

Counting Carrds: Investigating Personal Disclosure and Boundary Management in Transformative Fandom

Kelly Wang
wang.kelly1@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

Dan Bially Levy
dbiallyl@macalester.edu
Macalester College
St. Paul, Minnesota, USA

Kien T Nguyen
knguyen2@macalester.edu
Macalester College
St. Paul, Minnesota, USA

Ada Lerner
ada@ccs.neu.edu
Northeastern University
Boston, Massachusetts, USA

Abigail Marsh
amarsh1@macalester.edu
Macalester College
St. Paul, Minnesota, USA

ABSTRACT

The privacy practices of transformative fandom are of interest to HCI researchers both for the community's high proportion of queer members and for the community's sophisticated privacy norms and behaviors. We investigated fans' use of single-serving websites on Carrd.co ("Carrds") as personal profiles linked from Twitter accounts. We scraped Twitter to gather 5252 Carrds from fans in a variety of fandoms, which we analyzed using a combination of keyword searches and hand-coding. Fans' Carrds frequently disclose queer identity, and articulate a complex system of community values and boundary management. Inspired by how these findings aren't well-explained by individual theories of privacy, we articulate first steps towards a theory of collective privacy based in a communal process of values construction, trust building, and personal disclosure that we believe helps us to understand the sophisticated nature of fans' observed behaviors.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; • **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms**; Computer supported cooperative work; Social media; Social content sharing.

KEYWORDS

privacy, theory, fandom, norms, quantitative methods, Twitter, Carrd.co

ACM Reference Format:

Kelly Wang, Dan Bially Levy, Kien T Nguyen, Ada Lerner, and Abigail Marsh. 2024. Counting Carrds: Investigating Personal Disclosure and Boundary Management in Transformative Fandom. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3613904.3642664>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05.
<https://doi.org/10.1145/3613904.3642664>

1 INTRODUCTION

The HCI community is increasingly turning its attention to the concerns of vulnerable populations of users, recognizing that "users" are not a monolithic group and that designing for the imagined prototypical user often fails to serve the needs of other groups. Usable privacy researchers now acknowledge that privacy concerns often stem from a user's identity or the context of technology use, which we must understand in order to ensure the privacy of a diverse population. In this paper, we highlight one such community, transformative fandom as it existed on Twitter in 2022, and present a study of fans' disclosure habits and boundary-setting with the aim of engaging with privacy decision-making as a community endeavor.

Previous work in HCI has explored the privacy behavior of transformative fandom [14], a community that is difficult to fully profile due to its decentralized online nature, but which has consistently self-identified in community surveys as both predominantly female and predominantly queer [28, 44]. Due to the tension between wanting to participate in fan activities and, often, wishing to keep their involvement secret from non-fan relationships, fandom has developed sophisticated notions of privacy that both HCI and security and privacy researchers may be able to learn from [14].

In this study, we investigate the practice of online self-disclosure in transformative fandom, particularly the phenomenon of fans describing their identities and interaction boundaries through single-serving websites on Carrd.co ("Carrds"). Carrd.co is the observed preferred website for this behavior, and is in use by fans across many sites where fans share and discuss fanworks, including Twitter, Tumblr, and Instagram. These Carrds serve as extended user profiles, given the character and link limitations of Twitter user bios (see Figure 1). Fans will often disclose extensive sensitive, personal, and/or potentially contentious information about themselves in such Carrds, which contrasts with prior work on the self-censorship behaviors of queer and other marginalized groups [11, 27, 31]. We argue that this contrast highlights the importance of understanding privacy decision-making through the unique dynamics of online communities as an ecosystem, in addition to studying individual users' holistic experiences of digital privacy across platforms.

We find that many of the fans in our sample disclose potentially sensitive information about themselves, including gender and sexual identities, disability status, and age, and we reinforce past findings that fandom communities are predominantly queer. Fans

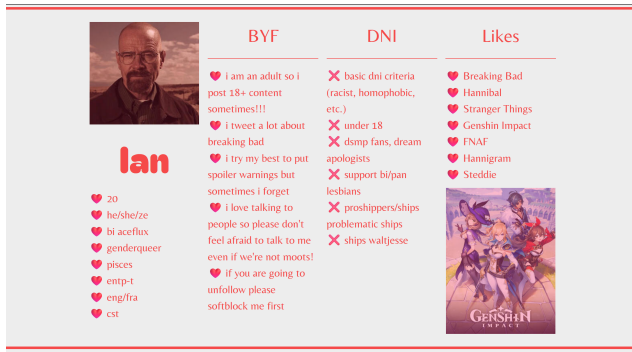


Figure 1: An example of a typical Carrd, created by the researchers. This demonstrates a typical layout and content, including a description of the user (the identity section), sections for “Do Not Interact” (DNI) and “Before You Follow” (BYF) criteria, and a list of the user’s interests. The image on the left is a still from the TV series *Breaking Bad*, while the image on the right is promotional artwork of the video game *Genshin Impact*.

also used two distinct concepts to communicate their interaction boundaries: “Do Not Interact” (DNI) and “Before You Follow” (BYF) lists, the components of which we categorized and counted in order to describe those boundaries.

In interrogating these results we find that many of these disclosure and interaction behaviors are challenging to explain using existing models of privacy. Noting that these challenges stem from the individual nature of existing theories, we propose first steps toward a theory of **collective privacy**. Inspired by our results and informed by existing literature from a variety of areas, including Waldman’s theory of privacy as trust [52], our theory seeks to explain how communities such as transformative fandom engage in a cyclic and communal process of values construction, trust building, and personal disclosure in order to reap the benefits of sharing the self in an environment of trust.

We make the following contributions:

- We present an analysis of 5252 single-serving websites (“Carrds”) containing Twitter users’ identity disclosures and boundary statements, identifying common types of disclosure and boundary-setting.
- We argue for the understanding of these disclosures and boundaries as privacy behaviors, as users negotiate their membership within Twitter fandom communities.
- We present first steps towards a theory of collective privacy inspired by our results, which expands on existing privacy theory by postulating a cyclic process of values construction, trust building, and personal disclosure.

2 BACKGROUND AND RELATED WORKS

2.1 Privacy Theory

Privacy scholars have long debated how best to conceptualize privacy, and their theories form valuable context and inspiration for our work. We hold with Solove in believing that privacy is likely

best conceived of a host of connected concepts, rather than a single unifying theory [48]. Many theories of privacy conceptualize it as a right to be let alone or a right to secrecy [25, 53, 54], ideas which capture certain key aspects of privacy while being vulnerable to critique as too narrow, too broad, or omitting others characteristics of privacy. One prominent idea that has become a key principle of such analyses is Nissenbaum’s theory of contextual integrity, which covers many situations by enabling flexible analysis of information flows in terms of contextual norms [36].

In contrast, other scholars have proposed almost counter-intuitive yet powerful reframings of privacy. For example, this work draws on Waldman’s theory of privacy as trust, which analyzes privacy through the trust which enables—rather than prevents—sharing with others [52]. Though perhaps surprising at first, this focus on sharing rather than hiding is intuitively privacy-related when compared with the idea of confidentiality, which involves trusting recipients to take appropriate care with sensitive information. Waldman highlights how beneficial sharing is to us as inherently social creatures, and thus focuses privacy on protecting those benefits by creating trusting environments in which sharing is safe. We integrate this theory into our analyses of sensitive disclosures by fans, reasoning about how that sharing may actually indicate strong trust within this community and how it may enable that trustful privacy through a process of constructing trust norms and promoting group solidarity.

