Social Norm Vulnerability and its Consequences for Privacy and Safety in an Online Community

BRIANNA DYM, University of Colorado Boulder, USA
CASEY FIESLER, University of Colorado Boulder, USA

Within online communities, social norms that both set expectations for and regulate behavior can be vital to the overall welfare of the community–particularly in the context of the privacy and safety of its members. For communities where the cost of regulatory failure can be high, it is important to understand both the conditions under which norms might be effective, and when they might fail. As a case study, we consider transformative fandom, a creative community dedicated to reimagining existing media in often subversive ways. In part due to the marginalized status of many members, there are strong, longstanding norms to protect the community. Through an interview study with 25 fandom participants, we investigate social norms that have been largely effective over time at maintaining member privacy and safety, but also break down under certain circumstances. Catalysts for these breakdowns include tensions between sub-communities and an increasing presence of outsiders, though most prominently, we identify a disconnect between the norms the community needs to support and the design of the platforms they occupy.

CCS Concepts: • **Human-centered computing → Empirical studies in collaborative and social computing**; **Computer supported cooperative work**.

Additional Key Words and Phrases: fandom; social norms; privacy; safety

ACM Reference Format:
Brianna Dym and Casey Fiesler. 2020. Social Norm Vulnerability and its Consequences for Privacy and Safety in an Online Community. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 155 (October 2020), 24 pages. https://doi.org/10.1145/3415226

1 INTRODUCTION

"Don't talk to strangers." "Look both ways before crossing the street." From an early age, we begin learning rules designed to keep us safe. We take similar rules with us online–for example, protecting our own privacy by not giving out our passwords [57]. We also learn how to keep other people safe–such as by being thoughtful about how and with whom we share photos of others [83]. These rules are rarely formalized, but more often come from our interactions with other people. They are social norms, the standards that govern or provide guidance to a group of people without necessarily scripting specific rules and punishments [24]. Within online communities, social norms that set expectations and regulate behavior can be vital to community welfare [48]–particularly for the privacy and safety of its members.

Relying on social norms to inform healthy community interactions has several benefits. Social norms often have more staying power than externally imposed rules [69] and can cultivate a strong

Authors' addresses: Brianna Dym, University of Colorado Boulder, Department of Information Science, ENVD 201, 1060 18th St. Boulder, CO, 80309, USA, brianna.dym@colorado.edu; Casey Fiesler, University of Colorado Boulder, Department of Information Science, ENVD 201, 1060 18th St. Boulder, CO, 80309, USA, casey.fiesler@colorado.edu.

2573-0142/2020/10-ART155 $15.00
https://doi.org/10.1145/3415226

sense of community identity [34, 47, 75]. Though norms hold powerful sway within communities, there are relatively few ways to enforce them, such as modeling desired behavior [74] or publicly shaming wrong-doers [12]. Therefore, even when social norms are strong within a community, they can fail to grant the same power as formal policy. Researchers and designers have an opportunity to better support online communities by understanding the conditions under which typically effective social norms might fail.

As a case study, we examine transformative fandom, a community that has a long history of using social norms to regulate behavior and promote positive community interactions [17, 34]. Transformative fandom is a subset of fan culture that remixes elements from original media (e.g. movies, video games, and books) into new media that is *transformative* of the original content, and is comprised largely of women and LGBTQ+ people [28, 49]. As a result of stigma attached to fandom as well as many members who are in vulnerable positions, the community has developed strong norms that help promote the safety and preserve the privacy of its members. However, as fandom grows to encompass new generations of members [29], there are situations where norms might break down or external factors such as platform design work against them.

As part of a broader investigation into privacy in content-sharing communities, we interviewed 25 participants about unspoken rules, behavior, and expectations around content shared in fandom. We found that though norms to protect privacy and safety are effective in many ways, we also identified reasons they break down: (1) value tensions and competing values within sub-communities; (2) the presence and intervention of outsiders (including researchers and journalists) who do not understand community norms; and (3) a disconnect between the norms the community needs to support and the design of the platforms they occupy. We conclude with suggestions for how community members can surface norms for discussion and education, how outsiders might approach learning norms when interacting with a new online community, and how platform designers might take action to better support important norms within a community—both for transformative fandom and communities with populations similarly vulnerable to privacy and safety risks. These suggestions aim to provide support for social norms that promote social benefits and well-being in online communities.

2 BACKGROUND

2.1 Transformative Fandom

In the broader context of online communities, *transformative* fandom (henceforth referred to as "fandom") is a place where people produce creative works derivative of original content. For example, members of *Star Trek* fandom might write a story that explores the off-screen relationship between Captain Kirk and Spock. Fandom is comprised of multiple sub-communities organized around a particular interest or "affinity" [40] such as a book series, movie, or television show. While these sub-communities form around affinities, they all belong to a broader community identity that shares some social norms and expectations [80]. For example, members of fandom share strong norms around copyright and attribution [34] and the community strongly encourages providing positive feedback toward one another's fanworks [17]. Fandom often serves as a safe and positive space for marginalized communities online [28, 38, 49], though it is not necessarily immune to toxicity and harm [30, 77, 79].

Because the community has existed since long before the internet [49], fandom is a technology-agnostic community that migrates and exists across an ecosystem of platforms rather than calling any one site home [35]. The values they have developed over time are so integrated into the community that they informed the design of the fan site Archive of Our Own (AO3), a platform built by members of fandom for their community [37]. A large portion of fandom uses AO3 for

archiving and sharing fanfiction, but social interactions take place across a variety of sites including Tumblr, Twitter, and Discord [35].

While AO3 was built by the community for their own needs, these other platforms were not designed with fandom in mind, though each is beneficial. For example, Tumblr's mix of text and visual media meets fandom's needs to share a variety of types of fanworks (e.g., art, writing, video) with each other, and their user-generated tags not only allow for labeling for multiple affinities, but also for flexible conventions around labeling gender identity and sexual orientation [25, 68], and the site's lack of a real-name policy provides some measure of safety for LGBTQ+ people that are closeted in other aspects of their lives [23]. Tumblr's default pseudonymity also supports fandom's strong norms around pseudonyms rather than "wallet names" to identify one another [16]. Tumblr's anonymous ask feature also allows people to participate in conversations without attaching their name to a question or comment they might be too hesitant to post otherwise. These characteristics compliment AO3's design, making Tumblr (at the time of our interviews) one of the most popular sites for social interaction among fandom. Since December of 2018, however, Tumblr introduced a ban on adult content that limited speech related to sexual content on the site, a topic important to LGBTQ+ groups and fandom [15, 45].

During Tumblr's heyday, fan communities used the site as a space to socialize for fandom in contrast to platforms like Facebook, which fans regarded as a space to talk with people they knew offline [46]. Tumblr lacked (and still lacks) methods for communities to enforce their own rules and norms on the platform. Other platforms like Reddit [36] or Discord [54] encourage communities to integrate third-party tools and applications that help communities enforce their own rules. On Tumblr, people can manage their own content by deleting posts, blocking other users, and deleting replies to posts. Users do not have access to sub-communities, moderation tools, or granular privacy settings on content they publish to the site. In lieu of these controls, fan communities rely on norms around privacy, safety, and secrecy to remain a safe and supportive space [16, 28, 30, 34].

2.2 Privacy, Safety, and Secrecy

This research explores what members of fandom perceive as threats to privacy and safety within the community. Here we conceptualize "privacy" with the contextual integrity framework introduced by Helen Nissenbaum [67], which suggests that information is contextual to specific circumstances and is appropriate to share in some spaces and inappropriate to share in others. A person might give their personal information to buy groceries and receive points from a loyalty program. The data will be used by the grocery store to recommend other purchases and offer discounts. Most of the time, people would agree that allowing a store to track their purchases for advertising is relatively harmless. However, people might view certain data as more private, or posing risks to privacy, if they are purchasing something sensitive [67], such as birth-control products. The purchaser might expect or more strongly desire that their consumer information be kept private in this case as opposed to just purchasing food.

Similarly, contextual integrity helps us understand fandom because of its history with personal- and community-level anxieties around information and where it is or is not appropriate to share. Even in pre-internet fandom, fans relied on social norms to manage appropriate information flows. Camille Bacon-Smith, for example, described how she only gained access to certain types of fanworks after she became integrated into the community, and how other members of fandom taught her what kind of behavior was appropriate to avoid harming members of the community. [9].

From both prior work in fandom that touches on appropriateness for content sharing [34] and our findings, we saw that attitudes toward privacy in fandom are closely linked to concepts of appropriate information flow, which is why we use this framework in our analysis, as opposed to

other ways of thinking about privacy (e.g., boundary regulation [2] or Westin's states of privacy [82]) that focus primarily on individual privacy negotiation. In fandom, both are critical, but tend to be negotiated at the community level. Bacon-Smith's experience illustrate these two axes on which fan communities consider privacy. First, privacy is important to each individual within fandom because there are negative consequences an individual might suffer if the information that they both share and access within fandom reach the wrong audience [16, 28, 34]. Second, privacy is important at a community level because it is something the community works to maintain as a group [9]. Therefore, privacy informs safety at both an individual and community level.

