

ABSTRACT

Title of Dissertation: REAL-WORLD SECURITY FOR AT-RISK USERS

Noel Warford
Doctor of Philosophy, 2023

Dissertation Directed by: Professor Michelle Mazurek
Department of Computer Science

Whether due to context or circumstance, “at-risk” users experience risk factors that augment or amplify their chances of being digitally attacked and/or suffering disproportionate harms in comparison to the so-called “average user.” Significant prior scholarship over the past decade has surfaced a variety of varying needs and concerns for different at-risk populations, from activists to survivors of intimate partner violence. However, when technology creators and researchers seek to develop tools to support these users, it can be difficult to reckon with the vastly varied, often contradictory, needs of different populations.

To address this challenge, this thesis develops a framework for understanding at-risk user research, designed to help researchers and technology creators better guide technology design to support at-risk users. The framework includes ten *contextual risk factors*—such as *prominence* and *marginalization*—that impact the digital safety of these users and three broad sets of both technical and non-technical *protective practices* these users currently employ to meet their needs. This thesis then analyzes the current *barriers* at-risk users face when seeking to address their

particular digital safety challenges.

In order to test the utility of this framework, this thesis explores the needs and concerns of two distinct at-risk populations through interview studies—library IT staff and entertainment journalists. Library IT staff work under severe resource constraints to serve patrons who often do not have access to computing resources outside of the library, or may need the particular protection of intellectual freedom libraries seek to provide. I observe that privacy and utility currently work in opposition to one another, but library IT staff in fact prioritize privacy over utility in accordance with strong library values.

As another case study, entertainment journalists—those who cover movies, TV shows, sports, and other pop culture topics—work at the intersection of prominence and marginalization. Due to their profession, they must maintain a public presence, but that presence often exposes them to severe harassment, particularly if they also experience marginalization. These journalists conceptualize harassment as an unavoidable cost of their profession, but they perceived taking any action other than attempting to ignore it as largely fruitless. In line with prior scholarship, this thesis shows that institutional support—both from their employer as an organization and other colleagues—was essential to managing harassment response.

This thesis demonstrates that understanding and accounting for particular contextual needs is crucial for improving digital safety for at-risk users. By working towards generalizing their needs and concerns, it will become easier to design both technical and non-technical solutions for both immediate and long-term challenges for at-risk populations.

REAL-WORLD SECURITY FOR AT-RISK USERS

by

Noel Warford

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Advisory Committee:

Professor Michelle L. Mazurek, Chair/Advisor
Professor Ian Miers
Doctor Sunny Consolvo
Professor Jordan Boyd-Graber
Professor Tudor Dumitras

© Copyright by
Noel Warford
2023

Acknowledgments

This dissertation is based in part upon work supported by DARPA under grant HR00112010011, as well as a gift from the Google Award for Inclusion Research program. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Government, DARPA, or Google.

No dissertation is completed by one person alone, and this one is no exception. Many thanks are included below, and if I forgot someone, please be assured that this error does not diminish my appreciation.

To my advisor, Dr. Michelle Mazurek: thanks for taking a chance on this liberal-arts-college graduate. Without your guidance and support, this simply would not have been possible.

To my family—Mom, Dad, Clara, John, Charles, Carolyn, and everyone else: thank you for your continued love and support during this many-year journey (even when I unexpectedly returned home for seven months in 2020).

To my friends and colleagues in SP2 (current and former)—Dan Votipka, Elissa Redmiles, Wei Bai, Kelsey Fulton, Omer Akgul, Nathan Reiting, Nathan Malkin, Wentao Guo, Alan Luo,

Phoebe Moh, and all the rest: thank you for your support, commiseration, collaboration, and cheerleading across the years.

To my other friends in the graduate program, especially Erica, Noemi, Makana, Lillian, MG, and the many folks I've been honored to study alongside: thanks for your constant support and companionship during this strange phase of life we call graduate school.

To all my friends and collaborators at other institutions, especially Rosie Bellini, Collins Mundayendo, Yasemin Acar, and Miranda Wei: thanks for helping me be a better researcher and for being a friendly face at conferences!

To all my awesome coauthors, especially Tom Ristenpart, Nikki Dell, Adam Aviv, Nora McDonald, Rachel Greenstadt, Kurt Thomas, Patrick Gage Kelley, Tara Matthews, Sunny Consolvo, and Manya Sleeper: the work in this thesis would simply not exist without all of your hard work and thoughtful feedback.

To my committee: thank you very much for your support, guidance, and everlasting patience around the bureaucratic process that is scheduling a thesis defense.

To the staff that make UMD run, especially Tom Hurst, Dana Purcell, and Petra Zapf: thanks for your support and making the bureaucracy easy!

To all the great folks in DC and elsewhere I've had the pleasure to play games of all descrip-

tions with—Avi, Josh, Ashley, Adam, Trevor, Noah, Sam, Jacqueline, Erin, Ben, Ryan S., Ryan L., Justin, Cate, Sylvie, Jay, Jim, Zoe, Brando, Cam, Ciero, Dave, Charles, all the Capital Clocktower crew I don't have space to mention: thanks for the good games and the good times.

To the Crate and Crowbar community, especially my good friend veerserif: y'all rock, and thanks for listening, everybody.

To my family at St. Michael's Episcopal Church, especially Beth Franklin, the Kelleher family, and my choir members: your support and encouragement during my time serving as music minister was truly unmatched, and I am eternally grateful.

Table of Contents

Acknowledgements	ii
Table of Contents	v
List of Tables	viii
List of Abbreviations	ix
Chapter 1: Introduction	1
Chapter 2: Background and Related Work	6
2.1 Foundational concerns of at-risk user security research	6
2.1.1 Defining the at-risk user	6
2.1.2 Structural challenges	9
2.1.3 Working with at-risk users	10
2.2 Libraries and digital safety	11
2.3 Journalists and digital safety	15
2.4 Online hate and harassment	17
2.5 Reflexive thematic analysis	23
Chapter 3: A Framework for Unifying At-Risk User Research	26
3.1 Introduction	26
3.2 Methods	27
3.2.1 Paper selection	27
3.2.2 Codebook development	28
3.2.3 Full analysis	29
3.2.4 Limitations	30
3.3 Contextual risk factors	31
3.3.1 Societal factors	34
3.3.2 Relationships	37
3.3.3 Personal circumstances	39
3.3.4 How do contextual risk factors interact?	41
3.4 Protective practices	43
3.4.1 Social strategies	44
3.4.2 Distancing behaviors	47
3.4.3 Technical solutions	48
3.5 Barriers to protective practices	52

3.5.1	Competing priorities	53
3.5.2	Lack of knowledge or experience	55
3.5.3	Broken technology assumptions	57
3.6	Discussion	59
3.6.1	Research	60
3.6.2	Technology development	62
Chapter 4: How Library IT Staff Navigate Privacy and Security Challenges and Responsibilities		66
4.1	Introduction	66
4.2	Methods	68
4.2.1	Recruitment	68
4.2.2	Intake survey	70
4.2.3	Interview protocol	70
4.2.4	Participants	72
4.2.5	Thematic analysis	73
4.2.6	The capabilities approach framing	73
4.2.7	Limitations	75
4.3	Library IT staff roles and responsibilities	76
4.4	Findings	80
4.4.1	Goals of library IT staff	80
4.4.2	Threats faced by library IT staff	87
4.4.3	Tensions among goals and responsibilities	91
4.4.4	Constraints of library IT staff	95
4.5	Discussion	101
4.5.1	Sen's capabilities applied to library IT	102
4.5.2	Privacy needs and moral capabilities	103
4.5.3	Values in conflict with resources, decision-makers, and other institutions	105
4.5.4	Future work	106
4.5.5	Connecting to the at-risk framework	107
Chapter 5: How Entertainment Journalists Manage Online Hate and Harassment		108
5.1	Introduction	108
5.2	Methods	110
5.2.1	Recruitment	110
5.2.2	Protocol	111
5.2.3	Limitations	113
5.3	Analysis	114
5.4	The context of harassment	114
5.4.1	The inevitable price of admission	115
5.4.2	Hurt people hurting people	119
5.4.3	White cisgender maleness is a shield	120
5.5	Strategies for dealing with harassment	123
5.5.1	"Just ignore it": the best worst option	123
5.5.2	Constant vigilance	127

5.5.3	External support is critical	129
5.5.4	IT hygiene	132
5.6	Discussion	133
5.6.1	Contextualizing with prior scholarship on harassment	134
5.6.2	Prominence and marginalization: Relating to the at-risk framework	135
5.6.3	An unequal distribution of labor	136
Chapter 6:	Discussion	138
6.1	Reflecting on the at-risk framework	139
6.2	How can technology creators and security researchers use the at-risk framework?	141
6.3	Understanding context is crucial	143
6.4	Non-technical support is just as important as new tools	144
6.5	Digital safety cannot be isolated from social conditions	146
Appendix A:	At-risk framework paper list	148
Appendix B:	Library IT staff protocols	156
B.1	Intake survey	156
B.2	Interview protocol	159
Appendix C:	Entertainment journalist protocol	168
Bibliography		175

List of Tables

3.1	Meta-analysis showing how the contextual risk factors from our at-risk framework apply to categories of populations in our dataset. A black circle means one or more citations indicated the contextual risk factor was relevant. Note that the absence of a black circle does not mean the population does not have that digital-safety risk. Rather, it means the set of papers cited did not discuss or explore that risk factor as defined in this framework.	33
4.1	Participant information including the type of institution they were affiliated with, work experience in IT (in a library setting and overall) the number of IT staff at their institution, and the estimated number of patrons their institution serves per month. Items are left blank when participants did not know or preferred not to answer.	69
4.2	Patron populations reported by our participants in the intake survey.	81
A.1	A list of the 95 papers included in the dataset for the analysis in Chapter 3	156

List of Abbreviations

ALA	American Library Association
CCPA	California Consumer Privacy Act
HCI	Human-computer interaction
IPV	Intimate partner violence
LGBTQ+	Lesbian, gay, bisexual, transgender, queer, and other non-cisgender and non-heterosexual identities
NGO	Non-governmental organization

Chapter 1: Introduction

Anyone can experience attacks related to their security, privacy, or safety online (i.e., *digital safety*), but some at-risk users experience risk factors that increase their chances of being digitally attacked, suffering disproportionate harms, or both. In addition to the already-existing difficulties of general comprehension and use of secure systems, these users experience challenges and constraints beyond the “average user” often assumed by digital-safety research. A common recommendation from researchers in this space is to consider at-risk users during the technology creation process (e.g., [23, 55, 129, 202, 221, 245]). For technology creators, however, it can be bewildering to consider dozens of different at-risk populations, each with disjoint and sometimes-contradictory digital-safety needs.

Prior work has studied a variety of at-risk populations which face different challenges and needs. Survivors of intimate partner violence have an attacker with intimate access to their accounts and devices, making defense challenging and breaking many assumptions about account security [42, 68–70, 87, 129]. Journalists and activists face highly-resourced adversaries who may even control telecommunications infrastructure [10, 55, 63, 101, 114, 123, 136–139, 208]. Children and teens are still developing an understanding of risks and safety on the internet, and so may make decisions that put them in danger [74, 74–76, 76, 94, 104, 105, 111, 140, 141, 143, 149, 237–239, 245]. Naturally, this work is almost universally focused on the specific at-risk population in

question, which can lead to difficulty applying the lessons learned to other groups. Generalizing solutions for at-risk users requires considering a huge variety of needs and concerns, which is challenging. What works for activists in the Sudanese Revolution [55] may not be applicable for foster teenagers [21], due to dramatically different constraints, circumstances, resources, and threat models.

In this thesis, I contribute to the above scholarship understanding and addressing the digital-safety needs of at-risk populations, framed by the following research questions:

1. What factors—such as a person’s situation in society, relationships, or personal circumstances—contribute to digital-safety risks for at-risk users?
2. Are there any commonalities across at-risk populations that can help researchers and technologists better account for these populations in their work?
3. What decisions do at-risk users make when faced with real-world constraints and goals? How do they make tradeoffs between digital safety and other goals?

First, described in Chapter 3, I (with many collaborators) rigorously developed a framework to help solve this problem [229]. By analyzing 6,534 papers from 2015 to 2021 from major security and HCI conferences, we created a dataset of 95 papers that focus on at-risk users—defined as having risk factors that augment or amplify their chances of being digitally attacked and/or suffering disproportionate harms. We developed ten *contextual risk factors* that describe the various ways in which digital-safety risks are increased: three *societal factors*, influenced by the user’s position in their society in culture, three *relationship factors*, emphasizing the particular personal relationships that can put users at risk, and four *personal circumstances*, dependent on who the at-risk user is, their personal activities, or their profession. Many of the populations we

analyzed were impacted by more than one of these factors, creating intersectional interactions between them that create unique and difficult digital-safety challenges.

We also identified three categories of *protective practices*—groups of common strategies that at-risk users *already* employ to deal with their unique challenges. These spanned *social strategies*, that relied on their network and personal connections, *distancing behaviors*, where they simply stepped away or did not engage in digital life, and *technical solutions*, although these technologies did not always meet their needs. At the intersection of these factors and practices, we also describe some *barriers to adopting protective practices*: strongly *competing priorities*, a *lack of knowledge or experience*, and many *broken technology assumptions* based on the myth of the “average user.” This framework allows researchers and technology creators to scalably consider the needs of at-risk users in their work, rather than needing to develop expertise on dozens of disparate at-risk populations.

In Chapters 4 and 5, my collaborators and I applied this framework to two at-risk populations: library IT staff and entertainment journalists. The framework allowed us to better position our work to specifically address these populations’ unique challenges in a systematic way, rather than needing to reinvent the wheel of at-risk user research.

Chapter 4 describes an interview study with 12 library IT staff [118]: the employees of libraries who are responsible for maintaining their library’s technical infrastructure and resources. Libraries serve as a critical source of digital resources for people who are *resource or time constrained* and/or experience *marginalization* – two of the *contextual risk factors* identified in Chapter 3. Library IT staff make decisions about privacy and utility that directly impact these users.

When making decisions about library IT infrastructure, our participants’ key goal was protecting patron privacy from threats outside their walls. Using Amartya Sen’s capabilities frame-

work, we describe how these digital-safety decisions can come at the cost of limiting patrons' ability to accomplish their other goals; fully erasing computers after use, for example, limits the ability of patrons without computers at home to do longer-term persistent work like developing a resume. This tradeoff—driven by *resource and time constraints*—illustrates the impact of real-world social and political constraints on digital safety. In some cases, particularly involving *marginalization*, maintaining digital safety comes at the meaningful cost of utility—if solutions do not account for this tradeoff, they risk making these disadvantages more severe.

In Chapter 5, I conducted another interview study with a different at-risk population: entertainment journalists. While prior scholarship on journalists and digital safety mostly focuses on investigative journalists with nation-state adversaries [63, 135–139], in this study, I focus on journalists who cover entertainment (movies, TV, video games, sports, and other popular media) and their experiences dealing with online harassment.

Online harassment is a prevalent and growing problem [210] that disproportionately affects women, transgender people, and people of color. Journalists who cover popular media and sports often face extreme harassment when covering entertainment, particularly when writing about the intersection between their entertainment beats, politics, and social issues. In this study, we interviewed nine entertainment journalists who described their experiences of harassment and their strategies for managing it. We find that harassment is simultaneously ephemeral and dangerous; although any individual instance of harassment may be temporary, harassment can escalate to real, physical harm. Our participants' main strategy for harassment management was to find ways to ignore it and move on. *Prominence* and *marginalization* work in combination with the threat of harassment to drive these users to reduce participation and self-censor, trading free expression for physical and emotional safety.

This thesis shows that contextual understanding is critical for addressing the digital-safety needs of at-risk users. By providing a framework for understanding these needs, we make it easier to account for these varying needs at scale, so that everyone’s digital-safety needs can be met. However, this work also points to the factors *outside* the realm of digital-safety research and technology development—we argue that advocacy and social change are just as important for addressing these issues as new tools and technology.

Chapter 2: Background and Related Work

I discuss related work across four key areas: at-risk user security research, libraries and digital safety, journalists and digital safety, and online hate and harassment. I also provide key methodological background for *reflexive thematic analysis*, the primary analytical technique used for this work.

2.1 Foundational concerns of at-risk user security research

In this section, I define the at-risk user and summarize current scholarship on various at-risk populations. I additionally describe some key methodological concerns for ethically working with at-risk populations.

2.1.1 Defining the at-risk user

I use *at-risk user* as an umbrella term for anyone with risk factors that increase their chances of being attacked digitally, suffering disproportionate harms from an attack, or both. I refer to groups of at-risk users as *at-risk populations*. Due to a lack of consensus in the literature on how to refer to such users or populations, I chose these terms with the goal of drawing focus to external risks these users face.

Below, I quote an illustrative example from Havron et al. [87] describing the challenges

in identifying technical interventions for survivors of intimate partner violence. This is not any particular person's story, for privacy reasons, but illustrates the complex environment people in this situation must navigate.