2.2 Online Identity Disclosure

Identity disclosure is increasingly recognized as area of particular interest for privacy researchers. An early user study examined the relationship between self-censorship and audience targeting, demonstrating that with a greater ability to target a relevant audience, users would share more on Facebook [47]. This can be understood as a foundational idea for later examinations of specific identity groups, especially queer users of social media.

DeVito et al.’s study of queer social media users observed that queer users were modulating what they shared about their queerness through audience awareness [11]. They highlighted that information visibility—what we discuss in this paper as disclosure—was a valuable strategy for managing online safety and privacy among their participants, and that decisions about visibility were made according to the subset of the platform which the participants viewed as their community, not the user base of the platform as a whole. In a study of transgender social media users, Lerner et al. also identified self-censorship as an important safety technique [27]. However, both studies examined queer users’ experiences broadly, including in spaces where they may not have been out as queer, and indeed found that these users often altered their identity presentations to reduce visibility. In contrast, we studied a specific context—transformative fandom—in which users’ disconnection from their legal identities allowed them to more freely perform identity exploration. This can be related to Andalibi et al.’s study on sexual abuse-related disclosures on the Reddit social media platform, where they found that the platform’s anonymity facilitated intimate disclosure and support seeking [2].

In order to understand these ecosystems in full, we must also look towards studies of language and disclosure in queer online

communities, which mirror patterns we see in our dataset. Studying queer users—primarily young people using Tumblr—Schudson and van Anders emphasize the value of creating labels for non-queer identities (e.g. “cisgender” as opposite “transgender”) as a means of challenging hegemonic discourses [46]. Supporting how disclosures of queerness are normalized on Tumblr, Cavalcante discusses the perception of Tumblr as a queer “utopia,” which provides an important space for young queer people to exist without the need to explain the meaning of their identities [7].

These works also highlight the importance of platform design in users’ sense of safety, which multiple features of Tumblr seem to provide: Tumblr allows users to maintain pseudonymous identities [46], uses unstructured profiles that enable users to self-select what they want to disclose [37], and does not automatically highlight comments on posts [7]. Contrast this with Facebook which requires users to use names by which they are addressed in “everyday life,” [16] a policy that was cited as disproportionately targeting transgender users and drag performers [56].

Thus, as Schudson and van Anders observe, while a platform does not need to be designed *for* the queer community, it must still be able to facilitate that community’s goals [46]. Tumblr was and largely continues to be used by fandom and the queer community, but platform changes pose an ongoing risk to stability—a risk which was partially realized when, in late 2018, Tumblr banned explicit adult content on the platform [12]. We examine the Twitter transformative fandom community as a result of this ban, which shifted many parts of fandom and queer community towards Twitter, as a platform which was willing to host adult content.

2.3 Transformative Fandom

We understand transformative fandom (also called transformational fandom and media fandom) as indicating the community of fans, largely female, whose activities are oriented around the creation, consumption, and discussion of fanworks (fanfiction, fan art, fan videos, etc.) [38, 39]. This community has been a subject of study for over three decades [3, 23], and recognized as a space predominantly occupied by women throughout that span. In more recent attempts to survey the community, fandom has also been recognized as a community with a largely queer membership [28, 44]. While surveying a large, decentralized online community like fandom poses challenges to validity, both of these surveys illustrate that there are many people who identify as queer within fandom. When combined with the relative youth of many fans as seen in those surveys and the frequent appearances of queerness and explicit sexual content as subjects in fanworks [34, 35], fandom can be understood as an important space for some fans to explore their gender and sexuality, which indeed has been documented in prior work [34, 43].

As a meeting point for queerness and sexually-explicit material, fans are particularly privacy-sensitive with respect to those outside of the community. Earlier work has discussed fans’ concerns about being misunderstood or harassed by outsiders who are not fans, especially when contentious comments spread beyond the relative safety of one’s own fan community [6, 15, 24], with fans citing potential real-world consequences such as job or relationship loss as a result of their membership in fandom being publicly discussed. As

a largely female and queer community, fans worried about the risk of their digital identities and legal names being connected through doxxing [14]. The collection of social media content for research or journalism can also be an unpleasant surprise for communities [4, 18]. We are careful in this work not to make it possible to re-identify users whose Carrds were included in our analysis (see Section 3.3).

Comparatively less work explores fans’ attitudes towards within-community privacy violations, only noting that privacy norms may be shifting over time to include more disclosure of personal information [13]. One of the primary contributions of our work is to examine those identity disclosures within a specific context—transformative media fans using Twitter as a platform to share and discuss fanworks—and thus enable a discussion of how those disclosures enable boundary-setting and community formation.

3 METHODS

We examine identity disclosure and personal boundaries of fans on Twitter by analyzing the contents of publicly available, single-serving websites, called Carrds, linked from Twitter bios, discovered via crawling the Twitter friend graph from significant fan accounts.

3.1 Data Collection

Twitter serves as a gathering space for many online communities, including transformative fandom. Fans on Twitter discuss fannish interests, share fanworks (fan art, fanfiction, and more); post live updates as they consume the original source material of their fandom, such as “livetweeting” their thoughts while watching an episode of a television series; connect with other fans over shared interests; and organize community events. Twitter currently limits user bios to 160 characters and one link associated with their profile, though users can share other links in bios. As a result, fans who want to share detailed information about themselves often use Carrd.co, a free single-serving website platform for developing and hosting simple HTML and CSS websites with text and image content.

From June to August in 2022, Twitter still had a publicly-available API. We used this API to perform breadth-first searches of depth 2 starting from starting point accounts. We chose as starting points users with mid-sized followings (1,000–10,000 Twitter followers) in six media categories of fandoms: books, video games, Eastern or Western live action video, and Eastern or Western animation. Starting points were identified by searching hashtags for specific fandoms active during the first half of 2022. We identified all accounts that followed any of the starting points (depth 1), as well as all accounts that followed those followers (depth 2). We then scraped the contents of each account’s Twitter profile, collecting all URLs contained in the bio or link fields, and filtered for those hosted on Carrd.co. We crawled all those Carrd.co websites, automatically clicking on all internal anchor links to view subpages, and collecting the contents of each subpage as plain text, raw HTML, and as screenshots. In total, we collected 5370 Carrd URLs, of which 5252 resolved to a website.

3.2 Analysis

We took an exploratory data analysis approach [51], using both qualitative and quantitative methods not to generalize our findings beyond our sample but to understand interesting and surprising

characteristics of our dataset that challenge our current understandings of privacy.

To quantify identity characteristic disclosure, we initially performed keyword searches over Carrd contents, but found false positive rates as high as 80% for some identifiers, (e.g., “man,”). Thus we manually coded the entire dataset for these identity characteristics. Four authors individually read the plain text version of each Carrd, marking the section of the text (“identity section”) where a person described their identity, if present, and coding within that section for age, gender, pronouns, and sexuality. Other identity characteristics (e.g., disability) less prone to false positives were quantified using keyword searches informed by our manual identification of terms in identity sections and existing lists of disabilities [8, 40].

Some Carrds contained multiple identity sections, because they were used by multiple people or because the author noted having multiple alternate personalities, and we coded all distinct identity sections. Because these categories of multiplicity were not easily distinguishable and they were rare (174 Carrds), we elected to leave them out of our final analysis.