We understand "safety" as freedom from emotional, physical, and social harm [71]. We draw our definition from Scheuerman et al.'s work with transgender communities because it closely mirrors the safety concerns our participants expressed as part of this study, as well as what we see in previous work on fandom as an LGBTQ+ support space [28]. As in Scheuerman et al [71], fans express concerns with safety from both purposefully abusive behavior and from online content that may be unintentionally harmful. For example, while harassing behavior such as "doxxing" (releasing personal information to an unintended audience without that person's consent), name-calling, or issuing death threats can be understood as abusive, posting content for an intended audience that others might find upsetting is not necessarily abuse (e.g. graphic depictions of violence against LGBTQ+ people [71]).

Safety can be regulated online through a variety of means. For example, platforms can apply top-down methods like banning hateful content across the platform [21] and automated systems can detect when specific content might be susceptible to harm, such as sensitive disclosures [62], or is aiming to cause harm, such as bullying or harassing comments [50]. Any perceived risks to a person's safety threatens their overall enjoyment of an online space [70], meaning that encouraging a sense of safety among individuals is important for encouraging long-term involvement in online communities. Preserving someone's safety online, however, is not simply a matter of identifying and removing harmful content. Vulnerable community members belonging to marginalized groups face heightened risks to their safety, especially from privacy violations like their personal information appearing elsewhere without their consent [30]. As a result, effectively managing privacy is closely tied to maintaining safety for marginalized populations in online communities [28, 58, 71]. These risks are especially prevalent when marginalized groups access an online community for social support, which requires a certain degree of personal disclosure to access supportive communities, mentorship, and other resources [7, 28, 73].

For this reason, deciding whether or not to disclose or respond to someone else's disclosure is affected by concerns around privacy among other variables [6, 8]. For certain populations (e.g. people either marginalized, stigmatized, or made vulnerable by broader social issues), online communities might be the only viable space to access social support and mentorship in times of life transitions [28, 66]. Despite needing social support, vulnerable populations must often consider the consequences of disclosing personal information and might be hindered by traditional platform mechanics like persistent real-life identities and profiles [84]. Vulnerable populations might choose to manage their identity presentation at a more granular level. For example, LGBTQ+ people often manage what aspects of their identity are visible across different social media platforms to avoid disclosing aspects of their identity to certain audiences before they are ready [18, 27, 44]. Prior research shows that anonymity can help individuals better manage social expectations and their identities while also learning about themselves in lower-risk scenarios for seeking support [31].

Anonymity is also a key normative behavior that relies on secrecy to minimize potential threats to individual privacy and safety [34, 58, 84]. Anonymity can better enable people to discuss certain challenges or sensitive topics tied to their offline lives, such as parenting struggles [73], pregnancy loss [5], or sexual abuse [7]. Similarly, the community members of fandom employ secrecy as

a primary method for preserving their privacy and safety [16, 30, 34]. Rather than being fully anonymous, fan communities rely on pseudonyms tied to persistent profiles and identities that are detached from their offline lives [16]. While the stigma associated with fandom has lessened over the years, publicly associating with fanworks still has risks around copyright infringement, job loss, and unintentionally disclosing identity in the wrong context [16, 28, 34]. As a result, people in fandom use secrecy norms to reduce the risks associated with participating in fandom [34].

2.3 Social Norms Online

Social norms, as the standards within a group of people that govern or provide guidance on appropriate ways to act within a community [24], are particularly important in fandom, where they are strongly adhered to by community members [34]. Social norms can effectively regulate online community behavior and promote a community's overall well-being. For example, longstanding community members can model desired behavior for newcomers, thus reducing punitive actions needed [55, 74]. Social norms can also discourage undesired behavior when community members react negatively to norm violations [55]. They can also provide an implicit feeling of what is and is not appropriate to do in a social setting, such as knowing when and where to show social support and when inappropriate humor might be appropriate [1, 53, 60]. Normative commitment, or fostering a sense of obligation to a community through social norms, can strengthen community ties and group identity [10, 56, 75]. Young people have also demonstrated a complex understanding of social norms around privacy, relying on normative expectations to determine how to interact with personal information online [13].

However, despite their potential benefits, social norms can also lead to unhealthy situations. For example, if multiple communities share the same platform, their social norms can come into conflict and cause norm violations, thus leading to harassment or encountering upsetting content [39]. Social norm enforcement might also disadvantage groups that are not represented or involved in the creation and enforcement of group norms [65], or even actively harm community members if they encourage destructive behavior [12, 20]. Furthermore, normative expectations around identity presentation and behavior can be damaging for people who do not fit that ideal, causing normative conflicts within a community [81].

Despite these challenges, social norms still play an important role in helping people preserve their privacy and safety online. Youth understand norms around what information is and is not acceptable to share to the public [57]. Transgender people form normative behavior around disclosing their status as trans early and often on dating apps to avoid encountering threats to physical and emotional harm [33]. Social norms encouraging pseudonymity online help people safely explore potentially stigmatized topics related to parenting, gender, and sexuality [4, 23, 28].

Norms can become part of a community's official rules [64], enforceable through moderation tools integrated into platforms such as Reddit [22, 51]. In contrast, many of fandom's social media sites lack infrastructure to enforce social norms, contributing to a tenuous sense of safety among community members [16, 30]. Though Reddit's tools therefore might be useful for fandom, a recent large-scale survey revealed that it is not a platform used by transformative fandom [35]. While this same survey revealed small pockets of fandom using platforms like Discord, Tumblr is still the most prevalent social platform by a large margin. Therefore, here we focus on how social norms work (and when they don't work) to preserve privacy and safety in an online community even when there are no moderation tools available.

3 RESEARCH METHODS

In recruiting for this interview study, we posted to the major social platforms associated with fandom, specifically Tumblr and Twitter, seeking fandom participants who were at least 18 years

old, considering that the typical secrecy associated with participating in fandom might place undue stress on a minor when obtaining permission from their parents. These interviews were part of a broader investigation of ethical and privacy issues, which was how the study was described in our recruitment materials. We tagged posts with common keywords associated with fandom (e.g., "fanfiction," "fandom") and these posts were shared far beyond our own immediate networks, with the recruitment post on Tumblr receiving hundreds of reblogs. The call for participants was also shared through a popular fandom podcast with thousands of weekly listeners, broadening the reach of the call for participants. Both authors are actively involved in fandom and disclosed this fact, fostering trust with the community. We responded to participant volunteers on a first-come, first-serve basis, recruiting 25 people total. The first author conducted all of the interviews, and only knew (as recent acquaintance) one participant, whose interview did not meaningfully diverge from the others.

Interviews occurred in the summer of 2018 via voice chat or instant messenger depending on the interviewee's preference. Interviews were semi-structured [76] and the interviewer opened by asking participants to describe how they were involved in fandom, then asked participants to describe how they managed their privacy, what they perceived as ethical behavior within fandom, and to describe any strong sense of etiquette they noticed people adhering to (or they adhered to themselves). Participants were also asked to describe any safety risks if their data were shared outside fandom. We conducted interviews until we noticed participants describing mostly the same ideas around privacy concerns (around 20 participants) and conducted several more interviews to ensure there were no major developments or new trends [43, 61]. During data collection, both authors reviewed transcripts and met and discussed themes emerging from the interviews on a weekly basis.

Because these interviews were part of a broader inquiry on privacy and ethics, following transcription we conducted an initial thematic analysis [14] that surfaced a number of broader themes on these topics. The first author led open coding, with the second author reading several transcripts and working with the first author to outline themes that developed from the data. Following this analysis, we re-visited the transcripts and the first author conducted a second open coding that focused on previously noted themes of rule and norm violations. We found that participants described a set of social norms that fell into two broad categories: norms for privacy and norms for safety. We then identified three primary themes tied to how participants described those norms breaking down or being violated: value tensions between sub-communities, outsiders and newcomers misunderstanding norms, and platform design and policy clashing with community norms.

3.1 Demographics

Close-knit communities run the risk of being able to accurately identify people from a few specific demographics [63]. In order to obscure participant identifiability, we describe our participant demographics in aggregate. We asked participants their age, ethnicity, gender identity, sexual orientation, and country of residence. We emphasized before collecting demographic information that each participant was free to not provide any information for whatever reason. As a result, we do not have demographic information on every dimensions for each individual participant. For gender, 5 participants are trans, nonbinary, or questioning. 2 participants are cisgender men and 18 are cisgender women. 4 of 25 participants are heterosexual. The age range was 18 to 40, though most participants were in their 20s and 30s. 17 participants reside in the U.S. and 15 participants are white. Only 3 participants informed us they were people of color. These demographics are typical of English-speaking fandom in that a majority of members tend to be (and historically were) white, US-based, LGBTQ+, and women [9, 28, 49].

In terms of their involvement in fandom, 4 participants had been active in fandom for over 20 years, 5 participants for 15-20 years, 4 participants for 10-15 years, 4 participants for 5-10 years, and 8 participants for at least 2-5 years. 21 of our participants created fanworks of some kind, with fanfiction being the most popular, and the remaining 4 participants described themselves as consumers of fanworks. All participants were active on Tumblr at the time of interviews, but participants who had been active in fandom for longer than 10 years all described previous platforms they had left, such as LiveJournal, message boards, and personal web archives. We know from previous work that, for some groups in fandom, Tumblr was a common space for fan communities to gather after other platforms went defunct [35].

