

Privacy Norms of Transformative Fandom: A Case Study of an Activity-Defined Community

ABBY MARSH, Macalester College, USA

ADA LERNER, Northeastern University, USA

Transformative media fandom is a remarkably coherent, long-lived, and diverse community united primarily by shared engagement in the varied activities of fandom. Its social norms are highly-developed and frequently debated, and have been studied by the CSCW and Media Studies communities in the past, but rarely using the tools and theories of privacy, despite fannish norms often bearing strongly on privacy. We use privacy scholarship and existing theories thereof to examine these norms and bring an additional perspective to understanding fandom communities. In this work, we analyze over 250,000 words of “meta” essays and comments on those essays, reflecting the views and debates of hundreds of fans on these privacy norms. Drawing on Solove’s theory of privacy as an aggregation of different ideas and on a variety of other academic theories of privacy, we analyze these norms as highly effective at protecting the integrity of fannish activities. We then articulate the value of studying these sorts of diverse “activity-defined” communities, arguing that such approaches grant us greater power to understand privacy experiences in ways that are specific, contextual, and intersectional yet still generalizable where possible.

CCS Concepts: • **Security and privacy** → **Social aspects of security and privacy**; • **Human-centered computing** → *Computer supported cooperative work*; *Collaborative content creation*; *Social media*.

Additional Key Words and Phrases: qualitative methods, privacy, norms, fandom, intersectionality

ACM Reference Format:

Abby Marsh and Ada Lerner. 2023. Privacy Norms of Transformative Fandom: A Case Study of an Activity-Defined Community. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 111 (April 2023), 29 pages. <https://doi.org/10.1145/3637388>

1 INTRODUCTION

Transformative media fandom (hereafter “fandom”) is a space understood from within and without by its *activities*: creating fanworks (fiction, art, videos, and more), talking about media properties (called canons) with other fans, and engaging in dozens of other social, community-shaping behaviors involved in being a fan of those canons [3, 16, 37, 42, 53, 54]. Despite the existence of thousands of subcommunities focused on specific canons and the diversity of activities that form the boundaries of fandom, it is united by a sense of “you know it when you see it.” The concrete and social evidence of its existence as a well-defined community includes entire fan-run platforms such as the Archive of Our Own (AO3), Dreamwidth, and Pillowfort, as well as a shared landscape of fiercely-held social norms [19, 26]. As privacy researchers, we have observed that while many fannish norms bear heavily on privacy, the word privacy is rarely used to describe them, and they have not often been studied from a privacy angle or using privacy’s tools and theories. In this work, we bring a privacy lens to examining fandom’s norms.

Authors’ addresses: Abby Marsh, amarsh1@macalester.edu, Macalester College, 1600 Grand Avenue, St. Paul, Minnesota, USA, 55105; Ada Lerner, ada@ccs.neu.edu, Northeastern University, 360 Huntington Ave, Boston, Massachusetts, USA, 02155.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

© 2023 Copyright held by the owner/author(s).

ACM 2573-0142/2023/4-ART111

<https://doi.org/10.1145/3637388>

Fandom has a number of properties that make it an invaluable subject for study at CSCW. These properties include the diversity of its community [43, 60], the transgressiveness of its activities (including explicit and queer content [53] and use of copyrighted materials [73]), and its long-lived nature, having existed for decades both offline and across a multitude of online platforms [25]. “Privacy” isn’t often the primary word that fans—and the scholars who study them—have used to define, describe, or debate the social norms and boundaries used within fannish communities. As privacy scholars, we use this work to argue that many of those community-defining values can be viewed as manifestations of privacy attitudes and responses to privacy needs. We do not claim that this is the only or correct way to analyze these norms. Rather, we believe that privacy offers a valuable additional perspective into fandom. Applying a privacy lens, we will demonstrate, develops new insights into how fans sustain community together. Simultaneously, it allows privacy scholars to develop our understanding of privacy in practice and theory through the observation of highly sophisticated norms and behaviors that have proven effective over decades.

Importantly, these fannish norms are far from monolithic, and neither are they isolated opinions of individuals. They are debated and refined constantly, including through the format of widely shared essays about community structure, termed “meta,” and comments on those essays. To examine these norms, we report on a qualitative analysis of meta essays posted by fans to fandom community spaces. We analyze over 250,000 words across over 60 documents, comprised of both meta essays and comment threads in response to those essays, representing the views of hundreds of fans. In consideration of the privacy risks of using public data, we collected meta essays which were already high-profile after being shared through community newsletters and aggregators, though we do not report on the specific identities of the documents and we paraphrase our quotations to protect individual writers from re-identification.

Fans discussed practices surrounding identity and re-identification of pseudonyms to real-world “wallet” names—a term used to connote names used on government-issued IDs. They explained the norms of participation that originated due to the features of different social media platforms. They extensively discussed the treatment of fanworks as public artifacts, from feedback and engagement to appropriate practices for sharing works to the end-of-life expectations for works that had been deleted by their original creators.

We derive two primary themes from these results. First, we find that as fans discussed privacy across these varied domains, they rarely used the language of academic privacy research, yet their norms and practices were varied and highly contextual. These norms were most effectively described by drawing on a wide variety of theories of privacy, which we consider an indication of sophisticated and practical perspectives on privacy as an aggregate of multiple related ideas, as theorized by Daniel Solove [67]. Our analyses shed new light on how fandom’s norms have enabled it to persist and strengthen over decades. Second, we define the notion of an *activity-defined* community, and argue for the value to privacy scholars of focusing research on such communities. Fandom perfectly illustrates such communities in the way that it incorporates many different activities under the umbrella of “fandom”, and we draw on our results, as well as building on past work on other diverse, at-risk communities, to present evidence for the value of such studies as opportunities to capture contextual and intersectional privacy experiences strengthen privacy scholarship as a whole.

In summary, we make the following contributions:

- We perform an analysis of discourse in the transformative media fandom community, identifying sophisticated privacy-relevant norms and practices.
- We argue for the value to both CSCW and privacy studies of applying a privacy lens to understanding fannish norms which typically aren’t described using the term “privacy”.

- We analyze these practices through the lens of academic privacy theories and articulate how fans negotiate tensions of value and risk.
- Using fandom as a case study, we argue for the scholarly value of analyzing privacy at the level of diverse, activity-defined communities.

2 RELATED WORKS

2.1 Theories of Privacy

Scholars across disciplines have long debated definitions and theories of privacy. Some of the theories on which we draw, such as Nissenbaum’s contextual integrity [55] or Waldman’s privacy as trust [76], were proposed by individual scholars. Many others are built on common understandings of privacy as seclusion, separation in a private space, control over information, in relation to intimacy, and in terms of a variety of positive benefits enabled by privacy, which have been admirably summarized [67, 76]. We discuss these in relation to our results in detail in Section 5.1.

Most powerfully for our analyses, Daniel Solove proposed that having many theories of privacy suggests that privacy is not a single idea, but an aggregation of related ones [67]. In the absence of a unifying theory, Solove argues for understanding privacy in relation to privacy threats and the activities those threats endanger. Under this model, the purpose of privacy is to protect the integrity of activities. We apply this idea by examining fannish privacy as an aggregation of norms that collectively protect safety, trust, and intimate spaces, enabling fandom community to exist.

2.2 Privacy Practice as Visibility

This study in part argues that visibility of online content, including practices to control and reactions to that visibility, is a privacy-relevant topic. Prior work in CSCW provides a baseline for this connection: many privacy research papers discuss deletion as a strategy for preserving privacy. For instance, the ephemerality of Snapchat is cited as one aspect of user agency around privacy management on the platform [80]. Further, many specific user populations whose privacy and security practices have been studied have used deletion as one of those practices. This can be observed in survivors of intimate partner violence [30], economically disadvantaged users [74], and people who have broken up with romantic partners [58].

In her body of work understanding self-presentation on social media, Michael Ann Devito and her coauthors explicitly identify ephemerality (defined as “regularity of content erasure”) and editability (defined as “includ[ing] deletion”) as privacy-relevant, in order to control visibility of content when using social media to represent oneself [12]. This line of research has subsequently discussed how these self-presentation techniques are relevant to the privacy behaviors of LGBTQ+ individuals on social media [13]. Transgender people in particular have been studied within the privacy community, establishing that online visibility as a transgender person can constitute a privacy risk [41]. Navigating that visibility is a necessary and difficult risk for many transgender users [11].

We argue that visibility of and engagement with a person’s online presence is usefully construed as privacy behavior, with insights for both privacy and fan scholarship emerging from applying the privacy lens to these practices. This is elaborated upon in our results, a significant portion of which center around a set of fannish behaviors—including engaging with (Section 4.3), sharing or linking to (Section 4.4), and deleting fanworks (Section 4.5)—that are critical for exploring the sophisticated privacy norms of these communities.

2.3 Privacy in Fandom

While a substantial body of work (discussed in Section 3.2) has touched on privacy in analyzing the dangers that outsiders (e.g., journalists and researchers) bring to fandom, remarkably little has analyzed fandom's norms and practices themselves through an explicit lens of privacy.

As part of a larger body of work, Dym and Fiesler highlighted that the many women and LGBTQIA individuals in fandom expressed concern about doxxing, the connection of their digital identities to their wallet names [17]. Of particular concern was outing, as many of their participants were not out as LGBTQIA in their offline lives and used fandom as a space to explore and build their queer identity. In their study, fans discussed primarily self-censorship as a strategy to prevent this risk, and Dym and Fiesler suggested that granular privacy controls would help address their concerns. Our work, in addition to supporting these findings, documents within-group privacy norms to add depth to our understanding of how privacy-sensitive populations regulate themselves in order to manage risk.

Dym and Fiesler further identified fandom privacy norms as changing over time, with many fans now more willing to connect their wallet names to pseudonyms, especially in response to an increase in activist efforts within fandom [16]. They highlight a key reason that we believe fandom is a fertile ground for privacy research: "Online fandom is a space that is at once both very personal and separate from all things personal." The complexity of fan identity, predating the internet but growing and adapting to over 30 years of digital communities, provides a space wherein users have multiple nuanced understandings of privacy, even if they do not use the word "privacy" to describe these norms.

One of the key contributions of this paper, then, is to demonstrate the value of applying a privacy lens to studying fannish norms. The value is two-way, in that not only can we better understand and support fannish activity by applying the scholarly tools of privacy to its study, but also because fans' sophisticated approaches to privacy (even if they do not use the word!) can teach us much about privacy itself as it manifests in diverse communities.