We considered but did not attempt to quantify racial and ethnic identities. Some Carrds included them, while others included only related concepts, such as languages spoken, country of residence, and time zone. Since racial and ethnic identities are highly contextual to specific cultures, we chose not to attempt to describe them under a single set of categories (e.g., American racial categories).

We also analyzed two types of boundary management sections. One was typically labeled “Do Not Interact” or “Do Not Follow,” (DNI or DNF), and included criteria indicating who should not interact with the Carrd’s owner. The other type of section consists of warnings and information for potential followers, labeled “Before You Follow,” (BYF). All five authors individually read randomized samples of approximately 10% of the total Carrd dataset, noting the DNI criteria and BYF contents if any. We then quantified the prevalence of these criteria and contents using regular expressions and keyword searching. We performed separate analyses for DNIs and BYFs. Through subsequent discussion, we grouped these criteria into the categories reported in Section 4.

3.3 Ethical Considerations

We did not recruit human participants, but we recognize that fans may consider their Carrds private despite their existence on public websites [4, 15, 18], and may worry about harassment, doxxing, or other social harms, as noted in previous research on fandom [6, 14, 15, 24]. We carefully considered the ethical implications of this research and followed best practices established by prior work on research ethics in online fandom.

Prior research has found that fans rely on obscurity as a means of privacy, and recommend that data is obscured to prevent re-identification when user permission is not acquired [15, 18]. We did not seek users’ permission to use the data because of the large size of our dataset, however we took several measures against re-identification. In reporting our results, we refrain from including names, usernames, URLs, direct quotes, or identifiable screenshots of Carrds, as well as the names of the fandoms or accounts used in constructing the dataset. Additionally, as we consider identity categories to be potentially identifying in aggregate, we do not list

specific combinations of identities (such as a fan’s age, sexuality, and pronouns) and we do not specify the exact count of Carrds that contained identity labels or phrases that could be reidentifiable.

We also recognize the potential harm of aggregating public data, as this increases visibility. This is an explicit concern for fans, who have discussed the importance of seeking permission from fans before linking to their posts [30]. In consideration of this, our dataset will not be released publicly.

3.4 Positionality

We acknowledge that both queer identities and membership within fandom are sensitive topics which demand consideration from researchers studying these communities. We present our identities in order to contextualize our experiences of queerness and fandom. All five authors identify as queer. Three authors identify as fans within transformative fandom, something that has been noted in prior fan studies work as important to understand the community norms when studying this space [15].

3.5 Limitations

Our research team was only fluent in English, and we do not include Carrds written in other languages in our analysis.

We used keyword searches to identify some aspects of identity and for all parts of our DNI and BYF analysis, and thus we did not capture context such as the valence associated with a term, whether a term appeared as a substring of something else, or whether “18+” described an age boundary instead of the nature of their content. We did examine random samples of each keyword search and determine false positive rates, however, to improve our search expressions and minimize context-related errors. Similarly, our analysis does not provide a full qualitative analysis of all concepts used as either DNIs and BYFs, as we could not reliably capture some ideas with the limitations of keyword search—for instance, DNIs or BYFs phrased as full sentences.

We parsed Carrds’ DNI and BYF sections using a script, which could have skipped or truncated viable data that was delineated in unexpected ways. We estimate the upper bound of false negatives to be 41 Carrds for DNI and 126 Carrds for BYF sections.

4 RESULTS

We collected 5370 total Carrds, 118 (2.2%) of which returned a “page not found” message. We analyzed the remaining 5252 manually and automatically, as described in Section 3.2. As we hand-coded, we noted that nearly all fell into one of three major categories: (1) Profile pages for a single user (“**personal Carrds**”), (2) Carrds advertising art commissions¹ (“**artist Carrds**”), and (3) Carrds representing fanzines² or other organizations or groups (“**organization Carrds**”). We identified artist Carrds via inclusion of the word “commission,” which yielded 1154 (22.0%) of the 5252 total Carrds. Organization Carrds were rare and we could not reliably and automatically identify them, so we did not analyze them separately from personal Carrds, which make up most of the remaining 78%

¹Fanartists used Carrds to host galleries of past work and pricing information for potential clients.

²Fanzines are non-professional magazines produced by and for fans of particular media.

of the dataset. Most of our results focus on the personal Carrds, but we first describe the dataset across these groups.

	All Carrds	Personal Carrds	Artist Carrds
Mean	1,611	1,433	2,241
Minimum	0	0	45
Maximum	47,723	47,723	28,383
25th Percentile	519	520	504
50th Percentile	1,023	960	1,473
75th Percentile	1,801	1,586	2,942

Table 1: Summary statistics describing the character length distributions for all accessible Carrds, the set of Carrds that did not contain the term “commission” (“Personal Carrds”), and the set of Carrds that did contain the term “commission” (“Artist Carrds”).

Collected cards varied from 0 (no content) to 47,723 characters, with a length distribution that is long-tailed and highly right-skewed with an average length of 1611 characters and median length of 1023 characters. Artist Carrds were on average longer than personal Carrds (Table 1), since they often included longer sections outlining terms of use, pricing, rules for commissions, and commission examples. As we hand-coded the data, we observed that longer outliers were typically but not exclusively organizational Carrds.

Many personal Carrds followed similar structures, often including sections disclosing personal identity characteristics (“**identity section**”), “DNI” (“Do Not Interact”) sections (“**DNIs**”), indicating who the user doesn’t want to interact with, “BYF” (“Before You Follow”) sections (“**BYFs**”) listing information the user thinks that others should read before following them, and sections listing the media, characters, and/or ships³ the user enjoys. Given the commonness of identity sections (appeared in 73% of Carrds), DNIs (50% of Carrds), and BYFs (43% of Carrds) and their relevance to our topic of study, the remainder of our results expand on these three types of sections.

4.1 Disclosure in Identity Sections

Of 5252 total Carrds, 3835 (73.0%) contained at least one identity section and 3661 (69.7%) contained exactly one identity section. Carrds containing multiple identity sections included ones where a single person described their identity on multiple pages of their Carrd site, projects (such as fanzines) with multiple organizers each describing their identities, friends and partners sharing one Carrd site, and people describing their alternate personalities. As we could not automatically distinguish between these varied cases and multiple identity Carrds were a small portion of the data, we omitted those 174 Carrds from the following analysis.

We analyzed disclosures of age, queerness, and disability status in particular, as we observed a large number of fans decided to disclose these characteristics despite the vulnerability such disclosures might carry. Below, we discuss our findings for each of these three

³Ship” is used within fandom to denote character relationships, whether romantic or platonic. A fan who likes a particular relationship “ships” those characters.




Figure 2: The distribution of ages disclosed through Carrds. The oldest age disclosed was 52, the youngest was 12, and the most common age given was 18.

categories of disclosures, and then report common combinations of disclosures (i.e. co-occurring in a single Carrd) in Section 4.1.4.

4.1.1 Age. Out of the 3661 Carrds with exactly one identity section, 2985 (81.5%) disclosed their age. 1243 age disclosures indicated being younger than 18 (41.6% of age disclosures), 32 disclosed age ranges that made it ambiguous whether they were 18 or older (1.1%), and the final 1710 clearly disclosed that they were adults (57.3%). While many of these strategies obscure the exact age of the user, a large majority (78.7% of age disclosures) disclosed their exact age. Of those who did use obscuring strategies, the plurality (10.8% of age disclosures) declared minor/adult status without providing any specific age or range.