3.2 Limitations in Recruitment

Our sampling method as well as recent changes to fandom's online communities imposes some potential limitations on our findings. For example, because we only spoke to fandom participants who were 18 years of age or older, we may be missing some perspectives that younger participants may have on their online privacy and safety, though many of our participants have been part of fandom since their youth and were able to speak to their prior experiences.

Because fandom is a space typically dominated by white, middle-class Americans, we are also missing important perspectives from racial minorities [79] that might have intersected with issues of online safety and privacy. Finally, our recruitment materials were shared most widely on Tumblr. While Tumblr was at the time one of the most active social media sites for fandom communities [35], our data may dominantly reflect norms and values inherent to the Tumblr community as opposed to other platforms. Also, Tumblr's adult content ban, taking place after our data collection, has resulted in many marginalized communities drifting away from the platform [45].

4 FINDINGS

When we asked participants to describe etiquette or rules in fandom, they described their own strategies and experiences that, across all participants, represents a set of norms. Descriptions of norms were surprisingly consistent across participants, regardless of length or type of involvement in fandom. Many participants described how these norms were successful and how they appreciated the overall positive experience that being part of fandom provided them. In order to fully describe these norms, however, participants often focused on the breaking points where norms were violated. Through our analysis, we identified three primary modes through which these breakdowns occur: (1) value tensions across sub-communities; (2) outsiders misunderstanding or neglecting to learn community norms; and (3) platform design disrupting norm enforcement. Before exploring these breaking points, we describe widely shared social norms among the community. It is important to understand these commonly held norms to understand their vulnerabilities. We then explore the ways these norms come into conflict.

4.1 Norms Toward Privacy and Safety in Fandom

Participants described norms that focus on maintaining privacy through using pseudonyms and adhering to what some participants described as the number one rule of fandom: "don't talk about fandom." Participants also described norms that supported community safety through encouraging self-responsibility in viewing content and discouraging harassment toward others. These norms generally serve two major functions in fandom. First, they protect a person's privacy by informing community members what information is and is not appropriate to share across platform, especially information belonging to other individuals much in the same way Nissembaum describes appropriate information flows working [67]. Second, these norms preserve safety for both individuals and

broader communities by both discouraging harassing behavior and helping community members avoid undesirable content.

4.1.1 Norms for Privacy. The privacy norms our participants described mirror prior research into privacy and fandom [16]. For example, our participants described using and respecting other people's pseudonyms as a way to avoid unwanted conflict in their offline lives. Only a few of our 25 participants had experienced severe privacy violations, though most participants had cautionary tales about friends or someone from their community being outed for what they enjoyed in fandom.

Participants described a general sense that the golden rule of fandom, or "first rule of fandom," as P17 called it, was that information shared for the purposes of fandom stayed within the fan community, especially personal information. This rule was described as a mutual form of trust between people by many of our participants as well as a "double-edged sword" (P3). Participants wanted to share with one another in fandom, but feared that bad actors or people from outside the community might misinterpret or maliciously use the information they gave to the community. For example, P2 knew that certain people in fandom had her legal name and mailing address because they had to mail her fan art, but she was not worried they would publicly release her information because she trusted the artists and did not think they had any reason to act maliciously toward her. She described keeping her fandom identity separate because of the sexually explicit content she explored in fandom, some of which was fan art that she commissioned from those artists:

> My blog says, "Not safe for work." There is erotica. There is porn. There is smut...Do I want my face on there? I just got a new job where my face will literally be on the company website. You'd really have to work hard to put two and two together, but it might be done, and that's the sort of thing that worries me. (P2)

P2 is not necessarily afraid of the artists betraying her trust, but is generally afraid of her offline life being connected with the content she shares in fandom. Because content in fandom can be intimate and sexually explicit, people needed fandom to stay separate from other parts of their life. Participants described different challenges to separating fandom profiles from their offline life, which social media can make difficult when accounts overlap or use Facebook and Google log-in credentials. For example, one participant (P21) had to remove any links to her Spotify account within fandom because the account was connected to her Facebook and therefore her legal name. As the author of a popular and somewhat contentious fanfiction, she worried that someone might weaponize her personal information from within the community if they disliked her story. Pseudonymity, therefore, involves employing a certain measure of secrecy, maintaining a distinct separation between fandom and other spaces.

> I think by default a lot of people remain under their pseudonyms to protect their work identity...I don't want my work and my personal life co-mingling purely because of the people that I work with. It could lead to awkward conversations...I imagine that for some people the things that they create are still taboo. You know, just because somebody is into tentacle porn doesn't necessarily mean that it should affect their life outside of fandom. (P17)

Beyond using and respecting other people's pseudonyms, participants also expressed the feeling that fandom content of any kind, not just their own, should not be shared outside of fandom.

> My first inner thought is, things should not be spoken of outside Fandom. It's kind of like fight club. First rule of Fandom. (P17)

Participants felt that fandom and its contents should stay within fandom because of the highly personal and vulnerable nature of the creative works shared there. Fanworks exploring sexuality or other socially taboo topics might cause more harm than good to be shared outside of fandom. In

particular, specific fandoms or fanworks might associate someone with a stigmatized identity. A participant explained their concerns:

> I don't think being outed as a fan has impacted anyone, unless they're a fan of something horrific. But, I definitely know that being outed [to your offline world] as what you are and how you act in a fandom in particular has caused problems. (P12)

Here, P12 is referencing people exploring their sexuality and gender identity, two themes that participants regularly use fandom to engage with, considering a majority of them identify as LGBTQ+ and fandom's increasing LGBTQ+ presence [28]. For many of our participants, fandom operates as a space entirely separate from the rest of their social lives:

> [Fandom is] not for you. And by "you," I mean the public. It is for a very specific set of people. [Fanfiction is] written with a certain audience in mind, and it's put in the place where yes, it is public. Anyone can go onto my AO3 and print it out and read it...Perhaps it isn't quite obvious that there's a certain sort of understanding or unspoken agreement about AO3. It's stuff for fandom. (P2)

P2's insistence that fanworks are "for fandom" speaks to the fact that personal data shared in fandom belongs to that community as part of its appropriate information flow [67]. When personal data is shared with fandom, it is willingly shared for that particular space and nowhere else, despite technically existing as "public data." For example, one participant explained that, within fandom, they used a different name and pronouns because they felt more comfortable that way, but did not want that name or pronouns used offline at work or with family because of negative perceptions toward LGBTQ+ people.

Participants cited threats to safety as well as privacy as reason to be protective of personally identifying information. While participants often cited threats outside of fandom as a concern, they also talked about backlash and harassment from within fandom as a common concern.

4.1.2 Norms for Safety. Previously, we defined safety as the freedom from emotional, physical, and social harm caused by abusive behavior [71]. Overwhelmingly, fandom is a safe space for people to explore sensitive topics [11, 28], with prior work demonstrating that vulnerable populations feel more comfortable in spaces where fandom operates, such as Tumblr [23]. Participants spoke about the anxieties and fears they felt over something potentially going wrong in the future rather than recounting experiences when their personal safety was violated. Participants were aware of possible threats to their safety, even if they were not actively experiencing any.

Our participants reported observing intentional and unintentional violations to safety norms through witnessing harassment or seeing other community members being doxxed. Participants also described a set of norms related to keeping themselves safe through avoiding unintentionally upsetting content. For example, many participants described a self-policing rule known as, "don't like, don't read," implying that community members are expected to ignore or look away from content that they might disagree with. Participants described that, in general, this practice was easy enough to maintain through the fandom community's infrastructure for tagging content:

> It can be difficult to deal with certain subjects, even in fiction. And [fanfiction], as opposed to other fiction, is something you can tailor to yourself in terms of reading. I can read a book blurb, I can look at reviews, I can ask friends or a store associate, but there's nothing that says "on page 18 the hero's dog dies," and that might be fine. But I can set AO3 searches to exclude any instances of dog death pretty easily. (P19)

Participants emphasized that this norm encouraged people to curate their own experience online because AO3 makes that process easy. Another participant stated that supporting this norm was about supporting people's right to creative processes as a resource:

> While I may not agree with [writing about underage characters in relationships with adults]...writing is an excellent way to process emotions and experiences. I think that fandom can be a really great outlet for that. (P22)

P22 explains above that "don't like, don't read" is important because even if they disapprove of the fanwork, the person writing it might need that process as a way to recover from difficult experiences. Putting the responsibility on the individual to regulate their own experience also plays into a closely related social norm that discourages people from harassing others: "ship and let ship."