2.4 Fandom's Activities as a Gift Economy

One such specific norm which defines fandom community is the idea of the gift economy. Fandom's activities have been described—by fans and by scholars of fandom—as a gift economy, wherein the giving, receiving, and reciprocation of gifts builds and maintains the community [35, 61]. Fanwork of all stripes, including fiction, art, songs, videos, graphics, meta essays, and even comments are the gifts being exchanged, under this paradigm. They are given without expectation of compensation, except perhaps that of community belonging.

Tisha Turk elaborates on this by observing that because fan works are copyable, digital artifacts which are posted publicly, gifting in this culture is one-to-many, and she theorizes the fan community itself as not just created by gifting but as the recipient of gifts [72]. Turk also observes that fanworks are not the only gift that fans give to each other and the community, and names commenting, writing meta, aggregating links, organizing and moderating communities and events, and much more as "effort gifts" ([57]) which contribute to and strengthen the community [72]. This sense that fanworks have been given to (and thus belong to) the community as a whole likely contributes to fans' norms and feelings around sharing and deleting fanworks, and thus has likely contributed significantly to the norms we analyze in this work.

2.5 Ethical Research Methods

To ground our specific ethical research choices (described in Section 3.2), we draw upon long traditions within both fan studies and the field of human-computer interaction.

The Association of Internet Research (AOIR) defines three main points of tension when conducting research with online subjects: whether it is human participant research, whether the data is public or private, and whether any particular piece of data collected is text or is a person [45]. By asking these questions, the AOIR suggests, a researcher must develop an individualized ethical practice that is cognizant of the specific risks of their work. For this, methods such as paraphrasing data might be appropriate [44], as the risks outweigh the benefits of direct quotation. We turned to the realm of fan studies in order to conceptualize those risks.

Foundational fan-studies texts base their approach to the field in ethnography [3, 37], and more recent work builds on those foundations, exploring and clarifying boundaries in an era of increasingly-online fandom. In suggesting a “goodwill approach” to studying fandom, Kelley states, “A goodwill approach to fan scholarship should take the time to consider and negotiate fans’ privacy concerns and make research findings fully available to fans [and] to analyze, engage with, and question fan texts just as fans do with popular culture texts” [39], but notes that asking individual permission of each fan can be an unrealistic burden, a view repeated in Natasha Whiteman’s thorough investigation of online research ethics [79]. Whiteman takes the approach that ethical considerations must be negotiated within individual contexts (“localizations”), rather than relying solely on generalized ethical precepts [79].

Particularly in the context of research using public data [8, 18], fannish privacy considerations have focused on the many and varied risks faced as a result of outsiders interacting with the community. Dangers of outside influence include within-fandom conflict, via something with wide publication connecting disparate parts of fandom and resulting in harassment [18], as well as the risk of academic or journalistic influence surfacing contentious viewpoints that otherwise would not have spread past the community to which they were originally posted, something also documented in literature about anti-fandom [38]. Social media users are also frequently surprised that what they post publicly might be collected for research [6, 27].

A recent example of similar research to what we present here takes a highly overlapping approach to our own: Dym et al. [19] examine public posts on the Pillowfort website, choosing to anonymize and paraphrase the specific posts and contributions of individual commenters, but not to reach out to those commenters for consent for their words to be paraphrased and included in the research. These methods make sense in the context of the subject matter (public discussion of platform policy) and the number of individuals (likely hundreds of commenters over a dozen posts), and the choice to anonymize and paraphrase specific contributions demonstrates that Dym et al. carefully weighed their ethical responsibilities and the privacy risks to fans in their reporting of results. We contrast this with the assumption often made in humanities-oriented fan studies work that researchers are studying, including, and linking to specific fictional fanworks [8], such as videos, fanfiction, or art, rather than discourse and community conversation within fan spaces. While re-identification of pseudonymous fans remains an important concern in our context of studying community norms and structures, important ethical steps in one context (e.g. asking author permission before excerpting identifiable quotes from a named piece of fanfiction) might not apply to all versions of working with fan or online communities.

3 METHODOLOGY

We investigate fandom community norms about privacy through a qualitative analysis of online discourse. We used meta essays, produced by fans about fandom, which discussed privacy either directly or through related concepts. We qualitatively analyzed these texts, including comments left by other users. Our research did not recruit participants, and was determined by the Institutional Review Board to not be human subjects research, but we still considered the risks inherent to studying digital artifacts and present the ethical measures taken to minimize them.

3.1 Meta Essays

Meta essays (“metas”) are defined as “a fan-authored piece of non-fiction writing discussing any aspect of fandom, fanworks, or the source text”¹. They represent a way for fans to talk about their experiences of fandom, and in the particular case of metas about fandom communities as comprised our sample, they are often explicit in both speaking to fans from multiple different fandoms and in seeking an audience of fans interested in a generalized discussion of fandom. We began our collection of meta by selecting essays that had been aggregated through meta-focused communities and mailing lists within fandom. All of the metas that we collected were publicly viewable by anyone on the internet, and we took additional steps to ensure ethical engagement with the source material and preservation of anonymity for its authors (see Section 3.2). As the authors read these metas, we evaluated the contents of each meta and any associated comment threads for relevance to privacy norms and online boundary-setting by asking the following questions:

- (1) Did the meta contain an explicit use of the word privacy?
- (2) Did the meta discuss privacy in proxy language, such as visibility (of fanworks, persons, or similar); sharing, deletion, and archiving of fanworks; etc.?
- (3) Did the meta discuss community boundaries, such as statements of community norms and what factors influence them, the sense of belonging, and platform use?

If a meta linked to another meta, we visited that link and evaluated its contents as well. We saved each meta as a document. Due to the design of journal sites like Dreamwidth and LiveJournal, when comments were split into multiple pages, we saved multiple documents in order to capture all of the pages of comments. We expanded each comment subthread as well so that automatically-collapsed comments were not missed in our analysis.

In total, we retrieved 62 relevant documents comprising 264,000 words and 1609 comments, published across 8 websites, dating between 2006 and 2022. Of these, 40 were journal posts, 17 were Tumblr posts, and 5 came from other sources. They represent fans who participate in multiple fandoms, including those for books, television, movies, and comics, among others, and many conversations involved fans explicitly discussing experiences which crossed different fandom communities. We selected metas for relevance to our research questions and captured discussion where it existed. We also made an effort to capture attitudes about privacy norms and boundary-setting over the course of many years, in order to examine (a) whether and how these norms have presented differently over time, especially as different platforms were the primary hub of fan activity at different times, and (b) many different events which have prompted fans to discuss and reconsider their norms.

A topic organization emerged from reading the collected metas. As we read and evaluated metas for inclusion, we took note of the main topics of each meta essay’s primary focus using memos. After our full data set was collected, we revisited these short-form descriptions and categorized them into a final set of topics: anonymity and identity; archiving fanworks; deleting fanworks; platform design; platform migration; sharing and linking norms.

First, one author and three research assistants coded the documents using two deductive code categories: privacy norms and platform features (defined below). We applied these codes to each document individually and met daily to review each other’s codes and discuss concerns until the process was complete. The second author met with the coding team weekly in order to give feedback on the coding process and the findings as they developed. All coders read and coded all documents. As our codes are not the primary product of our research and all documents were read and coded by everyone involved in the process, we do not calculate inter-rater reliability [4, 48].

¹<https://fanlore.org/wiki/Meta>

The two coding categories allowed us to focus on areas of primary interest: how fans discussed privacy both as a socially-mediated boundary and as a technically-mediated one. We defined them as follows:

- **privacy norms:** social or community norms of behavior, which influenced a fan's actions
- **platform features:** the technical capabilities of a given platform, which shaped fan behaviors.

Second, inductive codes based on the document categories previously mentioned were then applied by one author to the deductively-coded data, and both authors continued to meet weekly to discuss themes as they developed, with the second author contributing oversight to the inductive coding process.

3.2 Ethics

Our analysis relies on publicly-available data which was not intended for research, and which additionally comprises many people spanning nearly two decades. We carefully considered the ethical implications of collecting this data: people posting online do not expect that their posts will be used in this context [6, 27], and fans may have additional privacy concerns beyond that of the average online commenter [8, 18], which was discussed by fans within our collected data. We have made decisions in line with best practices in online research in order to shield the fans in our data set from risk and to respect their privacy. As we understand that our data represents perspectives across nearly two decades and many different contexts, we are particularly sensitive to the risk of re-identification for the fans involved. As researchers, our interest is in *how* conversations around privacy norms develop, not in attributing immutable opinions to specific individuals who may no longer hold them, and this is strongly echoed by fans talking about these same issues. We believe that presenting a list of sources would run counter to community trust, and so choose not to present one here. We took the additional following steps in order to maintain ethical standards for the data in our results:

We have elected not to report the features of individual documents which could allow those posts to be re-identified, such as the publication platform, date, or the number of words or comments.

We collected only posts which were viewable on the public internet: that is, the posts did not require an account or membership within a specific community in order to access. We additionally collected only posts which had been aggregated by within-fandom mailing lists or interest groups, or linked from such posts, indicating a high level of visibility; we did not search individual fans' social media profiles to find posts.

When reporting results in this paper, we paraphrase all quotes. Instead of providing direct quotes and attribution, which might be identifiable as specific individuals, we have rewritten the quotes so as to preserve the sentiment, in line with ethical best practices [44].

3.3 Position Statement

We acknowledge that as researchers conducting thematic analysis of qualitative data, we bring the context of our individual identities, experiences, and values to the work [7]. We describe our identities in order to provide that context to the reader, given that they relate to the topics discussed in this study. Both of the authors identify as queer. Additionally, one author has years of experience as an active participant in fandom, which has been noted in prior work as something to which fans are generally receptive when fandom is a topic of study [18].

4 RESULTS

4.1 Fannish Identity

In order to understand how fans construct digital privacy boundaries, it is first necessary to see “fan” as existing as its own identity, with many possible specific definitions but broadly understood by the fannish communities discussed in this paper. As fans put it:

“Fandom creates a home wherein fellow fans aren’t quite strangers—we are united by our fannishness.”

The multiplicity of fan identity is too large to discuss in detail here, but should be taken as a baseline which fans reference implicitly and explicitly when discussing boundaries. Someone who is a fan can be trusted to interpret another fan’s actions and interests within the norms of fandom, whereas someone who is not a fan cannot:

“People not in fandom don’t understand what fandom means to us. I feel judged when trying to explain fandom to outsiders.”

The perceived threat of judgment, leading to possible repercussions ranging from social embarrassment to job loss, led to many fans expressing that anonymity or pseudonymity was an essential way to interact with fandom, relative to attaching one’s wallet (legal) name to their fannish presence.