Looking only at those who listed an exact age, our sample ranged from 12 years old (one year younger than the minimum required age for making a Twitter account) to 52 years old. Figure 2 shows the age distribution, which is slightly right-skewed with a mean of 19.2 and median of 19. Note that this skew is likely impacted both by the bounds of age as a category and by the minimum age of use as stated in the Twitter Terms of Service.

Our sample is thus younger than Twitter as a whole, whose most common age range and median was 25-34 years old in 2021 (making up 38.5% of the sample) [9]. Our sample also appears to be younger than typical for fandom, based on comparisons with surveys of Tumblr and Archive of Our Own users, two platforms that largely serve fans. The bulk of Tumblr users are in the 18-29 (39%) and 30-39 ranges (32%) as of 2022 [49], while the bulk of Archive of Our Own users are in the 18-24 (36%) and 25-34 ranges (43%) as of 2022 [44]. Since these surveys did not poll users below 18 years old, our ability to compare is limited, but we can confirm that our sample contains proportionately fewer adults over the age of 25 or 30 (respectively).

4.1.2 Queer Identities. Consistent with past work in transformative fandom [28, 44], our sample was predominantly queer, with 2901 (79.2% of Carrds with exactly one identity section) disclosing identifiers associated with the LGBTQ+ community.

Pronouns. Out of the 3661 Carrds with exactly one identity section, 3183 (86.9%) disclosed their pronouns, with 2489 (68.0%) using pronouns other than he/him or she/her and 945 (25.8%) using

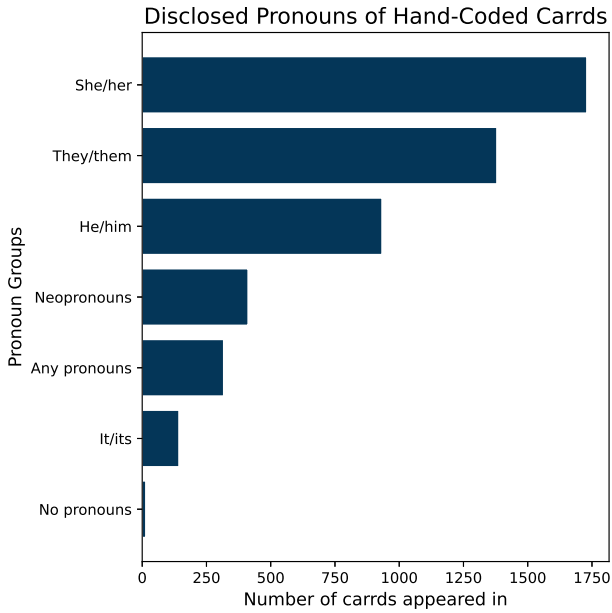


Figure 3: Counts of pronouns used by fans, as disclosed on Carrds. Labels were grouped according to similarity, where applicable—for example, “she/her,” “she,” and “she/her/hers” belong to the same group. Counts are non-exclusive, and a Carrd which listed multiple pronouns is counted in each applicable total.

multiple pronoun sets (e.g., “she/they”). Notably, the use of neopronouns (e.g., “xe/xem”) was common in our dataset, with 404 (11.0%) using neopronouns, much higher than recent estimates of 4% neopronoun use among queer youth [50]. Figure 3 summarizes this data.

Gender. Figure 4 summarizes the 1080 Carrds (29.5% of those with one identity section) that explicitly disclosed information about gender, assigned gender at birth, or transgender status. Such information includes terms that indicated gender (e.g., man, woman, non-binary) and terms that indicated relationship to assigned gender at birth (e.g., transgender, cisgender, intersex, perisex). We list the specific labels that went into those groups in Table 2, in order to emphasize the diversity of these categories.

Confirming recent work [50], we found that fans identified with a wide variety of gender labels. 965 (26.4%) gender terms were definitely associated with the queer community, and 688 (18.8%) disclosed a gender identity under the non-binary umbrella, with 362 solely listing the term “non-binary” (or a derivative word) while the other 326 included some other more specific label (e.g., genderfluid, agender, demigirl). While the plurality of Carrds specifying gender belonged to non-binary people, this doesn’t show that our overall sample was plurality non-binary, since many Carrds did not disclose gender.

Sexual and Romantic Orientation. Figure 5 summarizes the 1752 Carrds (47.9% of those with one identity section) that disclosed information about their sexual and/or romantic orientation, with




Figure 4: Counts of gender labels used by fans, as disclosed on Carrds. Labels were grouped according to similarity, where applicable; groupings which exceeded the chart space are indicated with an asterisk and listed in Table 2. Counts are non-exclusive, and a Carrd which listed multiple pronouns is counted in each applicable total.

Gender Labels in Grouped Categories	
Genderfluid	genderfluid, genderfaun, genderfae, genderdoe, genderfaunet, demifluid, cadensgender
Man	male, man, boy, guy, dude (both cis- and transgender)
Woman	female, woman, girl, gal (both cis- and transgender)
Genderflux	girlflux, boyflux, femflux, demiflux
Xenogenders	catgender, vampgender, stargender, candygender, canisgender, moongender, clowngender, thanatogender, pupgender, jinxgender, pikagender, strawbcagender, zombiegender
Other	feminine, femme, masculine, masc-aligned, mingender, gendervoid, gender apathetic, cassgender, autigender, bordgender

Table 2: Grouping categories for gender labels used in Figure 4. Gender identity definitions were confirmed by web search when not explained or previously known to the researchers; all labels we observed had existing definitions.

1729 (47.2%) disclosing labels associated with the queer community.




Figure 5: Counts of sexual orientation labels used by fans, as disclosed on Carrds. Labels were grouped according to similarity, where applicable; groupings which exceeded the chart space are indicated with an asterisk and listed in Table 3. Counts are non-exclusive, and a Carrd which listed multiple orientations is counted in each applicable total.

Orientation Labels in Grouped Categories	
Asexual Spectrum	asexual, demisexual, demi, aspec, aroacespec, aceflux, graysexual, acepec, aegosexual, autochorissexual, cupiosexual
Aromantic Spectrum	aromantic, demiromantic, demi, aspec, aroacespec, arospec, aroflux, grayromantic, cupioromantic, platonioromantic, lithromantic, recipromantic
Other	abromantic, nblnb, abrosexual, polysexual, androsexual, allosexual

Table 3: Grouping categories for sexual orientation labels used in Figure 5. Gender identity definitions were confirmed by web search when not explained or previously known to the researchers; all labels we observed had existing definitions.

Notably, the asexual and aromantic communities were highly represented in this data, with some partially obscuring the details of their orientation by saying “acespec,” “arospec,” “aspec,” or “aroacespec,”⁴ but many choosing to disclose the specific label(s) they identify with under the asexual or aromantic spectra (listed in Table 3, which breaks down how we grouped these for Figure 5).

⁴“Spec” is short for “spectrum,” and is used as part of some queer identity labels to denote membership in a wider group of similar identities.




Figure 6: Counts of the number of Carrds which disclosed specific types of disabilities. Disability statuses were counted through keyword search. Categories are non-exclusive; Carrds which included multiple disability statuses are counted in each category.

