview remote participants means that individuals who volunteered to participate in this study were comfortable speaking with us remotely over phone or other communication channels. We accommodated participants' requests for encrypted communication and provided assurances of both confidentiality and anonymization; however, potential interviewees with acute privacy concerns would not likely be comfortable participating in this study. For example, complaints appeared in comments on the Tor Project blog that we mentioned using unencrypted channels like Skype in our recruitment materials and that this prohibited participation by Tor users. Additionally, some Wikipedia editors recommended interviewing editors in Iran and China and even forwarded our recruitment materials to specific individuals, but no-one living in Iran or China volunteered to participate (although we interviewed one Iranian citizen who lived in the United States at the time of the interview.) Laws in Iran and other places specifically forbid use of Tor and other privacy-enhancing technologies and China generally blocks Wikipedia so these are examples of voices that are known to be missing from our account of privacy concerns in open collaboration.

FINDINGS

Our findings are organized in four main sections:

- **types of threats** perceived by contributors to open collaboration and other online projects,
- **perceived sources of threats**,
- **experiences of those who had few or no concerns** about their participation in online projects, and the
- **strategies that people used to mitigate perceived risk**, including *modifying participation* in projects and *enacting anonymity*.

Types of Threats

By design, participants in our study had diverse experiences of contributing to open collaboration projects. In nearly all interviews, participants described being wary about how aspects of their participation in open collaboration projects could compromise their privacy or safety. Many participants described crisis experiences of their own or of someone they knew as antecedent to their model of threat in online projects. In this section we explain the five types of threats that surfaced most frequently. None of them exist in isolation from others and dependencies frequently surfaced.

Table 2 reports the prevalence of the five most prevalent types of threats reported by participants. All threat types

were discussed by at least one Wikipedia editor and one Tor user. Still, the two participant types differ in some ways. As reported in the table, Tor users were more likely to discuss the threat of surveillance and Wikipedia users more likely to discuss threats of harassment and intimidation. Tor users' concern with surveillance could be anticipated; perhaps more surprising is the extent to which Wikipedia users found themselves in disagreements and conflicts as a part of their encyclopedia writing work that led to acute fears, threats, and concomitant privacy concerns.

Threat	# Participants		
	Total	WP	Tor
Surveillance/ Loss of privacy	12	3	9
Loss of employment/opportunity	10	4	6
Safety of Self/Loved Ones	9	5	4
Harassment/Intimidation	9	8	1
Reputation Loss	6	4	2

Table 2: Types of Perceived Threats

Threat 1: Surveillance/Loss of privacy

The most common concern voiced by participants was a fear that their online communication or activities may be accessed or logged by parties without their knowledge or consent. This loss of control over their personal information can also be seen in several of the other threat types, but most frequently surfaced as a general concern. In some cases, especially among Tor users, this concern reflected a general desire for online actions to be private-by-default, public-by-effort. Explained one Tor user from Europe, "You know in my country there's basically unknown surveillance going on... and I don't know what providers to use so at some point I decided to use Tor for everything." Several Tor users viewed the use of personal information for targeted advertising as gratuitous and invasive.

Some contributors to Wikipedia commented that their edit histories contain sensitive information. Wikipedia logs a public record of every edit made by every account. Much as Internet search histories are a record of a person's interests and needs, edit histories provide a log of Wikipedia editors' interests and expertise. One Wikipedia editor pointed out that he didn't want his edit history to be linked with his real identity. Part of it, he explained, "is just frustration that I don't have a choice." But, like others, he felt that his edit history could reveal aspects of his identity to unintended audiences that he would rather not engage, for example he explained, "Employers, I wouldn't want them to see my editing history. But it's mostly that I don't want it mandated that this stuff is linked to my identity online."

Threat 2: Loss of employment/opportunity

The potential for participation in open collaboration projects to reveal aspects of contributors' identities that could compromise their professional or educational goals surfaced repeatedly. In some cases, this was due to potential

for discrimination against specific aspects of participants' activities or identity. One Tor user remarked,

I am transgender. I am queer... my boss... would rant for hours about this kind of person, that kind of person, the other kind of person, all of which I happen to be. And I decided that if I was going to do anything [online] at all, I had best look into options for protecting myself because I didn't want to get fired.

In other examples, Wikipedia editors described editing articles on sexual health or drug abuse because they believed it was important for people to have access to sound information on these topics, but were wary about the potential for these edits to be seen and misinterpreted by professional colleagues or potential employers.

