



Privacy Research with Marginalized Groups: What We Know, What's Needed, and What's Next

SHRUTI SANNON, University of Michigan, USA

ANDREA FORTE, Drexel University, USA

People who are marginalized experience disproportionate harms when their privacy is violated. Meeting their needs is vital for developing equitable and privacy-protective technologies. In response, research at the intersection of privacy and marginalization has acquired newfound urgency in the HCI and social computing community. In this literature review, we set out to understand how researchers have investigated this area of study. What topics have been examined, and how? What are the key findings and recommendations? And, crucially, where do we go from here? Based on a review of papers on privacy and marginalization published between 2010–2020 across HCI, Communication, and Privacy-focused venues, we make three main contributions: (1) we identify key themes in existing work and introduce the *Privacy Responses and Costs* framework to describe the tensions around protecting privacy in marginalized contexts, (2) we identify understudied research topics (e.g., race) and other avenues for future work, and (3) we characterize trends in research practices, including the under-reporting of important methodological choices, and provide suggestions to establish shared best practices for this growing research area.

CCS Concepts: • Security and privacy → Human and societal aspects of security and privacy; • Human-centered computing;

Additional Key Words and Phrases: marginalization, vulnerable groups, discrimination, literature review

ACM Reference Format:

Shruti Sannon and Andrea Forte. 2022. Privacy Research with Marginalized Groups: What We Know, What's Needed, and What's Next. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 455 (November 2022), 33 pages. <https://doi.org/10.1145/3555556>

1 INTRODUCTION

People who face marginalization in society—whose needs and experiences are overlooked and who have limited resources and power due to some facet of their identity [18]—can have unique privacy-related needs and behaviors that must be recognized by researchers and designers of technology. Marginalized groups can experience disproportionate harms when their privacy is violated. Consider the case of a person living with HIV, who may risk social stigma, discrimination, or the dissolution of relationships should their HIV status become publicly known without their consent [44], as compared to the impact of a health information leak for a person who does not have to navigate a marginalized and often stigmatized identity. In response to these elevated privacy threats, marginalized individuals have to navigate complicated decisions around identity management and disclosure, which extend to their technology use [104]. At the same time, marginalization can also constrain privacy-protective behaviors; for example, people in economically disadvantaged communities often have limited access to privacy literacy resources [70]. Understanding the new

Authors' addresses: Shruti Sannon, sannon@umich.edu, University of Michigan, Ann Arbor, MI, USA; Andrea Forte, aforte@drexel.edu, Drexel University, Philadelphia, PA, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2022/11-ART455 \$15.00

<https://doi.org/10.1145/3555556>

dimensions and consequences that privacy issues take on in marginalized contexts is vital to designing more equitable technologies that respect the privacy of all users rather than a select few.

Given the need to represent marginalized voices in the design of technologies, and the potential privacy harms that stem from excluding them, there has been growing interest in this research area [111]. Wang suggests that we are currently seeing a third wave in privacy research that centers what he terms as “inclusive privacy” that encompasses “different human abilities, characteristics, needs, identities, and values,” as compared to prior waves that focused on the technical and usability aspects of privacy [114, 115]. Concurrently, privacy researchers in this space have begun to convene via panels [22] and workshops, such as the 2017 CSCW workshop on privacy for vulnerable populations [76], and the 2015-2022 SOUPS workshops on inclusive privacy and security (WIPS), pointing to a clear interest in this research area.

What remains unclear is how the work being conducted by researchers in this area coheres together, the bounds of communal knowledge in this space, and what gaps still need to be filled. Literature reviews can be useful to assess and guide new, emerging areas of research [26]. Literature reviews have been used to this end in computing research, including at CSCW and CHI (e.g., [26, 27]). Given the rapidly growing interest and research in this space, we see this as a key moment in time for a literature review that could help spur new research in this area while also providing an understanding of our cumulative knowledge and practices thus far.

To this end, we conducted a literature review of privacy research on marginalized populations, using a broad and open-ended definition of “marginalized adults” as adults facing any form of social exclusion or discrimination due to some facet of their identity [18, 48], such as along the lines of race, sexual identity, gender identity, socioeconomic status, and immigration status, among other factors. We examined papers published between 2010–2020 across a broad range of venues representing three inter-related disciplines: HCI and Social Computing (e.g., Proceedings of CSCW, CHI, and Ubicomp), Communication (e.g., Journal of Communication and New Media and Society), and Privacy-focused venues (e.g., Proceedings of the Privacy Enhancing Technologies Symposium).

After manually reviewing 2,823 privacy-related papers that were published across these disciplines between 2010-2020, we identified 88 papers that focused on the intersection of privacy and marginalization. Based on an analysis of these 88 papers, we offer the following contributions:

- *Key Research Areas and Findings.* We describe the major contexts of marginalization studied and the kinds of findings and recommendations that papers in our dataset put forward. In doing so, we introduce the *Privacy Responses and Costs* framework, which outlines 10 privacy responses people have to privacy threats, and the costs and consequences of these responses faced by marginalized groups.
- *Descriptive Norms for Research Practices.* We characterize the literature in terms of the practices that researchers report using to study privacy and marginalization.
- *Gaps and Implications.* Having surveyed the literature, we identify understudied areas and questions (e.g., race), lack of diversity in methods, and under-reporting of important methodological choices, and we make concrete recommendations for establishing shared best practices for this growing research area, such as designing protocols that minimize harm.

2 RELATED WORK

Given our focus on marginalization, we begin by explaining what we mean when we use this term. Then, we discuss the use of literature reviews in computing research broadly and privacy research both within and outside of HCI and social computing (for brevity, we refer to this field as just “HCI” in what follows). We end by discussing the focus of the current review.

2.1 What Do We Mean by Marginalization?

Although the terms “marginalization” or “marginalized” are sometimes used in studies of technology use and privacy concerns, the term is commonly left undefined in HCI and social computing work. In our findings, we used a grounded approach to identify the ways that researchers engage and discuss marginalization in the privacy literature; however, to begin the task of identifying research about marginalization, we needed to define precisely what we mean when we use the term.

Marginalization is a complicated concept with varying definitions and dimensions. We drew from several sources to construct our frame for “marginalized contexts.” Marginalized populations are defined as “persons who are peripheralized based on their identities, associations, experiences and environments” [48, p. 25]. As a result, these groups are “excluded from mainstream social, economic, cultural, or political life” [18]. People can be marginalized based on several factors, including race, disability, gender identity, sexual orientation, socioeconomic status, and immigration status.

The impact of marginalization can be seen not only in the ways groups are peripheralized, but also in social responses to this exclusion. Marginalization has been acknowledged through legal protections around the world through laws prohibiting discrimination in housing, employment, and other areas of life (e.g., the Fair Housing Act in the U.S. [86], the Racial Equality directive in the European Union [20], and the Rights of Persons with Disabilities Act in India [43]).

In this review, we focus on research about marginalized adults, many of whom face discrimination that these legal responses were set up to address. Although we were undoubtedly influenced by these readings and familiar *a priori* categories of marginalization commonly represented as legally protected classes, we conducted an open-ended analysis: if a paper described a group as being stigmatized or marginalized, or referred to privacy and marginalization more broadly, we opted to include it. As a result, our review includes a wide swath of marginalized identity characteristics, such as disability, sexual orientation, socioeconomic status, immigration status, and race, as well as experiences like sex work and human trafficking.

Although the studies we review in this paper are included because they study populations that have the characteristic of being “marginalized,” marginalization is not simply a feature of a population or individual, but a dynamic social process of exclusion that marshals the power of social norms, institutions, and interpersonal dynamics to render some people as privileged and others as inconsequential. As privacy researchers, we use this process orientation to understand privacy risks and violations as features of marginalization that change over time and in different contexts. As social computing scholars, we pay special attention to the role that technology design plays in initiating, continuing, accelerating, stalling, or obstructing processes of marginalization.

2.2 Why do a Literature Review? Reviews in HCI and Social Computing

Literature reviews are not uncommon in HCI and social computing venues, and can provide insight into both the knowledge built by a research community as well as the practices used to construct this knowledge. First, literature reviews are a useful way to identify research trends and unearth new directions for research communities, and have been used to this effect to further research on the sharing economy [26], HCI for development [23], and sustainable HCI [27], among other areas.

Second, literature reviews can also help develop clarity around methods and direct best practices for a research community. To this effect, literature reviews have provided insight and instruction around the use of reliability measures [78], the reporting of compensation [89], the standards for sample size [12], and anonymization practices [1] in social computing research.

Our goal is for this review to serve both of these purposes: (1) to understand the existing knowledge generated by researchers on privacy issues in marginalized contexts while pinpointing

new areas of exploration, and (2) to understand the methods through which this knowledge generation has occurred with an eye towards guiding future research practices.

2.3 Reviews of Privacy Research

Literature reviews are also not new to privacy research, both within and outside the HCI research community. Some of these reviews have explored conceptual or theoretical approaches in privacy research, such as how design relates to privacy in HCI research [120], how researchers have theorized about the privacy paradox [7], and how privacy has been conceptualized in HCI [6].

Alternatively, several privacy-related reviews have focused on the privacy considerations of specific technologies or contexts, such as eye-tracking [64], cryptocurrencies [53], Internet of Things (IoT) infrastructures [5], big data [85], or electronic health record systems [72, 98]. These reviews have put forth several implications for privacy research, including identifying privacy threats that still need to be addressed [53], and identifying a dearth of research on privacy protection in relation to the specific technologies under study [5, 64].

2.4 Our Focus

It is clear that literature reviews can be useful in evaluating nascent computing fields and subfields and unearthing future directions for research communities. We searched the ACM and other databases to determine whether a review of privacy research on marginalized contexts had been conducted, and to our knowledge, no such review exists. Given both the importance of and growing interest in this subfield of privacy, we aim to provide a review that will tie existing work together, highlight key similarities and differences in the subfield's knowledge and research practices, and provide recommendations for future research.

Having established the need for a review of research on privacy and marginalization, we drew on existing literature review practices to appropriately scope our inquiry. First, many literature reviews published in HCI and social computing venues (e.g., [1, 78]) focus on research published by the ACM. Since our primary audience is the HCI and CSCW community, we also chose to focus on HCI research published by the ACM. However, we also wanted to acknowledge the blurry nature of interdisciplinary boundaries. This led us to review a sample of publications from adjacent disciplines as well. Toward this end, we also examined papers published in a selection of high-impact Communication journals like *New Media and Society* and Privacy-focused venues like *IEEE Privacy & Security*.

Second, many social computing literature reviews are scoped by time period. We found considerable variation among the publication periods examined by reviews, ranging from three years [78] to ten years [26]. We scoped our review to papers published from 2010 through 2020 to capture a broad range of papers and measure activity in this research area over time.

As discussed in Section 2.2, reviews can synthesize knowledge in a domain, identify areas ripe for exploration, and develop best practices for a research community. Accordingly, our review is guided by a few overarching questions to achieve these goals. The first set involves the breadth of research in this area and key findings:

What research has been done on privacy in marginalized contexts? What can we learn from this cumulative knowledge about how marginalization relates to privacy, and what are the opportunities for future research directions in this space?

The second set of questions involves the practice of doing research on this topic:

How has this research been done? What can we learn in terms of best practices for conducting privacy research on marginalized contexts?

3 METHOD

Our method for conducting the review was made up of four main stages: 1) collecting a corpus of privacy-related papers in multiple fields and venues, 2) filtering this dataset to only include papers that focus on both privacy and marginalization, 3) conducting a quantitative analysis to identify descriptive trends (e.g., common publication venues), and 4) conducting a qualitative thematic analysis to understand themes in the papers' approaches, findings, and recommendations.

In what follows, we explain the steps we took at each stage of this process. We also present an overview of our method using the Preferred Reporting Items for Systematic Reviews and Meta-Analysis (PRISMA) schematic of study flow [81] in Figure 1.

3.1 Data Collection

To reflect the diversity of research on privacy, we sampled papers from three distinct but interdisciplinary and overlapping fields: HCI, Communication, and Privacy-focused venues. Our goal was to review a broad set of papers that represent research in these fields, rather than attempting to collect the complete set of papers published on privacy and marginalization.

Data collection took place in 2021. Per guidelines for effective literature reviews [65], we developed selection criteria to identify relevant papers for inclusion in the review: (1) full research articles (i.e., not posters, abstracts, panels, and so forth), (2) published between 2010 and 2020, (3) with a focus on both privacy and marginalization. In this section, we describe our process for identifying search terms for the review, and then detail our data collection procedures for each of the three fields.

3.1.1 Identifying Search Terms. We began by exploring search terms for identifying relevant papers. Since our aim was to identify papers that were both about privacy and marginalization, we considered using search terms for both of these foci. However, we quickly discovered that many relevant research papers on marginalized contexts do not use common keywords (such as "marginalization" or "vulnerable") and would thus be excluded from our dataset. While we considered developing a list of marginalized contexts to use as *a priori* search terms (e.g., "disability"), we decided that this would place artificial limits on the breadth of our data.