Example scenario, Carol's experience: Carol's now ex-husband subjected her to several years of increasing physical, emotional, and technology abuse before she obtained an order of protection, physically moved out, and filed for divorce. They are in a custody battle over their two children, ages four and ten, who live with the ex-husband part of the time.

Carol knows that he installed spyware on at least one of her devices, because she found the purchase of mSpy on their joint credit card statement. Additionally, he had access to her private photos that he then posted on Facebook. He would also routinely, over the period of a year, "hack" into her online accounts, posing as her in efforts to further alienate her from her friends and family. He even locked her out of her Gmail account by changing the recovery emails and phone number to his, which was devastating to her career in sales because it contained her business contacts.

Carol currently has five devices: a new Apple iPhone that is her primary device, two Android phones used by her children, an Apple iPad tablet bought for her children by her ex-husband, and a several-year-old Apple iPhone originally bought for her by her ex-husband. She routinely uses Facebook, a new Gmail account (since her old one was stolen by her ex-husband), and a variety of other social media apps that are important for her work in sales. [87]

Carol's digital life spans several devices and accounts, and is complicated by her relation-

ship with her ex-husband, especially considering their shared custody of young children. General security advice, like “use a password manager” or “keep your devices and software up to date” may not address the particular concerns she has based on context and circumstance.

Over the past several years, there has been a wide range of scholarship on various at-risk populations. As in the example above, survivors of intimate partner violence (IPV) have particular concerns around privacy, as their abuser may have unfettered access to their accounts and devices, making separation and safety paramount concerns over and above typical digital safety needs [42, 69, 70]. The Clinic to End Tech Abuse (CETA) is a research-based intervention working in collaboration with New York City’s Family Justice Centers that provides individual, clinical support to survivors of IPV, provides multimodal support to prevent abusers tracking, harassing, and otherwise harming their victims [68, 87, 217].

Children and teenagers also experience elevated risk, often due to their relationship with and reliance on third parties. Parents may desire to use tracking and surveillance software to protect their children online, but children experience this as a severe privacy invasion [74, 75]. More effective techniques for protecting children and teenagers online while respecting their autonomy include developing mutual trust [237–239] and including children in the design process of protective applications [76, 141]. Managing this balance can be very challenging, as young people may also lack full understanding of certain digital-safety risks, depending on age, experience and education [94, 104, 111, 245]. Young people who are at higher risk due to other factors, like foster teenagers with abusive relatives, may additionally have severe consequences if their digital safety is compromised [21].

Several scholars have additionally studied the particular needs of various at-risk users in non-Western countries and regions; strong social norms and different standards for technology

use mean that models for understanding digital safety do not necessarily apply across cultural lines [221]. Across a variety of cultures, countries, and regions, the assumption of one account and device for each individual is commonly broken [3, 5, 9, 178, 183, 184]. Women, in particular, are often expected to share devices and account credentials with male family members, as in Saudi Arabia [9] and South Asia [183, 184]. Western concepts of privacy may not even be useful constructs in other cultures, as in the case of the strongly community-oriented culture of Cambodia, where translating the concept of *privacy* on social media from English to Khmer was challenging [93].

I expand on the challenges of people with low socioeconomic status, marginalized individuals, and journalists in more detail in Sections 2.2 and 2.3. Further details and analysis on commonalities and differences between these and other at-risk populations can be found in Chapter 3.

2.1.2 Structural challenges

The challenges experienced by at-risk users can be inordinately complex, reflecting broader, societal challenges; McDonald et al. call for privacy research to take into consideration “structural inequalities and social norms” [131]. These inequalities, which vary globally, mean that particular care is required to integrate at-risk users’ experiences and identities into the technology creation process. Keyes et al. call for an “anarchist HCI,” calling HCI researchers to orient their work towards maximizing human autonomy, especially in relation to existing entrenched power structures [97]. Walker et al. provide a set of heuristics for performing research with at-risk populations, emphasizing in particular the needs of the participants, rather than solely focusing on the benefit to the research team [225]. This thesis seeks, in part, to respond to these calls to ethically

conducting research with and for at-risk users by focusing their needs and concerns, especially acknowledging that these people are experts on their own experience.

2.1.3 Working with at-risk users

Working with at-risk populations requires careful consideration in order to ensure they are active participants in the research, protected from undue harm, and benefit from the research itself. As Bellini et al. outline in work I collaborated on, both researchers and participants can face harms when this research is being conducted [25]. Currently, the digital-safety community lacks a standard set of guidelines to conducting at-risk user research, but Bellini et al. outline 14 research risks and 36 safety practices, which come from analysis of prior work in the digital-safety research community and elsewhere [17, 28, 84, 161, 186]. These practices aim to help researchers do this work well, safely, and for the benefit of all.

The human-subjects research in this thesis was conducted with a *trauma-informed* lens. I use the following definition of trauma from the Substance Abuse and Mental Health Services Administration (SAMHSA):

Individual trauma results from an event, series of events, or set of circumstances that is experienced by an individual as physically or emotionally harmful or life threatening and that has lasting adverse effects on the individual’s functioning and mental, physical, social, emotional, or spiritual well-being. [207]

Many at-risk users have experienced trauma in relation to digital safety, so extra care is needed to ensure that the research process does not create more harm than benefit, above and beyond standard ethical requirements. In order to account for this, I used Chen et al.’s trauma-

informed computing framework for computing [46] and Wong’s guidelines for trauma-informed care in qualitative research [241], which both provide recommendations for conducting user-centered research. In particular, Chen et al. recommend proceeding with extreme caution to avoid retraumatization of participants when discussing challenging events. I paraphrase some key guidelines below:

- *Avoid retraumatization*: Seek ways to minimize the potential harms of conducting this research, especially when asking participants to relive traumatic experiences.
- *Develop trust*: Be aware of the power dynamic between you and the participant. Build rapport with the participant and provide them time, space, and resources (as you are able) to deal with the issues they face.
- *Foster collaboration*: Participants should be directly involved in decisions regarding their care and wellbeing. Researchers should not assume they know the best solution, but work closely with participants to understand and address their specific concerns.

2.2 Libraries and digital safety

Libraries in the U.S. [13] and in other Western countries [20] are staunch guardians of their patrons’ right to free access to information, which has historically required protection from governments and other entities.

Librarians’ role as guardians of patrons’ privacy and intellectual freedom has, at times, been in tension with their civic role to defend public safety and security [38,204]. The challenges librarians face today are even more complex, spanning government institutions as well as the

third-party entities that provide digital resources to patrons, not to mention the vast space of digital resources and media that originate beyond library walls, but which are accessed using library computers.

In recent years, librarians have had to contend with privacy incursions from the Department of Homeland Security (DHS), who sought to bar libraries' use of privacy technology like Tor [16], and the FBI, who used the National Security Letter (NSL) provision of the Patriot Act to subpoena patrons' records. At each point, librarians fought back through legal means, but also directly at their libraries. Perhaps the most well known example is related to NSLs, when some libraries posted signs saying that the FBI had not been there but to watch closely for the removal of the sign to indicate that the FBI had requested patron records [77].

More recently, library scholars have focused on the way in which surveillance capitalism and big data policing [35, 66, 108, 247] implicates libraries that rely on third-party products. These products amass patrons' personal data, which may then be sold to law enforcement entities, including Immigration and Customs Enforcement (ICE) [108]. Consolidation of the major publishers has left librarians little choice but to be dependent on these vendors.

This creates a dilemma for libraries who want to both provide free access and protect their patrons' privacy. According to Lamdan [108], libraries may indeed see their role as being on the front lines, acting on behalf of their patrons where policy has failed. The American Library Association (ALA) has recently weighed in on these types of dilemmas [14]. The ALA has staunchly opposed efforts by top academic publishers like Elsevier, Springer, and Nature to clamp down on illegal use of academic material through the use of surveillance software, including biometric data collection [145]. According to a 2021 resolution, the ALA "stands firmly against behavioral data surveillance," citing the "Intellectual Freedom Manual" [72] to ensure that librarians protect

patrons from misuse of their data, including by the library itself [14]. The resolution also calls on librarians to uphold their role as advocates for their patrons' privacy and confidentiality, and to continue to educate themselves to ensure they can fulfill this role. Libraries aim to hold themselves to privacy standards defined in the Library Code of Ethics, first published in 1939 [12]; when vendors do not meet those standards [107], it is not entirely clear how they reconcile this tension.

In addition to threats from governments and vendors encroaching on protecting patron privacy, libraries must also consider privacy concerns relating to state and national law (e.g., the California Consumer Privacy Act, or CCPA) [52], as well as patrons' privacy on social media [109]. Further, many patrons rely on libraries' digital resources to seek health information or look for jobs [91], making libraries a space where sensitive data is transmitted on public computers.

A survey conducted by the Pew Research Center in 2016 found that 29% of patrons 16 years or older used library computers [92]. Among those patrons, 61% use them for research for school or work, 38% use them to get health information from online resources, and 26% use them to take online classes or complete online certifications. An older report from 2012 breaks down patron PC usage slightly differently, reporting that some 41% of patrons used public PCs to visit government websites or get information regarding government services, and 36% of patrons used them to look or apply for jobs. While these numbers may have changed in the intervening time, many patrons likely still use library resources this way [246].

While the literature points to efforts being made to study and inform librarians' privacy practices (e.g., [119]), there has been little systematic research focused on how librarians perform privacy stewardship, what their priorities are, and what technologies they use to reflect their priorities. In particular, there has been little study of those serving the role of library IT staff. We

begin to address this gap in Chapter 4.

Low-income marginalized communities and library privacy. Low-income marginalized¹ individuals face a “matrix of vulnerabilities” stemming from primary dependence on mobile technologies for internet access, further hampered by lack of use or familiarity with internet privacy tools and strategies [121]. Not only are low-socioeconomic individuals more reliant on public computers [92], they may lack strategies for keeping their information safe online [120].

Vitak et al. show how librarians struggle with providing assistance to patrons who rely on library computers for sensitive transactions, especially when these patrons lack concern or tools for protecting their data and/or trust librarians implicitly with their personal data (which extends to online sites) [223]. Research has also shown that libraries may lack formal privacy strategies for protecting their patrons’ privacy, leaving librarians to take an ad hoc and often intersectional approach to protecting their patrons [133].

In their study of security and privacy practices of refugees in the United States, Simko et al. find that refugees, having little or no experience with certain online threats, are at higher risk for scams. Moreover, the authors find that refugees tend to rely on public libraries, community centers, non-governmental organizations (NGOs), teachers, and case managers for support with sensitive data required by employers, but this trust can be complex [200].

Immigrants in the U.S. also have to contend with pervasive surveillance. Kalhan points out that immigration governance in the U.S. has rapidly (and invisibly) been transformed into a surveillance “technology-driven enterprise” [95], with U.S. Immigration and Customs Enforcement (ICE) operating a kind of “dragnet,” and that fear of surveillance is a deterrent for “im-

¹“Marginalized” here refers to the disadvantages experienced by those who are prevented from obtaining full membership and participation in society, often because of their race, ethnicity, gender, religion, sexual orientation, and/or sexual identity [162], which can overlap with lower socio-economic status.

migrants and their families from participating in a broad range of activities necessary for health and well-being not only of individuals, but of the communities of which they are part” [227]. We (echoing the library IT staff participants we interview) consider access to information and computing resources that allow one to access social services, find health information, or search for employment as some of these necessary activities.

These concerns echo beyond the library context for people with low socioeconomic status, who experience marginalization, or both. Low-income Black communities are more often subject to policing and surveillance in their neighborhoods [36, 66, 113]. Low socioeconomic status makes it more difficult to purchase private personal devices, or keep devices up to date [64, 202, 221, 224]. Attackers might also target a variety of marginalized individuals due to their background or identity, such as people who are LGBTQ+ [29, 30, 114, 188, 230, 231] or undocumented immigrants [82].

2.3 Journalists and digital safety

Journalists have particular digital-safety needs due to their profession and prominence. Investigative journalists need to be able to securely share and distribute relevant documents with trusted sources, but without the subject of investigation finding out [63, 138]. In the case of the high-profile Panama Papers investigation, a customized system for document sharing and collaboration was highly successful in controlling the flow of information on that investigation until time of publication, across many different journalists, organizations, and countries [138], but not all journalistic efforts received that high level of resources or attention.

Another key digital-safety task of journalists is protecting their sources, which is particu-

larly challenging in the digital age [110, 135, 136, 168, 187]. Journalists' priority when communicating with sources is to use a method that is "most convenient for the source, including the platform on which [the] source is most likely to respond" [136]. This leads to a natural conflict between the critical need to communicate and the secondary priority of preserving the security of this communication, especially when potential consequences of a security breach could be as severe as imprisonment (for the source) or serious reputational harm (for the journalist). Entertainment journalists also have sources that may need protection from retribution, especially when covering topics like harassment [99], labor rights and employer practices [100, 192], and industry sexual misconduct [81]. For example, Klepek writes in a piece about noncompete contracts: "An engineer at Sony [...] asked to remain anonymous to protect their job security." [100].

An important element of journalists' management of digital safety threats is organizational support. As McGregor et al. describe, organizational stakeholders and practicing journalists have different—and sometimes conflicting—priorities [137]. Journalists and organizational stakeholders (like editors and IT staff) share core concerns like protecting sources and preventing reputational harm. However, sometimes competing goals put these two groups in conflict. Journalists may, for example, strongly value being able to collaborate through externally managed cloud services like Google Docs, but organizational stakeholders may worry about those external services being subpoenaed and revealing sensitive information [137, 139]. Journalists are highly motivated to accomplish their digital-safety goals, especially in high-stakes reporting scenarios like the Panama Papers [63, 138], so long as the value of these security practices is made clear and the system involved is usable enough for all parties.

Most prior literature focuses on journalists who are protecting sources from nation-states or similarly resourced adversaries. However, entertainment journalists also face direct attacks in the

form of harassment from groups of otherwise-ordinary individuals. Their attackers use freely-available services of the modern internet and social media in order to target and harass these journalists. I discuss online hate and harassment in the next section.

2.4 Online hate and harassment

Danielle Citron defines online harassment as “a persistent and repeated course of conduct targeted at a specific person, that is designed to and that causes the person severe emotional distress, and often the fear of physical harm.” [48]. This type of attack has become more common over recent years, due in large part to “unintended applications of widely accessible technologies” [210]. Although cybercrime uses some of the same tools and techniques—spamming, creating fake accounts, and compromising sensitive information—online hate and harassment is not profit-motivated. This means that techniques to disrupt the financial ecosystem [96, 115, 171] underpinning cybercrime that uses these tools are ineffective, since there is no dependency chain that one can disrupt.

Thomas et al. taxonomize several important features of modern online hate and harassment, categorizing attacks based on the intended audience, the medium, and the capabilities required [210]. They define the threat model for online hate and harassment as consisting of a *target*—used rather than “victim” in order to empower those who face abuse—and an *attacker*, whose goal is to “emotionally harm or coercively control the target, irrespective of other side effects.” In their accompanying survey measuring the prevalence of harassment in the global population, Thomas et al. also show that harassment is particularly prevalent and severe for young adults (ages 18-24) and LGBTQ+ people.

Attacks are differentiated according to three criteria: the intended *audience* of the attack, the *medium* through which the attacker reaches the target, and the necessary *capabilities* for the attack to succeed. This taxonomy is platform-agnostic, focusing on broad themes in order to remain stable regardless of changes to the relevant technologies.

In this taxonomy, a harassment attack will either need to be seen by the *target* or an *audience*, distinguishing undetected attacks like covert stalking from intentionally leaking personal information, like a person’s address. The *medium* of the attack refers to the harasser having a way to disseminate text, images, or other media to the target or audience, as compared to attacks like *overloading*, wherein attackers flood notifications on a platform to render it unusable (but do not need to directly communicate). The four *capabilities* outlined in this taxonomy further differentiate the method by which the attack is executed: an attacker may need to *deceive an audience*, to *deceive a third-party authority*, to *amplify* their attack via coordinated group action, or to have *privileged access* to information, an account, or a device.

Thomas et al. then further outline seven categories of attacks identified in prior literature, organized by the above taxonomy. For this thesis, three are particularly relevant. *Toxic content* “covers a wide range of attacks involving media [...] that attackers send to a target or audience [...], but without the necessity of more advanced capabilities.” This includes attacks such as bullying, threats of violence, and sexual harassment. Another primary attack is *overloading*, where “an attacker forces a target [...] to triage hundreds of notifications or comments via amplification [...], or otherwise makes it technically infeasible for the target to participate online due to jamming a channel.” Finally, *false reporting* “broadly captures scenarios where an attacker deceives a reporting system or emergency service [...]—originally intended to protect people—to falsely accuse a target of abusive behavior.” The impact of all of these attacks are discussed in relation

to entertainment journalists in Chapter 5.