“Outing someone as a fan could cause backlash for students, for people with jobs, for all sorts of reasons.”

And fan communities were cited as being a place where many fans could safely explore queer identity in a trustworthy environment. Some first came out in fandom. And some created work which focused on queer characters and stories, as opposed to more overwhelmingly cisgender and heteronormative narratives in mainstream professional media.

“I know a lot of people who first came out as queer under a fandom pseudonym.”
“Fanfiction can center queer characters in whatever way the creator wants. But try explaining that to non-fans—they just see slash as weird.”

Fans almost universally expressed that non-fans would often see fan activities as weird, strange, or unusual. Because of this, they typically felt that either a strict separation of fandom and “real life” was necessary to avoid repercussions from being a fan, or they allowed for only very selective sharing of one’s fannish identity with trusted friends and family. They also recognized that different fans would have different relationships to this danger. This led to commentary on what some fans described as an increasing trend away from anonymity:

“I see a lot of fans who connect their wallet names and pseudonyms. It’s really distressing, because that’s the opposite of how I had to act when I was new to online fandom.”
“I hope that it becomes less devastating for fans to have their real name and author name connected as fandom grows more accepted.”

4.2 Platform Design

Fans were highly aware of the relationship between the technical capabilities of social media platforms and the nature of community norms. We contrast the two most-commonly named platforms, Tumblr and journaling sites like Dreamwidth or LiveJournal, in this section. We chose to highlight features within sites, rather than grouping similar technical features across different sites, because the site-specific cultures influenced the use of these features in ways which could not be divorced from that context.

4.2.1 Tumblr. The primary platform under discussion in the meta essays we collected was Tumblr. Many of the metas were direct responses to Tumblr's decision, in 2018, to ban adult content [14]. This event prompted fans to look for new platforms to serve as gathering spaces, and in the years since it took place, has served as a cautionary tale within fandom about over-reliance on a single social platform, much as LiveJournal's Strikethrough and Boldthrough incidents had done before it [23]. The result is a rich cataloging of Tumblr's features and the privacy implications thereof.

A defining, unique trait of Tumblr for many of its users, including fans, is the way that hashtags work on the platform. Tags serve as an organizational tool within one's own blog, similar to their function on journal sites, and also provide a way to find content with that tag across the site, similar to Twitter and Instagram. However, Tumblr tags have an additional use case: they serve as a "whisperspace" for users to make commentary on a post with reduced visibility.

Tumblr allows up to 30 tags on a post, each with a 140 character limit. These are frequently co-opted for human-readable commentary. Critically, tags are not copied when another user reblogs the post. This impermanence with respect to reblogs has been adopted as a version of granular privacy controls, given that Tumblr has not had the ability to change privacy and sharing settings on an individual post basis for the majority of its history. Tumblr users describe this across multiple metas in our dataset:

"Tumblr tags are a space for digressions, a "whisperspace" outside of the post body. There's no limit on the number of unique tags an account can use, unlike Dreamwidth, so you can't run out."

"I never used tags as a whisperspace before Tumblr. I miss it on other sites! It was a clear break between what was meant to be social and what was meant for public consumption."

"Tumblr tags as a private discussion space were a bad option, but they were the only one on the platform."

A bug with the iOS version of the Tumblr app in Spring 2022 temporarily resulted in a user's tags being duplicated as part of a reblog, propagating them beyond that initial context [70]. Many users of the platform reacted negatively to this bug, as can be seen in the responses to the official staff post—"Thank god it's a bug" was a common sentiment—emphasizing the widespread nature of this use case. The whisperspace was and is a point of pride for Tumblr users:

"Purposefully misusing the site features as a whisperspace was like a little act of resistance."

The lack of granular privacy controls was also noted for its negative outcomes, such as the potential for a post to spread beyond its intended audience, disagreement and conflict that might arise from that, and the lack of features aimed at enabling more intimate conversation. These concerns were raised across multiple metas.

"Reblogging a post to argue brought it to your followers, who agreed with you. You got better engagement by being brutal, rather than having a sincere discussion."

"There are no privacy controls so you can't prevent a post from being reblogged and seen by the entire website."

"If you wanted to say something in response to a Tumblr post, you had to reblog. There weren't direct comments, so it wasn't like talking with the original poster."

"Anyone on Tumblr for any length of time has reblogged a post that you weren't meant to see. That's just how the site works."

Reblogs sat at the heart of this: reblogging allows users to copy a post to their blog, with or without additional commentary, and is a core feature of the site. This concerned many fans, because

each reblog is its own post, wholly controlled by the user who reblogged it, and the original poster has no ability to remove content from blogs that have reblogged it. Deleting the post from one's own blog only deletes that copy, not any of the other copies that might exist on Tumblr, and the original poster's name remains attached to the reblogs. Some fans highlighted that deleting one's account was the only way to fully sever connection to a viral post. And though Tumblr allows users to de-index their blog from search engines, de-indexing doesn't prevent copies of their posts which were reblogged by a user with different privacy settings from being indexed by those search engines.

"Once your post has been reblogged, it's out of your control. You can't go back and private it and you can't delete it from the site, only delete your own copy."

"Tumblr posts stick around even if your account is gone, so your posts might circulate with your name attached even if that blog is defunct."

"Most people don't turn off Google indexing on Tumblr, so content they reblog from you will end up searchable even if you've changed your settings!"

De-indexing as a privacy measure came with downsides as well, as with one fan who was frustrated by the difficulty of searching Tumblr for fannish content:

"I'm in rare fandoms, and that means I can't find fellow fans easily because they've all turned off search indexing. Boo!"

Tumblr has historically had few privacy options, and users have been limited to a binary of extremes, with the potential for reblog fueled virality on one hand and the quiet, improvised, low-impact commentary of the tag whisperspace on the other. This may be changing—Tumblr added the ability to disable reblogs in Summer 2022 [71]—but for the period in which most of our dataset was created, no such controls existed.

"Reblogs take your content and put it into my space, under my control."

"Tumblr discourages conversation: you can shout or you can whisper, but nothing in-between."

"Posting to Tumblr was like shouting into a void. I had no idea whether anyone liked what I shared."

Another feature of platforms that related to fannish norms was the ease of posting different types of content. As compared to its journal site predecessors, Tumblr was better at hosting multimedia content, and thus served an important role for fannish community.

"Tumblr has a lot of flaws, but it made fandom way more accessible to non-writers than the journal sites, which privileged text."

This demonstrates one reason for fans' continued use of a platform which poses so many known obstacles to privacy and community-building: it is accessible and supports fan activities, and so fans are there. In contrast to the journal sites which will be discussed next, Tumblr is full of active communities and fans for nearly any media fandom, and so fans continue to use the platform despite the noted privacy concerns.

4.2.2 Journal Sites. The present (2023) state of the social web is such that journal sites, which include LiveJournal and Dreamwidth, are not as widely used as social media like Tumblr, discussed above, or Twitter. Our collected meta spans from about the time that sites like Tumblr and Twitter were being created to the present day, and much of what fans say about journal sites is in the context of those sites seeing a decline in use or having reached a steady state of niche use. This results in a reflective tone to many of the ways that fans examined their design: these sites are acknowledged as unlikely to resume the prominence they had in fandom during the early and mid aughts. However, the influence of journal sites was reflected in community norms produced under

the influence of journal site features and carried with the community as it has shifted toward more popular current platforms. Indeed, the specific strengths of journal platforms resulted in them continuing to be hubs for meta and the discussion of norms throughout the studied period.

Much of this conversation contrasts what journal sites provide as compared to Tumblr. For most fans, the strength of journal sites is that they offer granular privacy controls at the comment, post, journal, and community levels. This allows for a wide variety of interaction methods to coexist, from locked individual blogs to communities where fans broadcast fanworks—their own and those of others—to thousands of potential viewers.

“LiveJournal lets you lock or delete posts, delete comments, turn off anonymous commenting, track IP addresses, and more. And LJ communities offer a space for chatting with people who have similar interests, without feeling totally exposed to the public.”

The high degree of control over individual posts, including the ability to moderate one’s own comments section, was cited as leading to threaded conversations in comments as a normative practice on journal sites.

“In moving to Dreamwidth, it was helpful to learn that people expect chatty, digressing comments. I like that this site is about conversation!”

In contrast to Tumblr and Twitter, whose design models more easily led to virality as discussed above, these comment controls were viewed as making it relatively easy for a fan to recover from conversations gone awry by locking down the comment thread, journal, or post, as appropriate.

“If a conversation goes awry on Dreamwidth, you can just freeze or delete the thread. The only option on Tumblr is to delete your post, and that doesn’t stop the reblog copies already out there.”

“Because there’s no freezing, sometimes the only thing you can do to stop a bad conversation on Tumblr or Twitter is delete your account.”

However, that same low degree of discoverability was a hindrance when it came to building community, especially in an era when journal sites are not one of the primary social platforms for fans or online socialization more broadly. For all the downsides of the possibility of virality on Tumblr, the ease of finding relevant content for one’s fandom meant that joining the network of fans around it was simpler than on a journal site, potentially allowing it to better enable the core activities of fandom.

“There’s no automatic feed for seeing content on Dreamwidth, because you can’t track tags like on Tumblr. You have to use communities and hope people find them and post to them. It’s frustrating for niche content, and I miss parts of the Tumblr way!”

Whether in 2018, when users searched for an alternative to Tumblr in the wake of the ban on adult content, or again as we write, in 2022, as Twitter users consider alternatives, journal sites offer an old solution to a newer issue:

“Journal communities are more curated than Tumblr or Twitter tags. You don’t see a bunch of posts hating on the thing you’re looking up. I prefer communities, all said.”

4.3 The Inherent Tension of Online Engagement

Fans shaped their community spaces and social norms according to the values of fandom and in response to platform features in different degrees at different times. The comments reported in this section do not directly bear on privacy, but we include them because they form a key infrastructure

in which fannish privacy norms exist. Privacy is not the only framework for analyzing these comments, and we build on other frameworks (such as gift economy in fandom) to understand them, arguing that a combination of these perspectives are useful ways to understand fans' emotional reactions to engagement. In particular, approaching these comments from a privacy angle allows us to understand them through the idea of visibility, which we find to be a key element underlying the norms discussed here and which plays a role in the wider privacy literature (Section 2.2).