Thus, we decided to take a more expansive (albeit labor-intensive) approach: to collect all privacy-related papers published in our selected venues to arrive at a dataset of privacy research between 2010–2020, and then manually sort through these papers using our selection criteria to identify papers that also focused on marginalized contexts. The venues we selected to gain a broad overview of HCI, Communication, and Privacy venues are listed in Table 1. To identify privacy-related papers, we searched for the term "privacy" in papers' titles, abstracts, and/or keywords. We considered using a more extensive set of keywords related to privacy (e.g., surveillance, tracking, disclosure), but found that these terms were often overly broad or narrow, and thus introduced significant noise in the dataset (e.g., "tracking" brought up numerous results about eye sensor tracking that were not related to privacy). Since we were also searching fields outside of HCI where these terms could take on conflicting dimensions, we decided to use privacy as the sole keyword.

3.1.2 HCI. We searched the ACM Digital Library (ACM-DL) for papers with the term "privacy" in either the title, abstract, and/or keywords. We used the ACM-DL search options to restrict the results to (1) papers published between 2010–2020, and (2) papers classified as "Research article" (thus excluding posters, workshops, abstracts, and panels).

We then selected the "sponsored by SIGCHI" option to identify papers published in the proceedings of HCI conferences (e.g., CSCW, CHI, DIS, etc.). This search produced 623 papers. We also collected search results for the following journals: the Proceedings of the ACM on Human-computer

Interaction (PACM) (62 papers), Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT) (75 papers), and Transactions of Computer-Human Interaction (TOCHI) (12 papers).

We downloaded a dataset composed of 772 HCI papers based on the above procedure. We cleaned this dataset by removing 14 duplicates and 75 non-research articles (e.g., posters and doctoral consortium abstracts that were incorrectly tagged in the ACM-DL). Our HCI dataset after data cleaning contained 683 privacy-related papers.

3.1.3 Communication. There is no central database for Communication papers, and literature reviews on communication and technology typically examine a predetermined selection of journals (e.g., [32]). Following this approach, we selected a sample of top journals in Communication that publish papers on technology: Journal of Communication (JOC), Journal of Computer-Mediated Communication (JCMC), Communication Research (CR), International Journal of Communication (IJOC), New Media and Society (NMS), and Social Media + Society (SM+S).

We searched for “privacy” on each individual journal’s website, filtering for full research articles published during 2010–2020. This resulted in a total of 978 papers (JOC = 58, JCMC = 155, CR = 63, IJOC = 79, NMS = 380, and SM+S = 243). However, some journals’ search interfaces did not allow us to exclude full-text searches (unlike our ACM search), and thus many papers were irrelevant to privacy as a research topic (e.g., papers that mentioned privacy once in the Methods section when describing a consent process). Thus, to stay consistent with our ACM search, we manually filtered the Communication dataset to only include papers that contained privacy in their titles, abstracts, and/or keywords (when available). After this filtering process, the Communication dataset consisted of 208 privacy-related papers (JOC = 8, JCMC = 13, CR = 4, IJOC = 79, NMS = 58, SM+S = 46).

3.1.4 Privacy-focused venues. We sampled four Privacy venues: the USENIX Symposium on Usable Privacy and Security (SOUPS), IEEE Security and Privacy (S&P), Proceedings of the Privacy Enhancing Technologies Symposium (PoPETS), and USENIX Security. Since these venues are focused on privacy, we did not have to run a search to identify privacy-related papers. Instead, we included the entire corpus of 2010–2020 papers published at SOUPS, IEEE S&P and USENIX Security in our dataset; PoPETS began its proceedings in 2015, and we included all PoPETS papers from 2015–2020 in the dataset. The dataset from these venues consisted of 1,932 papers (SOUPS = 232, PoPETS = 306, USENIX Security = 813, IEEE S&P = 581).

3.2 Data Filtering

Once we had a corpus of 2,823 privacy-related papers from HCI, Communication, and Privacy-focused venues, we manually coded this dataset to identify papers that focused both on privacy and marginalization. We coded each paper for inclusion/exclusion according to the following codes: (1) relevant (the paper is about privacy in the context of marginalization), (2) not relevant (the paper

Venues	Journals and Proceedings
HCI	CSCW, CHI, DIS, MobileHCI, PACM CSCW, PACM GROUP, PACM IMWUT
Communication	CR, JOC, IJOC, JCMC, NM&S, SM+S
Privacy	SOUPS, PoPETS, IEEE S&P, USENIX Security

Table 1. Venues represented in the dataset comprising 8 conference proceedings and 9 journals from 2010–2020. Note that CSCW was a proceedings through 2016 and a journal thereafter.

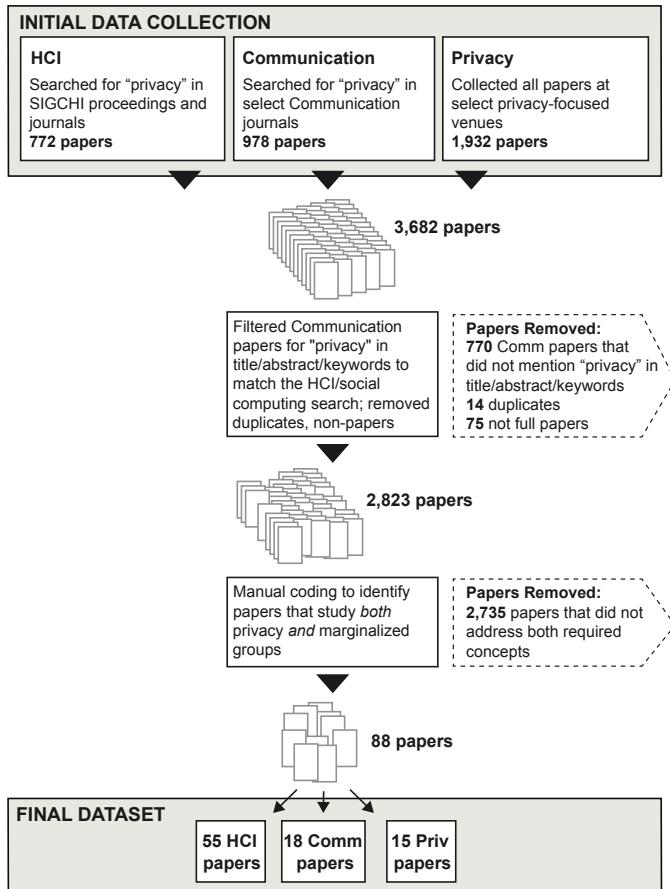


Fig. 1. The Preferred Reporting Items for Systematic Reviews and Meta-Analysis (PRISMA) schematic of study flow, illustrating the data collection and filtering process of publications in selected venues published during 2010–2020

does not cover both privacy *and* marginalization), and (3) unclear. We coded each paper based on its title and abstract, only referring to the full text of the paper if the coding decision was unclear.

We operationalized marginalization through an iterative process of reading and discussion—in addition to reading in the areas of HCI, Communication, and Privacy, we investigated the history of the term and read widely in journals from other fields, books, and websites of governmental agencies and NGOs. Definitions of marginalization in these readings were not homogeneous and sparked lively discussions in our research group about inclusion criteria. In cases where a group's marginalized status was unclear, we chose to only include papers that used language to plainly situate such populations as marginalized.

To ensure a clear and consistent interpretation of the concepts "privacy" and "marginalization", we used an iterative process of discussion and testing. Two researchers independently coded a subset of 10 papers each for relevance and discussed their initial impressions. Then, they coded 25 more papers and met to discuss and reconcile disagreements; at this point, their inter-rater reliability was good (Cohen's $\kappa = 0.74$). Following this discussion, they coded 25 additional papers; inter-rater reliability for this round of coding was strong (Cohen's $\kappa = 0.83$). Having established

close agreement over the coding scheme, one coder then independently coded the remaining papers in the dataset for relevance. Out of an abundance of caution, 61 papers were coded as “unclear” and flagged for further examination; through discussion, 13 of these papers were re-coded as relevant for inclusion in the final dataset.

Ultimately, 88 of the 2,823 privacy-related papers were coded as relevant for inclusion in the final dataset (HCI = 55, Communication = 18, Privacy = 15). All 88 papers were downloaded from their respective repositories for analysis. A list of the 88 papers is available as a supplementary file to this paper on the ACM digital library.

3.3 Data Analysis

During analysis, both authors read papers in the dataset; between us, we read all 88 papers. We conducted two main analyses: a quantitative analysis to identify descriptive trends across papers, and a qualitative analysis to identify themes in papers’ findings and practices, as detailed below.

3.3.1 Quantitative Analysis. We extracted descriptive data about each paper and compiled these into a spreadsheet. These data included each paper’s publication year, publication venue, methods, marginalized population/context, and technological focus, if any. We also noted whether each paper discussed ethics, compensation, and positionality; we extracted the content of these statements—when present—into the spreadsheet. We used these data to calculate descriptive statistics about the dataset (e.g., how many papers were published at CSCW, how many focused on LGBTQ+ individuals, how many provided information about compensating participants).

3.3.2 Qualitative Analysis. As we read through each paper, we created open codes to identify concepts and themes that cut across the dataset. We focused on study rationales, findings, implications, and recommendations of each paper, as well as how authors discussed the relationship between technology, privacy, and marginalization. We met regularly to discuss patterns across papers. For example, we found that several papers highlighted key privacy-related tensions in relation to marginalization (such as the need for social support versus the need for privacy); through discussion, we categorized these as four key tensions that we present in Section 7.2.

The primary goal of the qualitative analysis was to inductively generate concepts and themes, not to consistently identify examples of pre-defined concepts; as such, we did not calculate inter-rater reliability but instead used ongoing, iterative discussion to achieve consensus, aligned with best practices described in McDonald et al. [78]. This is in line with several other reviews that focus on surfacing themes in research areas (e.g., [27, 120]).

4 AN OVERVIEW OF THE DATASET

We begin by providing a snapshot of the work done at the intersection of privacy and marginalization. We describe when and where papers in our dataset were published, and then identify the foci of these papers, both in terms of the types of marginalization contexts and the types of technologies they focus on. To describe our findings, we use the following notation for each main type of venue: HCI (SIGCHI venues), C (Communication), and P (Privacy-focused venues).

4.1 Publication Venues and Years

Our final dataset represented 88 papers from 17 journals or conference proceedings. The majority of these papers were in HCI (55, 63%), followed by Communication (18, 20%) and Privacy-focused venues (15, 17%). Figure 2 illustrates the number of papers published between 2010–2020 by venue.

Within HCI, the majority of the 55 papers were published at CHI (24, 44%) and CSCW (including both the conference proceedings and the PACM journal, 24, 44%), suggesting that CHI and CSCW are the key venues that HCI researchers target for work on privacy and marginalization and that

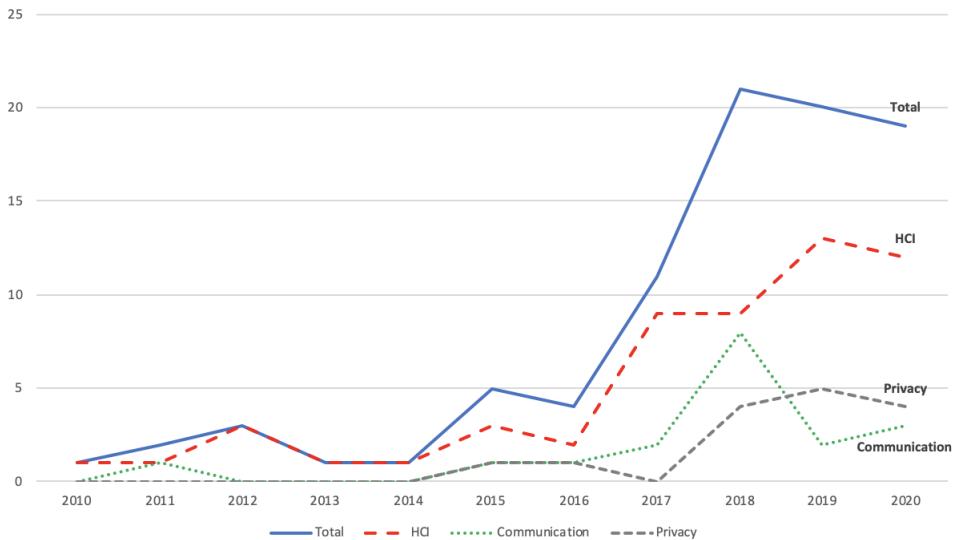


Fig. 2. Publication counts from 2010 to 2020 in HCI, Communication, Privacy-focused venues, as well as the total number of papers across venues over time

these two communities have an orientation that makes them receptive to the work. A few were published at other HCI venues: DIS (4, 7%), PACM IMWUT, GROUP, and MobileHCI (1 paper each).

In Communication, a little over half of the 18 papers were from the International Journal of Communication (10, 56%), followed by New Media and Society (4, 22%) and Social Media + Society (2, 11%). Journal of Communication and Communication Research included one paper each.