4.1.3 Disability Status. Out of the 3661 carrds with exactly one identity section, 401 (11.0%) disclosed disability, including information about physical and mental disabilities, chronic physical conditions, mental illness, and neurodivergence, in addition to disclosing whether they are able-bodied or neurotypical. We found that the large majority of those who disclosed disability status were neurodivergent users, and specifically Autistic (172, 4.7%) and ADHDer (165, 4.5%) users. This represents a relatively high presence of autism in our dataset compared to the 2.8% prevalence estimated in the United States in 2020 [20], and a lower than expected presence of ADHD in our dataset compared to the 9.6% prevalence estimated in the United States for 2016-2019 [19].

4.1.4 Disclosure Patterns. Table 4 summarizes the frequency of disclosure of each type of identity characteristic we studied. Age and pronouns were most commonly disclosed (over 80% of Carrds with exactly one identity section), followed by orientation (47.9%), gender (29.5%), and disability (11.0%). This pattern matches the frequencies with which Carrds included combinations of disclosures, with the three most common combinations being age + pronouns (1126, 30.8%), age + sexuality + pronouns (704, 19.2%) and age + sexuality + pronouns + gender (666, 18.2%).

While disability was overall more rarely disclosed, its disclosure tended to accompany the disclosure of other characteristics. Among those disclosing disability, the most common practice was to disclose all other studied characteristics (162, 4.4%), and well over half of all Carrds that disclosed disability status disclosed at least three other characteristics. We hypothesize reasons for these trends in the discussion (Section 5).

Identity Type	Prevalence
Pronouns	3183 (86.9%)
Age	2985 (81.5%)
Orientation(s)	1752 (47.9%)
Gender	1080 (29.5%)
Disability	401 (11.0%)

Table 4: Frequency of disclosure of our major categories of identity characteristics. Fractions are out of the 3661 Carrds with exactly one identity section.

4.2 Do Not Interact (DNI)

Fully half of the Carrds that we collected (50.0%, 2625 out of 5252) contained lists of criteria setting hard boundaries around interaction on Twitter under the header “Do Not Interact,” “DNI,” “Do Not Follow,” or “DNF.” The majority of these DNIs appeared in personal Carrds (2461, 93.7%), and the remaining 164 DNIs (6.3%) appeared in artist Carrds.

We display the most prevalent terms appearing in these DNIs in Figure 7. Most terms referred to groups of people, such as those holding certain prejudices (e.g. racists, homophobes, TERFs, sexists, ableists), those of certain ages, or people with specific identities (e.g. multi-spectrum attraction [mspec] lesbians). Specific types of fan content were also commonly used in DNIs: content that was not safe for work, “problematic,” contained incest, contained romantic or sexual relationships with an age gap (“pedophiles”), or related to Minecraft YouTube personalities (“MCYT+”). These content-related DNIs represented distinct within-community debates around allowable engagement with fandom which have only recently begun to receive attention from cultural commentators [41, 57] and academics [1, 42]. We present them here not with the aim to codify a definition of proshipping or otherwise weigh in on an active, ongoing matter of community discourse, but to show that this discourse affected how fans drew interaction boundaries.

Within DNIs, we observed many fans listing “basic DNI criteria” as the first item (1247 Carrds). This suggests that in these fan communities, there was an implicit understanding of what type of qualities were unwelcome. Some of these “basic DNI criteria” DNIs explicitly enumerated those criteria (283 Carrds); the most common terms were: racist (86.57%), etc (85.16%), homophobic (72.44%), transphobic (40.99%), ableist (23.32%), and sexist (15.55%). “Basic” DNIs focused on prejudicial beliefs, creating a picture of community norms in which prejudices were unwelcome and signalling one’s opposition to prejudices was an important part of aligning with community values.

We also found that many DNIs (1133 Carrds) listed age restrictions for followers. In Figure 8, we analyze the 778 Carrds that both disclosed their exact age and specified exact age(s) for age-related DNIs. The interaction formed a trend: fans most wanted to connect with users who were similar in age to themselves (darker regions), with the starkness of age restrictions being strongest for minors and tapering off after age 18.




Figure 7: Terms which were commonly present in “Do Not Interact” (DNI) sections of Carrds. Counts are non-exclusive; Carrds which used multiple terms are included in each total.




Figure 8: Comparison of disclosed age (x axis) and age restriction in DNI (y axis) across the 778 Carrds which both reported a specific age and had an age-related DNI. Darker regions indicate a higher percentage of individuals who allowed interaction with someone of that age. Note that there are fewer users above age 20 than below, consistent with the age distribution of our overall data.

4.3 Before You Follow (BYF)

Nearly as many Carrds contained lists of criteria for potential followers under a “Before You Follow” or “BYF” header as contained




Figure 9: Terms which were commonly present in “Before You Follow” (BYF) sections of Carrds. Bar color represents our categorization: content being tweeted (blue), Twitter platform features (magenta), or interaction with others (orange). Counts are non-exclusive; Carrds which used multiple terms are included in each total.

DNI sections: 2266 Carrds had BYFs, which was 43.1% of the total, in comparison to the 50.0% containing DNIs. As with DNIs, the majority of BYFs were in personal Carrds: 2122 (93.6%) were in personal Carrds and 144 (6.4%) appeared in artist Carrds.

In comparison to DNIs, BYFs were often framed as gentle warnings, useful information, or self-effacing apologies, and thus may serve a role similar to paratexts [21, 26] in terms of guiding a reader as to the intended way of reading the ‘text’ of a fan’s social media presence. In the remainder of this section, we highlight three categories of BYF terms: those discussing the content that a fan would tweet, those discussing Twitter’s platform features, and those describing interaction patterns with other fans. These are shown in Figure 9.

The most prevalent category of BYF terms related to content of a fan’s tweets. Fans often indicated use of trigger or content warnings, saying whether they warned for content at all, and if so, which specific warnings they used (e.g., “cw: blood”). Similar warnings appeared for actions like swearing, posting spoilers for media, not safe for work (NSFW) content, and the use of tone tags at the end of a tweet (e.g. “/hj” for “half-joking,” “/pos” for “positive”). These warnings usually indicated that this fan would post this type of content, but was aware that other community members might not want to see it or might want advanced notice about it when making a decision about whether to follow the account, implying an ecosystem in which fans reflected about the impact of their own presence on the community environment.

The second category of BYF terms contained references to Twitter’s features. In this category, fans discussed their expectations

around blocking and unfollowing (e.g. “block me to break our mutual following relationship”), when a direct message (DM) was welcome, and their tweeting and retweeting habits (e.g. “my timeline contains many retweets”). Though platform features might seem straightforward, as they behaved the same for any Twitter user, the frequency of these BYFs suggests that fans cared about the etiquette of when and how they were used. DM mentions typified this: DMs were often requested, either as invitations to meet new people or as a way to communicate potential issues with a fan’s actions, which ties into the following category.

This final category of BYF terms included those terms discussing interactions with other fans. As discussed, DMs sometimes appeared alongside these terms, such as requests to let the user know about issues (e.g. “DM to let me know if I say something problematic”). This category also included mentions of Twitter follower relationships in the form of “following back” and “softblocking” (quickly blocking and unblocking a user, forcing their account to unfollow yours), interaction (e.g. “I like interacting with followers”), discomfort (e.g. “Tell me if I make you uncomfortable”), and age concerns (e.g. “I’m a minor”). As age appeared in both DNIs and BYFs, it seemed to be a significant issue for these Twitter fan communities.