The word "ship" refers to the act of "shipping" or desiring certain fictional characters to be in a romantic relationship together, and is a staple of fandom [42, 49]. For example, most fanfiction focuses on the romantic relationship between two or more characters, with well over 75 percent of all fan fiction reading or writing focused on "romance" [19]. Shipping, used as a verb, implies the action of discussing and promoting a certain romantic pairing. Within fandom, shipping can become contentious when people disagree with one another over romantic pairings or believe that certain romantic or sexual pairings are unhealthy or morally objectionable, though there are still governing norms that provide a sense of what is and is not appropriate behavior when people disagree over shipping [42]. The norm "ship and let ship" serves to protect people from content policing. This social norm also encourages people to be more flexible of differences among community members. In other communities, differences in identity that deviate from a normative ideal can lead to conflict and harm toward vulnerable community members [81]. Violations of the norm "ship and let ship" tend to be examples of people enforcing normative expectations against the general spirit of fandom. One participant explained that some "ships" were viewed as unhealthy compared to others:

> It's like, dude, I know this chocolate cake's bad for me, but I'm going to eat two pieces of it anyway. It's like, all right, eat those two pieces of chocolate cake and enjoy yourself...Enjoy it all you want, but just admit that you're not eating vegetables. But these people take it a step further and they say, "No. I'm going to police what you can enjoy." That's not okay. (P23)

Participants drew a firm line at policing other people's involvement in fandom, with P23 referencing some people crossing that line to tell others what they can and cannot enjoy. "Ship and let ship" acknowledges that not all people in a community like fandom will share the exact same values, which is a concept that in part helped shape the community's tagging system [37]. Specifically, AO3 has four mandatory warning tags that content creators must use to label their fanfiction if it contains any four specific topics (content related to major character death, underage sex, non-consensual sex, or extreme violence). If a fanfiction contains none of those topics, the author can select "no archive warnings apply." If the author does not select any of these options, AO3 automatically labels the story "author chooses not to use archive warnings." These mandatory tags are a compromise over the tension of wanting to allow people to write what they want while also helping people avoid content they do not wish to see.

Safety norms in fandom serve two major purposes identified by our participants: (1) people are safer from online harassment like doxxing from within the community if their personal information is not easily accessible (e.g. by using a pseudonym); and (2) community members are also safe from being exposed to content they would otherwise avoid. For some participants like P2 and P17, these norms helped them maintain a professional online presence. For other participants, keeping their fandom content separate was a matter of preserving physical safety. For example, LGBTQ+ participants who were not out to their family described fears around losing their housing or being physically hurt by family members if their fandom content were discovered, with one participant referring to anything depicting LGBTQ+ characters as "contraband" (P15). These sorts of threats

stand as a very real concern to our participants because of the ways they described social norms around privacy and safety coming into conflict in fandom.

4.2 Reasons Norms Come Into Conflict

Because fandom exists across multiple platforms, community members are reliant on norms as a common set of rules for community behavior, with each platform containing different methods for interaction and regulation. Despite many of our participants describing and adhering to common social norms, they also described moments when these norms came into conflict or broke down. Participants talked about debates within fandom focused on how and when people should police fan content. Participants also described how outsiders entering fandom violated norms unintentionally by sharing fan content where it was not appropriate to share, thus amplifying fan content to the wrong communities. Furthermore, participants emphasized that certain platforms like Tumblr are not designed to support the privacy needs of fan communities, generating conflict around enforcing norms. Here, we describe three ways that our participants said norms came into conflict.

4.2.1 Value Tensions Between Sub-communities. We know from prior work that when communities in fandom agree upon norms, those norms are well-enforced by those communities [34]. Fiesler and Bruckman found that people in fandom, across different sub-communities, generally agreed on the social norms surrounding copyright and attribution in regards to creative works [34]. However, the social norms tied to privacy and safety we examine here are debated across fandom, causing value tensions between community members or entire sub-communities. For example, ten participants described witnessing internal community norm violation and community members disagreeing with others over norm enforcement. These value tensions are not necessarily new to fandom, considering that AO3 was designed with ways to mitigate those value tensions [37].

While we have previously described mandatory content warnings, another design feature that mitigates value tensions is the site's policy toward "orphaning" fanworks. Fan communities place a strong value on archiving fanworks for others to read. However, the community also values personal privacy and control over one's information. As a compromise, AO3 designed the "orphan" feature to disconnect a content creator from their fanwork, removing the associated author name and trace data. As a result, people can opt to remove their name from a work rather than deleting it, thus preserving the archive's collection [37]. Outside of AO3, social media platforms are not designed with the same care to diffuse value tensions specific to fandom. When describing norm violation *within* fandom, participants often pointed to intense community debates over what values fandom should hold and how those values should be enforced across the community.

> You have people posting these blanket statements of...almost always it relates to underage sex. So we can't write about teenagers having sex because then they'll think it's okay...There's all these conversations about how it needs to be not written about at all. Like if you write this or read this you're a terrible person. Or it's...where you have people who get hateful comments, or they get dragged on Tumblr for writing about abuse, or an older/younger person relationship. Things like that. (P22)

The debates that P22 describes conflict with the social norms participants described to us, particularly "don't like, don't read." P22 is describing a negotiation of content management meant to preserve community safety (e.g. making sure minors do not see inappropriate content), but notes that the debate leads to harm. Those "blanket statements" lead to people perceiving harassment toward someone writing objectionable content as justified, which can encourage people to take harmful actions toward others they might not otherwise [12]. The resulting consequences of that harassment have real-world impacts on community members.

> I see young people stumbling and then being shredded for something that was not done with ill intentions and could be a teachable moment. Instead, they are being treated as if they are the problem instead of somebody who is perpetuating the problem but in no way...could be the person who created the problem. (P1)

P1 describes witnessing general harassment toward others who unknowingly break rules in fandom. These people could easily benefit from a "teachable moment," P1 argues, but are instead punished harshly. Participants described a "steep learning curve" (P17) similar to other online communities highly structured by social norms [55]. While people are learning those norms, they might be exposed to situations like P14 describes below in which a young person received an influx of harassment for drawing fan art in a way that did not adhere to community norms:

> When young artists or authors are attacked by a force of [anonymous people], that just disrupt the safe space...I'm pretty sure there was an artist in the Steven Universe fandom who tried to commit suicide because of [anonymous] hate related to how she was drawing certain characters...I believe she survived, thankfully, but she was young and just some person, you know? (P14)

Participants noted that, even though they supported certain practices within fandom that encouraged people to let others create what they want, other people disagreed and felt that they had an obligation to police content. Many participants blamed the trend on "callout culture," the act of publicly shaming someone for doing something that is perceived as wrong or a group of people disagree with.

> I'm into ship and let ship. Callout culture seems to have gotten into fandom and I think it's unfortunate. I've been involved in several ships that are hashtag problematic. (P9)

> It's almost always blanket statements someone will post, like "All these people are terrible." Then you'll see pushback where someone goes "Well wait a second, it's okay as long as this," and then you'll have a fandom grandma saying "Wait a second, we've heard this before in periods of censorship and this is why communities like Archive of Our Own are created, and we need to remember that." It's a cycle. (P22)

Both P9 and P22 described callout behavior as though it is new to fandom, with P22 referring to a cycle of discussion that usually ends in the "fandom grandma" describing the history of why a certain norm exists. While Tumblr hosts many of these conversations, the archival history of fandom and documentation on why certain norms and tropes exist are hosted on other sites like the wiki Fanlore [1]. One participant described a moment where they were perceived as a newcomer to fandom–despite being actively involved across multiple platforms over the years–because their behavior did not align with the expectations of one sub-community:

> I came to Tumblr very late, and one of the first real interactions I had was one where a friend reblogged a third person who was [discussing] how much more amazing Tumblr is to LiveJournal. And my friend disagreed. I reblogged, agreed with her and then quite a few people jumped in...And the original poster freaked and threatened to dox me. And I was amused and bemused and freaked because the very reason I loved LiveJournal was the privacy settings and the reason I hated Tumblr was its lack. But this person (and their community) had created rules that I was unaware of, and felt I'd broken those rules and as such was a total newb who shouldn't engage in fandom. (P8)

Despite P8 being a long-term member of fandom (with over 20 years experience), their knowledge of fandom rules did not align with another person's knowledge of fandom rules. Other participants referenced that they felt "younger" or newer community members tried to police fan content

[1] https://fanlore.org

more than other community members. It is also entirely possible that, despite our participants sharing a strong sense of what the community norms are, that fandom norms have shifted over time or changed between sub-communities, much like we see normative clashes over how to appropriately express grief online [39]. Participants described how comments that were negative toward someone's fan content violated the norm "don't like, don't read."

> I don't know, it's jarring. If it's your preferred ship that someone's hitting on, it's like, "Well, this sucks." And if it's someone else's, you kind of feel bad for them. And why are you making someone feel that way? (P4)

P4 describes a sense that people in fandom should not criticize one another over their ships, mentioning a "jarring" sensation when noticing others violating the norm. P23 spoke to a fear of excessive regulation if people in fandom were allowed to police others for their fan content, describing that while some people might enjoy something in fandom that others did not, they still had the right to enjoy that content:

> Of course, we all have ships where we're like, "Okay, that's unhealthy," or "I think that's abusive." I think we've all collectively decided that Fifty Shades of Grey needs to burn on a pyre, but that doesn't give you the right to go and police other people because then it becomes a question of who watches the watchmen? Who is deciding what is pure and what isn't? (P23)

The tension our participants describe between policing undesired content and letting people enjoy what they want is in part the result of differences in values between sub-communities and individuals within fandom. Our participants felt as though it was inappropriate to contact or harass individuals for creating fanworks they objected to, though there is most likely a competing norm that compels the people perceived as harassers to take such actions. Prior work has been critical of that relies on anonymous shaming or callout posts to police behavior, noting that such norms impose certain aesthetic values on the community that restrict creative expression [26]. For example, P25 described a commonly-held value that middle eastern characters should be depicted with darker skin in fan art to increase representations of people of color. As a pale-skinned middle eastern person, she found that this value, while well-intended, was misinformed and limited her own creativity.