Wikipedia administrators in particular were concerned that their contributions might attract backlash that could eventually result in lost opportunities as a secondary effect. One Wikipedia administrator who was involved in resolving controversy with a disgruntled editor described her own experience and observed that another administrator

...was also involved in the discussion, and he actually got it worse than I did. He's in a position now where anyone who Googles him finds allegations that he... is this awful monster, and he's terrified of having to look for work now because you Google him and that's what you find.

The fear of losing professional opportunities surfaced repeatedly as people described their own experiences with harassment as well as others' experiences.

We have noted that we take a phenomenological approach to understanding privacy practices and threat models as constructions that are predicated on Internet users' subjective lived experience. Some interviews, like the one excerpted above, suggest that people who share characteristics (e.g. Wikipedia administrators) may construct understandings of threats and privacy practices based on observing *others* lived experiences as well as their own.

Threat 3: Safety of Self and Loved Ones

Some Wikipedians described threats of rape, physical assault, and death as reprisals for their contributions to the project. Although they sometimes stated that they didn't take such threats seriously themselves, the possibility of harm to their loved ones was much more serious. One Wikipedian in Europe said he didn't take it seriously when someone threatened a drive-by shooting at his home, but

I pulled back from some of that [Wikipedia] work when I could no longer hide in quite the same way. For a long time I lived on my own, so it's just my own personal risk I was taking with things. Now, my wife lives here as well, so I can't take that same risk.

Similarly, such fears led some political activists who used online spaces to organize and share information to use the Tor network. Explained one interviewee:

they busted [my friend's] door down and they beat the ever living crap out of him... and told him, 'If you and your family want to live, then you're going to stop causing trouble.'... I have a family. So, after I visited him in the hospital, I started—Well, at first I started shaking and went into a cold sweat, then I realized I have to—I started taking some of my human rights activities into other identities through the Tor network.

Threat 4: Harassment/Intimidation

People who contribute ideas and tools in a public forum open themselves up to criticism. One open source developer reported that he used Tor to protect information about his location because people “just wanna harass me, basically for a number of different reasons, mostly because I write software that they don't like, and provide tools for user bases and websites that they have issues with.” Similarly, a Wikipedia editor reported that to avoid being targeted by groups with a history of harassing Wikipedians, “when I'm reading Wikipediocracy or one of the Wikipedia criticism sites, because I know that they scoop up IP addresses, I use an IP obfuscator for that.”

Many Wikipedia administrators recognized the threat of harassment. We were surprised to learn how pervasive and dire the threat was perceived to be among people with central roles (like employees of the Wikimedia Foundation) and permissions like blocking editors or protecting pages (like administrators chosen by the community). One female administrator pointed out that “the fear of harassment, of real, of stalking and things like that, is quite substantial. At least among administrators I know, especially women.” Said another, “It's a lot of emotional work, and I remember being like 13 and getting a lot of rape threats and death threats and that was when I was doing administrative work.” Frequently, threats to both male and female Wikipedia editors stemmed from other editors who were angry about conflicts. Editors who took central positions like administrator or arbitration committee member found that additional authority and responsibility brought with it publicity and vulnerability.

Threat 5: Reputation Loss

We have already discussed the potential for contributions to controversial topics to be misinterpreted and result in lost opportunities. For some Wikipedians, the threat of reputation loss was unrelated to specific outcomes,

I know a couple of people who edit anonymously so that it doesn't impact their professional reputation. Not in that they're worried that there's going to be an article that “Local Scientist edits Wikipedia” and then a career in tarnish but more in the sense that they don't want someone to go on a vendetta against them and what's a volunteer hobby for them suddenly turns into something that affects their professional career.

Importantly, we learned that when people create content online, they aren't just concerned about the nature of that content reflecting on them, but about the social ties (or

appearance of ties) and affinity identity they create when they contribute. One Tor user explained that he contributes to online resources about pharmaceuticals where there are also contributors with “questionable morals” with whom he didn’t want to be affiliated. Explained another Tor user,

I do a lot of stuff on [8chan], sometimes I try to create highly valuable content... I wouldn't mind just being associated with the content. It's more like some people on there might be considered bad people, and I don't want to be associated with those kinds of things.

Other Threats

The above five types were not the only threats that surfaced in interviewees’ accounts, but they dominated threat-related accounts and appeared repeatedly across interviewees. Only two participants discussed the threat of legal sanctions for online participation. In one case, an interviewee noted that his country could revoke his access to the Internet if he was thought to be violating intellectual property law. In another case, an interviewee explained that because political dissidents have been arrested and beaten in Iran, it became common practice to use Tor to access political content and to rely on “braver” or non-Iranian friends to post content that could be deemed political.