In Privacy-focused venues, most of the 15 papers were in SOUPS (6, 40%) and USENIX Security (5, 33%), followed by PoPETS and IEEE S&P (2, 13.33% each).

Overall, the vast majority (71, 78%) of the 88 papers were published between 2017 and 2020, bolstering our view that this is a growing area of interest among researchers across HCI/Social Computing, Communication, and Privacy.

4.2 Types of Marginalization Contexts

We found that papers could be broadly categorized as focusing on the following types of marginalization contexts: (1) individuals and identities; (2) physical spaces and communities; (3) online spaces, tools, and communities; or (4) marginalization in general. We discuss each of these types below. Figure 3 presents a breakdown of paper distributions by context and publishing venue.

4.2.1 Individuals and Identities. Most papers (72, 82%) focused on specific marginalized identities (HCI = 46, C = 13, P = 13). Disability was the most popular topic (28, 32%) in the dataset, comprising more than a third of all HCI papers (21) and approximately half of papers at Privacy-focused venues (7). There were no Communication papers in this area. The most common disability topics included dementia and age-related cognitive impairments (HCI = 8, P = 4), visual impairments (HCI = 2, P = 4), and HIV (HCI = 5). In contrast, only three papers explored mental health conditions.

The second most common focus was individuals who identify as LGBTQ+ (14, 16%); within this category, most (8) papers focused on trans individuals. Several papers examined intimate partner abuse (9, 10%), low-income individuals (6, 7%), and the risks faced by women in patriarchal South Asian contexts (5, 6%). There were also several papers that focused on contexts that often (but not

exclusively) involve women, including sex work (1), sexual assault (1) and human trafficking (1), as well as women transitioning from incarceration (1).

Only two papers examined race and ethnicity as their primary focus; both in Communication. A few papers studied groups that were also often racial or ethnic minorities: undocumented immigrants (2), politically marginalized minorities (1), and refugees (1).

4.2.2 Physical Spaces and Communities. Rather than focusing on individuals, some studies (10, 11%) focused on marginalization at the level of communities (HCI = 6, C = 3, P = 1). Most of these papers (8) examined communities that are economically disadvantaged, while two examined other factors of marginalization: crime, income, and racial segregation across neighborhoods [57] and communities that are “geographically and culturally marginal” [58].

4.2.3 Online Spaces, Tools, and Communities. Two HCI papers in the dataset focused on tools that marginalized groups use to protect their privacy online or communities in which marginalized people participate: Tor, open-source software that enables anonymous communication [35], and online fandom communities [29].

4.2.4 Marginalization in General. Four papers did not focus on a particular context of marginalization: One non-empirical paper theorized about how to integrate vulnerability and marginalization in privacy theories and research [77]. Two explored how sensitive data about people can be identified

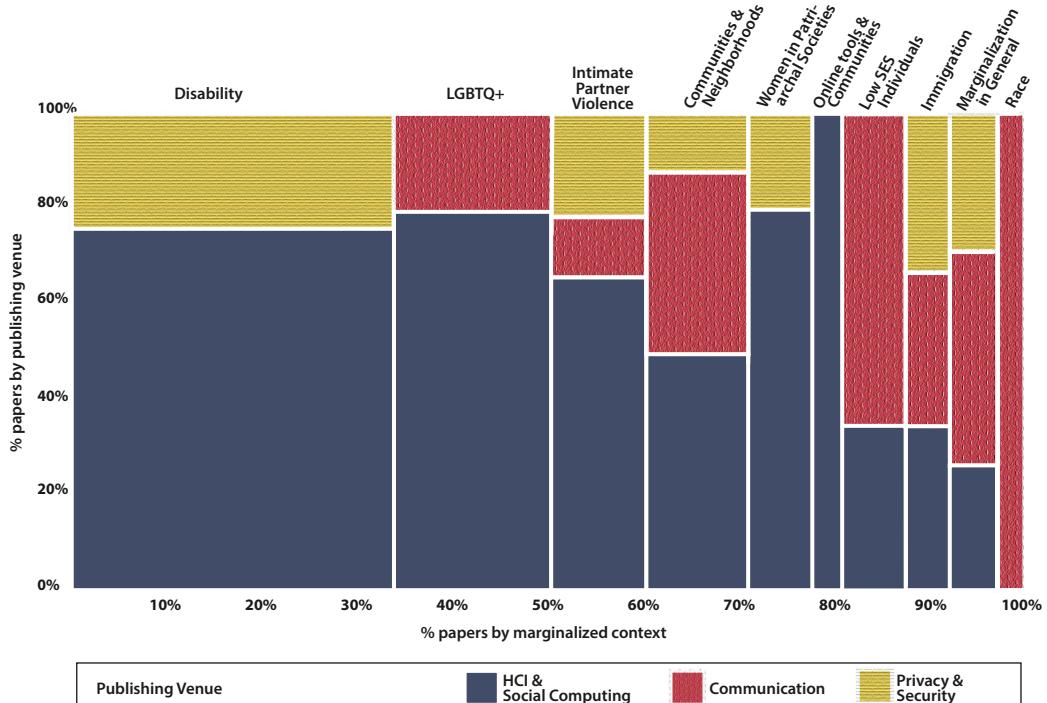


Fig. 3. Papers by publication venue and marginalization context. The width of each bar represents the percentage of papers in each marginalization context across all publication venues. The height of each color bar represents the percentage of papers in each publishing venue. Graph includes contexts represented by at least 2 papers.

through their digital traces [11, 73], and how these privacy violations can be particularly concerning for populations with stigmatized or marginalized attributes. One paper examined how the digital divide impacts marginalized groups more broadly [41].

4.3 Types of Technologies

In addition to identifying the marginalization context of papers, we identified the types of technologies they focused on, if any. Approximately one third of papers (31, 35%) did not focus on a specific technology but, rather, explored technology use more broadly. This approach appeared in all three fields (HCI = 20, C = 5, P = 6).

In terms of studies that focused on specific technologies, the most common focus was social media (23, 26%), particularly in HCI (HCI = 18, C = 3, P = 2); these included papers that examined a broad range of social media platforms or individual platforms, as well as different types of social media (e.g., online forums, dating apps, social networking sites, etc.).

Mobile phones were the second most common technology (8, 9%), followed by location-tracking technologies (5, 6%); the rest of the studies focused on many different types of technologies, including assistive technologies, automatic gender recognition algorithms, smart homes, and national identification systems. In a departure from the rest of the dataset, two Communication papers did not focus on technology or technology use at all, instead theorizing about the relationship between marginalization and privacy.

5 WHY STUDY PRIVACY AND MARGINALIZATION?

To study a phenomenon, it must be problematized in such a way that its importance is clear to the members of a research community. How researchers frame marginalization as a compelling area for privacy research not only sheds light on the research topic, but also on the scholarly community. In this section, we explore the rationales motivating work in this space.

5.1 Disproportionate risks and challenges

In many cases, studies problematized how marginalized groups may need to use technology—particularly in service of some identity-related goal—and simultaneously face elevated privacy risks from doing so. For example, survivors of intimate partner violence who shared custody of a child with an abusive ex-partner found that while electronic communication enabled remote contact between the child and ex-partner, it also opened up the possibility for intrusive access, such as through a live video feed of their home [38].

Several studies were motivated by the potential ramifications of privacy violations and their disproportionate impact on marginalized groups. Maris et al. observe that leaked URL data can disclose sexual interests, and note that “Those most likely to be impacted by online sexual privacy violations are traditionally marginalized and vulnerable communities, especially women, people of color, LGBTQ and other marginalized gender/sexuality communities” [73]. Data leaks about sexual preferences that are aligned with normative standards may prove embarrassing, but are not likely to lead to harassment or discrimination. However, such data leaked about sexual interests that are perceived as non-normative or stigmatized may lead to substantial harm.

While marginalized groups use a range of strategies to protect themselves from privacy risks, to resist technology, and to assert their agency [49], they can also face disproportionate challenges in protecting their privacy due to disparities in digital literacy, skills, technological access as well as linguistic and cultural barriers (e.g., [3, 109]).

5.2 Technological and societal exclusion

Researchers also cited technological and societal exclusion as reasons to study marginalized groups. Marginalized groups' access to resources and opportunities are often bounded by their access to (and willingness to use) technologies that can violate their privacy. Technologies can potentially exacerbate social exclusion because of the privacy risks they pose, and in this way, they can reify or accelerate the process of marginalization.

Some papers discussed limitations in various types of technologies, including phones and facial recognition software, with an emphasis on the fact that technology is often designed without consideration for the needs of marginalized groups. Ahmed et al. explain that for disabled persons, "there are also significant differences [in privacy concerns], due in part to their disability but also because of systems that were not designed with them in mind" [3].

Exclusion from the technology design process can create a host of privacy-related challenges for marginalized groups that warrant examination. For example, Dosono et al. uncover several barriers faced by people with visual impairments in using authentication to access websites [28].

Other papers called out structural challenges that marginalized groups face in society that complicate experiences in realms like education and employment that were not always tied to technology use; for example, in the case of children of undocumented immigrants to the United States, disclosures about immigrant status may be necessary for youth seeking support both online and offline, but carry the oppressive threat of deportation [61].

6 HOW DO RESEARCHERS STUDY PRIVACY AND MARGINALIZATION?

The majority of papers (97%) were empirical research (HCI = 54, C = 16, P = 15). Qualitative methods were most common across the board (HCI = 49, C = 12, P = 10), making up 81% of all papers. Here, we provide a field-specific breakdown of methods used in this space.

In HCI, several studies used multiple methods (e.g., interviews as well as focus groups or a survey); here, we present counts for each method used (i.e., some papers were counted more than once). Semi-structured interviews were the most used qualitative research method (45), followed by focus groups (9), participant observation (6); some papers also drew on textual content analyses (3), participatory action research (1), and ethnographic action research (1). Seven HCI studies used quantitative methods, including a survey (1), experiment (1), computational techniques (2), system evaluation (2), and quantified measures to supplement an interview study (1). In addition, seven HCI papers involved the design, development, and/or evaluation of prototypes or systems, such as a sound awareness system for deaf people [59], a monitoring system for dementia care [112], and a phone-based broadcast system to reach urban sex workers [101].

Communication studies in this space were primarily qualitative, using interviews (10), observational fieldwork (2), ethnographic methods (2), and/or focus groups (1). The four remaining empirical Communication studies were quantitative, comprising of 3 surveys and a quantitative content analysis.

Following the trend in HCI and Communication, all of the qualitative studies published in Privacy-focused venues drew on interviews (9). Computational methods were a little more common in Privacy-focused venues (4), such as natural language processing.

Three papers in the dataset were theoretical pieces that were not based on empirical research (HCI = 1, C = 2) (e.g., [96]).

In what follows, we discuss various aspects of researchers' methodological approaches and research practices, in terms of the types of lenses and frameworks they drew on, how they engage with marginalized groups, the ethical considerations of their research, decisions around compensation, and issues of positionality.

6.1 What lenses and frameworks do they draw on?

A minority of papers explicitly employed a critical lens or framework in their research approach. Critical lenses are those that consciously reflect on how social norms and cultures create systems of oppression. In HCI, six papers used some form of critical lens, such as a feminist, queer-Marxist, or intersectional approach (e.g., [16, 99]). In Communication, one paper provided a critical review of how privacy discourse has been historically used to further disenfranchise marginalized individuals along class and racial lines [96]. No studies in Privacy-focused venues employed these approaches.

6.2 How do they engage with marginalized groups?

There is a growing call to engage marginalized groups directly in research involving them, and to empower them by providing opportunities to co-develop and shape research activities, findings, and recommendations. However, building rapport, establishing equitable collaborations, and navigating ethical issues like establishing appropriate consent processes and providing appropriate compensation for research participants means that engagement can be a complex issue for researchers.

We examined the empirical studies in our dataset to identify how many engaged with marginalized groups as study participants. We found that participation or direct input from marginalized groups was fairly common across fields (HCI = 44, 80%; C = 15, 83%; P = 11, 73%).

In many studies, participants from marginalized groups took part in interviews, surveys, and co-design sessions, among other research activities. In some cases, researchers also interviewed other stakeholders, such as familial caregivers and professional care workers. In three studies, the researchers explicitly discussed difficulties in recruiting their population of interest due to cognitive impairments that compromised both the consent process and the nature of responses; in these cases, these difficulties were included in the researchers' rationale to interview other stakeholders instead of the marginalized group themselves.

6.3 How do researchers discuss the ethical considerations of this work?

Questions of research ethics have become increasingly complicated in HCI research given the wide range of methods at researchers' disposal [10]. At the same time, research ethics are a crucial consideration in studies involving marginalized groups. Research on sensitive topics can also engender challenges around "situational ethics"—ethical considerations that emerge in the process of conducting research that depart from formal procedural ethics requirements [83]. Munteanu et al. provide examples of situational ethics, such as discovering that some participants are not comfortable with sharing in focus groups, and altering the method to one-on-one interviews instead to respect participants' needs [83]. In this section, we explore reporting practices about ethical considerations.