5 DISCUSSION: TOWARDS A MODEL OF COLLECTIVE PRIVACY

In this section, we articulate first steps toward a theory of collective privacy. We found significant explanatory gaps when attempting to use existing, individual-focused privacy theories to explain some of our results, such as ostensibly risky disclosure behaviors. We offer our theory as a potential explanation of the sophisticated, beneficial, and inherently collective nature of these behaviors, and as a starting point which we plan to deepen and elaborate upon in future work with other groups and in other privacy contexts (Section 6.1).

5.1 Identifying the Need for a Theory of Collective Privacy

Fans frequently disclosed information like age, gender, pronouns, sexuality, and disability status (Section 4.1). Such information is typically considered sensitive and may sometimes be re-identifiable. Adversaries could abuse this disclosure: queer Twitter users may be exposed to harm from both in-group members and other Twitter communities [45], and fans face particular privacy threats from outsiders [6, 15, 24]. Additionally, many fans in our sample disclosed being minors, with some listing ages as young as twelve. The HCI research community has identified many salient risks to teenagers on social media, such as in the work of Wisniewski (e.g., [55]) and boyd (e.g., [5]). Fans in our Twitter sample skewed younger than fans in other surveys of transformative fandom [28, 44] while remaining predominantly queer. This combination of potential vulnerabilities makes the community’s continued norm of identity disclosure particularly remarkable.

Given this, we ask: how can we understand these disclosures? What role do they play in sustaining fandom’s longevity and remarkable inclusivity? Existing theories of privacy, we argue, struggle to describe what is happening here. Individually-focused theories of

privacy—even those that situate the individual in context—address only parts of the privacy dynamics here, missing a sense of privacy on the community level.

Building on Waldman’s theory of privacy as trust [52], we propose a new lens that understands these behaviors by situating them communally. In taking first steps towards a theory of **collective privacy**, we aim to add to and complement existing theories. This theory is not a theory specific to fans—rather, the robustly communal nature of fandom and fannish behavior acts here as a case study that challenges the explanatory capacity of existing theories and evoked the need for this proposal. Future work on this theory will refine it by exploring its applicability to a variety of communities and contexts.

5.2 What Are the Gaps in Existing Privacy Theory?

We motivate a new theory by illustrating how our results suggest a gap in existing approaches. Though these approaches are powerful and remain relevant, we find that even a combination of existing theories left significant aspects of the privacy behaviors from our findings unexplained.

Inness describes privacy as offering “control over decisions about intimate information, intimate access, and intimate action,” such as allowing access to a diary entry or love letter [22]. Applying this to personal disclosures in our findings, we might interpret fans as sharing sensitive information to express and develop intimacy. Even assuming that fans consider these disclosures intended specifically for the Twitter fandom community and not the general Twitter community [11], it is unlikely that they have an intimate relationship with every fan on the platform; indeed, BYFs make clear that fans write Carrds expecting them to be read by strangers. Therefore, privacy as intimacy is an insufficient explanation for our observations.

We also attempted to apply theories of privacy focused on the individual in context, including contextual integrity and networked privacy. Contextual integrity describes privacy violations as breaking norms of information flow [36], and could be applied to reason about norms articulated in DNIs. However, this is complicated by DNIs’ unenforceability. A literal reading of DNIs would suggest norms that rely on voluntary self-sorting into categories with which people rarely self-identify (e.g., “racists”), suggesting that DNIs play a signalling role beyond control of information flow as described by contextual integrity. Networked privacy highlights the fact that the flow of personal information is dependent on the privacy behaviors of others, and individuals must account for that when navigating social media [32]. This is true for the users in our data, as Carrds are public websites, and Twitter followers can further share the tweets of people they follow. However, while this perspective is useful, we found ourselves unable to use it to explain the reasons *for* these disclosure behaviors or to account for using unenforceable DNIs as a privacy measure.

In the next section, we discuss how our notion of collective privacy builds on Waldman’s theory of privacy as trust [52] and extends it to argue that, beyond Waldman’s *societal* privacy benefits, the environment of trust is deeply linked to the health of a *community*, and vice versa.

5.3 A Theory of Collective Privacy: An Iterative Process Rooted in Disclosure, Shared Norms, and Trust

We propose that our results can be explained by a new theory of collective privacy, containing the following cycle:

- (1) Personal disclosure and boundary setting by group members operate in a decentralized fashion to construct and update community norms and values.
- (2) Community norms and values operate to create a communal environment of trust.
- (3) A trusting environment creates safety to make personal disclosures and expose vulnerabilities via boundary setting, returning to the top of the loop.

This theory builds on Waldman, which reframes privacy away from concepts of secret information or hiding the self and towards conditions that *enable beneficial sharing*, namely trust. Trust is identified as the key factor that enables us to share with others in confidence, and it is this confidentiality that closes our loop, enabling personal disclosure and its benefits. However, informed by our data, we add to this model by proposing a recurring loop of behaviors that updates norms and maintains an environment of trust:

Step 1: Disclosure and Boundary-Setting Construct Norms and Values. First, personal disclosure and boundary-setting contribute to group norms and values. Personal disclosures on fans’ Carrds communicate that the community is queer, young, and inclusive. Using DNIs, fans codify violations of inclusive values (e.g. racist or homophobic behavior) as grounds for exclusion from the community through blocking and collective non-engagement. Though DNIs are worded as though directed outwards at bigots, they may also be understood as a public endorsement of community values. This leads us to conclude that privacy must adapt to the needs of each community member, and that individuals willingly take on that work for the sake of meaningful inclusion.

Step 2: Community Norms and Values Construct Trust. Next, shared norms and values establish a network of trust: new and existing fans can see Carrds articulating community values and get a sense of what is expected of them if they join the community, as well as feel assured that their safety will be protected by collective enforcement of these norms, even in the absence of effective platform moderation. DNIs act not only as an articulation of values, but also as a promise to enforce them. This promise must be complemented by an ethic of mutual care in order to be welcoming. This is achieved by BYFs, which use a softer tone to voice the caring dimensions of the community’s values, expressing readiness to adapt to others’ needs and confidence that others will do the same. Fans signal this adaptability in their BYFs by proactively asking others to “let me know” how they can adapt their behavior to accommodate the needs of others. Drawing on Möllering’s definition of trust as the “favourable expectation regarding other people’s actions and intentions” [33], we note that both DNIs and BYFs act as promises of future behavior—grounded in community values and mutual care—on which other members of the community can base their trust.

Step 3: Trust Enables Disclosure and Boundary Setting. Finally, drawing on Waldman’s idea of trust as privacy, the trust that others will exhibit care by upholding community norms creates the safety needed for people to share themselves with others via vulnerable personal disclosure. Trust also enables boundary-setting, since articulating boundaries often involves declaring vulnerabilities. The presence of trust gives confidence that those vulnerabilities will be cared for and protected, rather than abused. These two elements are the “inputs” to the first stage of the loop, closing the cycle and continuing privacy as an ongoing process mediated through collective construction of norms, individual boundary management, care for others, vulnerable self disclosure, and trust.