Finally, some Wikipedians discussed the threat that their efforts to contribute to the encyclopedia might be undermined. In most cases this threat was associated with holding an official role like being a Wikimedia Foundation employee or holding an elected position. One Wikipedia noted that in the language edition where he frequently edited, revealing features of his identity could influence how others perceived his contributions and potential biases: “Sometimes, your edit gets reversed immediately, just for the fact that you can be easily identified, affiliated to part of a religion or belonging to a religion or belonging to an ethnic group.” This interviewee described having his contributions challenged because his username revealed that he belonged to a minority ethnic group in his country.

Sources of Threats

We have explained the kinds of threats perceived by people who contribute to open collaboration projects. These threats were perceived to emanate from diverse sources, including the most commonly named sources of *governments*, *businesses*, and *private citizens* (Table 3).

Overall, Wikipedia editors and Tor users were equally likely to describe other individuals as a source of threats when they reflected on their online participation and need for privacy. Tor users were more likely to describe threats from governments and *only* Tor users described surveillance by businesses as a threat to their participation online. In some cases, perceived threats from governments and businesses were similarly predicated on the fact that contributors to projects didn’t know what information might be collected, for what purposes, or how information about their activities might be used in the future.

Source of Threat	# Participants		
	Total	WP	Tor
Governments	12	3	9
Businesses	4	0	4
Private Citizens	8	4	4

Table 3: Perceived Sources of Threats

A Wikipedia participant pointed out a case in which a French Wikipedia editor was pressured by French intelligence to delete an article about a military installation. The Wikipedia community responded by restoring and improving the article, translating it into several other languages and for a while it was the most viewed article in the French language Wikipedia. The incident highlighted the visibility of Wikipedians’ work to governments; our interviewee observed, “it's not just Iran, or Syria, or countries like that... here we are, we find out that editors in France have to be concerned of their own privacy.” In a few cases, participants perceived concrete threats from governments. For example, one participant described the likelihood of government sanctions of Iranian citizens for contributing to the Tor project.

Many interviewees viewed the interests of businesses and the state as potentially in conflict with their own interests. One Tor user explained his commitment to using privacy enhancing tools in terms of an escalating arms race:

We (private citizens) have to have these kind of capabilities in place, in particular in a world where the adversaries—whether that's national government, whether that's organized crime, whether that's an oppressive regime, whoever it may be—the adversaries are gearing up enormously strongly.

Participants spoke of threats from private citizens more concretely. They were afraid of things like threats to their families, being doxxed, having fake information about them circulated, being beaten up, or having their heads photoshopped onto porn because they experienced these things or saw them happen to others. In the context of open collaboration projects, these threats emanated both from other project members and from outsiders. In fact, nearly all of the harassment and intimidation described by those who took central positions in Wikipedia came from other editors. In multiple cases, Wikipedia contributors also described the threat of intimidation and harassment originating from organized groups whose members view doxxing Wikipedia administrators as a matter of policing the encyclopedia and providing a form of public service.

Lack of Perceived Threat

In some cases, although they may have been concerned about privacy in some contexts, participants conceded that they rarely perceive threats when contributing to open collaboration projects. In a few of these cases, participants pointed out that they enjoyed privileges due to their gender, nationality, race, or the scope of their interests.

Said one Tor user, “I come at it from a completely privileged position. I'm an employed white male, so I have no horse in the race. I have colleagues who get the death threats and the rape threats and all the rest of it” and a Wikipedia editor pointed out that, aside from demographic characteristics, “I'm in a privileged position of not being interested in any topics that would be of particular interest to, say, the NSA.” These interviewees perceived themselves as belonging to a privileged class online who enjoy freedoms because they are not vulnerable by virtue of their majority or socially approved status and interests. Despite perceptions of safety, some participants still use privacy enhancing tools like Tor because of a belief that privacy should be a default in electronic communications.

Mitigating Perceived Risk

The kinds of threats identified by participants are not specific to participation in open collaboration projects; however, their responses to these threats influenced their participation in ways that also affect the health and quality of projects. The experiences that led participants to perceive these threats and the ways they managed the threats were often influenced by their participation goals and the socio-technical systems that mediate their participation.