These opposing values lead to tense situations in which people are harassed, disrupting fandom as a safe space. That disruption can impact someone's life when their fan community is the only space they can safely be their true self [28]. One participant described purposefully doxxing a group of people in fandom as a means of punishing them for harassing her fandom friend:

> She...knew their fandom identities. And for example, some of them had written Real Person Fiction about this actor [they worked for]...I reported those employees for it in revenge for the way they had treated her...Both sides of their identity were known by this woman and because of that, she was able to get them fired from this job. (P23)

While P23 expressed that she regretted her actions and would not do something like this now, the story represents how norms like "don't talk about fandom" are easily overridden within the community when someone has the means and motive. Other participants like P1 and P7 self-described as "lurkers" and practiced non-use to avoid harassment. In addition to encountering threats from within the community, participants also described tension between participating in fandom and fearing that their content would be amplified to broader public spaces through outsiders entering fandom and misunderstanding community norms.

4.2.2 Outsiders and Newcomers Misunderstanding Community Norms. The social norms that govern rules around content in fandom are difficult to learn and not necessarily readily available to someone

just entering a fandom community. For example, outsiders and newcomers might be unaware of the privacy norms surrounding content. Participants described anxieties around people new to the space misunderstanding what fandom is for, or inappropriately commenting on fanworks.

In particular, fandom has specific norms around how people are supposed to interact with fan content, though these rules are not necessarily written down anywhere. Multiple participants described learning how to leave feedback and comment on people's fanworks appropriately, which involves being positive [32].

> I think the biggest rule that's stood out to me is how fans are supposed to respond to each other's work. So, it's not that critique is impossible; but, what I like is that there's a starting disposition of goodwill. You want to express your feelings and your love for this fandom and you want to share it with the community, and sort of taking that as a good thing and then moving from there to say most of our responses should therefore be positive. (P3)

Participants also voiced concerns over their fandom content, such as fan art or fanfiction, leaving fansites by being posted elsewhere without their permission, whether in a Buzzfeed article or shared to someone's personal blog. Participants described events where journalists, academics, or even harassers have taken fan content and put it elsewhere without asking permission. As P2 stated, fandom is not for the "public" but created for a specific community. However, fandom is situated in public spaces online, meaning that content can be relocated anywhere without the owner's knowledge. Because accessing fan communities is as easy as visiting a website, participants described people taking content from fandom without thinking about the consequences of displacing that content. One participant described such actions as parachute journalism: "people come into the community, they parachute in, they don't really look around, and they leave" (P5).

Parachute journalism can easily disrupt community norms and perpetuate harm. Returning to LGBTQ+ members of fandom, these participants described ways in which their content might accidentally reach the wrong people, even if none of their personal information did:

> There have been scares...I always worry about having something that gets extremely popular to the point of bleeding into something like a Buzzfeed article without my permission...For some people, it could be potentially dangerous if their identities are even accidentally linked and their life circumstances don't allow for that kind of openness. (P15)

P15, who writes fanfiction and illustrates fan art, worried that if one of their illustrations ever became popular enough, it might be relocated without their knowledge to other parts of the internet where their drawing style could be recognized by a family member and then followed back to their Tumblr account. P15 presented as openly queer on Tumblr but remained closeted at home for fear of homophobic reactions from their family. They described keeping their content within fandom as a matter of "survival." Despite there being so much at risk for some people, participants described how outsiders could better interact with fandom–by getting to know the community.

> Talk to people that are actually part of the community...If you come at it respectfully...it is a very open community that is willing to answer immense amount of questions that you have. If [outsiders] just took the time to actually contact enough people to get a good sense of things, and also just didn't portray things out of context, I think that'd be fine. (P12)

> If someone is going to feature a fandom creator, I think that basic courtesy [of asking permission] should be extended, especially if their content is going to be shown in a

> broader sense–art re-posted or reused, works copied and pasted, several posts compiled. (P15)

Asking permission to use fan content outside of fandom was the top recommendation participants gave for outsiders coming into the community. Participants actually encouraged people to write about fandom, citing that it was good to highlight the community as a normal space (as in, not a space to be stigmatized [77]). Our participants emphasized that the only way to know if a person is comfortable with their content being used elsewhere is to ask them directly. By encouraging others to ask permission, participants hoped that outsiders would understand that because certain content is openly shared in fandom does not mean it is public knowledge elsewhere in a person's life.

> I have had a personal friend who was outed by being featured in a local newspaper kissing her girlfriend when the paper did not get the permission of my friend nor her girlfriend. It was disastrous to say in the least. (P15)

The incident P15 describes above refers to when a news reporter took a photo from a fan convention's photo shoot page and ran it in a story. The example highlights the important distinction that fans emphasized about their fan content migrating elsewhere. Again, what might be appropriate information to share in one space might be inappropriate elsewhere [67]. What concerned participants most was the idea of people finding out what they did in fandom, not their status as a fan.

> What I do is for the fandom, people that are already in the pit...I don't mind people in real life knowing that I'm a fan and that I write fanfiction, because that is part of me...and I want people to know this, but it's not for [them], because unless you're in the pit already, you're not really gonna understand, and it might strike you as very bizarre. (P2)

Fandom is for fandom, as our participants stated numerous times. The norms built up within fandom are designed to prevent internal harassment as well as content leaving the platform. Unfortunately, not all platforms can support the privacy and safety norms of fandom.

4.2.3 Platform Design and Policy Versus Community Norms. Despite fandom existing independent of any one platform, the digital space it resides on can still influence the community. In many of the examples we have discussed, fandom's norms and the ways in which they break down represent gaps between the community's needs and the platform's affordances, especially concerning norms around preserving secrecy. Because Tumblr was, at the time, an important space for fandom's social interaction, many participants compared Tumblr's design to platforms that were previously popular, such as Livejournal.

> So back in the LiveJournal days, it was pretty much locking posts to certain friends, certain communities. On Tumblr, there's no privacy settings whatsoever, so it's just out there and so if you tried to track my name down, it would not be hard to find me. Here I am. (P21)

"Locking" posts refers to an ability people had on LiveJournal to restrict someone's ability to view posts, or entire journals, to certain people similar to the "friends only" setting on Facebook or circles in Google+ [52]. Tumblr does not have options for selective sharing. As a result, fandom participants cannot choose to make some of their content available only to a certain group of people. Other platforms posed privacy threats through single sign-on features, leading to P21 electing not to use her Spotify account to host any fandom-inspired playlists. Community members have come up with certain workarounds that point toward gaps between platform design and community needs that might easily be bridged.

> You don't really reblog personal posts. People will get onto Tumblr and go, "I had to deal with this guy who was harassing me and it sucked." It's sort of understood that you don't reblog it, you know? (P23)

> Some people are like, "please don't reblog" when they put stuff up. (P2)

Whether it is a sense that personal posts are not to be reblogged or explicit instructions on the post itself, the idea of *not reblogging* goes against Tumblr's design, which encourages people to spread content through reblogging and liking content. A design that encourages the spread of content can be good. Participants expressed that sharing and properly crediting fanworks throughout the community was an important part of fandom. However, our participants have also described a need for some content to stay strictly within fandom, to not talk about fandom outside of fandom.

Tumblr's design also makes following the norm "don't like, don't read" difficult. Participants talked about curating their experiences, but also talked about how Tumblr's infrastructure sometimes pushed unwanted content into their space.

> Tumblr is a site where tags are important for finding content and do not have a good on-site way of being monitored...If you are a user who doesn't have a blacklist [plug-in installed], then you are essentially seeing everything that goes on your dash, regardless of what it is. You can unfollow people, but sometimes people surprise you and randomly reblog porn in the middle of the day...So there's some aspect of Tumblr where people cannot curate everything. (P14)

This heightened visibility of content can fuel toxic interactions when Tumblr blogs with wide audiences specifically target or "call out" something in particular. P13 describes an incident in which two major people in fandom, also referred to as "Big Name Fans," amplified each other's presence and opinions to opposing sides of an issue and caused an influx of harassment for one another in the process.

> I've seen two people in a fandom pitchforking each other as to how a certain character's portrayal should be done by fan artists and fanfiction writers for it not to be racist. Given those two had big followings, both were harassed through anon asks. (P13)

Even when participants did not name specific platform mechanics as contributing to this behavior, they often contrasted how norms broke within fandom with a feeling that they worked differently on other platforms.

> It's just wild to me that people want to spend their time trash talking other people and other people's ships...I feel like it's something that's different from when I was last actually talking to people in fandom which was not on Tumblr. [It was] on LiveJournal...On LiveJournal, I was in specific separate communities. (P9)

The specific communities that P9 refer to represent separate journals, or sub-communities, within fandom that people could opt into. They would only see the content associated with that sub-community so long as they opted into it. On Tumblr, the site's design means that people will encounter anything that someone they follow reblogs, regardless of whether or not they are invested in that fan content.