5.4 Collective Privacy Helps Disclosure Benefits Outweigh Risks

Given the risks of disclosure mentioned above, why engage in this risky process rather stick with the neutral safety of obscurity? Fandom, we argue, may have leveraged the process described above to create a context—a community—in which the level of safety is high enough that the benefits of disclosure consistently outweigh the risks. The potential benefits of disclosure are huge: identity disclosure is a form of self-actualization and allows users to explore queerness [7, 37, 46], something which may not be safe or easy to do under their legal names [11, 27]. In recent years, queer online spaces have become even more important as physical spaces for queer community have dwindled [7], and this trend was only exacerbated by the ongoing isolation of the COVID-19 pandemic at the time of our data collection (mid-2022). Given the huge potential upsides, we shouldn’t be surprised to see some communities willing to take the risk in order to reap the benefits.

5.5 Identity Disclosure May Act as a Cost of Membership

In addition to our prior claims, we also hypothesize that disclosure may serve as a sort of “cost of membership”: as disclosure enables trust, a fan becomes or remains a trusted community member in part because of their disclosures, which show a willingness to take on risk for the sake of being in the community. We highlight two examples from our findings to illustrate these principles.

Age. Age was one of the most commonly-disclosed identifiers in our Carrds data, appearing in 81.5% of Carrds which disclosed any identity information. Age may be an important identifier because it enables fans to establish appropriate boundaries and care for their needs and the needs of others—fandom frequently includes content with mature themes (e.g., explicit sex [34] or gore and blood, common trigger warnings mentioned in BYFs), leaving openings for inappropriate interactions between people of different ages. Both younger and older fans shared their age as part of a collective process of mitigating that danger and protecting themselves/others from harm. Fans used hard boundaries (DNIs) to protect themselves, and built community norms around age-dissimilar interactions through BYFs, both of which contribute to a trustworthy community. This also enables fans to better build relationships by meeting people of similar age. Disclosing age may

help to enhance the quality of connections made, and thus helps people to better integrate into the community.

Pronouns. Pronouns likewise represented a significant proportion of identity disclosures, appearing in 86.9% of Carrds which disclosed any identity information. Pronouns indicate how to engage with someone, similar to the role of BYFs, which we found often signaled appropriate interaction patterns to others. Pronouns can’t be assumed or guessed and need to be explicitly shared in order to tell people how to use language to refer to someone, which is a major principle of respect in the queer community. Sharing pronouns constructs and reinforces a community norm that aligns fandom with this principle, and by association, with queer and inclusion-oriented communities in general.

6 FUTURE WORK

6.1 Further Explorations of Collective Privacy

Our research suggests that previous discussions of privacy are missing a key notion of privacy at the *community* level, and we have articulated first steps to a theory of collective privacy. As we explored collective privacy in response to how the Twitter fandom community was able to preserve the integrity of the community by dynamically negotiating an environment of trust, future work may study privacy behaviors within other communities that benefit from a trusting environment and investigate how those communities maintain that trust between members.

This may lead to new ways of conceptualizing privacy in HCI and other fields, such as law, where the goal of privacy protection necessitates an evolving understanding of privacy. Mulligan et al., for instance, call for technologists to embrace the instability of the definition of privacy, and propose that a system designed with different privacy models in mind will better support users [10]. By taking first steps to describe a collective theory of privacy, we also provoke scholars to consider what a technological system or law designed to protect collective privacy would look like.

6.2 Insight into the Twitter Fandom Community

Our findings illuminated a variety of online self-disclosure and boundary management behaviors within the Twitter fandom community, which future research can investigate further via qualitative studies. Such studies would also be positioned well to find out why fans might choose not to create Carrds, and how those fans might have different privacy needs, expectations, or attitudes from the users who made up our sample.

Additionally, Twitter has changed many of its policies and rebranded as “X” since we collected our data. Future work may investigate the state of transformative fandom on the platform after these changes and whether the community has migrated to new platforms [17]. Future research may need to use other methods of collecting data, however, as X has stopped making its API freely available as of the time of this publication. It is possible that the EU’s upcoming Digital Services Act [29] may require websites under its jurisdiction to make their APIs available to researchers, which would make it easier to conduct similar research, but its exact implications remain unclear.

7 CONCLUSION

In this paper, we begin the work of defining a **collective theory of privacy**, using the case study of transformative fandom on Twitter. Fans engaged in behaviors that could not be fully explained by existing privacy theories: despite the fan community being young and predominantly queer, both qualities which prior work has associated with a greater degree of risk in online spaces, fans disclosed personal identities and made themselves vulnerable by negotiating their social needs in the form of boundaries. We propose our collective privacy theory as an iterative loop: first, community members disclose and set boundaries to establish norms; second, community norms enable an environment of trust; third, this trusting environment makes members feel safe enough to disclose identities and be vulnerable in setting boundaries. This builds on Waldman's theory of "privacy as trust" and enables privacy researchers to understand privacy decision-making as a community process which is responsive to the needs and values of community members and essential for the existence of said community.

ACKNOWLEDGMENTS

We greatly thank Jenny Tang for her assistance with data collection, as well as Michael Ann DeVito for providing feedback on this manuscript. We also thank the Khoury Distinguished Fellowship, the Macalester Summer Research Collaboration Fund, and the National Science Foundation (grant 2334061) for supporting this research.