Strategies for mitigating perceived risks included two broad overlapping categories of activities: *modifying participation* in projects and *enacting anonymity*.

Notably, mitigation strategies have elements of negotiating temporal boundaries. One Tor user explained “Tor is one of those things where you want to have it before you need it, for obvious reasons, because if you're being censored, it's very hard to get.” Likewise, Wikipedians who took central positions in the community frequently described a shift in the way they viewed their privacy and participation as they took on more responsibility and become more visible.

Modifying Participation

One female Wikipedian explained that after experiencing harassment, she chose to divulge her real identity on her user page because the ease of obtaining her identity helped her avoid becoming a target:

I decided that my real life identity and my online identity were inextricably linked, and if I tried to hide them that would have been stupid, because that would have made it attractive for people to try and figure out who I was... It's just all about control, right? I made my information public on my own terms.

Editing under her real name gave her a sense of control; if she wasn't hiding, she couldn't be hunted. A Wikipedian with expertise in reproductive health described the effects of using her real name online on her editing activities:

I don't want to be getting flak when I'm applying to medical schools. Because I have friends who've been really badly hurt by that kind of thing... I avoid writing about sexual health.... I did some pokes at the abortion article and realized that I was gonna get

myself in a deep pile of shit if I kept going. Not because I was giving scientifically inaccurate information or anything—I had my giant obstetrics textbook right open next to me—but I just didn't want to wade in because I don't need backlash. So yeah, I definitely avoid saying things or editing about things because it's all connected to my real name.

In order to claim their identity online, then, Wikipedians may choose to give up the freedom to edit topics that they believe could be contentious or attract attention. Others described similar tradeoffs between revealing their identities to provide a public face for their Wikipedia work and being able to do the kinds of work they wanted to do:

I had a photo, my name and my private phone number on my user page... during that time I would not get into conflict with say, trolls or vandals, because since I didn't have any privacy, I felt limited in what I could do. I could still write articles, but blocking people and stuff like that was something I had tried to avoid.

One Wikipedian observed that surveillance “has a chilling effect on the way that we do business and on the ability which Wikipedia has, an enterprise, to continue. Because people are far less likely to engage with us, if they know that the American government is watching their every move.” Another Wikipedia editor offered corroborating evidence when he explained that “for the Edward Snowden page, I have pulled myself away from adding sensitive contributions, like different references, because I thought the name may be traced back to me in some way.”

Enacting Anonymity

Participants described efforts to enact degrees of anonymity either through **technical approaches** that circumvent observation (e.g. using Tor) or **operational approaches** that limit others' ability to connect activities with participants' real identities (e.g. maintaining multiple accounts). Each of these approaches circumvent different types of identity knowledge.

Tor provides a technical infrastructure for participating anonymously on the Internet; as such, all of the interviewees who used Tor used technical means of circumventing types of identity knowledge at least some of the time. However, because some service providers use aggressive tactics to prevent anonymous use of their services, some of our interviewees who use Tor were unable to participate in open collaboration projects or other online forums as freely as they would like. Most notorious among our interviewees, the web hosting service Cloudflare presents Tor users with CAPTCHAs to such an extent that:

It drives me insane that you solve one CAPTCHA after the next one. And even though it's technically possible to use Tor, it's just so much of a nuisance that I decided to just use... the normal browser. Just to make the CAPTCHA's go away.

Explained another Tor user, Cloudflare-hosted sites don't technically block Tor but "functionally" block Tor users by making it onerous for them to participate. The frustrations voiced by our participants about such approaches are reflected in measurement studies that document the frequency of such censorship events at the network level [22]. One interviewee told us that being blocked was problematic because the kind of identity knowledge he wanted to circumvent by using Tor was not his name, but his IP address. He explained that he does not care if blog owners know who he is when he comments, he cares that they are not able to hack or locate him.

Some websites go further and block Tor users from posting outright by blacklisting IP addresses that are known to be Tor exit nodes. One of these sites is Wikipedia.

A Wikimedia Foundation employee explained that "we do sometimes let people edit through them [tools like Tor]. I know that we have users in China coming through the Great Firewall and stuff." Wikipedia's policy is to block edits from Tor exit nodes unless a user requests special permission to edit via Tor. In order to obtain such an allowance, editors need to reveal perceived risks to employees of the Wikimedia Foundation to make their case.