LiveJournal was not without its own problems, our participants noted. The site's ban of explicit content prompted people to leave and build new communities elsewhere [37], demonstrating how platform policy can also severely impact a community, especially when that policy is in stark contrast to the community's norms. Tumblr's 2018 adult content ban might also prompt a similar exodus from Tumblr. The ban represents a policy decision that actively works against fandom's community norms. As we have previously explained, fandom is a safe space to explore topics related to gender and sexuality, which often involves "adult" content [15, 45]. The platform's policy,

therefore, cannot support the community with social norm enforcement tied to that content. When discussing platforms that supported community norms best, participants described AO3, a platform designed by fandom and for fandom, as a "step in the right direction" (P12), with accompanying policies that are supportive of the community's norms. Prior work demonstrates that aligning a platform's policies with community norms helps make those norms more enforceable [34]. When a platform is not able to support a community's social norms, those norms break down quickly. By understanding why these norms break down, we can better understand how communities interact with their platforms and consider ways to support norms from within a community.

5 DISCUSSION

These findings demonstrate that while fandom has strong social norms that help the community function, there are instances where a lack of means to enforce those norms can result in serious consequences for community members. Fandom's norms are only effective when people's values align with those norms and when there are ways to make those norms known. When those conditions are not met, disagreements between sub-communities or competing values can cause norms to come into conflict and break down. Fandom's social norms also cannot protect the community from outsiders stepping in and unintentionally violating them. In addition to these challenges, platforms like Tumblr are not designed in a way to support fandom's social norms. Because of these risks, vulnerable people in fandom are susceptible to harm depending on how their content is used and where it goes. Looking beyond fandom, we draw on these findings to consider how we might examine social norm vulnerabilities in broader online spaces and, if those norms are providing a benefit to the community, how we might better support them. Using fandom as a case study, we consider how to support norms as community insiders, outsiders, and as platform designers.

5.1 Supporting Norms within a Community

Supporting social norms from within a community can be a challenge when that community lacks formal moderation tools. Unlike platforms such as Facebook, Reddit, or Discord [3, 36, 54], Tumblr has no moderation controls that community members can use to establish or regulate sub-communities. People may delete their own posts or replies to their posts, but beyond their own content, they have no methods for regulating other people, barring reporting someone for policy violation. Therefore, with norms inevitably playing a larger role, disagreements over those norms led to tension between sub-communities. Participants described lengthy public discussions between community members about the purpose of certain norms, with "fandom grandmas" eventually stepping in to remind people why those social norms are in place.

The social norms in fandom are so effective because of the community's strong sense of group identity, as in community members behave a certain way to signal that they are part of the group [34]. However, our findings demonstrate a vulnerability where social norms might differ enough between sub-communities to cause tension. We can help people across different sub-communities by providing tools for surfacing a community's norms to its members. As an example, we might consider fan practices toward when it is and is not appropriate to discuss "spoilers" (e.g. plot-relevant details of a book, movie, or television show). Some sub-communities might display a "No spoilers!" message to let people know the community is not yet talking about new content related to whatever media is being discussed. The message sets expectations from within the community.

Still, miscommunication can happen. Prior work has recommended encouraging reintegrative norm enforcement, in which people are instructed on how to act the next time rather than being scolded for what they did wrong [34]. Communities should approach problem-solving through education over harsh sanctions when possible. For example, people who have never commented on

fanfiction might be unaware that positive feedback is prioritized, with constructive criticism only welcomed if asked for. If someone leaves the wrong kind of feedback, someone can reply to the comment and instruct them on what is expected in the community. If a person is never told what they are doing wrong, they never have the opportunity to correct their behavior. We see similar strategies working well with moderation tools on Reddit that not only delete inappropriate content, but also explain why the content was deleted. Explanations helped community members learn the social norms of a particular subreddit and improve their behavior in the future [51]. Our findings here support continued use of reintegrative norm enforcement across online communities.

5.2 Supporting Norms as an Outsider

When asked about outsiders entering the community, participants repeatedly emphasized that all are welcome in fandom as long as they are respectful and take the time to learn about the community. For example, the rule of *obtaining permission* for content sharing in fandom (as demonstrated by P8's experience of being called out for inappropriately reblogging) is important to members of fandom because it prompts a dialogue where an outsider or newcomer can better learn a person's life circumstances. As P15 noted, not every fan artist is excited to be featured in a Buzzfeed article, and may face harsh consequences if they are.

However, despite being generally welcoming, participants were still mistrustful of outsiders entering the community that did not intend to become part of fandom. For example, though they held more positive opinions toward researchers, community members feared journalists or mean-spirited trolls taking advantage of the community. This perspective might be fueled by the external stigmas participants were aware people held toward fandom. Participants assumed that, when coming to fandom, outsiders were more often than not looking for something of shock value or to quickly compile a list of memes, selected fan art, or fanfiction to feature elsewhere rather than linger with the community. One notable incident of outsiders flooding fan spaces involved a college professor assigning fanfiction as reading homework and students from the class leaving critiques for the fanfiction that did not adhere to community norms[2].

Of course, fandom is not the only space where newcomers or outsiders can disrupt a community. For example, large influxes of newcomers to subreddits typically mark a period of heightened rule violations [55]. Prior work has demonstrated that, despite an influx of new members, regular moderation paired with norm enforcement from the community can help newcomers learn norms quickly [55]. We return to idea of obtaining permission and asking questions of a community as a reliable way to understand community norms and reduce harmful norm violations. Fostering transparency and open communication between outsiders and community members can help decrease the chances of a journalist, researcher, or any other person unknowingly violating a norm.

5.3 Supporting Norms through the Platform

Because so many of our participants highlighted a gap in available tools for norm enforcement, we encourage platform designers to consider the communities that make use of their space and their needs. A website's interface will enforce certain normative behaviors on the communities using that site, whether or not those norms align with how the community can best use the platform. For example, Mel Stanfill [78] demonstrates through discursive interface analysis how certain fansites operated by corporate media properties actively discourage group formation by only allowing interaction through message boards. Stanfill argues that design choices like this encourage normative behavior toward fandom as an individual practice rather than a communal one. Common

[2]https://themillions.com/2015/03/from-the-internet-to-the-ivy-league-fanfiction-in-the-classroom.html

methods for eliciting values during design processes can aid designers in determining what a community's social norms are when building a site or adjusting it as its user base changes [37].

In contrast to AO3 which was built specifically for fandom, social media platforms such as Tumblr or Twitter were not designed with fandom's values or social norms in mind. Participants spoke about their frustration with other platforms' inability to do tasks they perceived as simple, such as making social media posts private to a select audience or managing the spread of viral content. Design and policy decisions can have a much stronger (and more negative) effect on marginalized groups using a platform even if that platform did not intend to drive off those users through their decisions [15, 45]. Considering the frustrations of our participants, we encourage researchers and designers to pay close attention to those pain points and how a platform might exacerbate a problem. Where those points of frustration are connected to a misalignment of a community's norms with the platform, we can ask whether or not changing the platform's design would encourage healthier community interaction. Asking what the norm is, what purpose it serves, and how the platform disrupts or supports a social norm allows consideration for how we might design to support social norms. Observing and understanding a community's social norms can help us design better platform mechanics that fit a community's needs without exposing those community members to undue harm. We might also consider ways to introduce flexibility in how communities manage themselves online, empowering people to regulate their own spaces through customizability and moderation.

For example, platforms seeking to support healthy social norm enforcement could allow community formation through sub-grouping, something that platforms like Facebook, Reddit, and Discord already encourage [3, 36, 54]. When a platform has sub-communities, not only can those communities form and enforce their own appropriate norms, but it allows users to know when they are in one community space versus another and decreases the possibility of inadvertently coming up against conflicting norms. Secondly, platforms can emphasize the importance of static information displays that can easily be accessed and referenced by community members who might need to reference norms or guidelines. Displaying rules in discussion boards has shown to help participants better adhere to the norms and expectations [64]. Our participants described a "cycle" of debate around social norms that could benefit from accessible, archival spaces integrated into the community. These sorts of spaces can keep records on community norms and expectations, not unlike the work certain wiki pages do for community records. Archival spaces can also help communities keep track of how their social norms might shift over time, or explain how expectations have evolved. Instances like P8 described, where norms might differ or change based on the sub-community, speak to a need to consider how platforms might encourage displaying rules and guidelines while also leaving room to show how those rules are amended and change with the community, or between sub-communities.

5.4 Cautions for Social Norm Support

Designing toward enabling communities to better enforce their own social norms has its challenges. First, not all norms are positive. For example, Blackwell et al. [12] have demonstrated how people can decide to direct harassment at other people online as a means of enacting justice on a wrongdoer. Social norms can also develop to enforce self-destructive behavior, such as in pro eating disorder communities where members encourage each other to adopt dieting behaviors that physically harm people [20]. Normative expectations can also discourage diverse expressions of identity or interest, causing harm to marginalized people in a community [81]. In these instances, we might instead consider how harmful norms could be disrupted by intervention from the platform.