Many other scholars have additionally examined the particular impact of harassment, especially in relationship to experiences of marginalization. Chadha et al. [40] describe how women employ a variety of strategies for dealing with harassment, which is often sexualized, that included simply normalizing it and self-censorship. Wei et al. [234] interviewed experts who provide advice for those facing harassment and found that generalizing advice for harassment is difficult, as it relies on the particular type of threat the target faces. Many experts echoed the advice described by the participants in Chadha et al. regarding self-censorship and limiting expression. Additionally, for people experiencing marginalization, these experts often added *extra* advice on top of pre-existing best practices, creating an unfair burden for those who face the most severe harassment. In other countries, especially those outside of the Western cultural sphere, harassment that damages one's reputation or the reputation of one's family was perceived as particularly harmful [191], and may escalate into direct physical or sexual violence [183]. The theme of restricting free expression as the only or best option in the face of harassment is a troubling one for researchers who seek to address this problem.

One specifically salient example of online hate and harassment for this dissertation is the #GamerGate controversy. In August 2014, game developer Zoe Quinn experienced targeted harassment due to a perception that their game *Depression Quest* “was lauded with awards, not because of excellent game design and execution, but because it symbolized the gaming world's movement to be more inclusive and progressive” [2]. Eron Gjoni, Quinn's ex-boyfriend, created the “Zoe Post,” alleging an affair between Quinn and games journalist Nathan Grayson, and shared it on the popular imageboard 4chan. This post served as the seed crystal for more severe attacks on Quinn. These attacks spread to journalists, like feminist games critic Anita Sarkeesian,

and game developers, like Brianna Wu, and eventually escalated to bomb threats and investigation by the FBI [2].

Crucially, this harassment demonstrates the impact of two *contextual risk factors* identified in Chapter 3—*prominence* and *marginalization*. Targets of #GamerGate experienced highly targeted harassment, due to their sudden and unexpected infamy amongst #GamerGate harassers. This newly focused set of attacks represented a significant change in circumstance, requiring new and often drastic measures to combat. The “average user” does not usually have to contend with focused harassment from an online group, so individuals who suddenly become *prominent* are often not equipped to deal with these threats. Additionally, this harassment was often “extremely offensive and derogatory” [2], employing graphic, misogynist threats aimed at women and minorities. Although it is true that not every target of #GamerGate was a woman—for example, Christopher Kushner, founder of 4chan, experienced targeted harassment when banning discussion of #GamerGate on the platform [106]—many of the attacks specifically relied on reinforcing or replicating prejudice against women and marginalized groups.

One key question of prior scholarship on #GamerGate is measuring how and why this controversy came to be. Trice examines organizational documents to attempt to reason about the actual size and scope of the #GamerGate audience, finding a contradictory set of community standards. As Trice writes, “The normalization of death threats within the organizational information represents an interesting pattern that emerged within in these documents, where largely normal instructions occasionally flourish unusual markers: cis white scum jokes, death threat normalization, be polite but accept nothing less then enforced radical reform” [214]. Chatzakou et al. [43, 44] describe Twitter users who engage in #GamerGate harassment as highly engaged (i.e., lots of social connections and activity) in comparison to other Twitter users, explaining in

part the controversy’s quick spread and enduring staying power.

Massanari describes how the features of Reddit’s platform policies, such as upvoting, the curation of Reddit’s homepage, and governance structure, might enable “toxic technocultures” [127]. At the time of that writing, Reddit had few systems set in place to systematically report hateful and abusive content, especially surrounding #GamerGate, with the best option being to report individual posts and users. However, since new account creation and new subreddit creation were very easy, even blocked users could easily reappear on the platform and continue to post hateful content. Salter [182] echoes this argument, using Marcuse’s framework of “technological rationality” [122] to describe how platform affordances encode certain values that permit or encourage online abuse. Chandrasekharan et al. [41] show that banning certain subreddits devoted to hate speech did reduce the level of hate speech on the platform overall, but some users may have simply migrated to other platforms that were willing to host that content. On Twitter, Trice and Potts show how #GamerGate activists were able to “turn the Twitter experience into an inescapable GamerGate experience” [215] by forcing targets to choose between either listening to their messages and suffering extreme online abuse, or leaving the platform entirely.

Games journalists themselves were also required to perform “paradigm repair” [163] in the face of this controversy; they were required to defend and uphold their standards of integrity against spurious claims of conspiracy while simultaneously defending the targets of this harassment. In an interview study with 17 participants conducted during the height of the #GamerGate controversy, Perreault and Vos describe games journalists giving themselves two primary roles: paternal figures who act deliberately in opposition to abusive behavior and traditional journalists who facilitated socially relevant discussion about video games. The paternalistic element here

refers to the practice of publicly criticizing gamers² for harassing female members of the industry. Games journalists in this study in fact acknowledged that *some* concerns about ethics in games journalism could be valid, particularly close relationships with industry, but distanced themselves strongly from practices like accepting gifts from game developers in return for favorable coverage. This practice sought to align games journalists with other good journalism, thus repairing the breach, spurious or not, to their reputation. This paper used discourse analysis to analyze interviews, as well as articles written by games journalists, to discuss how they defended their reputation. The authors did not examine specific technical strategies or other protective practices for harassment mitigation, differing from the work presented in Chapter 5.

#GamerGate has also been linked to larger cultural trends in online life. Feminist scholars link the rise of #GamerGate to concerns about feminism playing a greater role in discussion about videogames in the 2010s [78, 80, 148, 213]—although this is not a new topic in feminist scholarship [24, 61, 180, 236], it ironically rose to greater prominence in cultural conversations around digital games in part *due* to #GamerGate. Even more broadly, Bezio explicitly places #GamerGate as a precursor to the modern “alt-right” movement, especially through the controversy’s support by Milo Yiannopoulos and Breitbart [27].

This body of work emphasizes the difficulty of protecting oneself from a diffuse, cross-platform movement, and how modern affordances of social media permit the *overloading* and spreading of *toxic content* that was a core feature of #GamerGate harassment [210]. Although the platforms on which this harassment has taken place have adapted their policies over the intervening nine years, targeted harassment is still an ongoing problem. In this thesis, I use #GamerGate

²As Perrault and Vos point out, “gamers” is an unwieldy category, as 65 percent of American households played video games in 2023 [19]. I use it here for concision to describe the audience for games journalism.

and its influence on the landscape of online harassment to frame my work on entertainment journalists, many of whom were targets themselves during the height of that movement. In Chapter 5, I will discuss how my work, nine years later, builds on this scholarship.

2.5 Reflexive thematic analysis

The mode of analysis used in this dissertation is Braun and Clarke's *reflexive thematic analysis* process [33]. First described by Braun and Clarke in 2006 [31], in this work, I engage with their most recent perspectives on this methodology [32, 33]. This recursive, iterative analytical process takes place over six phases, which are paraphrased below, based on Braun and Clarke's 2021 textbook *Thematic Analysis* [33]:

1. *Dataset familiarization*: Deeply immersing oneself in the data, as well as thinking critically about how the data creates meaning
2. *Data coding*: Systematically working through the data and labeling discrete units of meaning.
3. *Initial theme generation*: Generating initial clusters of codes organized around a central organizing concept.
4. *Theme development and review*: Reviewing the initial themes for clarity, evidence, and importance.
5. *Theme refining, defining, and naming*: Figuring out the structure and flow of your argument through the process of revisiting and finishing the development of these themes.
6. *Writing up*: Communicating both thematic material and process.

This methodology posits that rather than excavating ideal truth from one’s data, researchers instead create meaning as an active, creative process through the work of interpretation. This interpretation is deliberately situated in the researcher’s inherent subjectivity, which is seen as a benefit and important part of the process, rather than a bias that should be removed.

Rigor is ensured, therefore, by describing the process and the particular situation of the researcher in relation to the work that they’re doing – the process and practice of reflecting on one’s assumptions and process and how they impact the research is *reflexivity*.

The phases outlined above are deliberately positioned as fluid, rather than a set of fixed steps. Analysis with a qualitative sensibility is inherently “constrained-but-open,” [33] meaning that it is driven by clear process, but that process is fluid and creative. Through constant reexamination of process and data, I aim to create and describe useful *themes* – patterns of shared meaning across the dataset that can illustrate the particular challenges and needs of at-risk users. For these projects, we span the spectrum of *inductive* to *deductive* analysis—in Chapter 3, we took a primarily inductive approach, focusing on the content of the data. The analysis in Chapters 4 and 5 included inductive themes and coding, but also used the framework developed in Chapter 3 as a deductive frame for understanding the needs of the at-risk populations studied in those chapters.

In writing this work, I take the position that while I bring a particular set of knowledge and experience as a computer security researcher, I do not have expertise on my participants’ experiences. Rather, the outcome of this work is to find the most productive union of the two – by combining the expertise and knowledge of the digital safety research community *and* the experiences and values of different at-risk populations, we can reach the most effective solutions for the unique issues they face, technical or not.

I advocate for increased focus on at-risk users' needs by technology creators and researchers during threat modeling, research, design, and development. Accounting for at-risk users can also elevate the digital safety of all users by making “more pronounced the need[s] that many of us have” [62]. Providing better digital-safety tools and guarantees can have far-reaching impact both to at-risk users and people in general. Additionally, providing choices and controls for at-risk users who know intimately the digital-safety threats they face can also benefit other users who may desire similar protections.

Each project in this thesis required a slightly different approach, so further methodological details and accountings of the reflexive thematic analysis process are found in each chapter.

Chapter 3: A Framework for Unifying At-Risk User Research

3.1 Introduction

This chapter describes a systemization of knowledge project developed by myself and my team of coauthors. The goal of this work is to take the past several years of scholarship on at-risk users, analyze it, and develop a framework for understanding this work to serve as an underpinning for further scholarship and technology development.

RQ1: *Contextual risk factors.* What factors—such as a person’s situation in society, relationships, or personal circumstances—contribute to digital-safety risks for at-risk users?

RQ2: *Interactions.* How do these contextual risk factors interact to elevate the risk or severity of digital-safety attacks for at-risk users?

RQ3: *Protective practices.* What protective practices are common across at-risk users when attempting to address their digital-safety risks?

RQ4: *Barriers.* What barriers do at-risk users encounter in protecting themselves from digital-safety risks?

3.2 Methods

We synthesize 95 research papers from a cross-section of computer science conferences. Here, we discuss how we identified and analyzed these papers.

3.2.1 Paper selection

Our dataset for this analysis was 95 papers describing digital-safety-related issues for various at-risk populations.¹ We collected papers from five years (2016–2020) of conferences spanning the security, privacy, and human-computer interaction (HCI) communities: CCS, CHI, CSCW, IEEE S&P, NDSS, PETS, SOUPS, and USENIX Security. We first gathered links to every paper from these conferences on DBLP.² From those links, we collected paper titles, abstracts, and publication dates, resulting in 6,534 papers.

To refine this list, three researchers independently read titles and abstracts for each paper and marked them as ‘relevant’ to our research questions or not. At this stage, we interpreted relevance broadly, selecting any paper even slightly within scope. Papers that no researcher marked as relevant were removed. Papers marked as relevant by only one researcher were reviewed by a fourth researcher and discussed. This process identified 127 potentially relevant papers.

Authors with extensive experience working with at-risk populations added 12 papers from other sources and/or from outside the target date range, in order to cover a broader range of populations, for a total of 139 potentially relevant papers.

¹The full list of papers can be found in Appendix A

²See <https://dblp.uni-trier.de/search>

3.2.2 Codebook development

Our goal was to identify contextual risk factors, protective practices, and other patterns discussed by the papers in our dataset. As a first step, we inductively built a codebook by analyzing, in detail, a subset of papers well-aligned with our research questions. Most of the core concepts in our framework were identified at this stage, although inductive refinement continued throughout our analysis.

To select this initial subset, we extracted from the dataset an initial list of populations (e.g., survivors of intimate partner abuse [129], refugees [200], activists [55], and children [245]). We also synthesized an initial list of risk factors: attributes of the population or the threats they faced that contributed to their digital-safety-related risks. We then selected our subset to ensure each population and risk factor on our list was represented, making sure to include some papers that combined multiple risk factors (e.g., low-income African American New York City residents [64] and foster teens [21]). This process yielded 27 papers.

We next analyzed these 27 papers and inductively built our codebook. We used the specific population(s) discussed in each paper³ as a unit of analysis [209]. One of four researchers read and summarized each paper. The full team used these summaries to iteratively build and refine our codebook [209], which included categories for risk factors, protective practices, and barriers to protection. We met throughout the process to develop, discuss, and refine codes. These detailed summaries also enabled us to examine relationships between codes and memo early ideas on themes [31].

³For example, Simko et al. [200] reported on interviews with refugees *and* associated caseworkers; we treated this as two populations.

3.2.3 Full analysis

Next, we used the codebook to analyze the remaining papers. Because our initial paper selection steps were deliberately inclusive, we continued to refine the dataset during this phase. Any researcher could flag a paper for possible exclusion if it did not address any of our research questions; determinations were made after discussion. Our final analyzed dataset included 95 papers (Appendix A).

We randomly selected 20% of the remaining papers, which two researchers independently coded, making minor updates to the codebook as needed. The researchers assigned codes to each paper and memoed additional relevant details and context [31]. After completing this 20% sample, the researchers calculated agreement using Krippendorff's α before discussing and resolving all disagreements. The researchers then repeated this procedure on another random 20%. Once sufficient reliability was reached, the researchers split the remaining papers between them to complete coding.

For the contextual risk factors and protective practices, we obtained $\alpha = [0.88, 1.00]$ after the second round of coding, and for the barriers, $\alpha = 0.83$ after the third round. As these are above the standard threshold of 0.8 for reliability [103], we then proceeded to individual coding. After coding all papers, the entire research team met to review the results, identify second-level codes and themes [31], and finalize our framework.

We conducted credibility checks of our framework and findings to verify they were accurate and clear. Seven experts who have worked with at-risk populations reviewed a version of this paper and met with us to discuss it. All experts found the framework sound and useful.

3.2.4 Limitations

Our 95-paper dataset is not exhaustive of all relevant papers published in the security, privacy, or HCI communities. However, given our systematic method of compiling papers, it should reasonably represent these communities’ published understanding of at-risk users. Also, reflecting the current state of literature from these communities, the dataset papers skewed heavily toward Western, and specifically U.S., populations. In addition, literature from other fields, especially the social sciences, could offer relevant perspectives on the digital safety of at-risk populations.

Biases in geographic representation. At-risk populations are found around the world, with social and structural circumstances that vary globally. The papers identified in our review skewed heavily toward Western, and specifically U.S., populations, but do include some studies with participants from other regions. About half of the papers in our dataset (48 out of 95 papers) included participants from the U.S. only, and 68% of papers included participants from Western countries only. Populations about which the literature reported perspectives from around the world (including most or all continents), included activists, journalists, NGO staff, crowdworkers, LGBTQ+ people, and survivors of trafficking. Other non-Western perspectives were reported for women in South Asia, older adults in Cambodia, people in developing regions (in Asian and African countries), refugees (resettled from Africa and Asia to the US), and people with visual impairments in India. The remaining populations listed in [Table 3.1](#) represent Western perspectives.

Researcher reflexivity. Though we do not know how the authors of dataset papers would identify their nationalities or describe their lived experiences, the majority appeared to be completed by

teams at institutions in Western contexts. Similarly, we work in U.S.-based organizations and have backgrounds in HCI, computer science, security, and/or privacy. These contexts influenced the papers we sampled and the analysis lenses used—both by our team and the researchers who produced dataset papers. While this SoK establishes a framework, we advocate for additional perspectives to enrich our findings.

At-risk user research is still evolving. As research methods and best practices for understanding at-risk users are still being developed, often differing from one community to the next, the papers in our dataset also often did not focus on the same issues or investigate to the same depth. As a result, our synthesis of the contextual risk factors, protective practices, and barriers covered in this paper may not reflect all the challenges the population in question experiences. Our coding is, instead, a reflection of the current understanding in the sampled literature on each population.

Despite these limitations, we believe this work serves as a critical first step towards recognizing contextual risk factors and protective practices that span at-risk users. We advocate for future work that builds on this framework by including broader literature and cultural perspectives.

3.3 Contextual risk factors

In our meta-analysis, we identified 10 contextual risk factors that augmented or amplified digital-safety risks. These risk factors, which form the first component of our at-risk framework, include: three *societal factors*, influenced by an at-risk user’s role in their society and culture; three *relationship factors* stemming from who an at-risk user knows or interacts with; and four *personal circumstances* dependent on who an at-risk user is or their personal or professional

activities. We note some common attacks for each risk factor, but do not consider being the target of a digital attack alone a risk factor.

We capture the presence of these risk factors within the papers from our dataset in Table 3.1. A black circle (●) indicates that we identified at least one prior study of an at-risk population that reported risks related to that factor. We caution that the absence of a black circle for an at-risk population in Table 3.1 does *not* imply that the risk factor is irrelevant to the population, only that it was not reported in our dataset.