Another core value was the idea of fandom as a gift economy, which shows up explicitly in multiple meta essays and their comments. Under the gift economy, fanworks, fannish efforts, and feedback on those were modeled as reciprocal gifts, with both one-to-one, one-to-many, and many-to-many variations. However, the tenets of this gift economy—specific actions surrounding necessary and appropriate forms of engagement—were a source of constant debate within our data set. All fans could agree that such an engagement etiquette existed, but not on one set of behaviors which fulfilled it.

Ultimately, violations of the expectations of engagement were violations of the sense of belonging to fannish community, whether that was a post spreading beyond its intended audience (e.g. through reblogs, discussed above) or a post not receiving a certain level of feedback. Both of these represented a transgression of expected visibility levels, resulting in fanworks being either more or less private than the creator wanted:

“Reader silence comes from some of us readers being lazy! I don’t expect responses as an author, because I know this, but it’s tough to post a story to crickets.”
“I think it’s reasonable for an author to ask if anyone is still reading their series. If there’s no response, maybe they want to put their writing effort elsewhere.”

Fans sometimes expressed different attitudes towards the necessity of particular actions, such as readers leaving comments on fanwork, as part of defining a fan community. Nevertheless, they broadly agreed that fan communities needed engagement of some form. For a gift economy to exist, there must be gifts and there must be reciprocity.

“If there’s no communication, there’s no community. Without interaction or feedback, people will eventually leave.”
“It’s not a lack of feedback for me—it’s a lack of creators. But I could see how it would be different for different people.”

In light of the necessity of engagement, fans speculated about obstacles. Much of our data set was written when AO3 was only a few years old, and concern about commenting rates on the new platform were common:

“Allow anonymous comments. Why are you shutting out potential readers?”
“Some custom journal layouts are really inaccessible. Where’s the reply button? Make it easy for me.”
“AO3 should position itself more like a journal community, not like a library of fanwork, if it’s going to get readers to comment more.”

More recent metas shifted the focus of this discussion to social platforms like Tumblr and Twitter, which have grown in popularity as fandom hubs since the start of the 2010s. These platforms offered participation options which could appeal to fans who preferred to read posts without actively contributing (lurking). However, while these options may have enabled participation at a visibility level comfortable for the consumer, they were potentially in tension with the visibility level desired by the creator.

“Twitter and Tumblr have options for passive users to participate—they can reblog or retweet other people’s content and might get their own following as an aggregator account.”

“Tumblr has likes, and a lot of people use those instead of commenting. It’s frustrating to post to the main tag and hope your content gets picked up, because mostly it doesn’t.”

Fundamentally, fans saw engagement, and particularly feedback to fanwork, as a primary condition of being active in fandom. This was an essential part of the definition of fan communities, a result reflective of but also sometimes in tension with past work that described reciprocity in fandom as not simply a one-to-one fanwork-to-comment exchange [72].

“The main reward for writing fic, after your own pleasure, is community feedback on your work.”

“We’re a gift economy, and the only payment we get is attention for our work.”

4.4 Sharing and Linking

A complex and varied set of norms influence the ways that fans share and link to content within fannish spaces. Even when the audience is presumed to be primarily or entirely fellow fans, fans expressed differing views of when the use of direct links was acceptable.

4.4.1 Public is Public. A common rule-of-thumb that fans used for when content is acceptable to share was that if a post is not private, especially on a platform which allows posts to individually be made private, then it could be linked.

“It’s fine to link something public. It’s public.”

“Something publicly posted within fandom is fair to share within fandom.”

“Posting something publicly on a site with good privacy controls means it’s okay to link.”

This view can be seen as aligning with notions of privacy as control, which put the responsibility on the poster to use settings to prevent linking. Phrases such as “on a site with good privacy controls” also suggest a social element to these rules. The community’s shared understanding of certain sites as having “good” or “bad” privacy controls informs inferences about the preferences of the original poster. The use of controls as a social signal not to share can be just as important as the technical effect of the control. Fans reflected this in discussing controls like friendslocking a journal entry so that only an approved subset of users could see it, emphasizing that the choice was to be respected.

“When someone posts something to their journal under a friendslock and asks people not to talk about it, you shouldn’t go share it with people who aren’t in the friendslock.”

“We know friendslock is the way to prevent our posts from being shared in someone else’s journal.”

This echoes theories of privacy which position it as sharing under conditions of trust [76]. Using a privacy control like friendslock sets an explicit social cue that the viewer is being trusted not to violate the poster’s privacy expectations, whereas leaving a post public can be seen as an implicit social cue that sharing that content—though perhaps only within fannish social circles—is not a privacy violation.

This philosophy runs into quandaries on platforms where the social semantics of privacy features are unclear. As previously mentioned, Tumblr only recently introduced controls to disable reblogs of one’s posts, in many cases years after the discussion in the meta posts we captured about sharing

and linking. Prior to that, users sometimes employed tags (or a lack thereof) to institute a socially enforced post-by-post public/private divide.

"Sometimes no tags on a post is a sign not to reblog it."

"It's confusing to me, but some people tag Tumblr posts with 'don't reblog.' People don't always listen."

The suitability of this came under question: why would any fan post publicly about something when they wished it wouldn't be shared? These posts were called out as a potential vulnerability, as the tags might make it easier for someone to harass the user. When commenting about a broadcast-oriented platform like Tumblr, many journal-site users expressed a lack of trust which seemed inherent to their conceptions of the platform.

"Tumblr has tag griefers who will search for personal, private, don't reblog tags and share those posts. I wouldn't trust fandom strangers to listen to tags."

This distinguishes between privacy threats emerging from the behavior of other fans, who are more likely to abide by socially-enforced privacy mechanisms, as compared to threats from outside the community. Additionally, it puts a spotlight on the use of apparently unenforceable privacy mechanisms, which we discuss in detail in Section 5.1.4.

4.4.2 Public is Complicated. However, publicity was not universally seen as permission, implicit or otherwise, for fans to share and link content. Many fans took the stance that a public post should *not* be taken as linkable-by-default. Fans argued the need to account for the emotional content, context, and age of posts when deciding whether to share:

"Sharing around a decade-old public post—especially a really emotional one—without asking first seems off. The internet was different back then, people had different expectations for how their ranty journal posts would spread."

Fans supporting these views can be described as emphasizing theories of privacy that use the intimacy of content as a signal for its sensitivity and thus for the appropriateness of sharing it. This may be particularly salient when viewing the community intersectionally, since marginalization may increase the danger of exposing intimate information, and multiple marginalization especially may create entirely new threats.

Some viewed older content as effectively being private, regardless of privacy settings, and that this social norm existed even for public posts back in the aughts. In contrast, past work examining feelings about old content on social media found that most participants didn't consider the privacy implications of post age until directly prompted to think of it [51]. This may reflect a greater focus on privacy within fandom compared to the general population. However, this was not unchallenged, as in the following comment and response, suggesting that a rich diversity of sharing and linking norms have likely existed for decades in online fandom. Each fandom or social circle potentially contributed individual expectations based on their experiences, positionality, or familiarity with the norms of other communities.

"Back in the LJ days, you could have an unlocked discussion that was private in practice because things were just smaller."

"I don't agree that there was a widespread expectation of privacy in the LJ days. That feels like rose-colored glasses. There was always a diversity of opinion—I wrote meta posts about this, which got linked by newsletters, and sometimes that got really awkward!"

Archiving fandom history was also cited as an important community value by other users, and thus explicit permission, versus a good-faith effort, might be an acceptable sacrifice:

“Linking to Open Doors² is fine because authors are contacted for permission first.”

“Fanlore is a resource maintained by and for the community—if you can’t link to it, it’s not valuable.”

“It’s unreasonable to partake in public discourse and then expect people to ask your consent before discussing you. Linking and archiving allow people to read about what already happened, so we’re not repeating the same things.”

Reflecting past work which identified the community as a whole as the recipient of fannish gifts [72], views like this one emphasize the notion of community ownership even of non-fanwork elements of fandom, and may de-emphasize individual privacy in order to sustain the core benefits of fannish activity.

4.4.3 Contextuality of Sharing/Linking. Beyond historical or communal value, additional forms of context were cited as examples of when something might have widely-understood social norms surrounding sharing and linking. For example, transformative fanwork, including fanfic, fanart, filk, gifsets, journal icons, fanvids, and more, were viewed as acceptable to link to because creators would want their work to be viewed:

“Why would you post to AO3 and not want people to link it?”

“I wouldn’t ask permission to link fanwork like fic, art, vids, etc., because I view them as created for consumption.”

“I can’t really imagine why you wouldn’t want people to link fic that you shared using your fandom pseudonym. It’s already connected to you.”

Asking and giving notice were raised as polite ways to interact with borderline content, such as posts by strangers or about sensitive topics:

“I usually tell strangers how I found their post if I comment on it.”

“I would ask before sharing meta with personal details in it.”

Journal communities were seen as public spaces like a town square, and thus anything posted to one was undoubtedly safe to link because that constituted a very strong signal of intent to disseminate beyond a friend circle:

“I think journal communities are understood as a public space.”

And, relatedly, fans argued that fandom newsletters and multi-fandom communities should operate under different sets of norms than individual users because of their larger influence and audience:

“The norm for the metafiction newsletter was to ask first because the audience was so big. A link in metafiction could send thousands of people to your journal, and that’s a lot to moderate.”

Spaces with many viewers, like communities and newsletters, were acknowledged as a potential vector for harassment.

“I think it’s important to let people choose whether their posts get shared by an aggregator. If you want fandom to not be hostile, you need to be careful about exposing newbies to attention they’re unprepared for.”

“Sometimes people are still developing. Maybe they say something racist on their own blog, and instead of taking the opportunity to help them learn from their mistake, the community just makes a public example of them.”

²<https://www.transformativeworks.org/opendoors/>

With hundreds or thousands of members—“metafandom”, a popular LiveJournal community which aggregated links across fandoms and had over 2,100 members [50])—these spaces were required to balance community interest with individual privacy and preparedness. Controversies could arise due to a clash in social norms between the community and the original poster or prejudices contained within the post (an example expressed in the quote above). The fans we observed understood the need to carefully consider the possibility of harassment or even fair but overwhelming negative responses after sharing or linking content.

4.4.4 Fan Response: Linking Policies. In response to the challenges of navigating so many different attitudes, some fans developed linking policies. These written statements were usually posted on a profile or pinned to the top of a social media account, and they expressed a fan’s preferences for linking or sharing their content, which could range from blanket permission to link public content to explicit statements of the types of implicit stipulations discussed above [22]. Such preferences can be seen as not only erecting personal boundaries, but also as participating in a community discourse about appropriate behavior more generally:

“We started adding notes about linking to our LJ profiles and sticky posts because people didn’t get it, they didn’t understand things we thought were obvious rules about how to treat people appropriately.”