There are also situations where social norm enforcement simply won't work. For example, social norms tend to have more sway when there is a strong group identity in a community [10, 75],

meaning that online spaces that lack a strong group identity can struggle to form strong social norms. It is also important to recognize that social norms can change over time, or that change might be necessary in the case of communities of harm. Finally, it is important to acknowledge that surfacing social norms as explicit rules, like on Reddit, can shift the burden of caring for those rules onto specific community members [41]. To this end, we offer our insights here as encouragement for deeper engagement with social norms in online communities. We encourage future research to not only examine how communities display and enact social norms (e.g. through rules lists, announcements, and active moderation) but also to examine where social norms come into conflict and fail, especially if those norms are benefiting the community.

6 CONCLUSION

Fandom is a unique community because of its longevity across different platforms and adherence to social norms for promoting community health and safety. Participants highlighted moments where these social norms do not always work, causing disruption to their privacy and sense of safety. The problems our participants face are not unique to fandom, either. Online communities comprised of vulnerable populations face similar privacy and safety concerns, and are often without the power to enforce their own beliefs against the grain of the platform. LGBTQ+ communities, communities of support, and communities devoted to sensitive or special topics all have specific norms that are difficult to enforce or are inadvertently sabotaged by platform design [7, 18, 27, 55, 59, 72].

This research demonstrates how analyzing a community's social norms can surface weak points in a platform's design as well as opportunities to better support online communities. Identifying social norms in other communities that rely on strong norms for regulation can provide researchers and designers with the opportunity to explore better ways to support those communities. Our research also highlights ways that online interactions can expose people to extreme risks. With so much at stake, why do our participants stay involved in fandom? One participant told us:

> [Fandom is] such a big part of what's gotten me through depressive periods...and it's hard to separate that...And I wish they could all exist in the same world, but I don't live in that world. So I have to draw lines. And the lines in the sand have shifted throughout my life. And they're going to keep shifting. (P2)

It is perhaps this deep connection that participants felt to the community that fuels the intense debates over adherence to norms and how to best keep the community safe. Within fandom, participants viewed social norms as an effective way to maintain a healthy community. When social norms work to preserve the privacy and safety of its community members, any conflict or breakdown in those norms can have serious consequences for the affected people. Identifying those norms, and their breaking points, helps us understand how to better support social norm enforcement and encourage healthy online communities.

7 ACKNOWLEDGMENTS

We would like to thank our participants for sharing their stories. We also would like to thank the members of the Internet Rules Lab (IRL) for their support throughout the research process. In particular, we would like to thank Aaron Jiang for his helpful feedback in the revision process and his help resolving formatting errors. Thank you as well to Katie Gach, Mikhaila Friske, and Morgan Scheuerman who provided close readings of paper drafts. This work was funded by NSF awards 1704369 and 1936741.

REFERENCES

[1] Kimberley R Allison, Kay Bussey, and Naomi Sweller. 2019. 'I'm going to hell for laughing at this' Norms, Humour, and the Neutralisation of Aggression in Online Communities. *Proceedings of the ACM on Human-Computer Interaction*

3, CSCW (2019), 1–25.
[2] Irwin Altman. 1975. The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. (1975).
[3] Tawfiq Ammari and Sarita Schoenebeck. 2016. "Thanks for your interest in our Facebook group, but it's only for dads" Social Roles of Stay-at-Home Dads. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1363–1375.
[4] Tawfiq Ammari, Sarita Schoenebeck, and Daniel Romero. 2019. Self-declared throwaway accounts on Reddit: How platform affordances and shared norms enable parenting disclosure and support. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
[5] Nazanin Andalibi and Andrea Forte. 2018. Announcing pregnancy loss on Facebook: A decision-making framework for stigmatized disclosures on identified social network sites. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
[6] Nazanin Andalibi and Andrea Forte. 2018. Responding to sensitive disclosures on social media: a decision-making framework. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 6 (2018), 1–29.
[7] Nazanin Andalibi, Oliver L Haimson, Munmun De Choudhury, and Andrea Forte. 2018. Social support, reciprocity, and anonymity in responses to sexual abuse disclosures on social media. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 5 (2018), 1–35.
[8] Nazanin Andalibi, Margaret E Morris, and Andrea Forte. 2018. Testing waters, sending clues: Indirect disclosures of socially stigmatized experiences on social media. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–23.
[9] Camille Bacon-Smith. 1992. *Enterprising Women: Television Fandom and the Creation of Popular Myth*. University of Pennsylvania Press, Philadelphia.
[10] Cristina Bicchieri. 2002. Covenants without swords: Group identity, norms, and communication in social dilemmas. *Rationality and Society* 14, 2 (2002), 192–228.
[11] Rebecca Black, Jonathan Alexander, Vicky Chen, and Jonathan Duarte. 2019. Representations of Autism in Online Harry Potter Fanfiction. *Journal of Literacy Research* 51, 1 (2019), 30–51.
[12] Lindsay Blackwell, Tianying Chen, Sarita Schoenebeck, and Cliff Lampe. 2018. When online harassment is perceived as justified. In *Twelfth International AAAI Conference on Web and Social Media*.
[13] Danah Boyd and Alice Marwick. 2011. Social steganography: Privacy in networked publics. *International Communication Association, Boston, MA* (2011), 93.
[14] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Pyschology* 3 (2006), 77–101.
[15] Carolyn Bronstein. 2020. Pornography, Trans Visibility, and the Demise of Tumblr. *Transgender Studies Quarterly* 7, 2 (2020), 240–254.
[16] Kristina Busse and Karen Hellekson. 2012. Identity, Ethics, and Fan Privacy. In *Fan Culture: Theory/Practice*, Katherine Larsen and Lynn Zubernis (Eds.). Cambridge Scholars Publishing, 38–56.
[17] Julie Ann Campbell, Sarah Evans, Cecilia Aragon, Abigail Evans, Katie Davis, and David P Randall. 2016. Thousands of Positive Reviews: Distributed Mentoring in Online Fan Communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 691–704.
[18] Matthew Carrasco and Andruid Kerne. 2018. Queer Visibility: Supporting LGBT+ Selective Visibility on Social Media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 250.
[19] Centrumlumina. 2013. AO3 Census Masterpost. http://centrumlumina.tumblr.com/post/63208278796/ao3-census-masterpost
[20] Stevie Chancellor, Andrea Hu, and Munmun De Choudhury. 2018. Norms matter: Contrasting social support around behavior change in online weight loss communities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
[21] Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. 2017. You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 31.
[22] Eshwar Chandrasekharan, Mattia Samory, Shagun Jhaver, Hunter Charvat, Amy Bruckman, Cliff Lampe, Jacob Eisenstein, and Eric Gilbert. 2018. The Internet's Hidden Rules: An Empirical Study of Reddit Norm Violations at Micro, Meso, and Macro Scales. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 32.
[23] Alexander Cho. 2018. Default publicness: Queer youth of color, social media, and being outed by the machine. *New Media and Society* 20, 9 (2018), 3183–3200. https://doi.org/10.1177/1461444817744784
[24] Robert B Cialdini and Melanie R Trost. 1998. Social influence: Social norms, conformity and compliance. (1998).
[25] Avery Dame. 2016. Making a name for yourself: tagging as transgender ontological practice on Tumblr. *Critical Studies in Media Communication* 33, 1 (2016), 23–37. https://doi.org/10.1080/15295036.2015.1130846

[26] Ruth A Deller. 2015. Simblr famous and SimSecret infamous: Performance, community norms, and shaming among fans of The Sims. *Transformative Works and Cultures* 18 (2015).

[27] Michael A DeVito, Ashley Marie Walker, and Jeremy Birnholtz. 2018. 'Too Gay for Facebook': Presenting LGBTQ+ Identity Throughout the Personal Social Media Ecosystem. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 44.

[28] Brianna Dym, Jed R Brubaker, Casey Fiesler, and Bryan Semaan. 2019. "Coming Out Okay": Community Narratives for LGBTQ Identity Recovery Work. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–28.

[29] Brianna Dym and Casey Fiesler. 2018. Generations, Migrations, and the Future of Fandom's Private Spaces. *Transformative Works and Cultures* 28 (2018).

[30] Brianna Dym and Casey Fiesler. 2018. Vulnerable and online: Fandom's case for stronger privacy norms and tools. In *CSCW '18 Companion*.

[31] Nicole B Ellison, Lindsay Blackwell, Cliff Lampe, and Penny Trieu. 2016. "The Question Exists, but You Don't Exist With It": Strategic Anonymity in the Social Lives of Adolescents. *Social Media+ Society* 2, 4 (2016), 2056305116670673.

[32] Sarah Evans, Katie Davis, Abigail Evans, Julie Ann Campbell, David P Randall, Kodlee Yin, and Cecilia Aragon. 2017. More Than Peer Production: Fanfiction Communities as Sites of Distributed Mentoring. *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (2017), 259–272. https://doi.org/10.1145/2998181.2998342

[33] Julia R Fernandez and Jeremy Birnholtz. 2019. "I Don't Want Them to Not Know": Investigating Decisions to Disclose Transgender Identity on Dating Platforms. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.