As shown in Table 2, papers varied widely in whether and how they discussed ethical considerations. Almost half (49%) discussed research ethics in some form—either in a dedicated section or integrated within a broader Methods section. However, this choice varied sharply by field, with

	HCI	Communication	Privacy	Total
Includes ethics discussion	30	3	10	43
No ethics discussion	25	15	5	45

Table 2. Papers that discuss ethical considerations

most computing papers including such discussions (HCI = 30, 55%, P = 10, 67%) in contrast to Communication (C = 3, 17%).

Research ethics discussions ranged from cursory to extensively detailed. Researchers most commonly reported anonymizing data (e.g., by removing participant identifiers) and obtaining consent from participants; papers that had short ethics discussions invariably only reported these procedures. Some papers reported more extensively on planned or emergent ethics-related procedures. Although some procedures were uncommon, we discuss some key themes around privacy, trust, and support to illustrate the breadth of approaches to ethics in this research space.

6.3.1 Protecting Privacy. Some papers went beyond anonymizing participant identifiers to protect participant privacy, such as by paraphrasing quotes taken from online data sources to reduce their searchability [60] or not recording audio during interviews [101]. One study took the uncommon approach of not asking for demographic information from participants, instead choosing to use existing demographic data from participants' neighborhoods as a proxy [110].

6.3.2 Building Trust. Some studies involved developing trust with communities and/or community allies and soliciting their guidance during the research process. Some researchers reported developing relationships with community members over significant periods of time by attending gatherings and volunteering (e.g. [52]). Others consulted with community allies or non-profit organizations to determine appropriate compensation [38] and how to recruit their participants in a privacy-protective way [45]. Another study used member checking practices, where the researchers shared a draft of their paper with participants for feedback to ensure they had accurately represented their experiences and to correct any misunderstandings [52].

6.3.3 Prioritizing Support. In some studies, researchers described how they prioritized the dignity of participants and supported them during the research process, particularly in the case of interviews about sensitive experiences. These research choices included stopping recording interviews when participants became emotional [100], having a therapist present for interviews with survivors of intimate partner abuse [68], and appointing same-gender interviewers for women who have faced abuse [99]. Beyond participants, some papers discussed the steps taken to also support the research teams themselves because of exposure to difficult interview content [51].

6.4 How do researchers navigate compensation with marginalized participants?

Carefully navigating participant compensation is particularly important when studying marginalized groups, many of whom may be economically disenfranchised or otherwise vulnerable. While providing participants with compensation recognizes and respects their time and effort, financial incentives may also coerce people into participating in the research out of economic need [31]. Thus, whether and how to compensate participants requires careful consideration and review by institutional ethics boards on a case-by-case basis [89]. Reporting and explaining decisions around compensation can help readers evaluate the research ethics of any given study (e.g., the potential for undue influence) and its methodology (e.g., the impact of compensation on recruitment) [66].

In our dataset, fifteen papers did not involve marginalized populations directly (e.g., content analyses), and thus the question of compensation was irrelevant (HCI = 7, C = 4, P = 4). The remaining 73 papers involved some form of research activity with participants, and we examined these to understand trends and rationales around compensation; an overview is in Table 3.

Overall, 52% (38) of the 73 user studies reported that participants were monetarily compensated (HCI = 25, C = 5, P = 8). Compensation ranged from \$5 to \$100 USD in cash, gift cards, or vouchers, and was typically provided for interviews. In a few studies, researchers determined appropriate compensation by consulting with expert organizations and communities (e.g., [38]), which can be a

	HCI	Communication	Privacy	Total
Provided monetary compensation	25	5	8	38
Provided small gifts or snacks	0	3	0	3
Explicitly did not provide compensation	2	0	0	2
Did not report compensation	21	6	3	30

Table 3. Compensation practices for the 73 studies involving participants

useful strategy when studying marginalized contexts [31]. In contrast, 41% (30) of the 73 papers provided no information about compensation (HCI = 21, C = 6, P = 3). In addition, two HCI papers stated that they did not compensate participants, and three Communication papers stated that the researchers provided snacks and small gifts in lieu of financial compensation.

6.5 How do researchers navigate positionality?

Researchers' positionality—their positions in society based on identity factors such as class, gender, and race, among others—inevitably influence the research process in several ways, from the types of questions that are asked to how the work is completed [30, 88]. Considering positionality is particularly important when conducting research with marginalized groups to ensure the research does not perpetuate the same marginalization that researchers seek to understand [88]. Positionality can also have implications for researchers' well-being. Insider research—when researchers belong to the same group as those being studied—in marginalized contexts can pose emotional risks for researchers, making self-care a critical part of such research [106].

Thus, we examined the papers in our dataset to identify whether and how researchers were engaging with issues of positionality. Overall, positionality statements were uncommon, with only 23% of papers reporting positionality across the dataset (HCI = 14, C = 6, P = 0). Positionality statements typically involved one or more coauthors disclosing that they identified as belonging to the marginalized population that was the focus of the study. It was comparatively rare for authors to use these statements to discuss their relative privilege or non-membership in the groups being studied. Overall, positionality statements were brief and served to communicate authors' identities (either marginalized or not), and most papers did not discuss how these positionality might inform or impact the research. Some studies in Communication used a less common practice that is worth noting. In these studies, when the authors did not belong to the population being studied, they enlisted population members to help inform the research process in terms of recruitment, feedback on questions, interviewing, and building rapport with the community.

7 WHAT ARE KEY THEMES IN THE FINDINGS?

In what follows, we discuss broad themes that reflect the ways in which marginalization, technology, and privacy were interrelated in the corpus of papers. Central to this analysis are the many tensions that arise when people choose responses to perceived privacy threats, the high-level tradeoffs involved in these choices, and the specific costs and consequences of different responses to threats.

7.1 Responses to Privacy Threats and their Costs and Consequences to Marginalized People

Many existing privacy frameworks overlook the impact of marginalization on people's behaviors [67]. In table 4, we present a *Privacy Responses and Costs* Framework that lists the ten main types of responses to privacy threats that we observed in the literature and the costs and consequences they

Privacy Response	Cost/Consequence	Select Examples from the Dataset
Apathy i.e. lack of response	• Exposure to risks	Undocumented immigrants felt government surveillance is inescapable, leading to inaction [45]; women transitioning from incarceration felt they have “nothing to lose” [105].
Non-use e.g. not using a technology, deleting an account	• Opportunity loss • Exclusion • Silencing • Isolation	Economically disadvantaged populations lose opportunities due to non-use of technologies [110].
Withholding disclosure e.g. self-censorship, information removal	• Restricts self-expression • Silencing	Low-SES youth [74], marginalized Cambodians [58], and political refugees in the U.S. [107] self-censored to avoid conflict and danger but were further silenced by doing so. Men who have sex with men may not disclose HIV status on dating apps but can inadvertently signal positive status [116].
Controlling disclosure e.g. compartmentalizing identity, multiple accounts, privacy controls, segmenting audiences	• Restricts self-expression • Labor-intensive • Social cost • Financial cost	Trans men crowdfunding top surgery used privacy controls to limit audiences [39], young Azerbaijanis maintained multiple accounts for political activism [90], and LGBTQ+ social media users managed identities across platforms [24]. Disclosure controls require extensive labor and restrict self-expression [8]. Complex privacy controls can be costly to access [95] and can be used incorrectly due to accessibility issues [3].
Privacy lies [102] i.e. providing false information	• Cognitive burden • Social/legal repercussions	Rural Appalachians provided false information as a form of vigilanteism [49]; South Asian women provided false information to protect themselves from online abuse [99].
Privacy-enhancing technologies (PETS) e.g., authentication, cloaking, encryption	• Social liability • Erasure of records	Women in patriarchal societies used private modes and locks on devices, which may be seen as incriminating and invite coercion to obtain access [100]. People who are financially insecure who lose access to trusted devices lose access to services that require two-factor authentication [108].
Physical workarounds e.g. hiding device, use of camera covers & headphones	• Limits environmental awareness • Vulnerable to physical coercion	People with visual impairments used headphones to avoid aural eavesdropping when using screen readers at the cost of physical safety [3].
Asking for help e.g. learning new practices, consulting network, websites, professionals	• Bad information • Involves risk/trust • Limited to help available	Professionals who provide support for survivors of intimate partner violence did not feel equipped to advise on identifying or coping with technology-enabled IPV [38]. People with visual impairments asked allies for help, but this risked trusting the ally with personal information [52].
Collaborative privacy practices e.g. shared guidelines, boundaries	• Loss of autonomy • Involves risk/trust	LGBTQ+ adults considered not only their own privacy boundaries but also those of their families, ex-partners, and children [8]. Families co-developed privacy guidelines for shared devices in Bangladesh [2].
Third-party protections e.g. parents removing devices, organizations destroying info	• Loss of autonomy • Outside of the person’s control	Art therapists removed identifying information from art created by persons with dementia to protect their privacy, but this also removed their voice [19]. Canadian government’s legal decision to destroy data documenting colonial abuses of indigenous people to protect their personal privacy also erased evidence of their abuse [40].

Table 4. Privacy Responses and Costs Framework: Types of responses to privacy threats observed in our dataset and their costs and consequences to marginalized people.

carry for marginalized people. Research on the privacy calculus has established that people weigh costs and benefits as they make choices about privacy [25]; our framework provides an account of costs and consequences involved in privacy decisions for marginalized groups. It is also noteworthy that weighing tradeoffs and making privacy decisions is challenging for people who are under stress [68, 75, 108], adding a new dimension of complexity to the privacy calculus. Sannon and Cosley characterize privacy management as a costly form of invisible labor that has an outsized impact on marginalized groups or those without power [103]. In what follows, we discuss each of these privacy responses and their associated costs in turn.

7.1.1 Apathy. In some cases, people can feel helpless and may do nothing in the face of privacy and security threats, a phenomena McDonald describes as “a sense of futility masquerading as apathy” [77]. For example, when exposed to online scams, people from economically disadvantaged communities may respond with feelings of resignation [110]. Similarly, undocumented immigrants describe the government as an all-knowing entity whose surveillance is impossible to evade, leading to inaction [45]. The major cost associated with this response is continued exposure to the disproportionate harms faced by marginalized groups.

7.1.2 Non-use. When faced with poor design and privacy threats, marginalized people may resist using technologies altogether, which can exacerbate the problem of exclusion. Economically disadvantaged populations in particular appeared to resist technology use in our dataset, which could limit their opportunities for jobs and social support [49, 56, 108]. For example, Vitak et al. described economically disadvantaged participants who were “hesitant to use technology or outright shunned it, preferring to use analog methods for submitting applications, forms, and payments whenever possible—even when that decision carried additional financial costs or took longer,” citing a participant who would not apply for jobs that required online applications [110].

7.1.3 Withholding disclosure. Self-censorship is a commonly discussed strategy for managing perceived risk, particularly of views that may give rise to interpersonal conflict [74] or political speech that may have negative repercussions [58, 107]. Marwick et al. observe that “choosing to self-censor and limit one’s participation is a choice to be rendered invisible” [74]. Self-censorship as a response to privacy threats can further silence marginalized voices. Moreover, withholding information may not be effective, with damaging consequences. Warner et al. describe how withholding HIV status can be ineffective, as people can also make inferences about undisclosed information [116, 117]. Efforts to withhold disclosure may also be fruitless in the face of intentional adversarial threats like the use of stylometry to identify content creators [35].

7.1.4 Controlling disclosure. Social media studies made up 26% of the dataset, and a common privacy-protective behavior, especially on social media, centered around people making efforts to compartmentalize their identity or choose their self-presentation in ways that would allow them to pursue their goals while shielding them from risk. Some ways to do this are to use privacy controls to restrict access to one’s information, to open multiple accounts on the same platform to keep various facets of one’s identity separate [46], and to select platforms based on the degree of privacy they afford [24, 62]. However, extensively locking down one’s online profile also means that one’s ability to express one’s identity is restricted [24]. The labor involved in controlling disclosure is also considerable. Fritz and Gonzales describe how one trans participant who was crowdfunding surgery “took an entire day to go through all his Facebook connections and block approximately 600 people who were connected to his family so they would not see when he promoted his fundraiser on Facebook” [39]. Pearce et al. describe the labor Azerbaijani young adults expend to segment audiences for political posts: “those with two profiles engaged in a great deal of labor to manage the two—defriending people on one, adding them to the other, inventing innocuous reasons why

a new profile was created, and so on” [90], and additionally describe the social cost of dissolving online ties to control disclosure. In a study of South African mobile privacy practices, Reichel et al. explain that using of privacy controls carried a financial cost associated with connectivity: “Nearly every time a participant expressed awareness of these privacy settings, they followed by explaining they had an inability to actually access the privacy settings, often mentioning data costs” [95].

7.1.5 Privacy Lies. Providing false personal information about oneself by telling “privacy lies” [102] appeared several times as a privacy strategy to deter threats. Sambasivan et al. describe the use of fake information as a strategy for South Asian women to protect themselves from abuse online [99]. Hamby et al. frame providing false information as a kind of privacy vigilanteism in rural Appalachia [49]. However, telling privacy lies requires cognitive effort to ensure the lies are not found out, and can also be risky, as being caught out in a lie can result in social or legal repercussions [102].