Provisional categories [130] of at-risk populations represented by our dataset are listed in the first column of Table 3.1. While these population categories flatten some of the demographic richness reported in individual dataset papers, where possible we unpack intersectional issues using our thematic analysis and examples that include more detailed participant descriptions throughout the paper. For example, the “activists” category in Table 3.1 describes activists across four continents, but examples note the specific geography when relevant. Further, these population categories are not exhaustive; instead they represent our dataset, demonstrating how risk factors can differ across categories and intersect for a single category.

Next, we describe each contextual factor, focusing on the nature of the risk and types of associated attackers⁴ and harms, as applicable. We also explore how risk factors may intersect to create more severe risks.

⁴We use *attacker* broadly to refer to anyone who introduces digital-safety issues for an at-risk user, regardless of the severity or intention.

Population Category	Citations	Societal factors			Relationships			Personal circumstances		
		Legal or political	Marginalization	Social norms	Relationship with the attacker	Reliance on a third party	Access to other at-risk users	Prominence	Resource or time constrained	Underserved accessibility needs
Children	[74, 76, 94, 104, 105, 111, 141, 143, 149, 245]		●		●				●	
Teens	[74–76, 140, 237–239]		●		●	●				●
Foster teens	[21]				●	●				●
Older adults	[71, 90, 142, 152]					●		●		●
Activists	[10, 55, 101, 123, 208]	●						●	●	●
Activists × Transgender people	[114]	●	●					●		●
People involved with political campaigns × US	[50]	●				●		●	●	●
Teachers	[105]					●			●	
Journalists	[63, 136–139]	●				●		●	●	●
Entertainment journalists	Chapter 5		●					●		
Sex workers	[23, 206]	●	●							
ER staff	[205]								●	●
NGO staff	[45, 112]	●				●		●		●
Library IT staff	Chapter 4, [118]	●				●			●	
Crowd workers	[243]				●			●		
LGBTQ+ people	[29, 30, 47, 85, 114, 188]		●							
LGBTQ+ people × With HIV	[230, 231]		●							
Marginalized racial group × US	[212]		●							
People with an illness	[164, 181, 190]		●							
Older adults × With cognitive impairments	[26, 51, 71, 134, 146]				●			●	●	
People with visual impairments	[6–8, 65, 88, 228]			●	●					●
People with other or multiple disabilities	[125, 167]		●							●
Non-Western culture × Women	[9, 60, 183, 184, 222]		●	●	●	●		●		
Developing regions × Older adults	[93]	●		●						●
Developing regions × Low SES	[3–5, 144, 178, 221]	●	●	●	●			●		
Developed regions × Low SES	[102, 174, 175, 202, 224, 232, 235]		●	●	●			●		
Developed regions × Low SES × Marginalized racial group	[64]		●	●				●		
Undocumented immigrants	[82]	●	●							●
Refugees	[200]			●		●		●	●	
People involved with armed conflict	[199]	●								
Survivors of sexual assault	[15, 156]		●		●	●				
Survivors of intimate partner abuse	[42, 68–70, 87, 129]		●		●	●		●		
Survivors of trafficking	[45]	●			●	●		●		

Table 3.1: Meta-analysis showing how the contextual risk factors from our at-risk framework apply to categories of populations in our dataset. A black circle means one or more citations indicated the contextual risk factor was relevant. Note that the absence of a black circle does not mean the population does not have that digital-safety risk. Rather, it means the set of papers cited did not discuss or explore that risk factor as defined in this framework.

3.3.1 Societal factors

The first set of contextual risk factors involve *societal factors* amorphously driven by cultures and institutions. Attacks related to societal risk factors tended to be *diffusely targeted*, i.e., directed toward anyone in a population or an entire at-risk population simultaneously, rather than at a specific person.

Legal or political. The government, political affairs, or laws of a country can contribute to at-risk populations experiencing heightened digital-safety risks, including potentially sophisticated attacks from government actors. A key theme associated with this factor was the power differential between government or quasi-governmental actors and the targeted populations.

Government or quasi-government actors may be able to intercept communications from at-risk populations in various ways, such as physically seizing devices or data [112, 136], impersonating trusted entities [123], coercing platform or telephony providers to bypass security measures [185], or preventing internet access entirely [55]. For example, in 2019, the Sudanese government shut off the country’s mobile data network to make organizing for activism as difficult as possible [55]. Similarly, the International Committee of the Red Cross—an NGO that collects information that could be used by armed groups for non-humanitarian intelligence—reported being obligated to physically surrender devices to meet with those armed groups [112].

Governments may also be able to enact surveillance, leading to real or perceived threats of monitoring. For example, undocumented immigrants in the U.S. reported concern about posting their activities on social media, due to perceived government monitoring [82]. More generally, residents of several countries with government-controlled internet surveillance have reported modifying their behavior [93, 184, 222].

Because of the power differential, threats associated with this factor may also escalate into offline harms, such as detainment, incarceration, or deportation [55, 82]. These harms may also have wide-ranging societal impacts, including restricted or self-censored speech [101, 208] or lasting damage to trust in public institutions or figures [50, 137]. For example, attacks on people involved with U.S. political campaigns were described as intending to undermine the institution of U.S. elections [50].

Marginalization. Pervasive negative treatment or exclusion at a societal level, due to an individual's identity attributes or life experiences, may also elevate digital-safety risk.

Attackers may target at-risk populations with online hate and harassment due to their beliefs, identity, or social status, such as people who are LGBTQ+ [29, 30, 85, 114, 188, 230, 231], undocumented immigrants [82], marginalized racial and ethnic groups [64, 212], or people with low socioeconomic status (SES) [64, 102]. These threats were typified by a broad set of potential attackers, and contending with hate and harassment led to emotional distress [30, 212, 222]. For this factor, past research has described the perception that anyone may be a potential attacker [188], exacerbated by the fact that attackers may hide behind anonymous identities [183, 212].

Harassment experienced by marginalized populations can be subtle or even unintentional. For example, in To et al. [212], people in marginalized racial and ethnic groups have reported experiencing widespread microaggressions pertaining to their race, via their interpersonal interactions with others on- and offline. Given the prevalence of this threat, some populations may be reluctant to fully participate online due to the risk of revealing a marginalized characteristic [15, 23, 70, 70, 156, 190, 230, 231].

Given that this factor can be associated with stigmatized information, attackers may also

coerce at-risk users by threatening to leak information that could be harmful. This can lead to reputation damage [222, 230] or sexual violence [82, 183]. For example, attackers may threaten to “out” transgender individuals [114] or in conservative regions, may threaten to leak chat logs between a man and a woman to damage the woman’s reputation [183, 222]. Algorithmic bias may also create or exacerbate digital-safety risks in this space. For example, automated gender recognition systems may misgender transgender people [85].

Social norms. Social norms are informal rules that govern behavior in society. Some norms can particularly restrict options for members of an at-risk population, leading to increased digital-safety risks when there is a mismatch between the expectations of technology creators and the lived experiences of that population.

For example, technology creators may assume that devices and accounts are personal and private. However, device-sharing norms mean that this assumption may not hold for a variety of at-risk populations, limiting the privacy protections afforded by private personal devices and accounts. For example, Sambasivan et al. and Ahmed et al. found that women in South Asia were expected to share devices or accounts with family members [3, 184]. Similar device sharing norms were also found among rural women in Greenland [235].

Technology creators may also assume that users understand implicit norms around when it is appropriate (or dangerous) to share personal information. However, some at-risk populations face changing norms, which can lead to unexpected threats. For example, alongside numerous other challenges, refugees must adapt to the technology norms and associated digital-safety risks of their new country of residence. For refugees in the U.S. specifically, certain types of information, such as Social Security numbers, were more sensitive than expected, placing refugees at

increased risk of financial harm until they learned when it was appropriate to disclose this type of information [200].

3.3.2 Relationships

The second set of risk factors is driven by the *relationships* of at-risk users, including direct relationships with an attacker and relationships with a third party. These risks tended to be associated with *focused targeting*, in which attackers pursue specific at-risk users (often with intense motivation).

Relationship with the attacker. A personal relationship with an attacker can lead to heightened risk of what Levy and Schneier [116] classify as “intimate threats.” It can be particularly difficult for individuals targeted by intimate threats to prevent and detect harm, because the attackers are likely to have physical access to the target’s devices and accounts, implicit or explicit authority over the target, and potentially detailed knowledge about the target they can leverage [116]. Intimate threats can involve a wide range of digital-safety attacks including surveillance, device or account compromise, destruction of data or devices, harassment, and more [116].

For example, survivors of intimate partner abuse [68, 70, 87, 129] were reported as often facing relentless attacks from abusers aiming to limit the survivor’s autonomy. These abusers may have physical or digital access to the survivors’ devices and accounts, including surveillance via coercive physical access [129] or spyware [42, 70]. Similarly, survivors of trafficking reported worrying about being recaptured because, prior to escape, the trafficker had full access to the survivor’s digital life [45]. Stalking attacks are also often characterized by a *relationship with the attacker*, as “the majority of stalkers are ...obsessionally focused on a specific person with

whom they have had some previous relationship,” often an intimate one [58]. Workers on online crowdsourcing platforms also experienced a severe power imbalance with requesters on those platforms, who could decide whether or not to pay workers based on their responses, and could gather detailed information on them via completed tasks [243].

Reliance on a third party. Individuals may be at risk due to their *reliance on a third party* for safety-focused care or help with essential tasks. Unlike intimate threats associated with a *relationship with the attacker*, the third parties in this context typically have helpful or safety-focused reasons for “privacy invasions” [116]. However, even with generally supportive intentions, these invasions can increase the individual’s digital-safety risks due to reduced privacy (from the third party), leading at-risk users to feel uncomfortable or that their autonomy is limited [146]. These types of risks can also increase the attack surface (e.g., through the third-party), or leave at-risk users vulnerable to attackers that impersonate the third party.

For example, children and teens often shared their devices and personal data with parents or caregivers [104, 141], or had applications or privacy settings enabled that allowed parental monitoring [75, 76, 237]. People with visual impairments sent photos to crowdsourced services for object identification, but also worried that there might be sensitive content in the photos (e.g., credit card numbers) [8]. Older adults with cognitive impairments depended on help from caregivers for digital tasks [71, 146]. Refugees relied on case workers for processing personal data, applying to jobs, or applying for support [200].

Access to other at-risk users. Having access to another at-risk person or population (the “primary target”) can also put someone at greater risk of focused, stepping-stone attacks that ultimately aim to harm the primary target. For example, children of survivors of intimate partner

abuse may be targeted by abusers in order to regain access to the survivor [129], journalists may be targeted to try to gain access to their sources [63, 136–139], and elementary school teachers may be targeted to gain access to their students [105].

3.3.3 Personal circumstances

Our third set of contextual risk factors cover *personal circumstances* that can increase risk, such as public prominence or socioeconomic constraints. All of these factors tended to involve *diffuse targeting* of groups, except *prominence*, which involved *focused targeting* of specific individuals.

Prominence. At-risk users who stand out in a population, because they are well-known publicly or have noticeable attributes (e.g., accomplishments, outspokenness, attractiveness, etc.), may face heightened risk associated with their prominence. *Prominence* can expose individuals to new attacks, like *focused targeting*. Celebrities, for example, may experience parasocial relationships with audiences, increasing the risk of stalking or personal information leaking [240].

Prominence may also exacerbate risks associated with *legal or political* factors or *marginalization*, a dynamic identified for prominent journalists [63, 136–138], politicians [50], NGO staff [112], and activists [55, 123]. In these cases, the level of prominence influenced the severity of risks. For example, people involved in a national U.S. political campaign were more at risk of focused attacks than campaigns receiving only state or local attention [50].

Resource or time constrained. Populations can experience elevated digital-safety risks and decreased ability to respond to digital-safety threats if they have limited access to technology or other resources (e.g., money, devices, connectivity), or limited ability to set aside time to cope

with risks. At-risk users experiencing this risk factor face constraints that go beyond what is typical. Individuals with low-SES, for example, may not be able to afford private personal devices, or may need to rely on old devices that can no longer receive security updates [64,202,221,224].

At-risk users can also face extraordinary time constraints. Political workers, for example, operated within extremely limited election timelines, which made it challenging to set up the security infrastructure needed to counter nation-state attackers [50]. Hospital emergency departments similarly “have strong availability demands ... and must provide services as quickly as possible” [205], leading to password reuse and use of unsecured personal devices in a sensitive work context.

Underserved accessibility needs. Some at-risk users have accessibility needs that are underserved by current technology, contributing to their digital-safety risks. In our dataset, this included accessibility needs due to a disability, neurodiversity, a language barrier, or developmental maturity. Inaccessible technology can cause anxiety about, and susceptibility to, potential attacks.

Members of populations experiencing this risk factor have described general anxiety about falling prey to “hackers” or vague bad actors, as well as worries about their ability to effectively protect themselves with existing technology. For example, some older adults reported asking trusted sources for digital-safety help because they did not feel confident protecting themselves [71]. This can stem from negative past experiences or from inaccessible language in online resources about digital safety [152]. Assumptions built into systems were unrealistic for older adults with mild cognitive impairments, who were sometimes unable to remember passwords and other crucial information [146]. In some cases, they had trouble remembering whether they made a particular purchase [134], making it difficult to differentiate an attack from a memory lapse.

Similarly, people with disabilities may struggle with a lack of accessible technology for some tasks [6–8, 65, 88, 125, 228]. For example, people with visual impairments reported to Wang et al. the conflict between the need to use screen readers in public and concerns about eavesdropping and safety [228].

Additionally, refugees with developing English skills reported struggling with language accessibility, finding it difficult to distinguish legitimate callers from scammers when the call was in English [200]. Zhao et al. found that children do not completely understand certain online privacy risks due to their age and development [245].

Access to a sensitive resource. Access to a sensitive resource (e.g., sensitive data, credentials, money) can increase the risk of attacks aimed at co-opting this access. In most cases, the at-risk users’ professional activities provided them privileged access to these resources.

For example, emergency department staff may be targeted for their access to patient medical data [205], journalists for their access to original source material, such as legal documents or financial records [63, 136, 137], and executive staff for their ability to approve wire transfers [59].

3.3.4 How do contextual risk factors interact?

Most at-risk populations in Table 3.1 experienced more than one contextual risk factor. Each factor contributed to risk on its own, but risk factors also combined to yield new digital-safety risks or amplify existing risks. Thus, we argue that technology creators and researchers should consider all of an at-risk population’s risk factors together when possible. This is related to prior work on intersectionality [54, 189], which considers how multiple marginalized identities or circumstances can combine to create unique modes of discrimination. We provide an illustrative

list of examples from our thematic analysis below, chosen because they represented experiences reported in multiple papers across populations and/or provided understandable illustrations of how risk factors could combine to yield new risks or amplify each other.

Prominence added focused targeting to other risk factors. When added to other risk factors, *prominence* appeared to make it more likely that an at-risk user would experience *focused targeting* that amplified their other risk factors. For example, political campaign workers had *access to a sensitive resource*, but *prominent* politicians were more likely to be subjected to *focused targeting* to access those sensitive resources [50]. Similarly, while many transgender people reported experiencing *marginalization*, the *prominence* of transgender activists resulted in highly targeted hate and harassment attacks [114].

Resource or time constraints made it harder to cope with other risk factors. Populations who are *resource or time constrained* were more likely to experience worse outcomes pertaining to their other risk factors, because they did not have the time, money, or other resources to effectively protect themselves or recover from attacks. For example, the primary risk factor for survivors of intimate partner abuse was typically their *relationship with the attacker*, but prior research has emphasized that survivors with low SES had particular difficulty protecting their (and their children's) digital lives [129].

Underserved accessibility needs and reliance on a third party often combined. In our dataset, at-risk users with *underserved accessibility needs* also tended to *rely on a third party* for care or help with technology, at least sometimes. This included children, teens, older adults, people with visual impairments, and refugees, as shown in Table 3.1. For example, older adults experiencing mild cognitive impairment often forgot important passwords and information (related to *under-*

served accessibility needs), which reinforced the need to share these passwords and information with caregivers (subjecting them to risks associated with *reliance on a third party*) [146].

3.4 Protective practices

At-risk users employed practices they perceived would help them prevent, mitigate, or respond to digital-safety risks. These *protective practices* form the second part of our at-risk framework. The risk factors presented above helped drive user decisions about which protective practices to use, but these practices were not always ideal or even effective. Protective practices involved tradeoffs and highlight barriers to technology use (which we explore in Section 3.5). Different users weighed the pros and cons differently, sometimes leading to seemingly contradictory choices. We catalog these imperfect practices to show what at-risk users currently do given their risks, and to set the stage for a discussion of barriers to protections in Section 3.5, both of which provide context for how to design technologies intended to support at-risk users.

Our meta-analysis identified three categories of protective practices: *social strategies* where at-risk users relied on their social connections to respond to threats; *distancing behaviors* where at-risk users distanced themselves from, or entirely abandoned, certain accounts and technologies; and *technical solutions* that involved leveraging technical tools and mechanisms to prevent or respond to threats. We found these strategies were not mutually exclusive, with at-risk users commonly relying on multiple strategies simultaneously.