“The podfic craze normalized consent statements in profiles. People included permissions for linking, quotes, etc. in that.”

Fandom newsletters may similarly develop policies explaining when and how they choose content to highlight. Many fans in our data set used their social media sites for a mix of fannish and personal posting, and these public statements allowed for a clear setting of expectations around sharing. As we discuss in Section 5.1.4, these policies are technically unenforceable, yet they likely function as social contracts and strengthen trust norms.

4.5 Deleted Works

One of the symptoms of fandom’s long history, which spans decades and many platforms [25], is that fanwork is sometimes taken down or made very difficult to access. Fans identified multiple reasons for this, the primary one of which was that the creator no longer wished for their work to be visible—which is often a practice used to preserve privacy (see Section 2.2). Other reasons were also named, including the publishing format being limited distribution (e.g. print zines), the original web host going offline or otherwise becoming inaccessible to new users, or a creator’s account being deleted by the platform (e.g. LiveJournal’s Strikethrough and Boldthrough; Tumblr’s 2018 adult content ban). The lenses of intellectual property, the community’s stake in fanwork, and plagiarism can also come to bear on understanding the complex relationship between fandom and works that have been deleted; privacy offers an additional perspective which accounts for the particular risks surrounding online harassment.

There was discussion of when a creator had the “right” to delete their work. Some comments hearkened back to the “public is public” views (Section 4.4), but the unique case here was whether a fanwork containing controversial or offensive content could be deleted. Some said that a creator always had the right to remove their work to protect themselves from continuing harassment, and that fandom history did not outweigh the concerns of an individual facing privacy and security threats. Others came down on the side of history and communal value. We discuss these fiercely held and contradictory opinions, both of which capture core values of fandom, in Section 5.1.5.

“Losing an important fic can be the end of a fandom.”

“It’s more important to ask whether archiving something would cause or continue hurt to the creator because hundreds of people can see it out of its original context than to ask if it’s historically important.”

“I don’t care if we lose our history—I’ve had people take screenshots of private entries so they could mock me and it sucks. Archiving this content prolongs the harassment.”

There was near-universal agreement that reposting someone’s work in full on a current archive site (e.g. taking a deleted work from LiveJournal and posting it as a work on AO3) was not acceptable and constituted plagiarism, even if the original author was named in the metadata (summary, author’s note, etc.). The Organization for Transformative Works (OTW), which also maintains AO3, runs a project called Open Doors. Through Open Doors, an archive that is closing its website can connect with the OTW to import all of the fanwork from the original archive to AO3. This policy involves an attempt to contact the creators of the fanworks via whatever contact information the original site has in its records; if a creator doesn’t respond or cannot be reached, the work is imported under a dummy account. Creators can claim their fanworks and own them under their own AO3 account by responding to this contact, or can ask for them not to be imported. This policy receives some criticism on the grounds of violating the no-reposting tenet established above: some fans feel that the creator has only given posting permission to the original archive and not to AO3.

“If a fanfic author has any rights, then distribution should be one of them. Reposting violates their right to control distribution, and attribution doesn’t change that.”

“It’s unreasonable to expect that people won’t pass on fic shared online, like in private conversations with fellow fans, which has been going on since the 1970s. If you don’t want that, state it when you post the fic.”

“I’m uncomfortable with Open Doors because maybe authors didn’t realize their fic was still up, or only wanted it on the original site.”

Beyond reposting, there was a varied set of opinions about other forms of sharing deleted or inaccessible works. Many fans said they would feel comfortable privately sharing PDFs, ebooks, or text files of a deleted story, such as with friends or with a fellow fan in the same community. Others said they would feel comfortable privately sharing if and only if the original creator gave them permission to do so, or if the original creator had a public statement affirming the same. One comment raised the point that what people would ideally do, and what they do in practice, probably differs.

“People want to be honorable and ask first, but if they really want to get their hands on something, sometimes they aren’t going to follow through on those good intentions.”

Some fans wished that creators would give notice before deleting fanworks, so that fans could download and save personal copies. The question of whether it was appropriate for a fan to have a personal copy of a deleted work invoked comparisons between fanworks and other types of media. For example, some fans made analogies to paper books, where there is no expectation that an out-of-print book would be destroyed. Others argued that fanworks were ultimately the property of the creator and it was respecting the creator’s ownership or wishes to delete the work when the creator deleted it—consumers had only a lending and borrowing relationship with creators. A wide range of perspectives on the rights of readers and authors showed up across multiple metas.

“I appreciate when authors give a short notice before deleting their works and tell readers what type of sharing they’re okay with.”

“Archives aren’t for individuals, they’re for the community. They’re for people finding the fandom years down the line and for preservation of these works, so that nothing disappears because a personal website goes down.”

“I save a copy of everything I like even a little.”

“What happens to a small archive when the only person working on it dies? How will its legacy live on?”

“Pro authors can’t delete their books from a library shelf.”

“Fanfic authors aren’t paid. They should have more distribution rights than paid pro authors.”

5 DISCUSSION

Fannish norms of privacy are remarkably varied, acting to protect the fans’ diverse values. These norms are sometimes enforced by technical systems, whether specifically designed by fans (e.g., orphaning [26]), used by fans (e.g., post privacy settings), or co-opted by fans and other users (e.g., Tumblr tag whisperspaces, Section 4.2.1). But at least as often, they are upheld by policies, whether centralized (e.g., the metafiction newsletter’s policies) or individual (e.g., linking policies on individual profiles), and socially via common respect for trust norms. These norms are strongly held because they help to define fannish identity and because they protect core values of fandom, including identity exploration, the gift economy, and social justice. And they are strongly debated, because they are often in tension and their enforcement is so reliant on trust rather than technical systems. This is why we believe that privacy offers a valuable framework for understanding fandom communities: describing these behaviors in terms of relevant privacy theory gives us insight into how they function.

Throughout this discussion, we emphasize the value of examining fandom as an *activity-defined community*. As the literature has long found and our results reiterate, fans have a strong sense of shared identity, and that identity revolves heavily around the activities of fandom (the creation and sharing of fanworks and surrounding activities). We aim to demonstrate the value of viewing fandom in terms of its activities by seeing how that lens interacts powerfully with theories of privacy.

In the first part of this section, we present the case for analyzing fannish norms using some lenses from privacy theory, demonstrating how they can be applied to both privacy and non-privacy fannish norms to more deeply understand fandom’s longevity and strength. We apply Daniel Solove’s theory of privacy as a collection of related ideas (rather than a single thing) and his proposal to analyze privacy in terms of the protection of the integrity of activities, interlocking with the idea of fandom as an activity-defined community. We then discuss how the apparent contradictory and unenforceable qualities of fannish privacy norms can be understood through these same lenses.

In the second part, we present lessons for privacy researchers. We explain *activity-defined communities* and argue for the value of studying privacy norms and experiences through them. This strategy, we believe, can increase the validity of findings while retaining significant generality, and can help to mitigate the siloing of identity characteristics via a holistic, contextual, and intersectional perspective.

5.1 Lessons for Understanding Fandom’s Strengths

As privacy researchers, we were inspired to conduct this study by observing interesting and unique features of fannish privacy norms. Our study of those norms allows us to conceptualize them

through a variety of theories of privacy, suggesting their strong sophistication in incorporating a variety of approaches to understanding privacy. Daniel Solove’s approach from “Understanding Privacy” gives us a powerful way to conceptualize these diverse privacy norms [67]. In this work, Solove argues that privacy is not best described as a single thing or through a single theory, but instead as “many different yet related things”. Observing that this approach runs the risk of making privacy too vague or diffuse to meaningfully study, he argues that a practical approach is to analyze privacy in terms of *important activities* and *the privacy threats to the “integrity” of those activities*. We believe that this framing of protecting the integrity of threatened activities can be generalized to explain the role of a variety of norms within fandom and other activity-defined communities.

In this section, we begin by producing a working definition of both fannish activities themselves and the properties those activities must possess to retain their integrity. We then draw on existing fan studies literature to identify non-privacy norms related to intellectual property and DIY control of technical infrastructure, describing how those norms help to sustain and protect the integrity of fandom from non-privacy threats. Finally, we describe how fannish norms of privacy also protect the integrity of fandom, and argue for the value of incorporating privacy lenses into broader discussions of fandom in order to understand how fandom has so enduringly managed to persist while upholding transgressive values.

5.1.1 Defining the Integrity of Fannish Activities. To discuss the integrity of fannish activities, we must first define those activities and the properties that are integral to them. We define fannish activities as those that center around the creation and consumption of transformative fanworks. This naturally includes the actual creation and consumption of those works, but also the surrounding ecosystem of activities ranging from forum moderation to cosplay to infrastructure development that build and sustain the sociotechnical spaces in which the creation and consumption of fanworks occurs. These activities often reflect the diversity of the fans who create and maintain them [43, 60], include sexually-explicit and queer content [53], and rely on a permissive interpretation of ownership and intellectual property law [63, 73].

5.1.2 Applying Integrity Protection to Non-Privacy Norms. Having built this definition of fannish activity as the ecosystem around creating and consuming transgressive fanworks, we can apply it to consider how various fannish norms act to protect the integrity of these activities. We do so by recontextualizing two examples of fannish norms from the literature—fandom’s gift economy and fandom’s value of build-your-own tech inherent in the creation of the Archive of Our Own. We then describe the ways in which fannish privacy norms work to uphold multiple aspects of that same integrity, and argue for the value both of taking this privacy-theory informed approach to considering fannish activity and the key (but by no means exclusive) role of privacy among other fannish norms in creating and sustaining fandom.

The gift economy model of fandom suggests that the means by which fans share fanworks, freely as gifts to other fans and the larger community, constitute a challenge to mainstream notions of intellectual property. Fannish activity runs the risk of being disrupted by copyright challenges via legal challenges to fanworks. This is mitigated by the community’s approach to sharing fanworks as gifts. The non-profit nature of gifting helps to keep fans within fair use’s protections for non-commercial works, making actual legal challenges less likely to succeed (and thus less likely to occur) [59]. The framework of the gift economy has also been used to argue that fanworks challenge capitalism by providing non-commercial ways for fans to engage with canons [35, 64]. Some scholars argue that fandom is inherently tied to capitalism [62], and the framework of the gift economy is therefore a contested model for understanding fandom, which we argue is a basis for the necessity of involving multiple theoretical perspectives in order to fully describe the actions

of the community. In this case, the idea of the gift economy allows us to understand the role of non-commerciality and gifting as vital to protecting the integrity of fannish activity.