7.1.6 Privacy-enhancing technologies (PETS). Papers in our dataset reported on a limited range of technological responses by marginalized groups to privacy threats. These largely involved private modes, encryption, and locking/adding additional authentication requirements. Despite the protections they offer, the use of PETs may incriminate the very people who need protection. Sambasivan et al. report in a study of women in India, Pakistan, and Bangladesh that “private modes are often associated with ‘secret’ activities, threatening participants’ values of openness as they performed culturally appropriate gender roles” [100]. Similarly Ahmed et al. observed in a study of shared mobile phone use in Bangladesh that “almost half of our participants reported that locking specific data or applications might also raise suspicions in the mind of their partners” [2]. Locking down devices and applications may also result in further coercion or physical harm in the context of intimate partner violence [37, 68]. Because of this cost, hiding the fact that privacy protections exist was surfaced as a design recommendation in multiple papers. Naseem et al. quoted a participant in Pakistan who applauded the design of secret PETs, “That way, at least one won’t come across as suspicious, especially since men in our society are very distrustful and suspicious” [84]. However, leaving no trace of activity carries additional consequences for people in abusive relationships in that it eliminates documentation of abusive behaviors [92]. Additionally, the use of two-factor authentication can contribute to a victim’s powerlessness if the victim loses access to devices required to authenticate, possibly due to interference by the abuser [108].

7.1.7 Physical workarounds. Hiding devices, covering cameras, using headphones to prevent others from overhearing screenreaders—these are all documented physical workarounds that members of marginalized groups have to privacy threats. Ahmed notes that many visually impaired participants used headphones to ensure privacy when using screenreaders; however, “since visually impaired people rely on hearing in order to sense the environment, headphones leave them more vulnerable to other privacy and safety concerns” [3]; thus the headphones might protect their information privacy at a cost to their physical safety. People in abusive relationships may also physically hide their devices to avoid unwanted snooping [75], but hiding devices comes with many of the same threats of physical or emotional coercion described in the above section on PETs.

7.1.8 Asking for help. Asking for help may involve learning new privacy practices from social contacts, or consulting privacy resources and professionals. When people rely on their social network, a practice found to be more common among lower SES individuals [94], the help they receive is only as good as the knowledge in their network. Poor advice can carry a high cost for those in vulnerable positions and even trusted professionals may not be sure what to advise. For example, professionals who support survivors of intimate partner violence don’t always know

what to advise [68]. Receiving help from others can also entail divulging sensitive information to third parties, which requires trust and can introduce a new privacy risk [79, 107, 110].

7.1.9 Collaborative privacy practices. Collaborative practices sometimes arise as a response to concerns about privacy threats within family units or other relationships. Practices like establishing shared rules and boundaries [2] provide a social alternative to PETs like app locks and encryption, but require high levels of trust and may compromise autonomous decision-making.

7.1.10 Third-party protections. In literature about marginalized populations with heightened vulnerabilities, it is unsurprising to find examples of third parties with power making decisions that impact marginalized groups. Examples in our dataset range from husbands making decisions for wives in patriarchal societies [62], governments taking action on behalf of vulnerable groups [40], and families and therapists taking responsibility for adults with cognitive decline [19, 82]. Costs of third-party protections include a loss of autonomy and control; when others make privacy decisions, the line between helpful and paternalistic can be difficult to see. Mantis et al. explore this tension in work with adults with cognitive decline and their partners [79]; they find that although partners aim to negotiate security decisions, the reality is often that decision-making ends up being one-sided by the caregiver partner.

7.2 High-level privacy-related tensions and trade-offs for marginalized technology users.

In addition to the granular costs and consequences described above, we identified several high-level privacy-related tensions in studies of marginalized groups' technology use:

7.2.1 Privacy vs. disclosure of identity. Most uses of technology entail some degree of identity disclosure, ranging from highly visible disclosures like creating a personal account with legal identifiers, to less visible disclosures like inadvertently sharing location via network data. For marginalized groups, the privacy threats associated with everyday use may be acute, even if the groups themselves are not cognizant of the threats [45]. In some cases, technologies create unique avenues for sharing among marginalized groups that come with elevated threats. For example, LGBTQ social media users may use platforms to connect and explore their identity [47] but also risk stigma from unintended audiences if their privacy strategies are compromised [24]. Marginalized groups engaging in activism [69] or participating in online collaborations [35] may disclose identity characteristics by virtue of these activities. In other cases, privacy breaches can be the result of secondary data use, for example by analyzing meetup data to infer LGBTQ identity [15] or social media posts to infer mental health status [118].

7.2.2 Privacy vs. support. Because marginalized groups by definition experience some form of stress, support seeking is particularly salient. Risking disclosure of vulnerabilities to receive support is not unique to technology-mediated contexts (e.g. [45]), but technologies can exacerbate the privacy threats associated with support seeking. Srinivasan et al. [109] observe in their research on identity infrastructures in India that low income and marginalized people often “are the people who most need benefits from the state, and to receive benefits, they must identify themselves,” which can introduce risk of stigma and persecution. In a different context, strikingly similar trade-offs arise: for trans men who crowdfund financial support for top surgery, support-seeking can entail highly public disclosures [39, 42]. Sambasivan et al. investigated the use of a phone messaging system intended to help deploy social and health services to urban sex workers who are vulnerable to physical violence and health issues, but for whom receiving such support could mean stigma or further physical threat [101].

7.2.3 Privacy vs. autonomy. Privacy is an important way of protecting individuals' freedoms (see [119]). As such, oversight of others' location, well-being, or activities is often paternalistic or invasive, but some of the research in our dataset discussed scenarios in which forms of surveillance were used to facilitate autonomy. One example of the challenges of navigating privacy and autonomy arises in the work of victim service providers (VSPs). These organizations support survivors of human trafficking. In order to help ensure that survivors of human trafficking are not revictimized and in a direct bid to protect survivors' freedoms, VSP shelters may surveil communications, particularly of minors, to enforce rules and monitor for risky behaviors [14]. Similarly, Mentis et al. identified threats, tradeoffs, and design considerations with respect to privacy of older adults with cognitive decline and their caretakers [80]; they note that while collaboration between adults with cognitive decline and caregivers is critical, supporting collaboration entails privacy concessions that can create opportunities for exploitation. Similarly, location tracking can be viewed as a safety measure that supports adults with dementia in retaining more freedom [82], but it can also be subject to potential abuse [21]. This handful of papers highlight important boundaries where social norms and cultural standards around the limits of privacy and acceptable privacy trade-offs for safety and wellbeing are negotiated.

7.2.4 Individual vs. collective. Although privacy theorists often weigh the interests of the individual against those of states and organizations (see [9, 119] for foundational examples in Western culture), in some cases, privacy concerns are collective. That is, more than one person may work together to protect the shared privacy interests of a group. The literature on marginalization and privacy included several examples of such framing. For example, LGBTQ+ parents on social media find that their privacy depends not only on their own disclosure decisions but also the disclosures and behaviors of their network. Further, their self-disclosures can also impact the privacy boundaries of those who make up their network, such as their children, partners, and former partners, and potentially expose the network to stigmatization as well [8]. In some cases, additional help may be needed in navigating privacy—for example, in the context of patients with dementia, Cornejo et al. describe how negotiating privacy is a process that is shared between the patients, family members, and care professionals [19]. People's privacy preferences may also be "socially-negotiated" rather than purely their own, as with people managing bipolar disorder whose family members and care team who would like them to regularly share their personal data in the form of a "check-in" [91].

7.3 Technologies and the way they handle privacy can either contribute to or impede the process of marginalization.

We identified two common narratives across the literature that were connected to the kind of mediating role that technologies can have as the relationship between privacy and marginalization plays out. On one hand, technologies can be viewed as a mechanism for greater equity and freedom, slowing or perhaps optimistically even reversing the process of marginalization for certain groups, such as technologies that provide assistive support [113] or safe spaces for disclosure and support seeking (e.g., [13, 54, 84]).

On the other hand, technology can be viewed as a mechanism that supports and strengthens processes of marginalization by disproportionately introducing privacy threats and harms or furthering exclusion. In some cases, the design of technologies is insufficient to protect the privacy interests of marginalized groups. For example, insufficient control over disclosures of HIV-positive status in dating apps can lead to "privacy unraveling"—particularly for men who have sex with men—which can further exacerbate their marginalized status [116]. Being from a marginalized group means that technologies may be designed in ways that exclude people from their use at all, as Rennie et al. explain, "social dynamics and obligations can prevent Aboriginal people from using

devices and settings in the way they were intended” [97]. Being prevented from using technology by virtue of poor design is yet another form of exclusion that marginalized groups experience. In some cases, poor design or the threat of privacy breaches may lead to non-use or resistance [49] which, while sometimes framed as a form of agency and empowerment, can also result in exclusion from digital public spaces [35]. Reichel goes further to examine how the very concept of privacy as a goal is embedded within and therefore yields social and technological structures that exclude marginalized groups [96] and Gangadharan [41] notes that inclusion efforts intended to assist marginalized groups may expose them to greater threats of surveillance and risk.

8 WHAT ARE KEY THEMES IN THE PAPERS’ RECOMMENDATIONS?

We organized the recommendations in the papers into three broad categories—conceptual, technological, and behavioral recommendations—that we discuss next.

8.1 Conceptual recommendations

8.1.1 Prioritizing autonomy and dignity in design. Several authors recommend rethinking the ways we talk about privacy and the concepts we use to understand and design privacy for marginalized groups. Akter et al. call for “humanizing” assistive technologies specifically because they found that “camera-based assistive systems were creating a lack of security in people’s daily lives—that is, these systems were serving to further marginalize their identities” [4]. Reichel rejects the concept of privacy, seeing it as inherently flawed in that it requires identification of a threatening “other,” and advocates instead to organize what is now privacy discourse around the more fraternal concept of dignity [96]. Fullenwieder and Molnar [40] similarly critique the notion of privacy as an individual right in their analysis of how privacy was leveraged to support the destruction of testimony about state-sanctioned violence against indigenous populations in Canada.

Designing technologies that do not threaten the privacy of marginalized groups also involves assessing which values are given importance in the design process, and considering how these values might impact marginalized groups. Values can be in conflict: for example, remote monitoring systems for people with dementia can prioritize the value of keeping them safe over preserving their privacy. In light of such conflicts, Dahl and Holbø recommend that value elicitation should be a necessary part of designing in such settings to evaluate technological biases and impacts [21]. Similarly, Wan et al. argue that flexibility should be built into technologies so that they can be adapted to fit potentially diverse needs in different organizational and family contexts [112].

Prioritizing people’s agency is one way of ensuring technologies do not harm marginalized groups. For example, automatic gender recognition software can engender numerous harms for transgender people; Hamidi et al. caution designers to consider whether gendering users is truly necessary and to exclude gender from their design if possible. In cases where such technologies are used, they recommend providing users with the agency to opt out of being gendered [50].

8.1.2 Recognizing the influence of power relations on technology use. Studying marginalization means studying power. As part of their focus on marginalized groups, many researchers call for designers “to consider the ways that technologies may not be one-size-fits-all” [33]. These researchers highlight the need for designers to pay more attention to how power relations structure the ways marginalized groups use technology and to tailor technologies accordingly. For example, locking a mobile phone may not be effective in a heavily patriarchal context where a woman can be coerced by her husband to unlock it [2]. Similar dynamics may occur due to familial and sociocultural power relations, especially in contexts where it is the norm for multiple people to share the same device [101], such as in South Asia [99]. Conducting digital risk assessments for technologies would also identify privacy risks and gaps in privacy protection based on different use

cases, such as assessing if certain technologies pose risks to survivors of intimate partner violence, given their unique threat model where they are avoiding a known other [38].

8.2 Technological recommendations

Most recommendations, particularly in HCI, centered around how technologies can be designed to better address the unique privacy needs and concerns of marginalized groups.

8.2.1 Providing greater control over information. Since many marginalized populations can face heightened risks from having their privacy violated, many studies stressed the need to afford people with greater control over their information (e.g., [13, 33, 42, 46]). To this end, a common design suggestion was that technologies should have granular privacy settings so that users can better control the visibility of their content, thereby avoiding the risks of context collapse. For example, Carrasco et al. discuss how LGBTQ+ social media users practice “selective visibility” by being more “out” to LGBTQ+ audiences but not cisgender or heterosexual audiences; they suggest that social media platforms that facilitate this form of selective sharing—for example, through supporting the use of multiple profiles—would give marginalized social media users more agency over their self-presentation (and by extension, their privacy and safety) [13].

Designing technologies that do not diminish the privacy of marginalized groups also requires designers to consider the “labor and risk involved in conveying sensitive information about the self to others”: for example, some groups—such as trans individuals—may not want to be seen by the broadest audience possible when using online dating platforms, as this could open them up to harassment, and instead would benefit from controlling the visibility of their profiles [33].