3.4.1 Social strategies

At-risk users frequently relied on social connections (in-person or online) to overcome digital-safety threats. This included relying on family or peers for trusted advice and support, vetting the identities of people they interacted with online, and controlling social interactions to minimize harms.

Informal help from trusted family and peers. A popular protective practice among at-risk users experiencing the *marginalization*, *social norms*, or *underserved accessibility needs* factors was to informally seek help from trusted family and peers. The at-risk user may be seeking direct help from someone perceived to be more knowledgeable or resourced, or may simply be seeking emotional support. For example, children and teenagers reportedly sought help from parents to understand security warnings [94] or deal with strange, scary, or confusing internet experiences [238, 239, 245]. Older adults who self-identified as having low technical understanding sought assistance from family members when addressing digital-safety concerns [71, 152]. At times, women in South Asia relied on family members for emotional support when harassed online [183, 184]. Informal support sometimes also came from anonymous peers, such as transgender individuals relying on social media to connect with other LGBTQ+ peers for emotional support when they experienced hate and harassment [114, 188], or people in armed conflict zones connecting to share critical information [199].

Formal help from trusted organizations. Public recognition of digital-safety threats facing at-risk users—particularly with respect to the *legal or political*, *marginalization*, *relationship with the attacker*, or *resource or time constrained* risk factors—has prompted trusted organizations to

offer support. For example, organizations that assist survivors of trafficking and intimate partner abuse, including NGOs [45] and government agencies [68, 70, 87], have provided assistance setting up digital-safety protections that are interwoven with continuous care [87]. Similarly, NGOs have assisted refugees with technology-required tasks, such as applying for work or submitting legal documents [200]. Public institutions, like libraries, have provided assistance and access to computers and connectivity to people with low SES [224]. While at-risk users may rely on these institutions for a broad variety of technical assistance [200, 224], digital-safety support was often explicitly reported as a key benefit. This assistance could incur additional digital-safety risks due to *reliance on a third party*, but the information and aid provided was often critical, both for digital safety and other basic needs.

Formal support may also stem from dedicated professional resources, particularly for those who have *access to a sensitive resource* or *access to other at-risk users*. For example, journalists have benefited from institutional support and accumulated best practices, as exemplified by the International Consortium of Investigative Journalists' (ICIJ) tools for digital safety in the Panama Papers investigation [138]. Elementary school teachers using technology in classrooms—who have a responsibility for ensuring the digital safety of their students—have sought assistance from school media specialists or librarians [105]. Political campaign workers in the U.S. have received training from organizations such as Defending Digital Campaigns and the Center for Democracy & Technology [50].

Vetting identities to avoid potential attackers. At-risk users who experience *marginalization* or *legal or political* risk factors have a documented need to verify that new people they encounter online are safe to associate with or to admit into a private space. For example, social media com-

munity moderators for transgender communities of color asked potential new members questions about topics relating to race and LGBTQ+ issues to help protect their spaces from potential attackers [114]. Women from non-Western cultures scrutinized online profiles to prevent men from infiltrating their private discussion spaces and subjecting them to sexual harassment [183]. Sex workers informally shared lists of abusive clients and aggressors to be avoided [23]. Undocumented immigrants restricted communication with untrusted parties until sufficiently vetted, due to concerns of impersonation by the police or immigration enforcement [82]. Activists used physical in-person meetings or mutually trusted peers to vet new individuals before adding them to sensitive group chats [55]. Unifying these practices is a lack of dependable digital signals of authenticity or identity, resulting in at-risk users relying on ad-hoc techniques to establish trust.

Preemptive disclosure for control. At-risk users experiencing the *marginalization* or *legal* or *political* risk factors sometimes proactively revealed sensitive personal information in a controlled manner to disempower attackers and preempt future emotional harm. For example, some men seeking men on dating apps opted to publicly disclose their HIV-positive status, both to simplify navigating sexual negotiations and as a path toward destigmatization [230, 231]. Some transgender people chose to “out” themselves to avoid the emotional stress of maintaining a secret identity and the risk of coercion or extortion [114]. Some people with disabilities—especially those with “visible” disabilities—chose to disclose these disabilities on dating apps, to reduce the chance of undesirable connections [167]. Some activists engaged in an “anti-surveillance tactic of openness” by collapsing public and private personas, rendering information they shared—like political beliefs—unusable against them in a separate context if leaked [185]. This practice is not applicable to all at-risk users, as it requires a willingness to be hyper-public, which carries

potential risks akin to those that *prominent* at-risk users face.

Social pleas. Some at-risk users who experience the *marginalization* or *social norms* risk factors reported reaching out directly to attackers with a plea to cease activities, such as content leaks or harassment. For example, some women in South Asia either independently, or with the assistance of family, asked attackers within their community to cease online sexual harassment [183, 222]. Similarly, some teens engaged with peers [239] or parents [149] to have embarrassing content removed from social media. This practice leverages the attacker’s empathy to help resolve digital-safety threats.

3.4.2 Distancing behaviors

The second category of protective practices in our framework involves *distancing*, or limiting use of technology. Distancing behaviors were prevalent in our dataset, used by most populations. The fact that limited or non-participation often felt like the safer choice may perhaps be a troubling signal to technology creators. Here we highlight two common themes.

Censoring online sharing. Some at-risk users carefully self-censored personal content they shared online. This was especially common for users experiencing the *marginalization* risk factor, who often did not feel safe sharing personal information or did not want to reveal stigmatized aspects of their lives to broad online audiences. For example, some LGBTQ+ parents refrained from sharing family or personal photos to try to avoid “outing” themselves beyond specific social circles [30]. In another case, some HIV-positive men seeking men reported not sharing their HIV status on dating apps to avoid unwanted stigmatization [230]. (As noted in the previous section, some members of this population did preemptively disclose their HIV status, demonstrating that

protective practices within a population can vary.)

Reducing one’s digital footprint. Some at-risk users dramatically reduced technology use to avoid attackers with extensive access, knowledge, and power, as was often the case with the *relationship with the attacker* and *legal or political* risk factors. For example, Chen et al. [45] found that survivors of human trafficking took relatively extreme measures to try to protect their new location, including abandoning devices, deleting social media accounts, and severely restricting online sharing. Survivors of intimate partner abuse employed similar measures when they suspected devices or accounts may have been compromised by their abuser [69]. As another example, activists in the Sudanese Revolution, concerned about the government confiscating their devices, used coded communications and regularly deleted sensitive data [55].

3.4.3 Technical solutions

At-risk users also employed technical solutions to protect their digital safety, such as specialized software or settings in common apps and services. Across these strategies, at-risk users did their best to protect themselves based on their knowledge and experience; however, the technical solutions they chose did not always provide the desired protections.

Secure communication and encryption. At-risk users experiencing the *legal or political* risk factor, such as journalists, activists, or politicians, commonly reported using encrypted messaging platforms. These at-risk users frequently wished to secure their communications against monitoring by a nation-state (or similarly resourced) attacker. For example, journalists working together to report on the Panama Papers used PGP as part of organizational security policies established by the ICIJ [138]. Similarly, journalists and activists in varied international contexts reported us-

ing encrypted chat apps to communicate over networks directly controlled by their nation-state adversaries [55, 185]. Transgender activists in the U.S., particularly those engaged in more *prominent* activities, similarly used encrypted chat to protect their communications [114]. NGOs, who had *access to other at-risk users*, sought to protect those at-risk users by using encrypted email for internal communications [45].

Beyond encrypted messaging, some at-risk users experiencing the *legal or political* risk factor reported encrypting files or entire devices to protect content from attackers—including nation-states and other focused attackers—who might gain physical access. At-risk users employing this technique included journalists, human rights organizations, and activists [112, 123, 136].

Strong(er) authentication. Strong authentication is commonly advised for internet users in general [157, 177], but can be especially important for users at higher risk of device or account compromises. While truly strong authentication practices were rare in our dataset, we identified a few authentication trends in subsets of select populations.

While frequently changing passwords is no longer considered a best practice generally [157], some at-risk users did this to address specific practical concerns. For example, at-risk users who have a *relationship with the attacker*, such as survivors of intimate partner abuse or trafficking, coped with attackers who may have ready access to their passwords (e.g., through coercion, co-presence leaks, or remote surveillance [68, 129]) by regularly changing passwords on accounts known to the attacker [45, 68]. Some users with visual impairments also chose to change passwords frequently in case others might have observed them inputting their passwords [6].

While not mentioned frequently in our dataset, our analysis revealed some cases of successful adoption of two-factor authentication (2FA). For example, in Matthews et al. [129], only

a small number of survivors of intimate partner abuse in their study reported using 2FA to protect their accounts, despite highly motivated attackers. Only a minority of participants involved with political campaigns reported adopting the strongest form of 2FA, despite commonly experiencing phishing attacks [50]. Journalists collaborating on the Panama Papers (*access to a sensitive resource*) were required by the ICIJ to use 2FA [138], though journalists are more broadly described as having limited awareness of 2FA’s benefits [137].⁵

Privacy settings and access control. We identified three categories of privacy settings (within widely used apps and services) that were discussed by multiple papers across a variety of at-risk populations and risk factors: location privacy [45, 60, 71, 129], social media visibility settings [45, 60, 71, 93, 129, 245], and blocking undesired contacts [60, 71, 129, 178, 183]. For users with a *relationship with the attacker* or *access to other at-risk users*, these privacy settings were commonly discussed as protecting highly sensitive information that, if obtained by an attacker, could compromise safety [45, 70, 129]. For example, some women in South Asia used platform controls to block unwanted contact from abusers on social media [183].

For at-risk users in professional settings with *access to a sensitive resource*—such as NGO staff, political campaign workers, and journalists—access control settings were also important for limiting who could access sensitive digital resources [50, 112, 138].

Online identity management. Some at-risk users reported attempting to hide their identity online via careful account management, including multiple and pseudonymous accounts.

Some at-risk users used multiple accounts or devices to maintain boundaries between different facets of their identities. *Marginalization* was one driver for this behavior; for example,

⁵We note that 2FA has become more commonplace since some of these papers were published.

sex workers (whose profession was stigmatized) reported keeping separate work and personal accounts [23]. Alternatively, survivors of sexual assault used throwaway accounts to seek support online without revealing their identities [15].

At-risk users also used multiple accounts or devices in response to potential account hijacking or surveillance. Survivors of intimate partner abuse, whose *relationship with the attacker* tended to give the attacker opportunities to access their accounts or devices, reported creating new accounts or purchasing new devices to avoid revictimization after escape [129]. Users experiencing *legal or political* risks, such as activists, reported using multiple accounts and devices to protect themselves from potential nation-state attackers. For example, activists in the Sudanese Revolution reported using SIM cards from other countries, creating fake U.S. phone numbers online, and asking relatives and friends overseas to verify social media accounts in an effort to thwart government surveillance [55]. Marczak et al. found similar evidence of an activist using multiple SIM cards to try to prevent governments from linking calls made in different countries and contexts [124].

In some cases—particularly cases of *marginalization*—rather than use entirely new accounts, at-risk users reported ad-hoc pseudonymity strategies for existing accounts. For example, in order to prevent personal photos being used for digital abuse, some women in South Asia chose neutral images, like flowers, as profile photos [183]. This strategy was shared by some transgender people, particularly activists (who also experienced elevated risk due to *prominence*) [114]. Low-income Black Americans similarly reported sometimes using emojis rather than contact names to protect their contacts’ identities on devices they thought might be compromised [64].

Network security. Only two at-risk populations in our review were reported to have used net-

work security tools—like Tor or virtual private networks—to disguise their web traffic. Certain activists [55, 123] and journalists [136] were aware of and used these technologies, but these tools did not appear elsewhere in our dataset—and not every user in those populations found these tools equally useful or usable [123, 137]. This could mean that other at-risk populations are largely unaware of these tools, that existing tools are not perceived as useful for the particular digital-safety concerns of other at-risk populations, or simply that researchers have not fully investigated this question with other at-risk populations.

Tracking and monitoring applications. Caregivers of at-risk users who *rely on a third party* sometimes used tracking or monitoring applications with the goal of protecting the at-risk user’s digital safety. Parents of foster teens reported using router settings to limit internet access at certain times of day and parental control apps to watch for potentially dangerous behavior [21]. These applications can create conflict between autonomy and privacy, but some children and teenagers indicated that they can be helpful, particularly when used as part of an ongoing dialogue with parents [74, 76].

3.5 Barriers to protective practices

In the previous section, we documented a range of protective practices at-risk users employed, motivated by the risk factors they experienced. In our thematic analysis, we also identified a variety of barriers that limited or prevented at-risk users from effectively adopting these practices. We apply the contextual risk factors and protective practices of our at-risk framework to discuss three categories of barriers: *competing priorities*, *lack of knowledge or experience*, and *broken technology assumptions*.

3.5.1 Competing priorities

Digital safety is not and cannot always be the top priority for all users; this is perhaps even more true for at-risk users with competing, often critical, needs. Some of the many competing priorities that appeared in our dataset included basic needs like food, income, or physical health and safety; social participation and compliance with social norms; and caring for others. We found that specific competing priorities were often associated with particular risk factors.

It is well understood that in general, users prioritize convenience, simplicity, and their tasks and goals over digital safety [1, 89]. Our analysis revealed a tendency for at-risk users to have layered or more severe conflicts between digital safety and other needs, making the leap to safer behaviors especially difficult.

Basic needs. People who are *resource or time constrained* in our dataset often prioritized other critical needs over digital safety, despite the potential for increased digital-safety risk. For example, Elliott and Brody [64] reported that low-income Black New Yorkers used apps to find cheaper food, despite suspecting the apps were insecure. Similarly, people experiencing homelessness described discomfort using public Wi-Fi, but used it anyway for critical needs such as applying for government assistance, housing, and jobs [202].

Relatedly, people who *rely on a third party* often gave up control of digital safety to accomplish basic tasks. Refugees, for example, shared account information, including passwords, with caseworkers to obtain social services or apply for jobs [200]. Some people who are visually impaired used crowdsourced assistive technology for tasks like identifying medicines correctly, despite the risk of exposing sensitive content in the background of the photos crowdworkers would evaluate [8]. Friik et al. [71] found that some older adults were willing to accept in-home

surveillance, giving up privacy to maintain some autonomy and independent living.

Participation and connection. As noted in Section 3.4.2 above, people who experience *marginalization* sometimes chose to distance or fully disconnect in order to keep themselves safe. Members of these populations who opted to engage were often reported as knowing that it increased their risk. For example, Blackwell et al. reported on LGBT+ parents implicitly “outing” themselves via “everyday” social media posts about their children and families [30].

We noted similar behavior related to *legal or political* risk, where the need to communicate with other activists, both to organize events and simply to be part of a community, motivated potentially risky modes of communication [185].

Similarly, maintaining good social standing in a family or community sometimes required people experiencing the *social norms* factor to deprioritize digital safety. Our dataset included, for example, women in South Asia and Saudi Arabia, as well as people in South Africa, sharing accounts, devices, and credentials with family members to meet cultural or community expectations [9, 178, 183, 184]. Some users in South Africa accepted Facebook friend requests from strangers, even though they were uncomfortable with the resultant potential for sharing, to avoid being rude [178]. Teenagers noted that they did not discuss or share risky digital-safety experiences with their parents to avoid awkwardness or a possible negative reaction [238].

Caring for others. Another example of competing priorities involved at-risk users taking on additional risk to help, care for, or support others. We noted this occurring frequently in connection with the *marginalization* factor, in part to reduce overall stigma via normalization. For example, Warner et al. found that some gay men disclosed their HIV+ status on a dating app in part to increase visibility [230]; similarly, separate studies found that LGBTQ+ parents and transgen-

der individuals posted on social media, despite possible risks, in part to increase visibility and reassure others that they were not alone [30, 114].

Our analysis revealed cases in which at-risk users prioritized others’ autonomy over their own digital safety, particularly in the case of *access to other at-risk users*. McGregor et al. [137] quoted a journalist who valued avoiding “impos[ing] any kind of burden on a source” over the journalist’s own digital safety. Similarly, NGO staff with *access to other at-risk users* reported the need to balance security with the autonomy of the survivors they worked with: Chen et al. [45] quoted a staff member working with trafficking survivors who tended to suggest that they turn off phone location, “But I do kind of leave it up to them. It’s not mandatory. . . . ’Cause we come from an empowering place. We don’t want to be telling them what to do.”

Caring for others can also lead at-risk users to prioritize efficiency in achieving their primary goals over digital safety, something we noted particularly in the cases of *legal or political, prominence*, or *access to a sensitive resource*, which frequently coincided with a profession or activity that benefited others. A human rights activist, for example, wondered, “Should I spend half a day figuring out digital security, or do work?” [123]. In several studies, journalists, political campaign workers, and emergency department personnel reported mixing personal and work data and devices for efficiency in their work [50, 137, 185, 205].