Fans have long built their own platforms for distributing fanworks and supporting communities, from the canon-specific fic archives and mailing lists of early online fandom to the platforms on which transformative fandom currently thrives, including AO3. The human-computer interaction community has documented this platform's creation and use [26] as a space motivated by feminist design values. It, and platforms like Dreamwidth, were built and popularized in the wake of mass censorship and account deletion events on other platforms, including then-popular Livejournal [23] and FanFiction.Net [21]. These purges disproportionately targeted sexually-explicit and queer content. This followed many earlier threats to the existence of transformative fanwork in the form of cease and desist letters, which led to even fanfiction hosts like FanFiction.Net disallowing fans from posting fiction by certain authors [20]. The impetus behind the initial creation of AO3, then, was explicitly laid out as providing a space for fans, created by fellow fans, which would always allow the publishing of any content which fit the definition of a fanwork—no matter how queer, explicit, or uncomfortable, and with the additional aim of providing legal advocacy to defend fanworks against legal threats around the world [2]. Dreamwidth, announced a year after AO3, was introduced with a similar commitment to freedom of expression [69], and continues its promise to host everything legal under US law [15]. Thus the build-it-yourself ethos that is present in transformative fandom communities protects the integrity of community activities by providing hosting for works and discussion which is not subject to the restrictive landscape of many other social sites or web hosts, and even directly advocates for the legal protection of such work at a global scale.

5.1.3 Applying Integrity Protection to Privacy Norms. Finally, we draw on our findings to describe how privacy norms protect the integrity of fannish activity by defending against privacy threats as well as by creating a safe environment. Here, we focus on the value of this analysis to fan studies, and we defer discussing the benefits to privacy scholarship until Section 5.2. We argue that the power of this approach is in the multidimensionality of privacy, which allows us to explore multiple factors involved in the integrity of an activity-defined community such as fandom. Essentially, Solove's theory, which posits privacy as multiple interrelated theories, serves as a basis for us to understand the ways in which both explicit and implicit expressions of privacy contribute to enabling and strengthening a variety of key aspects of fandom.

At its most simple, privacy norms protect against privacy threats. Such threats, as our results reiterate from past findings, include both individual and communal ones such as job loss, harassment, and unwanted press and public attention. We find that these threats are addressed in the community by a variety of norms, such as strong norms around connections between fannish pseudonymity and real life identities. However, our results show that fans don't think about privacy exclusively in negative, protective terms, instead often centering the positive value of fannish activity for people, fan communities, and broader society. Though in different words, fans are valuing the protection of the integrity of their activities—and the benefits that emerge from doing so—as the key goal of these norms. For example, in Section 4.4, one fan described how it would be unreasonable to prevent certain types of linking and sharing because “if you can't link to it, it's not valuable,” while another argued that public discourse must supersede some privacy protections.

Though the word privacy is not often used, fans incorporated a large number of views of privacy into their norms. This indicates a high degree of sophistication in protecting against complex privacy threats and navigating the positives and negatives of visibility. This multidimensional approach is enabled by Solove's theory of privacy as multiple interlocking things, and allows us to investigate how seemingly contradictory or difficult to enforce mechanisms benefit the

community in subtle ways. Beyond norms of pseudonymity, fans also discussed softer norms such as constraining the spread of fannish activity to fannish spaces in order to reduce the risk of virality. Similarly, fannish spaces define their boundaries by both creating new platform features (e.g., AO3's orphaning mechanism) and adapting existing ones (e.g., using Tumblr's tags as a whisperspace), bringing notions of privacy as related to seclusion into the conversation [78]. Additionally, fans were highly contextual [55] in assessing the boundaries of fannish space and when to share, delete, or archive content within and outside of that space. The contexts they consider included notions of privacy as intimacy [9, 66], such as where fans assessed the intensity of emotion and the age of a post to decide whether and how broadly to share it.

Fans sent and considered implicit and explicit signals to determine how to engage with others' content and speech. Implicit signals included emotions and age of content (as above), but also technical signals (e.g., not putting any tags on a post), while explicit signals included both social ones (e.g., saying "don't reblog") and technical ones (e.g., use of privacy controls such as friendslock). Theories of privacy that focus on individual control over personal information tend to center the direct effects of the binary release or non-release of information, and have often been broadly criticized as surveillance-friendly [77]. By contrast, fannish approaches to control are broader, subtler, and more sophisticated. Here, control acts not only as a direct black-and-white block on sharing but also signals socially-appropriate ways to engage.

We found that fans took this modeling of socially-appropriate engagement very seriously. They sought to act as trustworthy consumers and redistributors of others' works and commentary, and this use of both implicit and explicit signals to create a trustworthy environment that encourages sharing is just what is predicted by Waldman's idea of privacy as trust [76]. Here Waldman's and Solove's theories work in parallel to explain how strong and integrated privacy norms help to create an environment of trust, which enables sharing. As we have defined it, the whole point of fannish activity is to create and *share* fanworks, thus illustrating how the integrity of this sharing is deeply connected to privacy.

Privacy also affects the ability of fandom to maintain another aspect of its integrity: inclusivity. For example, our data shows fans discussing the phenomenon of people first coming out as queer under fandom pseudonyms. While other non-privacy-related features of fandom likely also contribute to protecting the inclusivity of fandom, privacy norms appear to play a major role here. Since we argue that fandom would not be what it is without this inclusion, we must attribute some significant role in fandom's long-term success to its ability to navigate privacy effectively.

This indicates one more way that the sophistication of fannish privacy norms is necessary for protecting fannish activity's integrity. Fandom's activities by their nature induce a tension related to visibility. The creation and sharing of fanworks requires—and its value is enhanced by—visibility. Fans were articulate in expressing how the ability to search for other fans, locate fanworks, and have their own fanworks discovered was critical to maintaining the community. On the other hand, as we have observed, privacy is critical for protecting fans from privacy threats and creating an inclusive space for those marginalized and therefore most likely to be harmed by privacy threats. Fandom's sophisticated approach to both privacy norms and its incorporation of those norms into technical features such as orphaning, we argue, has been critical to enabling fandom to continue and thrive over decades despite this tension.

In these sections, we have argued that taking a multi-dimensional approach to fandom as an activity-defined community provides a framework—the protection of the integrity of fannish activities—that can shed new light on how both privacy and non-privacy norms and practices (such as the gift economy and building your own platforms) have contributed to the persistence and success of fandom. In the next two sections, we explain how this perspective helps us to understand

some superficially confusing aspects of fannish privacy norms by looking at them in multiple dimensions.

5.1.4 Unenforceable: Multi-Purpose Approaches to Strengthen Trust. We found that fans sometimes use privacy strategies, such as linking policies on profiles, which are technically unenforceable and unlikely to be taken into consideration by outsiders who pose considerable privacy threats. What role do these strategies play? Why are they prevalent in a community whose privacy norms we analyze as sophisticated?

Contextual integrity describes privacy violations as flows of information that violate contextual norms [55]. Waldman’s theory of privacy as trust goes further, observing that norms constitute privacy by creating the environment of trust that makes people feel safe enough to share [75]. He followed this with data showing that social trust norms, which involve believing that others will behave favorably and predictably, predict sharing behavior [76]. This result was backed up by subsequent research that showed how trust mediates other important factors such as usefulness [65].

We hypothesize that these “unenforceable” strategies act to increase the salience of trust norms. In this model, we expect that a community member who reads a linking policy comes away with a stronger understanding of the community’s norms and a stronger belief that others will act according to these norms. These social signals encourage people to act favorably and predictably toward one another, and to believe that others will do so. The end result is to reinforce trust within the community, and thus encourage the sharing behaviors at the core of fannish activity. In preliminary follow-on work, we have also found that many of these apparently unenforceable strategies occur in situations where fans are seeking to create privacy from people who challenge the inclusivity of the community.

Future work should examine these “unenforceable” strategies to validate this hypothesis and describe their other purposes. For example, we expect that such policies also act as performances of community membership and adherence to community values for impression management purposes [31], and compatible with our preliminary follow-on work, that they may also signal support for other fandom values, such as social justice, furthering those values and making the community more welcome for diverse fans, and thus protect the inclusivity aspect of fandom’s integrity as well.

5.1.5 Contradictory: Tensions of Key Community Values. We found that deletion norms were strongly held yet frequently in conflict with one another. This tension emerged because of two critical fannish values: fanworks as communal artifacts, and the privacy and safety of fans. AO3’s orphaning feature is one answer to this. Fiesler et al. identified orphaning as reconciling the tension between history and creator control [26], and we build on this by further analyzing creator control in terms of individual and communal privacy and the existential value of privacy for fannish activities.

Some fans drew analogies between fanworks and professionally published print books, or between conservation of fanworks and mainstream efforts to preserve culture. Both analogies positioned fanworks as valuable to the fan community and perhaps to society in the way that art and history are inherently valuable. Others described archives’ purpose as preserving fanworks for future fans, or noted that entire small fandoms could collapse if one creator withdrew their works. We observe that these analogies are strengthened by characterizations of fan culture as a gift economy, in that gifts no longer belong to the giver once they are given. Additionally, Turk’s identification of the fannish community as the recipient of these one-to-many gifts can be seen as justifying these claims of communal ownership [72].

Other fans argued for the importance of leaving distribution of fanworks in the control of their creators, to protect creators’ privacy. Contrasting the idea of a gift economy granting communal

control, some fans argued that creators who are not paid for their works should retain *more* control than paid professionals. Others emphasized consent, potential harms of harassment, and out-of-context publicity, arguing that these concerns took priority given the value of anonymity and safety on protecting fandom as a safe place for sharing.

We can see this tension emerging because these positions are both critical to the integrity of fandom's activities. Arguments emphasizing the individual and communal value of works uphold the positive good—there would be no fandom without the sharing that explores transgressive identities, relationships, ideas, and joy. Yet those activities cannot exist if fandom does not remain a safe and protected place, one that inspires the trust needed for that sharing. The diversity of fandom means that one or the other value may be more salient for certain fans or in specific contexts.

5.2 Lessons for Privacy Research

Recent years have seen a proliferation of important research examining the privacy needs of specific marginalized and vulnerable populations, such as disabled people (e.g., [1, 34]), queer and trans people (e.g., [32, 33, 40, 41, 56]), sex workers [5, 47, 68], and survivors of intimate partner violence [29, 46]. We conclude our discussion with principles for future approaches to inclusive privacy work. Reflecting on these foundational works and inspired by the findings of this paper, we suggest that studying *activity-defined communities* may be a powerful and naturally intersectional approach that enables us to characterize the experiences of diverse people while strengthening the ecological validity of findings.