8.2.2 Facilitating management of communal and networked aspects of privacy. Several papers stressed the need for designers to consider communal aspects of privacy. In addition to managing their personal privacy boundaries, people often have to navigate collective boundaries as well, and granular privacy controls on online platforms would help such networked privacy management [8]. In cases where multiple stakeholders jointly negotiate privacy, such as in the case of people with dementia and caregivers, systems could support cooperatively setting privacy preferences [19].

People can also face privacy threats from others who have access to their devices, either with or without their consent. In such cases, a common design recommendation is to facilitate secrecy and the hiding of sensitive information. For example, mobile phones could be set up to allow an individual to hold multiple accounts that are kept secret if others access the device [2]. Similarly, enabling multiple users to create individual accounts on shared devices also helps maintain each individual’s privacy [68]. Devices could also enable on-demand information “hiding”, a feature that would be especially useful for groups who could face severe consequences if their information were accessed by the wrong parties, as in the case of undocumented Latinx immigrants [45].

8.2.3 Making privacy decisions easier. Another type of recommendation centered around making it easier for marginalized groups—and people in general—to make informed privacy-related decisions. One way of doing this is to improve the usability of privacy and security features and settings, including when people are under dire stress, as in the case of avoiding intimate partner abuse [75]. Another option is to design for transparency, so that they have more clarity around when their information is being collected and how it might be used. Sometimes this transparency is important with respect to how data will be used by social media companies and service providers [11], but also with respect to how information is shared with other parties. For example, in the context of an app that helps people with bipolar disorder continuously disclose mental health information to trusted others, managing this continuous sharing in a sensitive way requires interfaces to clearly indicate who can see what data at any given time [91].

8.2.4 Building in technical safeguards. Whereas many of the above-discussed technical recommendations hinged on user experience and interface features, several papers also discussed how technical infrastructures could be better designed to safeguard marginalized populations. For example, in the call for humanizing assistive technologies, Akter et al. [4] note that computer vision algorithms must be designed to detect not only objects, but features of context that matter to the humans who use them. Systems that provide infrastructure for anonymous online activity can protect people with marginalized identities in a variety of contexts, for example, survivors of abuse who wish to report abusers [14].

8.3 Behavioral recommendations

Some researchers stressed the need to help people engage in behaviors that keep them safe and align with their privacy needs. One way to do this is to provide people with education and resources. Access to privacy-related education can be particularly useful for groups such as undocumented immigrants, who may not be aware of the intricacies of the technological privacy and security threats they face [45]. Educating people about how to protect their privacy can help them continue to use technologies while mitigating the risks in doing so, as in the case of contributors to open collaboration projects like Wikipedia who have marginalized attributes [35]. In both these studies, the authors stress the need for such social interventions to occur in conjunction with technological solutions that focus on improving the technologies themselves.

In some cases, marginalized groups may be well aware of the fact that they face privacy and security threats, but feel limited in their ability to address these threats. Research on intimate partner abuse suggests that survivors need targeted instructional materials that will help them avoid their abusers, such as information on how to use security features like two-factor authentication [75]. In response to this need for access to resources, two studies in our dataset explored the usefulness of providing survivors of intimate partner violence with security consultations with a trained technologist, finding that these consultations were generally perceived as valuable and also uncovered security vulnerabilities that participants were not aware of [36, 51].

Despite multiple studies pointing to a need for improved access to privacy education, a challenge in improving privacy education among marginalized groups is that digital literacy programs are not always well-attended; in response to this challenge, Vitak et al. point out that stories are often an effective means of spreading information in low-income communities and could be a way of sharing privacy-related knowledge and resources [110]. Reichel et al. suggest that “lightweight privacy on-boarding interfaces” could help resource-constrained users, including making privacy settings available offline for those who have limited connectivity [95]. People may also be able to improve their digital privacy skills if they have unrestricted, private access to the Internet, but many disadvantaged communities rely on libraries and schools to access the Internet where their usage is time-limited; thus, addressing the digital divide remains crucial to increasing both the autonomy and privacy literacy of marginalized groups [70].

9 WHERE DO WE GO NEXT? CHARTING FUTURE DIRECTIONS

Our initial dataset included 2,823 privacy-related papers that were published between 2010 and 2020; of these, only 3% (88) focused on marginalized contexts. Although our dataset does not include all venues where such work might appear, this highlights how profoundly understudied this area of research remains. We also found that research in this area is growing over time, and we end this paper by discussing some potential avenues for future research in terms of focus, topic area, methods, and research practices.

9.1 Broaden How We Problematize Privacy and Marginalization

Although we identified several common themes that cut across studies, papers in our dataset did not clearly converge on a shared articulation of problems or solutions. We note that, although all the papers examined marginalized groups that are, by definition, marginalized as a function of social norms and structural inequities, only a handful of papers problematized privacy in the context of social structures, policies and laws. We are confident that nearly all researchers doing work on privacy at the margins understand that technologies, policies, and social structures intersect and inform one another; however, efforts to examine these intersections were rarely central features of papers. In contrast, empirical studies, which were the most common type of contribution in our dataset, tended to focus on people's experiences and behavior. These studies mark an important shift towards representing the voices of users who have been traditionally excluded from technology research. Alongside these much-needed contributions, we suggest that researchers leverage and build on critical lenses that highlight structural and systemic aspects of marginalization that underlie technology design and use. Correcting inequities in privacy is also a problem with potential policy implications. Our dataset yielded few if any policy recommendations, and we see a need for more interdisciplinary work that bridges technology design and policy.

9.2 Fill Understudied Research Gaps

While research on marginalization and privacy is growing rapidly, particularly in the past three years, this growth has not been uniform across research areas. As we coded our dataset of existing literature, it was quickly apparent that some research topics at the intersection of privacy and marginalization remain undeveloped.

The paucity of discussion around race and privacy in our dataset was conspicuous. Papers that examined issues of race and power as their central focus were few, and mainly within the Communication dataset [40, 96]. Additionally, some papers looked at contexts in which race was inextricable from the context under study, such as refugees and other immigrants, but for the most part, these did not explicitly reflect on the influence of race in the experience of marginalization. Additionally, race was sometimes mentioned as a factor to consider in contexts like crime prevention or low-income neighborhoods, but was not used as a lens through which to understand privacy concerns or as an analytical tool in empirical work. This is surprising for HCI in particular, given a trend toward considering technologies as potential instruments of oppression and marginalization and increasing attention to conceptual frames that address features of race, like critical race theory [87] and—at times controversially—intersectionality [93]. Although such frames were at times mentioned in discussion, they were generally absent from the body of empirical work we reviewed. To fill this important gap in the literature, we see a dire need for privacy research—particularly within HCI and Privacy-focused venues—that includes and centers race.

It is worth noting that this is a new subfield; even the most common topics are still relatively understudied and remain rich areas for further exploration. This is also true within various topic areas. For example, disability and LGBTQ+ issues were the two most common topics and together made up almost half of our dataset, but certain subpopulations were less studied than others, such as people with mental health conditions ($n = 3$) and queer women ($n = 1$), respectively. As such, our review illustrates the current breadth of work in this space but also shows that there is much room for future work across topics.

In order to fill these gaps, we see an immediate opportunity for conferences and journals to foster research in this area. Scholarly leadership can support these efforts by mentioning topics related to marginalization in CfPs, dedicating special issues, panels, and workshops, identifying keynotes and themes. Moreover, program committee members and reviewers must recognize the value of

studying marginalized groups independent of comparisons to or in the interests of designing for the majority population.

9.3 Diversify Methods and Questions

Qualitative methods—particularly semi-structured interviews—were by far the most common methods used by researchers. This makes sense given that the goal of much empirical work thus far has been to understand the privacy-related needs and experiences of marginalized user groups. While these methods are likely common because they are the most appropriate for the types of questions the community has been asking, our findings indicate the potential to leverage other methodologies to complement and build on existing research questions. Within the umbrella of qualitative research, ethnographies, textual analyses, and case studies are all examples of methods that are under-represented in this space but may still be well-suited to the questions being posed. There is also potential to leverage quantitative methods to a greater degree; for example, to measure privacy inequities and disparities, or to experimentally test technology designs intended to mitigate these disparities. Further, a minority of papers used participatory methods to co-design privacy-aware technologies with users, and the privacy community may benefit from greater use of these methods, particularly as mechanisms by which to involve marginalized people in design. While it is important for researchers to acknowledge that new tools may not be wanted or needed by the individuals they intend to serve, the few participatory studies in our dataset suggest that there is potential to work in conjunction with marginalized communities to co-develop tools or re-design existing technologies to better meet their needs and practices. To do so, we can also look to HCI research on other marginalized contexts, such as participatory action research in under-resourced neighborhoods [55], for guidance.

9.4 Develop Shared Best Practices

We found wide variation in both research and reporting practices across studies, such as the amount of detail provided in methods sections and decisions around ethical considerations, such as compensation. Particularly in the context of studying marginalization, we see an urgent need to consider, justify, and report study details, ranging from why a given recruitment strategy was chosen to potential harms from the research and how these were mitigated. These are challenging aspects of study design, and reporting the decision process underlying these research choices is important for research communities to engage in discussion/critique and to develop shared practices and ethical standards.

Here, we discuss some research considerations that researchers need to reflect on when crafting and conducting our studies and report on when publishing our papers as a first step towards developing shared best practices for studying privacy and marginalization.

9.4.1 Plan for and reduce a wide range of potential harms. Reducing harm and maximizing benefits are central tenets of international standards for ethical research [17, 34]. Categories of “vulnerable” populations highlighted by bodies like Institutional Review Boards (IRBs)—for example, prisoners, pregnant women, children—do not cover all possible at-risk groups. Marginalization by definition creates vulnerabilities. When working with marginalized groups, researchers need to take particular care to understand the potential risks involved in their participation. Whereas a large portion of our data corpus was focused on identifying privacy risks, these papers rarely discussed the privacy risks and precautions related to study design beyond anonymizing data or using pseudonyms. According to a review of anonymization practices at CHI, this is not atypical [1]. Moreover, some study designs in our corpus may have introduced novel risks to the population, for example through automated detection of marginalized identity features, without a robust discussion of the ethical implications

of such research endeavors. As a research community it is important not only to carefully consider, but also unambiguously report our considerations of potential harms or unintended consequences and how these risks are mitigated in study designs.

Another important consideration to reduce harms in research is in the language that researchers choose to represent marginalized people and experiences. Language that others, diminishes, or inadvertently stigmatizes marginalized groups can introduce harms through the research itself. Researchers should take care to learn and discuss the preferences of the groups involved in research and report if and how specific choices were made around labels and terminology where controversy or disagreement exist.

Other strategies used to mitigate harm in our dataset included having a therapist on hand to assist if interview participants experienced distress [68], along with other strategies already discussed in Section 6.3. Researchers who conduct interviews on sensitive topics or with groups facing adversity may also find Kasket's [63] protocol for responding to participant distress useful.

Finally, while studying marginalization, it is also worth considering potential harms to researchers themselves. Researchers can face burnout as well as physical safety risks when conducting research on sensitive topics [31]. A handful of papers in our dataset discussed the steps that were taken to ensure researcher well-being, and these discussions were useful to understand some of the challenges that may surface when conducting research in this space, as well as how these risks might impact the research process and the scope of the findings. For example, a study in our dataset that involved conducting interviews in unsafe neighborhoods noted that the interviews were strictly time-bound to ensure the safety of the researchers, sometimes at the cost of asking additional follow-up questions [95].

9.4.2 Report compensation. A review of HCI papers published by the ACM found that the majority of user studies did not report whether participants were compensated [89]. Our dataset suggests that researchers working at the intersection of privacy and marginalization may be more likely to report this, but still, about half of the user studies in our dataset did not include this information.

We echo Pater et al.'s call for researchers to report this information, including their rationales. When providing compensation, researchers could consider whether there are implicit biases in their compensation choices—for example, not all participants may want an Amazon gift card [89], and in the context of studying marginalized groups, there may be differences in access that make some forms of compensation more appropriate or desirable than others. A few studies in our dataset described consulting with community members or partners to decide on appropriate compensation, which is a practice that could be particularly well suited for researchers in this space to adopt.

9.4.3 Discuss positionality. Positionality is an important feature of research—researchers' identities, beliefs, experiences, and backgrounds directly inform their selection of methods, engagement with participants in the case of human-subjects research, and interpretations of data. How to go about reporting positionality is a complex decision—for example, positionality statements that involve identity self-disclosures may disproportionately harm marginalized researchers. We do not make prescriptive statements about whether all studies should explicitly include this information or not. However, considering positionality is a vital part of the research process, particularly in the context of working with marginalized groups. In our dataset, a minority of papers reported positionality, often in the form of disclosing researchers' own marginalized identities. We suggest, echoing Liang et al. [71], that positionality statements need not always include identity characteristics. Positionality statements do not confer legitimacy, and although they may provide useful information about researchers' insider status within a community of study, the goal of a positionality statement is to provide context that helps readers understand the research being presented. In some cases, researchers' political beliefs, epistemological commitments, and disciplinary training may be more

helpful in understanding how research was conducted than specific identity characteristics. For example, in the case of this paper, both authors have experienced marginalization in multiple dimensions, which may help readers contextualize our interest in and commitment to this topic, and we are both privacy researchers who approach research from a critical, interpretivist perspective. These are important things to know about us, since we are offering research advice, whereas specific disclosures around our race, sexual orientation, and other identity features may not be.