[34] Casey Fiesler and Amy S Bruckman. 2019. Creativity, Copyright, and Close-Knit Communities: A Case Study of Social Norm Formation and Enforcement. *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (2019), 1–24.

[35] Casey Fiesler and Brianna Dym. 2020. Moving Across Lands: Online Platform Migration in Fandom Communities. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW (2020), 1–25.

[36] Casey Fiesler, Jialun "Aaron" Jiang, Joshua McCann, Kyle Frye, Jed R Brubaker, et al. 2018. Reddit rules! characterizing an ecosystem of governance. In *Twelfth International AAAI Conference on Web and Social Media*.

[37] Casey Fiesler, Shannon Morrison, and Amy S. Bruckman. 2016. An Archive of Their Own: A Case Study of Feminist HCI and Values in Design. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. 2574–2585. https://doi.org/10.1145/2858036.2858409

[38] Diana Floegel and Kaitlin L Costello. 2019. Entertainment media and the information practices of queer individuals. *Library & Information Science Research* 41, 1 (2019), 31–38.

[39] Katie Z Gach, Casey Fiesler, and Jed R Brubaker. 2017. "Control your emotions, Potter": An Analysis of Grief Policing on Facebook in Response to Celebrity Death. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 47.

[40] James Paul Gee. 2005. Semiotic social spaces and affinity spaces. In *Beyond communities of practice language power and social context*. 214–232.

[41] Sarah Gilbert. 2020. "I run the world's largest historical outreach project and it's on a cesspool of a website." Moderating a public scholarship site on Reddit: A case study of r/AskHistorians. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2020), 1–27.

[42] Victoria M Gonzalez. 2016. Swan Queen, shipping, and boundary regulation in fandom. *Transformative Works and Cultures* 22 (2016).

[43] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field methods* 18, 1 (2006), 59–82.

[44] Oliver L Haimson. 2018. Social Media as Social Transition Machinery. In *Proc. ACM Hum.-Comput. Interact.*, Vol. 2.

[45] Oliver L Haimson, Avery Dame-Griff, Elias Capello, and Zahari Richter. 2019. Tumblr was a trans technology: the meaning, importance, history, and future of trans technologies. *Feminist Media Studies* (2019), 1–17.

[46] Serena Hillman, Jason Procyk, and Carman Neustaedter. 2014. 'alksjdf; Lksfd' tumblr and the fandom user experience. In *Proceedings of the 2014 conference on Designing interactive systems*. 775–784.

[47] Michael A Hogg and Scott A Reid. 2006. Social identity, self-categorization, and the communication of group norms. *Communication theory* 16, 1 (2006), 7–30.

[48] Christine Horne. 2001. The enforcement of norms: Group cohesion and meta-norms. *Social psychology quarterly* (2001), 253–266.

[49] Henry Jenkins. 1992. *Textual Poachers: Television Fans and Participatory Culture*. Routledge.

[50] Shagun Jhaver, Iris Birman, Eric Gilbert, and Amy Bruckman. 2019. Human-machine collaboration for content regulation: The case of Reddit Automoderator. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 5 (2019), 1–35.

[51] Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2019. Does transparency in moderation really matter? User behavior after content removal explanations on reddit. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–27.

[52] Sanjay Kairam, Mike Brzozowski, David Huffaker, and Ed Chi. 2012. Talking in circles: selective sharing in google+. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1065–1074.

[53] Anna Kasunic and Geoff Kaufman. 2018. "At Least the Pizzas You Make Are Hot": Norms, Values, and Abrasive Humor on the Subreddit r/RoastMe. In *Twelfth International AAAI Conference on Web and Social Media*.

[54] Charles Kiene, Jialun Aaron Jiang, and Benjamin Mako Hill. 2019. Technological Frames and User Innovation: Exploring Technological Change in Community Moderation Teams. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.

[55] Charles Kiene, Andrés Monroy-Hernández, and Benjamin Mako Hill. 2016. Surviving an eternal september: How an online community managed a surge of newcomers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1152–1156.

[56] Robert E Kraut and Paul Resnick. 2012. *Building successful online communities: Evidence-based social design*. Mit Press.

[57] Priya Kumar, Shalmali Milind Naik, Utkarsha Ramesh Devkar, Marshini Chetty, Tamara L Clegg, and Jessica Vitak. 2017. "No Telling Passcodes Out Because They're Private": Understanding Children's Mental Models of Privacy and Security Online. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 64.

[58] Jessa Lingel and danah boyd. 2013. "Keep it secret, keep it safe": Information poverty, information norms, and stigma. *Journal of the American Society for Information Science and Technology* 64, 5 (2013), 981–991.

[59] Alexis Lothian. 2013. Archival anarchies: Online fandom, subcultural conservation, and the transformative work of digital ephemera. *International Journal of Cultural Studies* 16, 6 (2013), 541–556.

[60] Danielle Lottridge, Nazanin Andalibi, Joy Kim, and Jofish Kaye. 2019. "Giving a little 'ayyy, I feel ya' to someone's personal post": Performing Support on Social Media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–22.

[61] Jacqueline Low. 2019. A pragmatic definition of the concept of theoretical saturation. *Sociological Focus* 52, 2 (2019), 131–139.

[62] Enrico Mariconti, Guillermo Suarez-Tangil, Jeremy Blackburn, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Jordi Luque Serrano, and Gianluca Stringhini. 2019. "You Know What to Do": Proactive Detection of YouTube Videos Targeted by Coordinated Hate Attacks. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.

[63] Annette Markham. 2012. Fabrication as ethical practice: Qualitative inquiry in ambiguous internet contexts. *Information, Communication & Society* 15, 3 (2012), 334–353.

[64] J Nathan Matias. 2019. Preventing harassment and increasing group participation through social norms in 2,190 online science discussions. *Proceedings of the National Academy of Sciences* 116, 20 (2019), 9785–9789.

[65] Nora McDonald and Andrea Forte. 2020. The Politics of Privacy Theories: Moving from Norms to Vulnerabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[66] Tsubasa Morioka, Nicole B Ellison, and Michael Brown. 2016. Identity work on social media sites: Disadvantaged students' college transition processes. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 848–859.

[67] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[68] Abigail Oakley. 2016. Disturbing hegemonic discourse: Nonbinary gender and sexual orientation labeling on Tumblr. *Social Media+ Society* 2, 3 (2016), 2056305116664217.

[69] Elinor Ostrom. 2000. Collective action and the evolution of social norms. *Journal of economic perspectives* 14, 3 (2000), 137–158.

[70] Elissa M Redmiles, Jessica Bodford, and Lindsay Blackwell. 2019. "Just Want to Feel Safe"A Diary Study of Safety Perceptions on Social Media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 405–416.

[71] Morgan K Scheuerman, Stacy M Branham, and Foad Hamidi. 2018. Safe Spaces and Safe Places : Unpacking Technology-Mediated Experiences of Safety and Harm with Transgender People. In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 2.

[72] Zachary Schmitt and Svetlana Yarosh. 2018. Participatory design of technologies to support recovery from substance use disorders. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 156.

[73] Sarita Yardi Schoenebeck. 2013. The secret life of online moms: Anonymity and disinhibition on youbemom. com. In *Seventh International AAAI Conference on Weblogs and Social Media*.

[74] Joseph Seering, Robert Kraut, and Laura Dabbish. 2017. Shaping pro and anti-social behavior on twitch through moderation and example-setting. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*. ACM, 111–125.

[75] Joseph Seering, Felicia Ng, Zheng Yao, and Geoff Kaufman. 2018. Applications of Social Identity Theory to Research and Design in Social Computing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW'18). ACM, New York, NY, USA.*

[76] Irving Seidman. 2006. *Interviewing as qualitative research: A guide for researchers in education and the social sciences.* Teachers college press.

[77] Mel Stanfill. 2013. "They're Losers, but I Know Better": Intra-Fandom Stereotyping and the Normalization of the Fan Subject. *Critical Studies in Media Communication* 30, 2 (2013), 117–134.

[78] Mel Stanfill. 2015. The interface as discourse: The production of norms through web design. *New media & society* 17, 7 (2015), 1059–1074.

[79] Mel Stanfill. 2018. The unbearable whiteness of fandom and fan studies. *A Companion to Media Fandom and Fan Studies* (2018), 305.

[80] John Tulloch, Henry Jenkins, et al. 1995. *Science fiction audiences: watching Doctor Who and Star Trek.* Routledge.

[81] Ashley Marie Walker and Michael Ann DeVito. 2020. "'More gay'fits in better": Intracommunity Power Dynamics and Harms in Online LGBTQ+ Spaces. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* 1–15.

[82] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.

[83] Pamela Wisniewski, Heng Xu, Mary Beth Rosson, Daniel F Perkins, and John M Carroll. 2016. Dear diary: Teens reflect on their weekly online risk experiences. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* ACM, 3919–3930.

[84] Svetlana Yarosh. 2013. Shifting dynamics or breaking sacred traditions? The role of technology in twelve-step fellowships. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* 3413–3422.

Received January 2020; revised June 2020; accepted July 2020