In some cases, our Tor-using participants treated Wikipedia with the same annoyance as Cloudflare sites that forced them onto the "open" Internet to make contributions. One Tor user explained that he still edits Wikipedia although he feels "kind of vulnerable not using Tor." Two of our participants stopped editing Wikipedia because of the prohibition of Tor edits. Wikipedia editors who use Tor, then, must either stop participating or strategize other ways of protecting their identities while they edit.

Operational tactics for enacting anonymity involve subverting identity knowledge by changing behaviors. For example, Wikipedia editors subverted potential efforts to identify their location by carefully managing their editing activity. Explained one interviewee, "these are small things but I would usually not edit things relating to my school or places near where I lived while logged in... It is actually weirdly easy to piece together someone's identity based on—say it was based on the location."

Likewise, interviewees observed that maintaining different accounts for different purposes can be helpful, but behavioral or linguistic patterns across multiple accounts or identities can allow observers to link them to the same person and learn a lot:

People can look at the edits of those [accounts] and the patterns of both and correlate them and realize, "Hey, wait a minute this is probably the same person behind these two." And if you look at the patterns of what people read, what people edit, and what people comment on, you can start to draw conclusions about the sorts of things that they as the person are interested in, why they're doing it, for what, what some of their

backgrounds might be... That cuts a little bit too close to the bone for my taste.

DISCUSSION AND IMPLICATIONS

We have presented findings about the kinds of threats that people perceive when contributing to open collaboration projects, why some people do not perceive threats, and what actions people take when to mediate threats. In the Related Work section, we discussed the tension between participation and privacy highlighted by Palen and Dourish [29]: in order to participate, the participant must be revealed. Nissenbaum's concept of contextual integrity stipulates that the ways that people manage those revelations is bound to the contexts in which they are made [27]. In many contexts, understanding information flow and the norms surrounding appropriate disclosures are robust; but this is not the case for open collaboration projects. We have shown that contributors reveal features of their identity before they have come to understand the social and technical infrastructures in which their contributions to projects are being made and as their participation changes, so too do their privacy concerns.

We discuss two sets of implications of these findings. First, we challenge the implicit assumption underlying many discussions about open collaboration projects – i.e. that human knowledge and skills are equally sharable. This assumption has led to the problematization of volunteer engagement almost exclusively as a function of motivation, incentive, and skills. Our findings suggest that safety and privacy concerns can be equally important for some potential contributors. Second, we examine the implications of our findings for technical and social interventions that can alleviate threats perceived by participants in open collaboration projects and offer a two-dimensional framework for thinking about such interventions.

Chilling Effects, Privilege, and Free Culture

The potential for thousands of people to come together online and create resources like encyclopedias or operating systems for the common good is exciting. Open collaboration is predicated on people's ability to self-select into projects for which they have the skills and interest to contribute. For example, Wikipedia aims to produce and make accessible the "sum of all human knowledge" through contributions of hundreds of thousands of volunteer editors who write about things they know or have an interest in learning. Engaging a diverse editorship is critical to this mission; yet, as we have demonstrated, some swaths of human experience may be excluded from this project if certain topics or people are associated with heightened threat. A *chilling effect* is what happens when the threat of prosecution, persecution, or other undesired effects associated with a behavior cause people to stop engaging in it. For example, Fiesler et al. found that confusion over copyright laws can have a chilling effect on online participation [17]. Our findings identify additional perceived threats to participation; in fact, the threat of

government surveillance alone has been found to have a chilling effect on Wikipedia traffic. In a recent study of English language Wikipedia, author Jonathan Penney found that traffic to Wikipedia pages on sensitive topics decreased after widespread revelations about U.S. government surveillance activities in June of 2013 [30].

Those of our interviewees who did not perceive threats self-identified as privileged in various ways. Conversely, those who expressed concern self-identified characteristics that created vulnerabilities; for example, being female, being from an ethnic minority, or being transgender. The work of female Wikipedia editors in coping with harassment and other emotional labor has been identified in related work as a contributor to gender disparity on the site [26]. If such voices are systematically dampened by the threat of harassment, intimidation, violence, or opportunity and reputation loss, projects like Wikipedia cannot hope to attract the diversity of contributors required to produce “the sum of all human knowledge.”

The problem then is not always motivation or incentive. Our interviewees participated in different forms of open collaboration and reported high levels of commitment to these projects. Yet, many encountered privacy concerns that caused them to modify their participation and some stopped editing Wikipedia completely. Open collaboration communities must go beyond attracting participants, to develop social and technical arrangements that support contributors’ needs for privacy—especially contributors who are not fully engaged and whose privacy requirements may not be well understood.