9.5 Limitations and Future Work

Literature reviews have limitations that revolve around the selection of conferences and databases, as well as the criteria used to search for and filter papers.

First, our sample was scoped to a large but still limited set of venues. This is because our sampling strategy was designed to uncover themes in a broad, interdisciplinary set of papers at the intersection of privacy and marginalization, rather than to identify a comprehensive dataset of all papers published in this space across fields. As a result, we limited our search to SIGCHI-sponsored venues, select Communication journals, and select venues that focus on privacy and security research. While our exploratory searches suggested that these were some of the most common spaces for research relevant to our focus, it also constrained our sample. We think there is potential for future work to examine other conferences and fields in greater detail, including venues that focus on historically marginalized groups, such as the ACM Conference on Computers and Accessibility (ASSETS).

Second, the criteria we used to search for and filter papers also influenced our sample. When searching for relevant papers in ACM SIGCHI-sponsored venues and Communication journals, we constrained our search to just one keyword (“privacy”). This allowed us to limit the size of our dataset and remain consistent across subfields; however, in doing so, we may have missed work on marginalization as it intersects with other topics related to privacy, such as trust and disclosure. Since we filtered papers based on whether they contained the search term “privacy” in the title, abstract, and/or keywords, we likely missed some relevant papers that are not focused on privacy but nevertheless have interesting privacy-related findings pertinent to our focus.

Further, “marginalization” is a term with blurry boundaries and multiple definitions that are open to interpretation. Manually filtering our dataset to identify papers on marginalization required us to draw on several definitions of marginalization that often differed from each other, as well as our conceptualizations that were necessarily constrained by our own readings and positionalities. Our initial step was to read widely about how the concept of marginalization has been developed and used in a variety of fields, and we engaged in ongoing discussions about marginalization within our research group. Ultimately, our decision about whether to include studies in our dataset was informed by our understanding of what different social groups experience, our interpretations of definitions, and colored by our own beliefs and backgrounds. While we chose to focus on the experiences of marginalized adults, we also see potential for future work to explore research on children, who are a vulnerable population with unique privacy risks.

10 CONCLUSION

The goal of this paper was to serve as a roadmap for current and prospective researchers working at the intersection of privacy and marginalization by providing a review of current knowledge in the field, the practices through which it has been generated, and to chart a way forward. Our review of 88 papers published between 2010–2020 in HCI, Communication, and Privacy-focused venues found that this is a fast-growing area of study with wide variation in topics and research practices. Many existing privacy frameworks do not account for marginalization [67], and in response, we

introduced the *Privacy Responses and Costs* framework to reflect the range of privacy responses people enact and the costs and consequences of these responses to marginalized groups.

We also uncovered topics that need further study, such as race, the structural aspects of marginalization, and the role of policy, as well as the potential to use more diverse methods in our practices, including quantitative and participatory methods. Finally, given our focus on marginalized groups, we see a need to discuss and report research practices in greater detail, particularly around the ethical considerations of our work, and we put forth some suggestions for establishing shared best practices to do so.

ACKNOWLEDGMENTS

The authors thank Mako Hill, Rachel Greenstadt, and Kaylea Champion for their feedback throughout this study. The work was supported by the National Science Foundation (Award #2031951). The first author was also supported by the NSF under Grant #2127309 to the Computing Research Association for the CIFellows Project. The findings and recommendations in this paper are those of the authors and do not necessarily reflect the views of the NSF or the CRA.

REFERENCES

- [1] Jacob Abbott, Haley MacLeod, Novia Nurain, Gustave Ekobe, and Sameer Patil. 2019. Local standards for anonymization practices in health, wellness, accessibility, and aging research at CHI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [2] Syed Ishtiaque Ahmed, Md Romael Haque, Jay Chen, and Nicola Dell. 2017. Digital privacy challenges with shared mobile phone use in Bangladesh. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–20.
- [3] 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.
- [4] Taslima Akter, Bryan Dosono, Tousif Ahmed, Apu Kapadia, and Bryan Semaan. 2020. “I am uncomfortable sharing what I can’t see”: Privacy Concerns of the Visually Impaired with Camera Based Assistive Applications. In *29th USENIX Security Symposium*. 1929–1948.
- [5] Fatma Alshohoumi, Mohammed Sarrab, Abdulla AlHamadani, and Dawood Al-Abri. 2019. Systematic review of existing IoT architectures security and privacy issues and concerns. *Int. J. Adv. Comput. Sci. Appl.* 10, 7 (2019), 232–251.
- [6] Louise Barkhuus. 2012. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 367–376.
- [7] Susanne Barth and Menno DT De Jong. 2017. The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior—A systematic literature review. *Telematics and informatics* 34, 7 (2017), 1038–1058.
- [8] Lindsay Blackwell, Jean Hardy, Tawfiq Ammari, Tiffany Veinot, Cliff Lampe, and Sarita Schoenebeck. 2016. LGBT parents and social media: Advocacy, privacy, and disclosure during shifting social movements. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 610–622.
- [9] Louis Brandeis and Samuel Warren. 1890. The right to privacy. *Harvard law review* 4, 5 (1890), 193–220.
- [10] Amy Bruckman. 2014. Research ethics and HCI. In *Ways of Knowing in HCI*, Judith S. Olson and Wendy A. Kellogg (Eds.). Springer, 449–468.
- [11] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. 2018. Unveiling and quantifying Facebook exploitation of sensitive personal data for advertising purposes. In *27th USENIX Security Symposium*. 479–495.
- [12] Kelly Caine. 2016. Local standards for sample size at CHI. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 981–992.
- [13] Matthew Carrasco and Andruid Kerne. 2018. Queer visibility: Supporting LGBT+ selective visibility on social media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [14] Christine Chen, Nicola Dell, and Franziska Roesner. 2019. Computer security and privacy in the interactions between victim service providers and human trafficking survivors. In *28th USENIX Security Symposium*. 89–104.
- [15] Taejoong Chung, Jinyoung Han, Daejin Choi, Ted Taekyoung Kwon, Jong-Youn Rha, and Hyunchul Kim. 2017. Privacy leakage in event-based social networks: A meetup case study. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–22.
- [16] Rachel Clarke, Peter Wright, Madeline Balaam, and John McCarthy. 2013. Digital portraits: photo-sharing after domestic violence. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2517–2526.

[17] Nuremberg Code. 1949. The Nuremberg Code. *Trials of war criminals before the Nuremberg military tribunals under control council law* 10, 1949 (1949), 181–2.

[18] Kay Cook. 2008. Marginalized populations. *The SAGE encyclopedia of qualitative research methods* (2008), 495–496.

[19] Raymundo Cornejo, Robin Brewer, Caroline Edasis, and Anne Marie Piper. 2016. Vulnerability, sharing, and privacy: Analyzing art therapy for older adults with dementia. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1572–1583.

[20] Council of European Union. 2014. Council Directive 2000/43/EC implementing the principle of equal treatment between persons irrespective of racial or ethnic origin. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32000L0043>.

[21] Yngve Dahl and Kristine Holbø. 2012. “There are no secrets here!” Professional stakeholders’ views on the use of GPS for tracking dementia patients. In *Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services*. 133–142.

[22] Sanchari Das, Robert S Gutzwiler, Rod D Roscoe, Prashanth Rajivan, Yang Wang, L Jean Camp, and Roberto Hoyle. 2020. Humans and technology for inclusive privacy and security. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 64. SAGE Publications Sage CA: Los Angeles, CA, 461–464.

[23] Nicola Dell and Neha Kumar. 2016. The ins and outs of HCI for development. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 2220–2232.

[24] Michael Ann 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.

[25] Tobias Dienlin and Miriam J Metzger. 2016. An extended privacy calculus model for SNSs: Analyzing self-disclosure and self-withdrawal in a representative US sample. *Journal of Computer-Mediated Communication* 21, 5 (2016), 368–383.

[26] Tawanna R Dillahunt, Xinyi Wang, Earnest Wheeler, Hao Fei Cheng, Brent Hecht, and Haiyi Zhu. 2017. The sharing economy in computing: A systematic literature review. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–26.

[27] Carl DiSalvo, Phoebe Sengers, and Hrönn Brynjarsdóttir. 2010. Mapping the landscape of sustainable HCI. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1975–1984.

[28] Bryan Dosono, Jordan Hayes, and Yang Wang. 2015. “I’m Stuck!”: A Contextual Inquiry of People with Visual Impairments in Authentication. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 151–168.

[29] Brianna Dym and Casey Fiesler. 2020. Social Norm Vulnerability and its Consequences for Privacy and Safety in an Online Community. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–24.

[30] Kim VL England. 1994. Getting personal: Reflexivity, positionality, and feminist research. *The professional geographer* 46, 1 (1994), 80–89.

[31] Josephine Ensign. 2003. Ethical issues in qualitative health research with homeless youths. *Journal of advanced nursing* 43, 1 (2003), 43–50.

[32] Sandra K Evans, Katy E Pearce, Jessica Vitak, and Jeffrey W Treem. 2017. Explicating affordances: A conceptual framework for understanding affordances in communication research. *Journal of Computer-Mediated Communication* 22, 1 (2017), 35–52.

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

[34] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1979. The Belmont Report. (1979).

[35] Andrea Forte, Nazanin Andalibi, and Rachel Greenstadt. 2017. Privacy, anonymity, and perceived risk in open collaboration: A study of Tor users and Wikipedians. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1800–1811.

[36] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. 2019. “Is my phone hacked?” Analyzing Clinical Computer Security Interventions with Survivors of Intimate Partner Violence. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.

[37] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. 2018. “A Stalker’s Paradise” How Intimate Partner Abusers Exploit Technology. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.

[38] 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.

- [39] Niki Fritz and Amy Gonzales. 2018. Not the normal trans story: negotiating trans narratives while crowdfunding at the margins. *International Journal of Communication* 12 (2018), 20.
- [40] Lara Fullenwider and Adam Molnar. 2018. Settler governance and privacy: Canada's Indian Residential School Settlement Agreement and the mediation of state-based violence. *International journal of communication* 12 (2018), 1332–1349.
- [41] Seeta Peña Gangadharan. 2017. The downside of digital inclusion: Expectations and experiences of privacy and surveillance among marginal Internet users. *New Media & Society* 19, 4 (2017), 597–615.
- [42] Amy Gonzales and Nicole Fritz. 2017. Prioritizing flexibility and intangibles: Medical crowdfunding for stigmatized individuals. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2371–2375.
- [43] Government of India. 2016. The Rights of Persons with Disabilities Act. https://legislative.gov.in/sites/default/files/A2016-49_1.pdf.
- [44] Kathryn Greene, Valerian J Derlega, Gust A Yep, and Sandra Petronio. 2003. *Privacy and disclosure of HIV in interpersonal relationships: A sourcebook for researchers and practitioners*. Routledge.
- [45] Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H Mhaidli, Kentaro Toyama, and Florian Schaub. 2018. Keeping a low profile? Technology, risk and privacy among undocumented immigrants. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [46] Oliver L Haimson, Anne E Bowser, Edward F Melcer, and Elizabeth F Churchill. 2015. Online inspiration and exploration for identity reinvention. In *Proceedings of the 33rd annual ACM Conference on Human Factors in Computing Systems*. 3809–3818.
- [47] Oliver L Haimson, Justin Buss, Zu Weinger, Denny L Starks, Dykee 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.
- [48] Joanne M Hall, Patricia E Stevens, and Afaf Ibrahim Meleis. 1994. Marginalization: A guiding concept for valuing diversity in nursing knowledge development. *Advances in nursing science* (1994).
- [49] Sherry Hamby, Elizabeth Taylor, Alli Smith, Kimberly Mitchell, and Lisa Jones. 2018. Technology in rural Appalachia: Cultural strategies of resistance and navigation. *International Journal of Communication* 12 (2018), 21.
- [50] 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.
- [51] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. 2019. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium*. 105–122.
- [52] Jordan Hayes, Smriti 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)*.
- [53] Lasse Herskind, Panagiota Katsikouli, and Nicola Dragoni. 2020. Privacy and cryptocurrencies—A systematic literature review. *IEEE Access* 8 (2020), 54044–54059.
- [54] Hwajung Hong, Jennifer G Kim, Gregory D Abowd, and Rosa I Arriaga. 2012. Designing a social network to support the independence of young adults with autism. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 627–636.
- [55] Julie Hui, Nefer Ra Barber, Wendy Casey, Suzanne Cleage, Danny C Dolley, Frances Worthy, Kentaro Toyama, and Tawanna R Dillahunt. 2020. Community collectives: Low-tech social support for digitally-engaged entrepreneurship. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–15.
- [56] Julie Hui, Kentaro Toyama, Joyojeet Pal, and Tawanna Dillahunt. 2018. Making a living my way: Necessity-driven entrepreneurship in resource-constrained communities. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–24.
- [57] Aarti Israni, Sheena Erete, and Che L Smith. 2017. Snitches, Trolls, and Social Norms: Unpacking Perceptions of Social Media Use for Crime Prevention. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1193–1209.
- [58] Margaret C Jack, Pang Sovannaroth, and Nicola Dell. 2019. “Privacy is not a concept, but a way of dealing with life” Localization of Transnational Technology Platforms and Liminal Privacy Practices in Cambodia. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–19.
- [59] Dhruv Jain, Kelly Mack, Akli Amrous, Matt Wright, Steven Goodman, Leah Findlater, and Jon E Froehlich. 2020. Homesound: An iterative field deployment of an in-home sound awareness system for deaf or hard of hearing users. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [60] Jazette Johnson, Rebecca W Black, and Gillian R Hayes. 2020. Roles in the Discussion: An Analysis of Social Support in an Online Forum for People with Dementia. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–30.