3.5.2 Lack of knowledge or experience

In our dataset, nearly all at-risk populations had less digital-safety knowledge than they needed to mitigate the risks they faced. While prior work has found that general users tend to have limited digital-safety knowledge [233], this was exacerbated for at-risk users who, by definition,

faced serious digital-safety risks and thus usually needed to understand and deploy more robust protections than a typical user. The at-risk framework helps us understand current patterns in how at-risk users experience digital-safety knowledge limitations.

Legal or political and prominence tended to attract sophisticated attackers, leading to stark knowledge asymmetries. It was difficult for at-risk populations—like activists, undocumented immigrants, journalists, or political campaign workers—to surmount the knowledge differential needed to counter the nation-state attackers that might or did target them [50, 55, 136, 137, 208]. For example, undocumented immigrants in the U.S. were often unaware that Immigration and Customs Enforcement might target them for surveillance via government requests for information to the social media platforms they used [82].

Relationship with the attacker involved intimate threats that required expertise to counter. Even in the absence of sophisticated attacks, these threats can require substantial, robust protective measures. For example, survivors of intimate partner abuse contended with a motivated attacker who may launch repeated attacks leveraging intimate knowledge about them, physical access to their devices and accounts, and relational power [50, 69, 116].

Some resource or time constrained users had limited technology experience. Examples included some people with low SES, people living in developing regions, and survivors of intimate partner abuse or trafficking, who did not have regular access to new or trusted devices or the internet, which limited their ability to gain technology experience and skills [45, 93, 129, 202]. Other populations, like journalists and people involved with political campaigns, commonly did not have the time to develop the technical skills to counter the digital-safety risks they faced [50, 136, 137]. Several studies reported that at-risk users who are *resource or time con-*

strained (emergency department workers [205], older adults [71, 90], refugees [200], and others) did not understand digital-safety settings that were, in theory, available to them.

3.5.3 Broken technology assumptions

The atypical threat models at-risk users face sometimes break assumptions that are built into secure system designs. Because of these assumptions, digital-safety best practices and technologies may be inaccessible, non-functional, or only minimally useful to at-risk users, often with magnified consequences.

One person per device or account. A common assumption of technology creators is that for every given device and account, there will be exactly one user who has access, despite prior work showing that convenience-based device sharing between trusted parties is common [128].⁶ This core assumption fails for several at-risk populations. At-risk users who *rely on a third party* may share account information with these trusted parties in order to accomplish important tasks [105, 200] or receive needed monitoring [146]. Different cultural privacy models have resulted in certain users sharing devices and accounts due to *social norms* [5, 184] or even the *legal or political* requirements of where they live [9]. Additionally, at-risk users who have a *relationship with the attacker* were frequently forced to give these attackers access to their devices under duress, directly breaking this assumption [42, 45, 68, 70, 87, 129].

Everyone has sufficient technology access. Some security techniques assume minimum levels of technology access that are out of reach for some at-risk users who are *resource constrained*. For example, 2FA that depends on a mobile phone (e.g., via SMS or a code-generating app)

⁶We note that since many of the papers in our dataset were published, there has been progress in addressing this issue, particularly in supporting child and family accounts. Nonetheless, we believe more can still be done to disrupt this assumption.

may not be available to at-risk users experiencing homelessness who are unable to consistently pay a phone bill or have an old device without space for new apps [202]. Some low-income Black Americans reported needing to stay on family mobile plans—which could have enabled surveillance from untrusted relations—because they could not afford separate service [64]. U.S. sanctions on Sudan entirely prevented Sudanese activists from enabling 2FA on a particular social media platform, since that platform’s 2FA system was prohibited from recognizing Sudanese phone numbers [55].

Physical and cognitive capacities are universal. Even for general users, advice about authentication tends to assume impossible cognitive capabilities, such as memorizing a unique, strong password for every account [157, 219]. This mistaken assumption is compounded in cases of *underserved accessibility needs*. For example, several papers have identified password memorization as a critical challenge for older adults with cognitive impairments [71, 146]. Password memorization can be difficult for people with disabilities that relate to alphanumeric comprehension, like dyslexia or aphasia [125].

Concepts, values, and experiences are universal. Other digital-safety paradigms rest on ideas, definitions, values, and morals that may not apply universally. Shortcomings in these assumptions can create digital-safety risks, especially for at-risk users with different *social norms*. For example, translations from English to Khmer of concepts like *privacy* within social media settings were hard to understand in the strongly community-oriented culture in Cambodia [93]. Similarly, some refugees who immigrated to the U.S. came from cultures where birthdays were not recorded; when these users were assigned a default value of January 1 in the U.S., authentication systems that relied on knowledge or entropic distribution of birthdays were less effective [200].

Separately, Barwulor et al. [23] found that moral codes and laws enacted by the U.S. government and enforced by U.S. companies made it difficult for sex workers in other countries, where their work is legal, to access safe payment and advertisement platforms, placing their physical and digital safety at risk.

Limiting digital-safety options is good for everyone. Digital-safety options that are too numerous can be hard to use [203]. But limiting the nuanced control enabled by digital-safety options to meet the usability needs of typical users may not always work well for at-risk users. Risk factors that increased the chances of focused targeting—which included *prominence*, *relationship with the attacker*, *reliance on a third party*, and *access to other at-risk users*—can lead at-risk users to have highly contextual digital-safety needs. For example, users protecting against a focused, intimate attacker, must account for nuances in the current relationship, the attacker’s mood, whether or not they are physically copresent, and more [116, 129]. Because of this nuance, at-risk users may benefit from additional options or enhanced transparency that they can deploy in specific alignment with their goals. When these options are not easily available or understood, it may lead at-risk users to fall back on *distancing behaviors* in which they try to stay safe at the cost of fully engaging with technology [55, 69, 71, 82, 114, 136].

3.6 Discussion

The at-risk framework can be used in multiple ways by researchers and technology creators, including guiding research and developing technologies to be inclusive of at-risk users.

3.6.1 Research

The framework, as applied to our dataset in [Table 3.1](#), can be used to help identify where knowledge of at-risk users is underdeveloped, sparse, or missing, giving researchers a way to prioritize their efforts. The framework can also be used to guide development of study designs and research questions—two examples of this in practice can be found in [Chapters 4 and 5](#).

Identify the *who* and the *what*. The digital-safety community would benefit from research about all at-risk populations and factors—this includes expanding knowledge about those that have already been studied and creating new knowledge about those that have not. In this complex and nuanced space, even after several papers have reported on a population, researchers may still have much to learn about the populations’ digital-safety needs.

Nonetheless, our meta-analysis makes clear that certain at-risk populations and risk factors have received particularly limited attention, at least recently and within the digital-safety community. We observed a general tendency to study participants from Western cultures, especially the U.S. Studying people from other regions and cultures could shed light on new risks related to *social norms*, un- or understudied interactions among risk factors, or even new risk factors. We advocate for more research involving geographically and culturally diverse at-risk participants, conducted by researchers who represent the world.

We also saw multiple studies exploring populations that experienced *marginalization*, focusing on a fairly narrow set of risk experiences (i.e., some populations’ only black circle (●) in [Table 3.1](#) was for the *marginalization* factor); future work could expand our understanding of digital-safety risks for these populations. Other groups often referenced as at-risk in popular media (e.g., real estate agents who have *access to sensitive resources* [\[218\]](#); celebrities who face

digital-safety risks due to *prominence* and parasocial relationships [240]) have not been studied in the research venues we analyzed. New or additional foundational research may be needed to inform the digital-safety community about these populations' experiences and needs.

Our framework also highlights how risk factors often combine and interact, but the literature in our dataset has not deeply explored this topic, precluding a thorough synthesis. We advocate *interactions* as a critical area of future research—one which our framework can support (e.g., enumerating the risk factors to consider for interactions).

Finally, in crisis situations, such as natural disasters or war, people may lose access to essential resources, their circumstances may change, or they may change their behaviors in exchange for critical services in ways that amplify their risk of attack or the severity of resulting harms. Recent papers on the impacts of the COVID-19 pandemic provide evidence that such changes can create new risks and amplify existing risks (e.g., [216, 244]). Future work on digital safety during various kinds of crisis situations could yield a new risk factor, or expand the definitions of existing factors.

Guide the *how*. The framework can also be used to help shape study designs and reporting, particularly interview or survey questions. For a population about which little is known, asking about all 10 contextual risk factors can ensure fairly comprehensive coverage of digital-safety concerns. For populations where some research exists, researchers can use the framework to explore how previously un- or understudied risk factors may (or may not) apply, or add depth on the impact of a specific, previously identified risk factor of interest. Researchers can also use the protective practices portion of the framework to explore more comprehensively how their participants currently protect themselves and why they choose their practices. We hope that using

the framework to guide research and reporting can enable better comparisons among studies, helping to uncover when disparate populations have overlapping (or distinctive) digital-safety practices and needs. Beyond this, it is important for researchers to carefully plan ethical methods when working with at-risk users, guidance for which is an emerging area of research [53, 117].

3.6.2 Technology development

Technology creators can use the at-risk framework to better support a wide range of at-risk users in their products.

Consider at-risk users at scale. Our framework does not replace direct engagement with at-risk users, and we advocate for such engagement when appropriate. However, doing so selectively and ethically is important, and there are still open questions about ethical methods for at-risk user research (e.g., how do researchers not overtax already stressed and resource-constrained groups?). Meanwhile, many papers in our dataset recommended considering the impact of a technology design on the at-risk population they studied, which is incredibly important but difficult to scale across populations without a guiding framework. Our framework simplifies the challenging but important process of thinking through potential risks and needs of multiple at-risk populations together. It does this by providing 10 *contextual risk factors* that organize patterns of risks and needs, and a set of *protective practices* at-risk users currently deploy to (sometimes ineffectively) cope. Using our framework can help researchers and technology creators prepare for user research, at interim points during multi-phased technology creation projects, or when user research across multiple at-risk populations is not an ethical option.

Each risk factor suggests specific, sometimes overlapping, technology needs. For *social*

norms, relationship with the attacker, and reliance on a third party, users often could not keep their devices and accounts private. These users might benefit from the ability to keep select technology use and data secret on shared devices and accounts, or perhaps to enable digital traces that are ambiguous or imprecise to support plausible deniability [18]. Users with *access to a sensitive resource* would benefit from robust protections for the sensitive resource (e.g., encryption, strong authentication). Users facing focused and/or sophisticated attackers (e.g., *legal or political, prominence, relationship with the attacker, access to other at-risk users*) could benefit from easier-to-use versions of strong protections (such as hardware-based 2FA, strong passwords, and encryption) coupled with guided set-up flows and education. Users with *prominence* or who experience *marginalization* would benefit from support for managing bulk or pervasive attacks from potentially anyone.

Our systematization of protective practices (Section 3.4) and barriers (Section 3.5) together highlight how at-risk users cope with their risks, sometimes in ways that are not completely effective or that introduce vulnerabilities. Notably, at-risk users commonly employed non-technical practices—such as a host of social strategies and distancing behaviors—and it is important for technology creators to understand and not disrupt the important role these practices play. The protective practices we identified also show that at-risk users are not commonly using some existing digital-safety solutions (at least as reported in the dataset), suggesting areas where technology creators could improve accessibility and/or usability. Further, we discuss broken assumptions that contribute to ineffective protections (Section 3.5.3), and encourage technology creators to consider these in new technologies.

Balance tensions. Those creating technology for at-risk users will have to contend with inher-

ent tensions: “perfect” digital-safety protections usually do not exist. At-risk users already use a variety of practices, technical and otherwise, to address their pressing digital-safety concerns (Section 3.4). These practices all have some protective value, but may also come with significant downsides, like reduced social participation, lack of agency, and loss of transparency. For example, *distancing* from technology (Section 3.4.2) was a common protective practice across our dataset, but it also reduced access to *social support*, which was another essential protective practice broadly employed by at-risk users (Section 3.4.1). Technology creators should understand that any technical intervention is likely to introduce benefits and drawbacks, which should be carefully studied to help ensure such tensions are understood, manageable, and net beneficial. Careful evaluation of potential designs using our framework can help technology teams reason about how to balance various risks and benefits.

Balance usability and options. In Section 3.5.3, we discussed at-risk users’ need for more nuanced digital safety options than typical users. At the same time, adding options and transparency must be balanced with a very real need for usability. At-risk users experience heightened stress associated with higher potential for digital harm, as well as stress from their particular risk factors (e.g., *resource constraints*, *marginalization*, etc.).

To balance the need for options that addresses the unique needs of at-risk users with the high bar for usability, we encourage practitioners to make transparency and controls actionable and manageable. For example, transparency features that make users aware of a threat can be more actionable if they provide clear next steps to fix the issue and then guide the user through relevant protective measures. Similarly, layered or directed designs can help users find the options and controls that best meet their needs. Equally important are easy-to-understand defaults that

minimize barriers to deploying protections and are carefully selected to support at-risk users with many competing priorities (Section [3.5.1](#)).

Chapter 4: How Library IT Staff Navigate Privacy and Security Challenges and Responsibilities

4.1 Introduction

Libraries serve as a central hub for technology access for a large number of people in the U.S. and in other Western countries [91, 169]. As people require technology access for things like searching for jobs, education, and accessing essential services, libraries provide this access to many users, particularly those who cannot afford personal computing devices. This well-studied phenomenon is sometimes referred to as the “digital divide” [220]. Further—in contrast to earlier notions of the digital divide—access to infrastructure is not the only or even most important barrier: digital equity also encompasses ease and quality of use [86].

As a free and public service, which for many users is the only or primary means of technology access, the security and privacy of library IT resources can be critically important, in multiple dimensions. First, when patrons¹ are reliant on libraries for all or most of their computing needs [223], their usage includes potentially very sensitive activities. For example, LGBTQ+ people [73], teenagers [238], people who have experienced intimate partner violence seeking out information resources that they must keep private from others in their households [129], or marginalized and older individuals filing taxes [133] may be more susceptible to threats. Second,

¹A “library patron” (hereafter, “patron”) is someone who uses a library service.

patrons may not be as familiar with assumed-to-be “standard” self-protective behaviors that many home computer users have learned [176]. For example, Madden [121] argues that low socioeconomic adults struggle to find resources on how to protect their online data, and are uniquely vulnerable to privacy risk, perhaps due to increased dependence on mobile internet technologies in the absence of home broadband connectivity. This increases the importance of security and privacy protocols adopted by librarians. Third, and relatedly, the types of risks that vulnerable and marginalized patrons perceive might be different from (or put them at odds with) what is commonly expected when we think of privacy literacy [126, 132, 229].

Library IT staff members can make security and privacy decisions that have a direct impact on these marginalized individuals, who may turn to the library for privacy information [120], or for essential computing resources [92, 170] that put them potentially at a higher risk for privacy violations [121].

This paper reports on interviews with 12 library staff members from the U.S who are responsible for technology management at their libraries. The participants span a variety of roles—from one of many on a large IT team, to library director who doubles as system administrator—at a variety of institutions—from small rural libraries to large library systems in major cities.

Research Question and Contributions. We approached this research with the broad research question:

What kind of privacy tools and strategies (if any) do library IT staff provide for their patrons, and why?

We frame our findings using Sen’s capabilities approach [196], which allows us to go beyond specious notions of access and usage to illuminate the tradeoffs that library IT staff consider

when they make decisions about privacy in the service of patron freedoms. We will show that library IT staff provide critical resources for supporting Sen’s notion of expansion of freedom, particularly in the sense of making information available [196].

We find that library IT staff are more concerned with protecting patrons from external threats (e.g., governments and institutions that sell their data to third-parties, which they believe infringe on patron freedoms [196]) than with protecting library systems from potentially malicious patrons. For this reason, librarians are most concerned with ensuring that patron data is not stored, even if it requires that they exert less oversight. At times, library IT staff must make difficult tradeoffs in enabling patron capabilities—for example, between supporting patrons in achieving their IT goals (availability and maximum utility of services) and protecting privacy (confidentiality). As might be expected, library IT staff struggle with lack of resources and training, as well as conflicts with decision makers and other organizations, resorting in some cases to creative or unorthodox solutions in order to uphold these patron freedoms.

4.2 Methods

Here, we describe our interview study and data analysis.

4.2.1 Recruitment

To recruit participants, we reached out to candidates directly via email using information available on public websites for libraries. We specifically emailed candidates who identified as information technology specialists for the libraries they worked at. We also posted advertisements

ID	Institution	Years in:		# IT staff	Est. patrons/month
		Library IT	Total IT		
1	Advocate organization	2	8	2	-
2	City library system	2	2	3	48K
3	University library	2	2	-	-
4	City library branch	1	1	17	300K
5	County Library system	10	22	14	100K
6	Association library	9	9	1	-
7	Local library system	7	9	1	-
8	Advocate organization	12	12	18	408K
9	Advocate organization	20	20	-	-
10	City library system	1	2	1	22K
11	Local public library	10	10	1	3K
12	City library branch	2	2	-	50K
13	City library system	8	14	2	20K
14	County library system	4	7	2	50K

Table 4.1: Participant information including the type of institution they were affiliated with, work experience in IT (in a library setting and overall) the number of IT staff at their institution, and the estimated number of patrons their institution serves per month. Items are left blank when participants did not know or preferred not to answer.

on public fora, including Reddit and Twitter.² We also reached out via email to library staff in our personal networks to ask for referrals to IT staff. As the study progressed, we also relied on snow-ball sampling to recruit additional participants. We asked participants if they had connections to other library IT staff that might be willing to interview with us. Participants were considered eligible if they were 18 or older, were employed at a public or university library, and their job included IT work (formally or informally). We also accepted participants currently working in a library advocacy organization who had prior IT experience in a library.