Designing along Two Dimensions for Privacy in Open Collaboration Systems

In this section, we describe potential social and technical interventions that address many privacy concerns voiced by our interviewees and could help create a more equitable arena for participation in open collaboration projects. These interventions represent a range of considerations along two dimensions: social/technical and internal/external. Social interventions include efforts like educating people about how to protect their privacy while contributing to projects whereas technical interventions involve changes to the infrastructures that support open collaboration projects to better align with participant privacy practices. Internal interventions may focus on technologies and social solutions that are implementable by community members by modifying policies and tools; whereas external interventions are appropriate where open collaboration projects depend on shared infrastructures and toolkits provided by service providers.

Educating about Operational Privacy Practices

Wikipedian interviewees were primarily concerned with being identified through the content of their edits, the content edited, or contextual factors (location/timing of edits). Although privacy-enhancing tools such as Tor can mitigate contextual factors (particularly location), this was

not how most editors describe being outed. Social intervention may be helpful in this case: projects could benefit by producing and publicizing guidelines for participants who wish to make anonymous contributions.

Respecting Temporal Features of Privacy and Participation

The most severe privacy threats associated with participating in Wikipedia were usually not realized until contributors had taken on leadership roles associated with content curation and managing conflict. When people begin contributing to an open collaboration project, they are often not even aware of community structures [11], and do not begin with the goal of taking a leadership role. When they were new editors, our interviewees systematically failed to consider the effect of time on their privacy concerns. For some, participation changed; for others, life changes such as starting a family or applying for a job/school affected their ability to handle privacy threats. One interviewee suggested a social intervention to help:

Until it happens to you, it doesn't occur to you that, "Well, why would anyone be Googling me? Why would they want my address?" And by the time you realize that all it takes is some kid being bored someday, it's a little late to hide it. So I would like there to be, and I don't even know how there could be because people don't really read things on the internet, but I wish we had a way to warn people, "No really, you should give this some serious thought."

Yet, as she observed, such warning messages as social interventions alone are not often effective. Where possible, participants in open collaboration projects could benefit from privacy enhancing technical defaults when they begin contributing. To support privacy practices, open collaboration systems could be designed to offer people the option of creating a new technical identity that is publicly unlinked to their past technical identities when taking on significant new responsibilities and roles.

Helping Providers Control Abuse

A tension exists between preventing abuse of online services (stopping Wikipedia vandalism, for example) and allowing the use of privacy-enhancing tools such as Tor (which can be used by vandals). The CHI and CSCW communities have spent decades discussing possible approaches to managing deviant behavior [10]; in this paper we focus on those that are most pertinent to maintaining high-quality information/product and ensuring the safety of contributors to projects.

Most of our Wikipedian interviewees did not use Tor regularly and they mainly considered the existing procedures for getting exemptions from the Tor IP block list (application to the Wikimedia Foundation) to be reasonable. Many of our Tor interviewees were also Wikipedia editors; in contrast, we often heard from this group that editing Wikipedia through Tor was extremely difficult if not impossible and they often weren't aware of

the protocol for exemptions. This difference represents the perennial problem of “internal testing”: although members of an organization may use a tool they are designing, they often have different requirements or practices than potential users outside the organization.

There exist some examples of social moderation of anonymous participation, ranging from community moderation (trusted German Wikipedia editors can approve edits from users lacking sufficiently reputable accounts), to flagging behavior from anonymous users so other users can choose how to interact with them [14]. However, applying these approaches has not had widespread success. Part of the challenge is knowledge transfer—each website deals with the question of anonymity in isolation. There are no standard tools for managing anonymous contributions even if site managers want to do so. In addition, sites face different problems that require different solutions: Wikipedia doesn’t want vandals to compromise content, whereas Yelp doesn’t want competitors to scrape its pages.

Technical mechanisms also exist that allow anonymous use of websites in ways that help service providers control abuse. Simple technical approaches include “you can read but you can’t post” (e.g. Wikipedia) or “you must log in to post.” Many sites use “flags,” a sociotechnical approach that allows community members to flag contributed content they deem socially or legally unacceptable, but only after it is posted. More complex approaches track the reputation of individuals and give them access to site features based on past behavior of the person rather than on past behavior associated with a network address. There are also rate-limiting approaches involving CAPTCHAs (e.g. Cloudflare), or bandwidth throttling for anonymous visitors.