REFERENCES

- [1] Samantha Aburime. 2022. Hate narratives, conditioned language and networked harassment: A new breed of anti-shipper and anti-fan-antis. *Journal of Fandom Studies* 10, 2-3 (2022), 135–155.
- [2] 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 Trans. Comput.-Hum. Interact.* 25, 5 (2018).
- [3] Camille Bacon-Smith. 1992. *Enterprising women: Television fandom and the creation of popular myth*. University of Pennsylvania Press.
- [4] danah boyd. 2010. Privacy and Publicity in the Context of Big Data. WWW, Raleigh, North Carolina, April 29.
- [5] danah boyd. 2014. *It's complicated: The social lives of networked teens*. Yale University Press.
- [6] Kristina Busse and Karen Hellekson. 2012. Identity, Ethics, and Fan Privacy. *Fan Culture: Theory/Practice* 38 (2012).
- [7] Andre Cavalcante. 2020. Tumbling into queer utopias and vortexes: Experiences of LGBTQ social media users on Tumblr. In *LGBTQ Culture*. Routledge, 77–97.
- [8] Cleveland Clinic. n.d.. Neurodivergent. <https://my.clevelandclinic.org/health/symptoms/23154-neurodivergent>. Accessed on June 13, 2023.
- [9] DataReportal, We Are Social, and Hootsuite. 2021. Distribution of Twitter users worldwide as of April 2021, by age group. <https://www.statista.com/statistics/283119/age-distribution-of-global-twitter-users/>. Accessed September 7, 2023.
- [10] Colin Koopman Deirdre K. Mulligan and Nick Doty. 2016. Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy. *Philosophical Transactions of the Royal Society A* 374 (2016). <https://doi.org/10.1098/rsta.2016.0118>
- [11] 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), 1–23.
- [12] Jeff D'Onofrio. 2018. A Better, More Positive Tumblr. <https://staff.tumblr.com/post/180758987165/a-better-more-positive-tumblr>. Accessed January 12, 2023.
- [13] Brianna Dym and Casey Fiesler. 2018. Generations, migrations, and the future of fandom's private spaces. *Transformative Works and Cultures* 28 (2018).
- [14] Brianna Dym and Casey Fiesler. 2018. Vulnerable and online: Fandom's case for stronger privacy norms and tools. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 329–332.
- [15] Brianna Dym and Casey Fiesler. 2020. Ethical and Privacy Considerations for Research Using Online Fandom Data. *Transformative Works and Cultures* 33 (2020).
- [16] Facebook. [n. d.]. Names allowed on Facebook. <https://www.facebook.com/help/112146705538576>. Accessed on September 10, 2023.
- [17] Casey Fiesler and Brianna Dym. 2020. Moving across lands: Online platform migration in fandom communities. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–25.
- [18] Casey Fiesler and Nicholas Proferes. 2018. "Participant" perceptions of Twitter research ethics. *Social Media+ Society* 4, 1 (2018), 2056305118763366.
- [19] Centers for Disease Control and Prevention. 2022. Data & Statistics About ADHD. <https://www.cdc.gov/ncbddd/adhd/data.html>. Accessed September 7, 2023.
- [20] Centers for Disease Control and Prevention. 2022. Data & Statistics on Autism Spectrum Disorder. <https://www.cdc.gov/ncbddd/autism/data.html>. Accessed September 7, 2023.
- [21] Alexandra Herzog. 2012. 'But this is my story and this is how I wanted to write it': Author's Notes as a Fannish Claim to Power in Fan Fiction Writing. *Transformative Works and Cultures* 11 (2012), 15.
- [22] Julie C. Inness. 1992. *Privacy Intimacy, and Isolation*. Oxford University Press.
- [23] Henry Jenkins. 1992. *Textual poachers: Television fans and participatory culture*. Routledge.
- [24] Bethan Jones. 2016. 'I hate Beyoncé and I don't care who knows it': Towards an ethics of studying anti-fandom. *The Journal of Fandom Studies* 4, 3 (2016), 283–299.
- [25] Sidney M Jourard. 1966. Some psychological aspects of privacy. *Law & Contemp. Probs.* 31 (1966), 307.
- [26] Maria Lindgren Leavenworth. 2015. The paratext of fan fiction. *Narrative* 23, 1 (2015), 40–60.
- [27] Ada Lerner, Helen Yuxun He, Anna Kawakami, Silvia Catherine Zeamer, and Roberto Hoyle. 2020. Privacy and activism in the transgender community. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [28] Lulu. 2013. AO3 Census: Masterpost. <https://centrumlumina.tumblr.com/post/63208278796/ao3-census-masterpost>. Accessed January 9, 2023.
- [29] Tambiana Madiega. 2022. Digital Services Act. [https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/689357/EPRS_BRI\(2021\)689357_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/689357/EPRS_BRI(2021)689357_EN.pdf). Accessed on November 27, 2023.
- [30] Abigail Marsh and Ada Lerner. 2024. Privacy Norms of Transformative Fandom: A Case Study of an Activity-Defined Community. *Proceedings of the ACM on Human-Computer Interaction* (2024). To appear in CSCW.
- [31] Alice Marwick, Claire Fontaine, and danah boyd. 2017. "Nobody Sees It, Nobody Gets Mad": Social Media, Privacy, and Personal Responsibility Among Low-SES Youth. *Social Media + Society* 3, 2 (2017), 2056305117710455.
- [32] Alice E Marwick and danah boyd. 2014. Networked privacy: How teenagers negotiate context in social media. *New Media & Society* 16, 7 (2014), 1051–1067. <https://doi.org/10.1177/1461444814543995>
- [33] Guido Möllering. 2001. The nature of trust: From Georg Simmel to a theory of expectation, interpretation and suspension. *Sociology* 35, 2 (2001), 403–420.
- [34] Lucy Neville. 2018. *Girls Who like boys who like boys: Women and gay male pornography and erotica*. Springer.
- [35] Lucy Neville. 2018. 'The Ten's Big Enough for Everyone': online slash fiction as a site for activism and change. *Gender, Place & Culture* 25, 3 (2018), 384–398.
- [36] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.
- [37] Abigail Oakley. 2016. Disturbing hegemonic discourse: Nonbinary gender and sexual orientation labeling on Tumblr. *Social Media+ Society* 2, 3 (2016), 2056305116664217.
- [38] obsession_inc. 2009. Affirmational fandom vs. Transformational fandom. <https://web.archive.org/web/20151101202710/http://obsession-inc.dreamwidth.org/82589.html>. Accessed on September 11, 2023. Archived on November 1, 2015.
- [39] obsession_inc. n.d.. Transformational fandom. https://fanlore.org/wiki/Transformational_Fandom. Accessed on September 11, 2023.
- [40] Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. "Accessibility Came by Accident": Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities (*CHI '18*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174033>
- [41] Aja Romano. 2023. Puritanism took over online fandom — and then came for the rest of the internet. <https://www.vox.com/culture/23733213/fandom-purity-culture-what-is-proship-antiship-antifandom>. Accessed on September 12, 2023.
- [42] Allegra Rosenberg. 2021. "Writing To Cope": Anti-Shipping Rhetoric in Media Fandom. In *Electronic Literature Organization 2021: Platform (Post?) Pandemic*. <https://elmcip.net/node/16301>.
- [43] Elizabeth Rosenblatt and Rebecca Tushnet. 2015. Transformative works: Young women's voices on fandom and fair use. *Egirls, eCitizens* 385 (2015).
- [44] Lauren Rouse and Mel Stanfill. 2023. Over*Flow: Fan Demographics on Archive of Our Own. <https://www.flowjournal.org/2023/02/fan-demographics-on-aos/>. Accessed July 7, 2023.
- [45] Morgan Klaus Scheuerman, Stacy M. Branham, and Foad Hamidi. 2018. Safe Spaces and Safe Places: Unpacking Technology-Mediated Experiences of Safety and Harm with Transgender People. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW

- (2018).
- [46] Zach Schudson and Sari van Anders. 2019. 'You have to coin new things': sexual and gender identity discourses in asexual, queer, and/or trans young people's networked counterpublics. *Psychology & Sexuality* 10, 4 (2019), 354–368.
- [47] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013. The post that wasn't: exploring self-censorship on facebook. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 793–802.
- [48] Daniel J Solove. 2008. Understanding privacy. (2008).
- [49] Statista. 2022. Distribution of leading social media platform users in the United States as of August 2022, by age group. <https://www.statista.com/statistics/1337525/us-distribution-leading-social-media-platforms-by-age-group/>. Accessed September 7, 2023.
- [50] The Trevor Project. 2020. National Survey on LGBTQ Youth Mental Health 2020.
- [51] John W Tukey et al. 1977. *Exploratory data analysis*. Vol. 2. Reading, MA.
- [52] Ari Ezra Waldman. 2018. *Privacy as trust: Information privacy for an information age*. Cambridge University Press.
- [53] Samuel Warren and Louis Brandeis. 1989. The right to privacy. In *Killing the Messenger: 100 Years of Media Criticism*. Columbia University Press, 1–21.
- [54] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.
- [55] 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*. 3919–3930.
- [56] Jillian C. York and Dia Kayyali. 2014. Facebook's 'Real Name' Policy Can Cause Real-World Harm for the LGBTQ Community. <https://www.eff.org/deeplinks/2014/09/facebooks-real-name-policy-can-cause-real-world-harm-lgbtq-community>. Accessed on September 10, 2023..
- [57] Sarah Z. 2021. Fandom's Biggest Controversy: The Story of Proshippers vs Antis. <https://www.youtube.com/watch?v=5OclDcg7UJw>. Accessed on September 12, 2023..

Received 14 September 2023; revised 12 December 2023; accepted 18 January 2024