Participants who responded were asked to complete an intake survey (to provide background information to the interviewer) and then scheduled for an interview.

The study was approved by the University of Maryland IRB. Participants were asked to

²For clarity, I use Twitter instead of X to describe the social media platform in this thesis. All data was collected before the rebranding.

complete separate online consent forms prior to the intake survey and the interview. The consent form described the purpose of the study, the researchers conducting the study, the study procedure, and how interview recordings, participant names, and contact information would be handled. Participants were also asked prior to the interview if researchers could record the audio of the interview for analysis, and were informed that they could stop the interview at any point or decline to answer any questions. Participants were also explicitly informed that recordings might be sent to a third-party transcription service.

4.2.2 Intake survey

Participants who expressed interest in interviews were asked to complete an intake survey including questions which characterized the size and type of library they worked at, their experience with IT in libraries, summary information about the computing resources provided by their library, and vulnerable populations their library served.

Information about technology use and staffing was used to guide interview questions—for example, if the participant described working with contractors to provide IT services in the intake survey, we asked about those contractors in the interview. The intake survey also provided context for the responses for each participant.

4.2.3 Interview protocol

One researcher conducted the 12 semi-structured interviews via video conference. These interviews were audio-recorded and automatically transcribed, supplemented by notes taken by the researcher. The consent form explicitly noted the possibility of third-party transcription when

participants were asked about audio recording. Interviews lasted, on average, about one hour, and participants were compensated with a \$35 USD Amazon gift card. The intake survey and interview protocol are found in Appendix B. We asked participants about a wide range of topics, including:

Role and responsibilities. We asked participants about their day-to-day responsibilities at their organization. We also asked how participants ended up in their roles and what they thought about those roles.

Staffing structure. Participants were asked to describe the structure of IT staffing at their organization, as well as who they reported to, and how many employees responsible for IT operations their organization employed.

Patron populations. We asked participants to talk about the patron population that their organization served, how many patrons they served, and the ways that patrons typically used the library services. We also explicitly asked about populations that are historically underserved or face higher-than-average digital-safety risks, and how their organizations were serving those patrons.

IT services and infrastructure. Participants were asked to describe the kinds of IT services their organization offered, such as public PCs, internet access, and printing services. We also asked about how those services were offered and managed, specifically whether the organization self-managed these services or if they contracted third-party vendors to do it for them.

Security and privacy policies and tools. We asked participants about the specific security and privacy policies and tools they regularly employed to protect their infrastructure and their patrons. We also asked whether they felt that those policies and tools were sufficient to meet the security and privacy needs of their organizations.

Security and privacy challenges. Participants were asked about specific security and privacy challenges that they regularly face at their organization, and whether or not they felt they were adequately addressing them at the moment. In some cases, we also asked about what participants felt would need to be done to resolve those challenges.

COVID-19 pandemic. We also asked about how the COVID-19 pandemic affected the typical operations of their organization and whether or not that impacted the security and privacy considerations of IT staff.

4.2.4 Participants

We recruited 12 participants in Fall and Winter of 2020. Participants worked at libraries ranging from rural to urban, with a wide range of experience and educational backgrounds. The particular positions these participants held varied—one participant was the director of a small, rural library in addition to being responsible for the IT concerns of this library, whereas multiple participants worked as part of an IT-specific team at large, metropolitan library systems in major U.S. cities as library IT staff. Three participants currently work for advocacy organizations rather than directly for a library, but all three had library IT experience—this experience, in fact, was what drove them to advocacy in the first place. Some participants had backgrounds in computer science or information technology, but many had library science backgrounds. Participant institution types, IT experience, staff size, and estimated patrons served per month are presented in Table 4.1.

4.2.5 Thematic analysis

Two researchers from the team conducted a reflexive thematic analysis (RTA) [31,32]. After reading through all the transcripts, they generated initial codes that were applied to a sample of interviews and discussed. These researchers met with the team to further discuss these codes and potential research themes. They then continued to iterate on and apply these codes to the remaining interviews and developed a set of refined themes. Codes were first categorized, then discussed and further refined with the research team over the course of several meetings, as themes were developed. The two researchers continued to code the remainder of the transcripts. After developing the themes, we assigned illustrative quotes. We ultimately generated four high-level categories (goals, responsibilities, tensions, constraints). Within these categories, we identified our themes (formatted in bold below). We applied Sen's capabilities approach framing to interpret our findings.

To develop our themes, we engaged with Braun and Clarke's most recent perspectives on RTA [32], developing themes flexibly and organically, while comparing and testing them for coherence and consistency. The resulting themes organize salient, meaningful segments of the data.

4.2.6 The capabilities approach framing

We used Amartya Sen's capabilities approach to interpret our findings. His framework centers freedom to achieve wellbeing as a moral imperative and posits that this freedom must be understood in terms of people's capabilities. We found that it usefully described the way that library IT staff are thinking about patron access in terms of what they are able to do.

Library IT staff experienced privacy tensions when providing internet resources. However, these tensions arose not simply in providing direct access to these resources, but in more nuanced navigation of capabilities (patrons using internet resources to e.g., do work or educate themselves), which were related to the specific context of their patrons and the structural barriers they faced. This mismatch between literal measures of “digital divide” success (i.e., internet access) and the nuanced needs of library patrons finds parallels in the capabilities approach. Both Sen [195] and Martha Nussbaum [155] (who has extended the framework) offer a critique of economic discourse focused on standard of living, arguing that the focus should be on well-being of a given community.

Sen first proposed the capabilities approach [195] as a way of measuring what people are capable of doing with the means provided them. It assumes a moral paradigm in arguing for the capabilities (also called “freedoms”) to achieve well-being understood in terms of people’s capabilities, which when enhanced increase their capabilities.

In Sen’s capabilities approach, “functionings” encompass what an individual may value “doing or being,” like being healthy or having a home and safe place to live. Capabilities describe the set of functionings a person has available to them, from which they can choose. In other words, capabilities represent the freedom of an individual to choose between sets of different functionings, according to their own values, in service of a “good life” [196].

Sen enumerated specific exemplar capabilities, including economic opportunities, political freedoms, social facilities, transparency guarantees (in terms of information and relationships between people), and protective security. At the same time, Sen emphasizes that the choice of which capabilities to value must be left to individuals and communities. Individual capabilities (e.g., freedom from surveillance) may be more highly valued in particular communities because

of their relevance to that community’s experiences and their relationship to other capabilities.

More recently, scholarship has applied Sen’s framing to ICT for development, focusing on access to information [158, 242]. Wresch considers that while the number of websites in developing countries falls far behind Western countries, perhaps the availability of “helpful” local web resources is what is most important [242]. Kleine applies Sen’s concept of capabilities development to ICTs in Chile, arguing that by giving individuals greater access to information, they contribute to one concept of capabilities [98].

While Nussbaum [155] further articulated specific capabilities (e.g., “life,” “bodily health,” “bodily integrity,” etc.), we follow Sen and allow participants to define relevant capabilities rather than relying on Nussbaum’s categories. While Nussbaum’s capabilities, which have recently seen considerable use in privacy and human-computer interaction fields, are useful in elevating the importance of human dignity across cultures, Nussbaum has been criticized for defining capabilities in a Western-centric way. We agree, in this context, with Sen that capabilities are more usefully defined by those who have insight into the day-to-day experiences of their patrons (and the nuances of what constitutes good access), including library IT staff. We believe this approach is more intersectional and appropriate, given that every community has its own matrix of power that makes certain identities vulnerable in certain specific ways [49]. We use this framing in our results and in our discussion.

4.2.7 Limitations

As is typical for qualitative studies, our goal is not to be fully generalizable, but rather to identify key themes and ideas related to our research questions.

Our sample size of 12 is relatively small, but is well within norms for human-computer interaction studies [37] and qualitative best practices [83]. To broaden our results, we recruited from many types of libraries in a wide variety of U.S. locations; however we acknowledge that other IT staff working in public libraries, especially in other countries, may have different experiences. Nonetheless, we believe that the themes we discuss here are highly relevant to the greater library IT space, even if they are not generalizable to all libraries.

Additionally, the researcher who conducted the interviews was different from the researchers who performed the thematic analysis on the interview transcripts. The researchers who performed the thematic analysis discussed the interviews with the interviewer and were careful to immerse themselves within the data sufficiently to understand it as a whole.

We also acknowledge that the researchers who performed the thematic analysis have a primarily security and privacy background, whereas most participants did not, and that this difference in experiences may have affected the interpretations the analysts generated. To address this, the researchers took special care to faithfully organize and describe the phenomena our participants described, and to conduct the analysis process with these differences in mind.

4.3 Library IT staff roles and responsibilities

Before describing the findings of our thematic analysis, we first report descriptively on the roles and responsibilities that library IT staff reported to us. Library IT staff provide a wide variety of digital services to their patrons, and the IT staff participants we spoke with are responsible for making sure those services are functional. Many participants described a responsibility to minimize the security and privacy risk these devices pose to patrons.

Public PC management. A near-universal responsibility of our participants was to manage public-facing PCs. Participants reported that patrons use these PCs for recreation, but also for serious and essential tasks, in line with prior findings(Chapter 2.2).

P11: "... they're just going to sit somewhere upstairs all day and watch YouTube and Facebook."

P12: "Our older patrons tend to use it to do research or read emails, print information. Print forms. A lot of people having been using them for tax information over the last couple of months."

In addition to traditional stationary PCs, some libraries also loan out laptops for individual use, particularly when impacted by COVID-19 closures. Library IT staff are typically also responsible for managing these laptops, which can come and go from the library.

Several participants describe setups where public computing infrastructures are, in fact, run by local or city government or even a vendor, leaving library IT staff with little direct control of how these systems are managed.

Print, copy, and fax management. Printing, copying, and faxing are crucial library services. The basic functioning and administration of these services is another key component of what our participants say they do from day to day. Notably, despite reliance on outside vendors, participants expressed less concern about patrons' privacy in this realm.

Libraries handle printing processes in a variety of ways. For example, some libraries use vendor-provided tools to allow automated printing and payment:

P10: "[Envision, a third-party vendor, has] a print vending system that ties into—we've got like a coin and bill acceptance machine for paying...so somebody could send a

document from home or from their phone whatever and come to the library and print it out.”

Others use more ad-hoc methods, like a public email address to which patrons can send attachments, which then requires a librarian to manually print (and therefore have access to) their document:

P12: “Right now, there’s only one way to print and what you can do is email your job as an attachment. . . to the email address we provide. . . and the subject line has to say print so that it will go into our print folder and get filtered into that.”

Managing access to digital resources. Libraries often contract with third-party vendors to provide digital resources, like ebooks or access to academic journals. This naturally saves time and effort, compared to developing these materials in-house. However, some participants described privacy concerns relating to the relationship between these vendors and the library’s patrons, which we will discuss in detail in Section [4.4.3](#) below.

Providing digital skills education. Participants say they occasionally provide digital skills education to their patrons, either through open office hours or scheduled events. These cover topics from using common software suites and Microsoft Office to basic online safety. P07 spoke about a more specifically privacy-focused class that ranges from basics to more advanced concepts, which was (perhaps surprisingly) the only class of this sort described by our participants:

“We have a much more in-depth class that we call “online self-defense” which we start out with the basics and then go all the way up to like, using a secure operating system and stuff like that.”

Assisting with day-to-day troubleshooting. Some participants say they are also accessible to the public on an as-needed basis. This was sometimes described not as a part of their official job description, but rather a responsibility that arose organically from the needs of their patrons. Circulation staff—who are publicly visible to patrons by the nature of their position—will often receive technical support questions, and when they are unable to answer those questions themselves, will refer the patron to library IT staff. Different participants’ libraries dealt with these support requests in different ways, ranging from deliberately arranging computers to be visible from the circulation desk (so that circulation staffers could keep an eye on both), to ad-hoc solutions, like simply flagging down the nearest IT staff member, regardless of what they were previously doing. In most cases, the official role of library IT staff is to maintain library infrastructure rather than provide public-facing tech support, but in practice, they are often asked to provide this support anyway.

P02 experienced both of these strategies, describing technical support questions interrupting their more formal IT work:

“At this point, the only time I have any interaction with patrons is when I am doing stuff in the lower-level tech center. . . last week I was imaging some computers down there that weren’t being used by patrons. And so, you know, maybe one of the patrons will try to have a conversation with me [for tech support], or one of the librarians that are stationed down there will take this as an opportunity to have me help them help a patron.”

However, this participant also indicated that before the pandemic lockdown, they used to run three to four hours a week of dedicated technical support desk time.

4.4 Findings

In this section, we report on library IT staff members' goals and values, and the threats they described facing. After outlining those individual factors, we describe tensions our participants highlighted between those factors, as well as the constraints that limit their ability to achieve their goals. We use Sen's capabilities approach to frame our findings about how library IT staff think about their patrons' use of library technology and the tradeoffs they make with regard to security and privacy.

4.4.1 Goals of library IT staff

Participants describe several high-level goals related to library IT, several of which concern the security and privacy of their patrons. Library IT primarily focus on ensuring the capabilities of their patrons to access information, social services, employment, etc. Doing so requires that they provide a trusted space for accessing economic resources and other information in ways that coincide with Sen's notions of capabilities to social opportunities (e.g., education, healthcare) and protective security (including protection from poverty and incarceration or deportation) [196]. Privacy is seen as primarily a way of shielding patrons from surveillance (or the fear of it) by government institutions [196].

Making libraries a safe and trusted space. We see this dedication to patrons' capabilities in our participants' goal of making the library a safe and trusted space u This accords with other scholars' views of ICTs as information resources ("making information available and facilitating interaction") that are fundamental to Sen's capability to choose [98]. In this case, providing a safe

Population	Count
LGBTQIA+ people	10
People with disabilities	10
Individuals without permanent residence	10
Low-income communities	9
Racial/ethnic minorities	9
Older adults	9
Individuals with a criminal history	8
Low-income single mothers	8
Undocumented immigrants	8
Political activists	5
Journalists	3

Table 4.2: Patron populations reported by our participants in the intake survey.

and trusted space supports this capability.

Our participants frequently explicitly identified libraries as the last remaining trusted public organization/resource in their community.

P04 says, “We are the most trusted institution of any [local government agency]. . . routinely, in poll after poll we are shown to be a beloved institution.” This belief is supported by studies from the Pew Research Center indicating that patrons in the United States view libraries as helpful for finding trustworthy information and making informed decisions [92].

A few participants emphasized the need to maintain the trust of patrons, especially those from populations who might experience higher digital-safety risk. Participants mentioned examples like LGBTQ+ youth, populations who are disproportionately targeted by the police, and people with low socioeconomic status, as patrons who, participants felt, particularly depend on the library for critical information and resources. Almost every participant described serving various such populations in their libraries, as detailed in Table 4.2. According to participants, protecting these at-risk users’ privacy is key to maintaining their safety as well as their trust.

P05: “There are a lot of LGBT youth who grew up in very conservative household[s]

who could have never searched for that information about who they are, without going to the library. [The library] provided them an avenue and a space to search for information without their parents, knowing or without anybody else, knowing what they were looking at.”

P08 described a different concern with retaining the trust of immigrants and undocumented people, where the library itself was seen as a potential adversary that threatened the capabilities of their patrons to live in the U.S. without fear of surveillance or deportation. Framed in terms of capabilities, this participant believes this mistrust (i.e., the chilling effects imposed by fear of being surveilled and deported) interferes with the capability of their patrons to achieve wellbeing through access to resources they use for information seeking and employment:

“[Patrons are] essentially saying ‘We’re not going to sign up for a library card because this information [their library activities] is going to be sent over to the city government and once it gets sent over to the city government, then law enforcement officials might have access to it.’ ”

This participant clarified that their library system did *not* share browsing data or other patron information with law enforcement (especially U.S. Immigration and Customs Enforcement (ICE)).

“We do not release information to government officials unless they have a signed judicial warrant, not an administrative warrant the ones that ICE likes to serve people. It has to be a court-issued warrant or subpoena.”

However, P08 also indicated that communication around this issue was a challenge, suggesting that it was not unreasonable for patrons to worry that the library, as part of city govern-

ment, might have an obligation to share data with law enforcement. Despite their efforts to guard against infringements of their patrons' capabilities, the bureaucratic configurations of the system had chilling effects.