Anonymous blacklisting systems provide a promising cryptographic alternative to blocking entire anonymity networks like Tor. In a one-time registration step, users of an anonymous black-listing system receive a credential that they can use to authenticate with participating services without revealing any information that might help the services link them to any past or future sessions. Every anonymous authentication results in a unique ticket, which is related mathematically to the secret part of the user’s credential but which appears random to the service provider and third parties. If the user commits an abusive act during its session, the service provider can place this ticket on a blacklist. In the most basic form of anonymous blacklisting, a returning abusive user will be unable to authenticate if any of its own tickets are on the blacklist. Some schemes support blocking policies that are more sophisticated, such as policies allowing multiple strikes [34] or policies that consider community-generated reputation scores [3].

CONCLUSIONS

When gathering privacy requirements, designers of open collaboration systems need to remain sensitive to a variety of threats. Many threats to Wikipedians came from individuals who were part of the project, sometimes

including individuals who held positions of responsibility. Designing for “communities” suffers from some of the same blindspots as designing for “families” or “neighborhoods” – although these words are often imbued with values that include mutual trust and protection, in some cases, neighbors and even children and parents can constitute a threat to one another. Likewise, Wikipedians and other open collaboration project members can pose threats to one another. Just as social and technical solutions can be external and internal to project infrastructures, threats can be external and internal. When considering the design of privacy enhancing technologies for large, open communities, efforts should be made to understand the potential for both internal and external threats.

Ultimately, any approach to handling anonymous participation will require that organizations and service providers value anonymous contributions. We have demonstrated that privacy concerns pervade open collaboration projects; risks are perceived both by individuals who occupy central leadership roles in projects and by rank-and-file contributors. Further work needs to be done to measure the effect of project policies, practices, and technologies that interfere with contributors’ efforts to protect their privacy and to evaluate the impact of potential interventions that shift privacy protection strategies from the individual to the community.

ACKNOWLEDGEMENTS

Our thanks to Tor Project and Wikimedia affiliates who helped our recruitment efforts; to Michael Dickard, Roger Dingledine, Tim Gorichanaz, and Susan Wyche for feedback on paper drafts; and to the participants who shared their stories with us. This work was partially supported by National Science Foundation grants #1253302, 1347151, and 1253418.

REFERENCES

1. Nazanin Andalibi, Oliver L. Haimson, Munmun De Choudhury and Andrea Forte. 2016. Understanding Social Media Disclosures of Sexual Abuse Through the Lenses of Support Seeking and Anonymity. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 16)*.
2. Alexander Ardichvili, Vaughn Page and Tim Wentling. 2003. Motivation and barriers to participation in virtual knowledge-sharing communities of practice. *Journal of knowledge management*, 7, 1. 64-77.
3. Man Ho Au, Apu Kapadia and Willy Susilo. 2012 of Conference. BLACR: TTP-free blacklistable anonymous credentials with reputation. In *Proceedings of NDSS 2012*.
4. Eric P.S. Baumer, Phil Adams, Vera D. Khovanskaya, Tony C. Liao, Madeline E. Smith, Victoria Schwanda Sosik and Kaiton Williams. 2013 of Conference. Limiting, leaving, and (re)lapsing: an exploration of facebook non-use practices and experiences. In *Proceedings of the SIGCHI Conference on Human*