5.2.1 Study Activity-Defined Communities. Though fandom is a predominantly female and disproportionately queer space [43, 60], with media fandom being associated with female fans since its earliest appearances in academic literature [3, 37], it is not defined by the identity characteristics of its membership. That is, queer and female identities are neither necessary nor sufficient to be a fan. Instead, fannishness is defined by a collection of activities, including engagement with fanworks; discussion of canons; community maintenance activities such as moderating fan-oriented spaces; and connected activities such as the exploration of queer identity [53] and activism and advocacy for social justice [16, 54].

These activities bring enormous benefits. They sustain well-being, root strong communities, make space for free thought and artistic expression, and enable advocacy for justice. Despite the diverse nature of these activities and the many individual fandoms for specific canons, fannish activity loosely unites people as “fans.” Our results illustrate this shared identity, with writers discussing how “Fandom creates a home wherein fellow fans aren’t quite strangers,” a property difficult to explain to “outsiders” but understood within the diverse but recognizable bounds of the community.

We refer to this type of community as an *activity-defined community*, and we argue that we can learn much by researching them. Activity-defined communities can encompass a diverse membership. Therefore, their privacy practices must evolve to be flexible enough to meet the needs of all of their members. And yet, by having coherent boundaries of activity, the community is not so vaguely-defined as to be illegible. As a researcher, conceptualizing a community this way encourages the researcher not to pre-suppose the limits the community or the identities of its members, but rather to observe them and their interactions throughout the work. Additionally, some notion of “success” is available for us to observe. If the activities of the community continue, that is evidence of privacy norms protecting their integrity. And if the membership remains diverse, that is evidence that those norms are protecting the needs of those who are marginalized. Such norms are promising in their ability to teach us how to protect privacy elsewhere.

Fans face significant privacy threats for several reasons. The activities of fandom are often transgressive. Fans interpret and iterate on popular media through new, often queer or female, perspectives [42, 53], producing commentary and work that is not a common focus of that popular media. The often queer or sexually explicit nature of this work marks it as unusual, and therefore controversial or problematic. Merely existing in such a space can require privacy protections. Additionally, fans often produce creative works based on copyrighted intellectual properties, placing fanwork at risk of legal action or intimidation. Such threats are a foundational concern of the OTW, the parent organization of AO3, which operates a legal advocacy team to defend fans against legal challenges to their fanworks [28].

These factors make fannish activities particularly vulnerable to disruption by privacy violations. As a consequence, fannish norms of privacy (as well as non-privacy norms, such as around copying and remix [24]) have been battle-tested for their ability to protect the integrity of these activities. We argue that such communities are particularly valuable subjects of study. Recent work focusing on privacy among activist communities, sex workers, and investigative journalists has demonstrated the value we articulate here. For example, sex workers have clearly defined goals and complex strategies for protecting their safety from privacy harms [47] and to be especially savvy in navigating online spaces [5], while investigative journalists have been found to employ sophisticated technology alongside operational security practices to maintain the privacy of extensive investigations such as in the “Panama Papers” project [49].

What properties do we gain by studying such activity-defined communities? In short, we believe they can help researchers to strike a sweet spot between *specificity* and *generality*. Specificity is needed to avoid broad findings that lack validity, while generality is needed to apply findings beyond specific participants, challenges, or time periods.

One approach to increasing specificity has gained prevalence in recent years: studying people in a particular marginalized category, such as a racial or gender minority. As researchers who conduct such work, we note that there is a danger that we treat all members of a group the same if we focus on one axis of marginalization. When the group contains multiply-marginalized people, their intersecting identities are easily erased and replaced with identities of greater privilege [10]. Focusing on activities may help incorporate all relevant experiences, rather than passing over those outside the frame of a specific axis of marginalization.

One threat to generality is the contextual nature of privacy [55]. Findings may only apply to the narrow context of the particular participants of a study at a given time. By studying a group like fandom, which has persisted over decades, we are capturing a wide set of contexts under one coherent label. The privacy norms which emerge are likely to be adaptable across these contexts and to accommodate many threat models.

5.2.2 Intersectionality and Activity-Defined Communities. Studying activity-defined communities can cross-cut contexts and avoid excluding identities that might be lost in the silo of a study framed by a single identity characteristic. We’ve also observed that communities like fandom, which have sustained a diverse membership over decades, may be particularly good subjects. Drawing on Solove’s theory [67], if we interpret privacy through the lens of its ability to protect the integrity of an activity, then the continued activity among diverse members exhibited by fandom is evidence that those privacy norms are successful. Thus studying the privacy norms of such a community helps us to understand both the needs of marginalized members as well as the nature of privacy practices that have been proven successful at protecting those needs.

Studying activity-defined communities avoids the trap of siloing subjects according to a single identity characteristic. As Crenshaw observed when defining intersectionality, systems often silo axes of marginalization, implicitly assuming privilege along other axes [10]. Thus privacy research

about one axis of marginalization risks centering the most privileged members of that group and erasing those who are multiply marginalized.

We hypothesize that the sometimes contradictory norms we observed (see Section 5.1.5) may in part reflect the experiences of multiply marginalized people. By studying the privacy norms of a diverse community, we identify practices that protect people with different experiences of privacy. They may share experiences along one axis of marginalization, but they also experience the additive and composed effects of marginalization, which can both amplify universal threats as well as create entirely distinct ones that are not able to be reduced to the sum of their parts [10]. Where the threats discussed are divergent, we may be seeing intersectional oppression at work. Where the norms have adapted in sophisticated ways, we may be seeing intersectional norms in action.

Because the meta we analyzed rarely foregrounded the identity characteristics of the writers, we do not validate in this study that our approach is observing intersectional experiences. We believe that our data is compatible with this hypothesis, however. For example, we saw participants who valued the same norms but experienced them differently, such as the multiple perspectives on connections between pseudonyms and real-life names, where some authors were distressed by those connections, while others were more hopeful for a future in which such links were ordinary. These compatible views may represent different individual experiences of the dangers of connecting names and pseudonyms, which could emerge from the greater vulnerability experienced by those facing greater marginalization, either additively (i.e., the same threats are more severe) or in composition (i.e., multiply marginalized fans experiencing different threats). We plan in our future work to examine datasets such as fans' profiles, which frequently disclose identity characteristics, and to engage directly with fans in user studies to examine these possibilities more rigorously.

The existing literature on activity-defined groups offers evidence for the intersectional benefits of this work as well. For example, like fandom, sex workers include many marginalized and multiply marginalized members [36, 52]. Also like fandom, its activities are not monolithic and sex workers often collaborate on mechanisms for protecting one another, such as the Bad Client and Aggressor List [68]. This diversity provides multiple contexts for researchers while preserving validity through the umbrella of sex work. The result is insightful work with deeply intersectional approaches, including exploring risk management language [52], safety and privacy practices [47], the ability to use protective mechanisms such as legal takedown requests [5], and more. These studies are united by their intersectionality, despite differing methods, types of sex work, and experiences captured.

We encourage future work to examine more of these activity-based communities, which we name in the hope that it becomes part of the privacy researcher's toolkit.

6 CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we conducted a qualitative analysis of meta essays that discussed transformative media fandom's privacy-related norms. Though fans and those who study them may not often use the word "privacy" or related academic terms, we argue in this work for the practical and theoretical value of analyzing these norms as sophisticated approaches to privacy governing a variety of complex domains. We examined these norms using Solove's model of privacy as an aggregate of related things, which characterizes privacy by the ways it protects the integrity of activities. Under this framing, we find that fandom's norms—both those privacy-related and not—appear to have been broadly effective at protecting fandom's activities, and have done so for multiple decades in which online fandom has existed. Alongside pre-existing theoretical frameworks, privacy provides us new insights into the function and structure of these norms.

Then, drawing on these results and on recent literature, we codified our approach by defining *activity-defined communities* (Section 5.2). We propose that framing privacy research around

complexes of threatened activities is a valuable technique that can increase ecological validity (by grounding work within real-life contexts defined by activities) and enhance generality (by observing multiple different contexts and intersectional experiences of the studied activities). Additionally, these communities provide a definition of successful privacy practices as activities continuing unimpeded, and lend themselves to intersectional analysis. Going forward, we plan to explore these ideas directly with fans by bringing our results to them and refining them through their feedback in interviews. We hope that our insights will help researchers to learn from the expertise of these communities and to translate the knowledge we gain from them into protections for others.

ACKNOWLEDGMENTS

We would like to thank the additional people who made this work possible: Emanuel Deleon Otazu, Teddy Gu, and Vy Lam, Macalester undergraduates who worked as research assistants in Summer 2022, and Dr. Michael Ann DeVito, who provided feedback on the draft version of the paper. We are also grateful for funding from 3M Gives, the Macalester Young Researchers in the Sciences Fund, and a Hearst grant for supporting our undergraduate researchers.

REFERENCES

- [1] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. 2015. Privacy concerns and behaviors of people with visual impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 3523–3532.
- [2] astolat. 2007. An Archive Of One's Own. <https://web.archive.org/web/20230126003023/https://astolat.livejournal.com/150556.html>. Accessed July 7, 2023.
- [3] Camille Bacon-Smith. 1992. *Enterprising women: Television fandom and the creation of popular myth*. University of Pennsylvania Press.
- [4] Rosaline S Barbour. 2001. Checklists for improving rigour in qualitative research: a case of the tail wagging the dog? *Bmj* 322, 7294 (2001), 1115–1117.
- [5] Catherine Barwulor, Allison McDonald, Eszter Hargittai, and Elissa M Redmiles. 2021. “Disadvantaged in the American-dominated internet”: Sex, Work, and Technology. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [6] danah boyd. 2010. Privacy and Publicity in the Context of Big Data. WWW. Raleigh, North Carolina, April 29.
- [7] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.
- [8] Kristina Busse and Karen Hellekson. 2012. Identity, Ethics, and Fan Privacy. *Fan Culture: Theory/Practice* 38 (2012).
- [9] US Supreme Court. 2015. *Obergefell v. Hodges*.
- [10] KW Crenshaw. 1989. Demarginalising the Intersection of Race and Gender: a Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics. In *University of Chicago legal forum*, Vol. 1. 138–167.
- [11] Michael Ann DeVito. 2022. How transfeminine TikTok creators navigate the algorithmic trap of visibility via folk theorization. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–31.
- [12] Michael A DeVito, Jeremy Birnholtz, and Jeffery T Hancock. 2017. Platforms, people, and perception: Using affordances to understand self-presentation on social media. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*. 740–754.
- [13] 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.
- [14] Jeff D’Onofrio. 2018. A Better, More Positive Tumblr. <https://staff.tumblr.com/post/180758987165/a-better-more-positive-tumblr>. Accessed January 12, 2023.
- [15] Dreamwidth. n.d. Diversity Statement. <https://www.dreamwidth.org/legal/diversity>. Accessed July 7, 2023.
- [16] Brianna Dym and Casey Fiesler. 2018. Generations, migrations, and the future of fandom’s private spaces. *Transformative works and cultures* 28 (2018).
- [17] 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.
- [18] Brianna Dym and Casey Fiesler. 2020. Ethical and Privacy Considerations for Research Using Online Fandom Data. *Transformative Works and Cultures* 33 (2020).