[61] Jennifer A Kam, Andy J Merolla, and Andrew C High. 2020. Latinx immigrant youth's indirect and direct disclosures about their family - undocumented experiences, received emotional support, and depressive symptoms. *Communication Research* 47, 4 (2020), 599–622.

[62] Naveena Karusala, Apoorva Bhalla, and Neha Kumar. 2019. Privacy, patriarchy, and participation on social media. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 511–526.

[63] Elaine Kasket. 2009. Protocol for Responding to Participant Distress. *Adapted Version for Telephone/Skype* (2009).

[64] Christina Katsini, Yasmeen Abdrabou, George E Raptis, Mohamed Khamis, and Florian Alt. 2020. The role of eye gaze in security and privacy applications: Survey and future HCI research directions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–21.

[65] Khalid S Khan, Regina Kunz, Jos Kleijnen, and Gerd Antes. 2003. Five steps to conducting a systematic review. *Journal of the royal society of medicine* 96, 3 (2003), 118–121.

[66] Robert Klitzman, Ilene Albala, Joseph Siragusa, Kristen N Nelson, and Paul S Appelbaum. 2007. The reporting of monetary compensation in research articles. *Journal of Empirical Research on Human Research Ethics* 2, 4 (2007), 61–67.

[67] Bart P Knijnenburg, Xinru Page, Pamela Wisniewski, Heather Richter Lipford, Nicholas Proferes, and Jennifer Romano. 2022. Modern Socio-Technical Perspectives on Privacy.

[68] Roxanne Leitão. 2019. Anticipating smart home security and privacy threats with survivors of intimate partner abuse. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 527–539.

[69] 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.

[70] Xiaoqian Li, Wenhong Chen, and Joseph D Straubhaar. 2018. Concerns, Skills, and Activities: Multilayered Privacy Issues in Disadvantaged Urban Communities. *International Journal of Communication* 12 (2018), 22.

[71] Calvin A Liang, Sean A Munson, and Julie A Kientz. 2021. Embracing Four Tensions in Human-Computer Interaction Research with Marginalized People. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 2 (2021), 1–47.

[72] Amjad Mahfuth, Jaspaljeet Singh Dhillon, and S Mohd Drus. 2016. A systematic review on data security and patient privacy issues in electronic medical records. (2016).

[73] Elena Maris, Timothy Libert, and Jennifer R Henrichsen. 2020. Tracking sex: The implications of widespread sexual data leakage and tracking on porn websites. *New Media & Society* 22, 11 (2020), 2018–2038.

[74] Alice Marwick, Claire Fontaine, and Danah Boyd. 2017. "Nobody sees it, nobody gets mad": Social media, privacy, and personal responsibility among low-SES youth. *Social Media+ Society* 3, 2 (2017).

[75] 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.

[76] Nora McDonald, Karla Badillo-Urquiola, Morgan G Ames, Nicola Dell, Elizabeth Keneski, Manya Sleeper, and Pamela J Wisniewski. 2020. Privacy and Power: Acknowledging the Importance of Privacy Research and Design for Vulnerable Populations. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.

[77] Nora McDonald and Andrea Forte. 2020. The politics of privacy theories: Moving from norms to vulnerabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[78] 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.

[79] Helena M Mantis, Galina Madjaroff, Aaron Massey, and Zoya Trendafilova. 2020. The Illusion of Choice in Discussing Cybersecurity Safeguards Between Older Adults with Mild Cognitive Impairment and Their Caregivers. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–19.

[80] Helena M Mantis, Galina Madjaroff, and Aaron K Massey. 2019. Upside and downside risk in online security for older adults with mild cognitive impairment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.

[81] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G Altman, and Prisma Group. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *PLoS medicine* 6, 7 (2009), e1000097.

[82] Claudia Müller, Lin Wan, and Dalibor Hrg. 2010. Dealing with wandering: a case study on caregivers' attitudes towards privacy and autonomy when reflecting the use of LBS. In *Proceedings of the 16th ACM international conference on Supporting group work*. 75–84.

[83] Cosmin Munteanu, Heather Molyneaux, Wendy Moncur, Mario Romero, Susan O'Donnell, and John Vines. 2015. Situational ethics: Re-thinking approaches to formal ethics requirements for human-computer interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 105–114.

- [84] Mustafa Naseem, Fouzia Younas, and Maryam Mustafa. 2020. Designing Digital Safe Spaces for Peer Support and Connectivity in Patriarchal Contexts. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–24.
- [85] Boel Nelson and Tomas Olovsson. 2016. Security and privacy for big data: A systematic literature review. In *2016 IEEE international conference on big data (big data)*. IEEE, 3693–3702.
- [86] The United States Department of Justice. 2021. Fair Housing Act. <https://www.justice.gov/crt/fair-housing-act-2>
- [87] Ihudiya Finda Ogbonnaya-Ogburu, Angela DR Smith, Alexandra To, and Kentaro Toyama. 2020. Critical race theory for HCI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [88] Laura Parson. 2019. Considering positionality: the ethics of conducting research with marginalized groups. In *Research methods for social justice and equity in education*. Springer, 15–32.
- [89] Jessica Pater, Amanda Coupe, Rachel Pfafman, Chanda Phelan, Tammy Toscos, and Maia Jacobs. 2021. Standardizing Reporting of Participant Compensation in HCI: A Systematic Literature Review and Recommendations for the Field. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [90] Katy E Pearce, Jessica Vitak, and Kristen Barta. 2018. Socially mediated visibility: Friendship and dissent in authoritarian Azerbaijan. *International Journal of Communication* 12 (2018), 22.
- [91] Justin Petelka, Lucy Van Kleunen, Liam Albright, Elizabeth Murnane, Stephen Voda, and Jaime Snyder. 2020. Being (In) Visible: Privacy, Transparency, and Disclosure in the Self-Management of Bipolar Disorder. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [92] Fanny A Ramirez and Jeffrey Lane. 2019. Communication privacy management and digital evidence in an intimate partner violence case. *International Journal of Communication* 13 (2019), 18.
- [93] Yolanda A Rankin and Jakita O Thomas. 2019. Straighten up and fly right: Rethinking intersectionality in HCI research. *Interactions* 26, 6 (2019), 64–68.
- [94] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2017. Where is the digital divide? a survey of security, privacy, and socioeconomics. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 931–936.
- [95] Jake Reichel, Fleming Peck, Mikako Inaba, Bisrat Moges, Brahmnoor Singh Chawla, and Marshini Chetty. 2020. 'I have too much respect for my elders': Understanding South African Mobile Users' Perceptions of Privacy and Current Behaviors on Facebook and WhatsApp. In *29th USENIX Security Symposium*. 1949–1966.
- [96] Matt Reichel. 2017. Race, Class, and Privacy: A Critical Historical Review. *International Journal of Communication* 11 (2017), 4757–4768.
- [97] Ellie Rennie, Tyson Yunkaporta, and Indigo Holcombe-James. 2018. Privacy at the margins| privacy versus relatedness: managing device use in Australia's remote aboriginal communities. *International Journal of Communication* 12 (2018), 19.
- [98] Fatemeh Rezaeibagha, Khin Than Win, and Willy Susilo. 2015. A systematic literature review on security and privacy of electronic health record systems: technical perspectives. *Health Information Management Journal* 44, 3 (2015), 23–38.
- [99] Nithya Sambasivan, Amna Batool, Nova Ahmed, Tara Matthews, Kurt Thomas, Laura Sanely Gaytán-Lugo, David Nemer, Elie Bursztein, Elizabeth Churchill, and Sunny Consolvo. 2019. "They Don't Leave Us Alone Anywhere We Go" Gender and Digital Abuse in South Asia. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [100] Nithya Sambasivan, Garen Checkley, Amna Batool, Nova Ahmed, David Nemer, Laura Sanely Gaytán-Lugo, Tara Matthews, Sunny Consolvo, and Elizabeth Churchill. 2018. "Privacy is not for me, it's for those rich women": Performative Privacy Practices on Mobile Phones by Women in South Asia. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. 127–142.
- [101] Nithya Sambasivan, Julie Weber, and Edward Cutrell. 2011. Designing a phone broadcasting system for urban sex workers in India. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 267–276.
- [102] Shruti Sannon, Natalya N Bazarova, and Dan Cosley. 2018. Privacy lies: Understanding how, when, and why people lie to protect their privacy in multiple online contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [103] Shruti Sannon and Dan Cosley. 2019. Privacy, power, and invisible labor on Amazon Mechanical Turk. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [104] Shruti Sannon, Elizabeth L Murnane, Natalya N Bazarova, and Geri Gay. 2019. "I was really, really nervous posting it" Communicating about Invisible Chronic Illnesses across Social Media Platforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [105] Hyunjin Seo, Hannah Britton, Megha Ramaswamy, Darcey Altschwager, Mathew Blomberg, Shola Aromona, Bernard Schuster, Ellie Booton, Marilyn Ault, and Joi Wickliffe. 2020. Returning to the digital world: Digital technology use and privacy management of women transitioning from incarceration. *New Media & Society* (2020).

- [106] Rhonda M Shaw, Julie Howe, Jonathan Beazer, and Toni Carr. 2020. Ethics and positionality in qualitative research with vulnerable and marginal groups. *Qualitative Research* 20, 3 (2020), 277–293.
- [107] Lucy Simko, Ada Lerner, Samia Ibtaasam, Franziska Roesner, and Tadayoshi Kohno. 2018. Computer security and privacy for refugees in the United States. In *IEEE Symposium on Security and Privacy (S&P)*. IEEE, 409–423.
- [108] Manya Sleeper, Tara Matthews, Kathleen O’Leary, Anna Turner, Jill Palzkill Woelfer, Martin Shelton, Andrew Oplinger, Andreas Schou, and Sunny Consolvo. 2019. Tough times at transitional homeless shelters: Considering the impact of financial insecurity on digital security and privacy. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [109] Janaki Srinivasan, Savita Bailur, Emrys Schoemaker, and Sarita Seshagiri. 2018. The poverty of privacy: Understanding privacy trade-offs from identity infrastructure users in India. *International Journal of Communication* 12 (2018), 20.
- [110] 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.
- [111] Ashley Marie Walker, Yaxing Yao, Christine Geeng, Roberto Hoyle, and Pamela Wisniewski. 2019. Moving beyond ‘one size fits all’ research considerations for working with vulnerable populations. *Interactions* 26, 6 (2019), 34–39.
- [112] Lin Wan, Claudia Müller, Volker Wulf, and David William Randall. 2014. Addressing the subtleties in dementia care: pre-study & evaluation of a GPS monitoring system. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3987–3996.
- [113] Ruolin Wang, Chun Yu, Xing-Dong Yang, Weijie He, and Yuanchun Shi. 2019. EarTouch: facilitating smartphone use for visually impaired people in mobile and public scenarios. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [114] Yang Wang. 2017. The third wave? Inclusive privacy and security. In *Proceedings of the 2017 New Security Paradigms Workshop*. 122–130.
- [115] Yang Wang. 2018. Inclusive security and privacy. *IEEE Security & Privacy* 16, 4 (2018), 82–87.
- [116] Mark Warner, Andreas Gutmann, M Angela Sasse, and Ann Blandford. 2018. Privacy unraveling around explicit HIV status disclosure fields in the online geosocial hookup app Grindr. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–22.
- [117] Mark Warner, Agnieszka Kitkowska, Jo Gibbs, Juan F Maestre, and Ann Blandford. 2020. Evaluating ‘Prefer not to say’ Around Sensitive Disclosures. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [118] Janith Weerasinghe, Kediel Morales, and Rachel Greenstadt. 2019. “Because... I was told... so much”: Linguistic Indicators of Mental Health Status on Twitter. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 152–171.
- [119] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.
- [120] Richmond Y Wong and Deirdre K Mulligan. 2019. Bringing design to the privacy table: Broadening “design” in “privacy by design” through the lens of HCI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–17.

Received January 2022; revised April 2022; accepted May 2022