- Factors in Computing Systems*, 3257-3266. <http://dx.doi.org/10.1145/2470654.2466446>.
5. Yochai Benkler. 2002. Coase's Penguin, or, Linux and The Nature of the Firm. *The Yale Law Journal*, 112, 3. 369-446.
  6. Yochai Benkler. 2006. *The Wealth of Networks*. Yale University Press, New Haven.
  7. Valerie Malhotra Bentz and Jeremy J Shapiro. 1998. *Mindful inquiry in social research*. Sage Publications.
  8. Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3, 2. 77-101.
  9. Michael Brennan, Sadia Afroz and Rachel Greenstadt. 2012. Adversarial stylometry: Circumventing authorship recognition to preserve privacy and anonymity. *ACM Transactions on Information and System Security (TISSEC)*, 15, 3. 12.
  10. Amy Bruckman, Catalina Danis, Cliff Lampe, Janet Sternberg and Chris Waldron. 2006 of Conference. Managing deviant behavior in online communities. In *Proceedings of CHI '06 Extended Abstracts on Human Factors in Computing Systems*, 21-24. <http://dx.doi.org/10.1145/1125451.1125458>.
  11. Susan Bryant, Andrea Forte and Amy Bruckman. 2005. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of Group*, 1-10.
  12. Zoë Corbyn. 2015. Bassel Khartabil: fears for man who brought open internet to the Arab world *The Guardian*.
  13. K. Crowston, Li Q. Wei K., Eseryel Y. and Howison J. 2007. Self-organization of teams in free/libre open source software development. *Information and Software Technology Journal: Special issue on Understanding the Social Side of Software Engineering*, 49, 6. 564-575.
  14. Roger Dingledine. 2014. A call to arms: Helping Internet services accept anonymous users. *Tor Project Blog*.
  15. Roger Dingledine, Nick Mathewson and Paul Syverson. 2004 of Conference. Tor: The second-generation onion router. In *Proceedings of Usenix Security 2004*.
  16. Yulin Fang and Derrick Neufeld. 2009. Understanding sustained participation in open source software projects. *Journal of Management Information Systems*, 25, 4. 9-50.
  17. Casey Fiesler, Jessica L Feuston and Amy S Bruckman. 2015. Understanding copyright law in online creative communities. *Proceedings of CSCW*, 116-129.
  18. A. Forte and A. Bruckman. 2006. From Wikipedia to the Classroom: Exploring online publication and learning. In *Proceedings of Int'l Conference of the Learning Sciences*, 182-188.
  19. Andrea Forte and Cliff Lampe. 2013. Defining, Understanding and Supporting Open Collaboration: Lessons from the Literature. *American Behavioral Scientist*, 57, 5. 535-547.
  20. Graeme Horsman. 2015. The challenges surrounding the regulation of anonymous communication provision in the United Kingdom. *Computers & Security*, 56. 151-162. <http://dx.doi.org/10.1016/j.cose.2015.06.005>.
  21. Ruogu Kang, Stephanie Brown and Sara Kiesler. 2013. Why do people seek anonymity on the internet?: informing policy and design. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 13)*, 2657-2666.
  22. Sheharbano Khattak, David Fifield, Sadia Afroz, Mobin Javed, Srikanth Sundaresan, Vern Paxson, Steven J Murdoch and Damon McCoy. 2016. Do You See What I See? Differential Treatment of Anonymous Users. In *Proceedings of Network and Distributed System Security Symposium*.
  23. Aniket Kittur and Robert Kraut. 2008. Harnessing the wisdom of crowds in Wikipedia: quality through coordination. In *Proceedings of CSCW*, 37-46.
  24. Shyong Tony K Lam, Anuradha Uduwage, Zhenhua Dong, Shilad Sen, David R Musicant, Loren Terveen and John Riedl. 2011. WP: clubhouse?: an exploration of Wikipedia's gender imbalance. *Proceedings of the 7th Int'l Symposium on Wikis and Open Collaboration*, 1-10.
  25. Gary T Marx. 1999. What's in a Name? Some Reflections on the Sociology of Anonymity. *The Information Society*, 15, 2. 99-112.
  26. Amanda Menking and Ingrid Erickson. 2015. The Heart Work of Wikipedia: Gendered, Emotional Labor in the World's Largest Online Encyclopedia. *Proceedings of CHI*, 207-210.
  27. Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review*.
  28. Oded Nov, Ofer Arazy and David Anderson. 2011. Dusting for science: motivation and participation of digital citizen science volunteers. *Proceedings of 2011 iConference*, 68-74.
  29. Leysia Palen and Paul Dourish. 2003. Unpacking privacy for a networked world. *Proceedings of CHI*, 129-136.
  30. Jon Penney. 2016. Chilling Effects: Online Surveillance and Wikipedia Use. *Berkeley Technology Law Journal*.
  31. Jenny Preece, Blair Nonnecke and Dorine Andrews. 2004. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, 20, 2. 201-223.
  32. Alfred Schutz. 1967. *The Phenomenology of the Social World*. Northwestern University Press.
  33. H. Jeff Smith, Sandra J. Milberg and Sandra J. Burke. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS Quarterly*, 20, 2. 167-196.
  34. Patrick P Tsang, Man Ho Au, Apu Kapadia and Sean W Smith. 2010. BLAC: Revoking repeatedly misbehaving anonymous users without relying on TTPs. *ACM Transactions on Information and System Security (TISSEC)*, 13, 4. 1-30.
  35. Zaher Yamak, Julien Saunier and Laurent Vercouter. 2016 of Conference. Detection of Multiple Identity Manipulation in Collaborative Projects. *Proceedings of the 25th International Conference Companion on World Wide Web*, 955-960.