- [19] Brianna Dym, Namita Pasupuleti, and Casey Fiesler. 2022. Building a Pillowfort: Political Tensions in Platform Design and Policy. *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–23.
- [20] Fanlore. n.d. Cease & Desist. https://fanlore.org/wiki/Cease_&_Desist. Accessed July 7, 2023.
- [21] Fanlore. n.d.. FanFiction.Net’s NC-17 Purges: 2002 and 2012. https://fanlore.org/wiki/FanFiction.Net’s_NC-17_Purges:_2002_and_2012. Accessed July 7, 2023.
- [22] Fanlore. n.d. Linking to Public Fan Sites. https://fanlore.org/wiki/Linking_to_Public_Fan_Sites. Accessed January 12, 2023.
- [23] Fanlore. n.d. Strikethrough and Boldthrough. https://fanlore.org/wiki/Strikethrough_and_Boldthrough. Accessed January 12, 2023.
- [24] 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.
- [25] 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.
- [26] 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*. 2574–2585.
- [27] Casey Fiesler and Nicholas Proferes. 2018. “Participant” perceptions of Twitter research ethics. *Social Media+ Society* 4, 1 (2018), 2056305118763366.
- [28] Organization for Transformative Works. n.d.. Legal Advocacy | Organization for Transformative Works. <https://www.transformativeworks.org/legal/>. Accessed January 15, 2023.
- [29] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. 2018. “A Stalker’s Paradise” How Intimate Partner Abusers Exploit Technology. 1–13.
- [30] Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. 2017. Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders. *Proceedings of the ACM on human-computer interaction* 1, CSCW (2017), 1–22.
- [31] Erving Goffman. 1959. The presentation of self in everyday life. (1959).
- [32] Oliver L Haimson, Justin Buss, Zu Weinger, Denny L Starks, Dyke Gorrell, and Briar Sweetbriar Baron. 2020. Trans Time: Safety, Privacy, and Content Warnings on a Transgender-Specific Social Media Site. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–27.
- [33] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–13.
- [34] Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang. 2019. Cooperative privacy and security: Learning from people with visual impairments and their allies. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 1–20.
- [35] Karen Hellekson. 2009. A fannish field of value: Online fan gift culture. *Cinema journal* 48, 4 (2009), 113–118.
- [36] Sandy James, Jody Herman, Susan Rankin, Mara Keisling, Lisa Mottet, and Ma’ayan Anafi. 2016. The report of the 2015 US transgender survey. (2016).
- [37] Henry Jenkins. 1992. *Textual poachers: Television fans and participatory culture*. Routledge.
- [38] 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.
- [39] Brittany Kelley. 2016. Toward a goodwill ethics of online research methods. *Transformative works and cultures* 22, 0 (2016).
- [40] Os Keyes. 2018. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–22.
- [41] 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.
- [42] Alexis Lothian, Kristina Busse, and Robin Anne Reid. 2007. “Yearning Void and Infinite Potential”: Online Slash Fandom as Queer Female Space. *English Language Notes* 45, 2 (2007), 103.
- [43] Lulu. 2013. AO3 Census: Masterpost. <https://centrumlumina.tumblr.com/post/63208278796/ao3-census-masterpost>. Accessed January 9, 2023.
- [44] Annette Markham. 2012. Fabrication as ethical practice: Qualitative inquiry in ambiguous internet contexts. *Information, Communication & Society* 15, 3 (2012), 334–353.
- [45] Annette Markham and Elizabeth Buchanan. 2012. Ethical decision-making and internet research: Recommendations from the AoIR Ethics Working Committee. *Association of Internet Researchers* (2012).
- [46] Tara Matthews, Kathleen O’Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. 2017. Stories from survivors: Privacy & security practices when coping with

- intimate partner abuse. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2189–2201.
- [47] Allison McDonald, Catherine Barwulor, Michelle L Mazurek, Florian Schaub, and Elissa M Redmiles. 2021. "It's stressful having all these phones": Investigating Sex Workers' Safety Goals, Risks, and Practices Online. In *30th USENIX Security Symposium (USENIX Security 21)*. 375–392.
 - [48] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
 - [49] Susan E McGregor, Elizabeth Anne Watkins, Mahdi Nasrullah Al-Ameen, Kelly Caine, and Franziska Roesner. 2017. When the weakest link is strong: Secure collaboration in the case of the Panama Papers. In *26th USENIX Security Symposium (USENIX Security 17)*. 505–522.
 - [50] metafiction. n.d. metafiction: interesting discussions in fandom. <https://web.archive.org/web/20140828092758/http://metafiction.livejournal.com/profile>. Archived on the WayBack Machine, August 28, 2014. Accessed January 12, 2023.
 - [51] Reham Ebada Mohamed and Sonia Chiasson. 2018. Online privacy and aging of digital artifacts. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. 177–195.
 - [52] Jessica D Moorman and Kristen Harrison. 2016. Gender, race, and risk: Intersectional risk management in the sale of sex online. *The Journal of Sex Research* 53, 7 (2016), 816–824.
 - [53] Lucy Neville. 2018. *Girls Who like boys who like boys: Women and gay male pornography and erotica*. Springer.
 - [54] Lucy Neville. 2018. 'The Tent's Big Enough for Everyone': online slash fiction as a site for activism and change. *Gender, Place & Culture* 25, 3 (2018), 384–398.
 - [55] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.
 - [56] Fayika Farhat Nova, Michael Ann DeVito, Pratyasha Saha, Kazi Shohanur Rashid, Shashwata Roy Turzo, Sadia Afrin, and Shion Guha. 2021. "Facebook Promotes More Harassment" Social Media Ecosystem, Skill and Marginalized Hijra Identity in Bangladesh. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–35.
 - [57] Erika Pearson. 2007. Digital gifts: Participation and gift exchange in LiveJournal communities. *First Monday* (2007).
 - [58] Anthony T Pinter, Jialun Aaron Jiang, Katie Z Gach, Melanie M Sidwell, James E Dykes, and Jed R Brubaker. 2019. "Am I Never Going to Be Free of All This Crap?" Upsetting Encounters with Algorithmically Curated Content About Ex-Partners. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
 - [59] Elizabeth Rosenblatt and Rebecca Tushnet. 2015. Transformative works: Young women's voices on fandom and fair use. *Egirls, eCitizens* 385 (2015).
 - [60] Lauren Rouse and Mel Stanfill. 2023. Over*Flow: Fan Demographics on Archive of Our Own. <https://www.flowjournal.org/2023/02/fan-demographics-on-ao3/>. Accessed July 7, 2023.
 - [61] Rachael Sabotini. 1999. The Fannish Potlatch: Creation of status within the fan community. In *The Fanfic Symposium*, Vol. 20.
 - [62] Cornel Sandvoss. 2005. *Fans: The mirror of consumption*. Polity.
 - [63] Aaron Schwabach. 2016. *Fan fiction and copyright: Outsider works and intellectual property protection*. Routledge.
 - [64] Suzanne Scott. 2009. Repackaging fan culture: The regifting economy of ancillary content models. *Transformative Works and Cultures* 3, 1.6 (2009), 1–11.
 - [65] Ashraf Sharif, Saira Hanif Soroya, Shakil Ahmad, and Khalid Mahmood. 2021. Antecedents of self-disclosure on social networking sites (sns): A study of facebook users. *Sustainability* 13, 3 (2021), 1220.
 - [66] Scott Skinner-Thompson. 2015. Outing Privacy. *Nw. UL Rev.* 110 (2015), 159.
 - [67] Daniel J Solove. 2008. Understanding privacy. (2008).
 - [68] Angelika Strohmayer, Jenn Clamen, and Mary Laing. 2019. Technologies for social justice: Lessons from sex workers on the front lines. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–14.
 - [69] Synecdochic. 2008. Announcing Dreamwidth Studios. Coming Summer 2008. <https://synecdochic.dreamwidth.org/221888.html>. Accessed July 7, 2023.
 - [70] Tumblr Changes. 2022. Ongoing incident. <https://changes.tumblr.com/post/679998365786947584/ongoing-incident>. Accessed January 12, 2023.
 - [71] Tumblr Staff. 2022. Introducing: Reblog Controls. <https://staff.tumblr.com/post/685964115655737344/introducing-reblog-controls>. Accessed January 12, 2023.
 - [72] Tisha Turk. 2014. Fan work: Labor, worth, and participation in fandom's gift economy. *Transformative Works and Cultures* 15 (2014), 113–18.
 - [73] Rebecca Tushnet. 2007. Copyright law, fan practices, and the rights of the author.
 - [74] Jessica Vitak, Yuting Liao, Mega Subramaniam, and Priya Kumar. 2018. 'I Knew It Was Too Good to Be True' The Challenges Economically Disadvantaged Internet Users Face in Assessing Trustworthiness, Avoiding Scams, and Developing Self-Efficacy Online. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–25.
 - [75] Ari Ezra Waldman. 2016. Privacy, sharing, and trust: The Facebook study. *Case W. Res. L. Rev.* 67 (2016), 193.

- [76] Ari Ezra Waldman. 2018. *Privacy as trust: Information privacy for an information age*. Cambridge University Press.
- [77] Ari Ezra Waldman. 2021. *Industry unbound: The inside story of privacy, data, and corporate Power*. Cambridge University Press.
- [78] Samuel D WARREN. 1890. The Right to Privacy. *Harvard Law Review* 4, 5 (1890).
- [79] Natasha Whiteman. 2012. Undoing ethics. In *Undoing Ethics*. Springer, 135–149.
- [80] Bin Xu, Pamara Chang, Christopher L Welker, Natalya N Bazarova, and Dan Cosley. 2016. Automatic archiving versus default deletion: what Snapchat tells us about ephemerality in design. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 1662–1675.

Received January 2023; revised July 2023; accepted November 2023