“Patrons rightly are not very sure what we exactly do. . . We are part of the city government, but the way that [the library system] works is that we are separate in terms of hierarchy, we answer to the library board and not straight to the mayor. Our IT systems are separate, meaning that the library IT department has control over the network, has control over the public computers, has control over the Wi Fi network. City IT does not get any of that. So they can't come in and. . . sniff traffic.”

Maintaining library values within larger IT systems. Several of our participants—particularly those who worked in smaller towns and cities (and thus in smaller library systems with fewer resources)—described being dependent on local government IT departments to maintain or manage library IT. A few described conflicts in this relationship. P07 describes how the city's IT department, not understanding the library's mission of staunch protection of “patron privacy and intellectual freedom,” would use monitoring software to monitor patrons:

P07: “Initially when I was hired, the city IT actually took care of library IT needs, but. . . there was a lot of conflict there. They did not have a good understanding of the library's mission and they would often do things that were contrary to our library's mission and values. . . They had monitoring software on the public computers. . . and they would sometimes observe what the patrons were doing because they thought that it was necessary to make sure that the patron didn't violate our policies. In a corporate IT setting that's perhaps reasonable, and corporate IT may need to monitor

employees to make sure they aren't violating corporate policies, but that is directly in contradiction with library values—patron privacy and intellectual freedom.”

This participant later noted that “It’s easier to teach a librarian IT skills than it is to teach an IT person library values,” referring to the need to preserve “patron privacy and intellectual freedom.” Library IT staffers’ dedication to privacy and intellectual freedom reflects their concern that patrons’ capabilities—freedom, in Sen’s framework, to make choices to advance their own wellbeing—are infringed upon when those patrons are surveilled (especially in ways that home computer users would not be subject to) reducing transparency and creating distrust.

Availability and usefulness of computing services. Another goal of library IT staff members was to support and improve patrons’ capabilities by ensuring that use of computing resources was easy, and that patrons could accomplish their goals using library resources.

One component of this goal is making access to public PCs and other infrastructure simple and painless. Most participants described systems wherein participants needed only a library card, or even a guest pass for those without library cards, to use library public computers. P10 described fully open access to their library’s computers—patrons only had to sit down at the computer to use it, without any login procedures. This fully open access was an exception, rather than standard practice across our participants.

This goal also manifests in a desire to help patrons accomplish their goals. This was mentioned explicitly in the context of formal and informal tech support, and described implicitly when discussing services like public PC use and printing, which exist to support patrons’ needs and capabilities.

Privacy of patron data and activities. Most participants described protecting the privacy of

patrons who use library IT resources as a core goal.

One near-universal approach was to ensure that patron data or activities are not retained on a public PC after the patron is finished using it, under the theory that if no data is kept, then it cannot be revealed or abused. Participants described wanting patrons to feel comfortable enough with privacy to use library equipment (supporting expansion of capabilities) and also expressed concern about the potential consequences of storing data that could be used against the patrons (limiting capabilities).

P06: “In principle, I like the ability for somebody to use the internet anonymously without a paper trail if they want. I think that that’s a good thing to offer.”

Mechanisms for preventing data retention vary across participants’ organizations, based on available tools, financial resources, and knowledge. In particular, participants cited vendor software, such as Deep Freeze—a software solution that allows computers to be automatically restored to a default image³—as well as OS-level features, like Windows 10’s unified write filter (UWF),⁴ as primary strategies for ensuring that public terminals are kept in their original state. P02 described using a FOG Project⁵ server to manage this need, since it is open-source and thus not a financial drain on the library. However, this method requires significant setup time and maintenance of physical infrastructure, as well as the knowledge to seek out and implement an open-source solution. Participants noted that these data retention protocols are broadly similar for both stationary computers and loanable laptops.

Whether or not patrons must log in (e.g., using a library card number), participants empha-

³<https://www.faronics.com/products/deep-freeze/enterprise>

⁴Unified Write Filter is a Windows 10 operating system feature that redirects writes to the drive to a virtual overlay. This virtual overlay is typically cleared during a reboot or when a guest user signs out.

⁵<https://fogproject.org/>

sized that they do require patrons to begin and end sessions, in order to protect data when the patron signs off. However, a number of participants lamented that patrons do not log out appropriately, and so in addition to messaging on the computer reminding them to sign out, they have staff provide additional monitoring to ensure that computers are wiped between uses:

P07: “When they’re done, there’s a button on the desktop that says ‘log out securely.’

We also have staff that walk around and if they see something up, then they will hit that ‘log out securely’ button, which basically wipes out the entire session—any documents, any downloads, any browser history goes away.”

Library staff thus ensure that patrons’ privacy is consistently preserved.

Security of library infrastructure. One might assume that protecting the library’s digital infrastructure—public computing but also other resources, like administrative computers for library staff or printing systems—would also be a key goal for library IT staff. In fact, few of our participants actively mentioned this goal as a high priority, as we will discuss further in Sections [4.4.2](#) and [4.4.3](#) below.

Indeed, P02 is an exception, noting that Deep Freeze serves two functions: protecting patron privacy, but also “lock[ing] down” the public-facing terminals to prevent tampering, suggesting that other patrons might pose a threat:

“We do have Deep Freeze on there to reset everything. There are very few permissions on . . . you can’t even change the wallpaper, so we try to lock it down as much as possible. We try to make sure everything is wiped off when they get off so that nobody has access to their information.”

4.4.2 Threats faced by library IT staff

We have just described the ways in which library IT staff support patrons' freedom to access information and resources. In this section we discuss several threats they faced when aiming to achieve their goals and responsibilities which we detail below, noting that the the biggest come from outside their walls. That is, library IT staff view the greatest threats to patron's capabilities as institutions like governments and third parties.

Attacks from patrons. When asked about security and privacy concerns, participants rarely mentioned their patrons posing any malicious or intentional security or privacy risk to the library or other patrons. Some pointed out that when patrons do create risks, it's most often unintentional.

P06 and P09, for example, noted that problems are rarely caused by malicious users, but that patrons can unintentionally disrupt library computers by accidentally downloading malicious or undesirable software:

P06: "I don't often deal with people here who I think would do anything maliciously, but it's like they screw up and download something that then slows the computer down or so[mething]."

P09: "I guess my biggest worry would be, you know downloads... They [patrons] click something bad, download something... that's probably what I would worry about the most, someone doing something bad unintentionally rather than someone coming in and plugging in a keylogger or whatever."

However, participants are still aware of some potential attacks—P12, for example, described watching for keyloggers as a possible threat. However, they also noted that they doubt

that it would be possible to run one on their computers:

P12: “You can’t download anything to the computer and save it there. So unless there’s a way to have a key logger running in the background of a computer that is not running... I don’t know how anyone would be able to do it.”

Ransomware attacks. One specific attack mentioned by a few participants was ransomware, which is a comparatively common attack on libraries [39].

As P09, who has done advocacy work with several libraries, said, “It’s probably every other week that libraries get hit by ransomware.” P09 also relates this to patron behavior, as ransomware is usually delivered in this context by clicking on a suspicious download link on the machine in question. However, only one of the five participants who mentioned this attack had actually experienced it, and their PC restoration system was able to solve the problem via a restart. Other participants who discussed this attack mentioned either hearing about it in the news or from colleagues from similar libraries.

Data collection by third-party vendors. Several participants expressed concerns about third-party vendors (e.g., ebook providers) having access to patron data. Participants typically described this as a general threat related to their responsibility to protect patrons and their capabilities, rather than citing any specific threats or past incidents with any particular vendor. Participants generally felt ill-informed and ill-equipped to fully understand or counteract threats from third parties.

For example, P01 describes their concern about patrons being required to give personal data to vendors, and how libraries are mostly powerless to enact change:

“With these vendors there’s various kinds of data sharing to like, big corporate enti-

ties like Amazon. It's a really wild environment and I think it is a great concern. We don't have any meaningful collective negotiating power around it to say no to these practices."

P12 expressed frustration that few vendors hold the same values that libraries do and that libraries are often forced to work with such vendors regardless of this mismatch because of the services they offer:

"We're [public libraries] trying to just provide the information you want and not pry, but a lot of the third-party apps don't seem to care about that. It's frustrating because they're more interested in how much money they can make and that's not really our objective. A lot of the time they'll put together records of what customers are interested in to make really good targeted advertising and are refusing to stop because it's not like there is a library-friendly corporation out there that wants to give us their ebooks. You have to work with a lot of these big monopoly companies and their software so they can do whatever they want."

P08 describes a similar sense of powerlessness: trying and often failing to better understand third parties' privacy practices and potential threats:

"I just read a lot about privacy policies. I read about privacy policies and vendors, their reviews. I also read about vendors' relationships with privacy and public computing, but I don't find that much information."

P08 lamented that not only are libraries unable to prevent existing third party vendor privacy incursions, but the influence (and associated surveillance) of third parties continues to increase:

“We as a profession really don’t have a handle of data privacy and security that we should have, especially with . . . this exponential increase and working with vendors to provide critical services. . . And those vendors consolidating and on top of that coming up with new ways to track and surveil our users.”

Data requests from law enforcement. A few participants defined the goal of patron privacy explicitly as defense against law enforcement data requests. Library IT staff perceive that protecting their patrons from law enforcement supports their capabilities to freely access library resources, whatever their specific goals.

P06, for example, explained their library’s data collection policy (with respect to patrons using public PCs) explicitly in terms of how their library’s goals (to protect the privacy of patrons) differed from those of law enforcement.

“Our highest concern is our patrons’ privacy, it’s not serving the police or solving crime, we have different goals. . . [When] they started saying that they could come in and get your list of patrons or take computers, is when public libraries generally stopped having lists. I wouldn’t argue that’s coincidental. I think librarians looked at the situation and were like, ‘Well, what can we do in this situation to protect the privacy of our users?’ The answer was to stop keeping lists because we can’t be forced to give them something we don’t have.”

This threat relates to the goal of making libraries a safe and trusted space (Section [4.4.1](#))—P02 described patrons worrying about law enforcement being able to gain records of their activities or information about them. Participants recognize that to support patrons capabilities, they must protect patrons from fear of surveillance, and that doing so requires transparency: they can-

not hand over records they do not have. This accords with Sen's notion of transparency as a means to build trust, and to proactively ensure that trust cannot be abused.

4.4.3 Tensions among goals and responsibilities

Our participants communicated several tensions among their goals and responsibilities, where library IT staff must make choices among competing priorities. We found that library IT staff almost always prioritize capabilities of their patrons over possible threats to the library itself (e.g., allowing anonymous logins to provide patrons with the freedom to search for information without surveillance, even at the risk of not being able to identify threats). The situation is more complicated when different types of patron capabilities appear to conflict with each other (e.g., the capability to access and use services fluently compared to the freedom to do so without surveillance). We find that library IT staff are making nuanced decisions about balancing these tradeoffs, grounded in their patrons' needs and experiences.

Patron privacy vs. security of library infrastructure. In some cases, ensuring patron privacy can be at odds with taking maximum precautions to protect the security of library infrastructure. For example, libraries could choose to install monitoring software in order to detect misbehavior, or enforce strong login requirements in order to attribute any problems to particular users, but monitoring and logging inherently poses a threat to privacy by tracking user activities.

Our participants almost always prioritized protecting patrons' privacy rather than protecting library infrastructure. P02 describes their library regularly checking hardware for tampering, as well as having security cameras to monitor the physical machines, but notes that such security cameras are unable to view the computer displays so as to preserve patron privacy.

“We have security cameras but they’re not fixed onto the displays of the computers and they’re not that detailed anyway. More for like a macro shot. The librarians after a session will see if somebody left anything in the computer, like a USB flash drive left in there. They’ll just take it out and put it in Lost and Found.”

Similarly, several participants said that their libraries allow anonymous (or near-anonymous) login to public PCs, potentially at the cost of not being able to attribute misbehavior to specific users. The means for implementing this access vary. While many libraries do require patrons to enter their library card number and a PIN in order to begin using workstations, some also allow visitors without a library card to obtain a “guest pass” with temporary credentials for PC usage.

Other participants from university libraries reported that they generally require patrons to log in via their university’s central identity management service to access most PCs, but that they do reserve a few PCs with limited services for general access without requiring university credentials.

Other libraries do not require any login at all, and simply present terms of service that users must agree to in order to use the public PC, as P07 describes:

“If it’s a desktop computer, they can just sit down and use it. We have no sign in process or signup process.”

Notably, these login decisions also prioritize availability and usefulness of computing services—another way of supporting patrons’ capabilities—over security of library infrastructure, by making it easy for people to use the public PCs without extensive signup or login procedures.

Patron privacy vs. availability and usefulness of computing services. In other cases, the goals of patron privacy and availability/usefulness of computing services come into conflict, mirroring

well-known tradeoffs between confidentiality and availability in other contexts [79]. Measures taken to protect patron privacy from third-parties and law enforcement can create new challenges for patrons, who are unable to store data for later reuse. For instance, one participant describes a patron losing all their documents because they did not realize that anything “saved” to the computer would be erased at the end of the session.

P04: “Patrons are not always aware of how obsessed with privacy library folks are, so what happens is they will save their resume to the to the desktop, not understanding that when their session ends [it is deleted]. . . I remember a situation where the person had just saved it there. They didn’t ask for the time extension before it ended, and so she lost everything so she had to start all over.”

This demonstrates that, if library IT staff’s efforts to ensure patrons’ freedom from surveillance are confusing or unclear to patrons, they may hamper other patron capabilities.

Relatedly, libraries often block software downloads to protect patrons’ safety, which means that patrons don’t get the same rich user experience that individuals who own their own PCs would:

P05: “People can’t download anything, they can’t add an extension to their browser, like if they were at home. It’s kind of locked down. . . it’s a subpar user experience.”

These restrictions can also limit patrons’ access to critical modern social and work-related activities. For example, P02 described a patron who wanted to use the developer software Xcode, presumably for work or education, but it was not supported on the library’s current version of the MacOS operating system. P02 was wary of providing an upgrade, in case the updated OS was not compatible with their security software, Deep Freeze:

P02: “A patron came up to me and they were using one of our iMacs and they wanted Xcode, but all the iMacs are running High Sierra and didn’t support the current version of Xcode off the App Store. You’d need to do a full update of the operating system and I wasn’t able to do it for them at the time, because of how much time it would take, but also because I’m not sure how much compatibility Deep Freeze would have with the newer versions of Mac.”

In each of the above cases, libraries opted for privacy and security over availability: limiting some capabilities associated with work, education, or flexible use of computing resources. In these cases, library IT staff placed higher value on freedom from surveillance. Several participants noted that while patrons do express frustration with this situation, they accept it as a necessary condition of using library resources.

In other circumstances, library IT staff make the opposite choice. For example, the print-via-email-attachment mechanism described by several participants (Section 4.3) prioritizes availability and usefulness over privacy. As P04 points out, sensitive data is sent to their branch email address.

P04: “[A branch email address is] set up to get attachments from the public. And then we print them. . . it’s usually pretty sensitive, you know, like a woman yesterday had tax returns.”

Despite the sensitivity of these printed materials, P04 indicated that they did not consider *themselves* as a potential source of privacy violation:

P04: “I don’t want to be like the doctor who has seen enough naked bodies, but we see a lot of tax returns. It doesn’t really faze me anymore.”

This tension between privacy and availability is particularly acute in cases where patrons do not have personal computers at home, or cannot safely use them for sensitive tasks. If these patrons cannot usefully take advantage of library resources, their capabilities will necessarily be limited. Choosing between privacy and availability risks reifying privacy as a “luxury commodity” [160].

Relatedly, P05 worries that simply not having a computer at home—thus making internet use less routine—limits patrons’ ability to anticipate or guard against risk:

“If your only computer experiences at the public library... it’s more likely that you don’t have a computer at home. Or if you’re there using the WiFi regularly, you might not have internet at home. And so it’s just probably likely that you have less tech knowledge. And so you’re not really necessarily thinking about the danger mechanisms or threats.”

Other participants reported struggling to inform patrons about risks to their privacy, despite the tools and strategies they had in place, especially when patrons have important competing priorities. For instance, P04 described patrons’ needs to complete essential tasks, despite potential privacy risks:

“It gets very tricky, because like if you’re doing your taxes... what I consider sensitive, [the patron] may not.”

4.4.4 Constraints of library IT staff

Participants report that their ability to achieve their various goals in support of patron capabilities is also hampered by several constraints that are common across institutions, including

reliance on external vendors; limited resources, staff, and training; and outside decision-makers. Here the desire to protect patron's capabilities is undermined by library IT staff members' lack of political and perhaps technical clout.

Library staff rely on external vendors. Vendors hold a significant amount of power when it comes to the relationship between vendors and libraries. As P01 notes, libraries are not in a position to negotiate with vendors who, for example, provide funding support for the American Library Association:

“We don’t have meaningful collective negotiating power . . . to say no to these practices, and part of the reason we don’t is that the American Library Association is, like, captured by the vendors, so they are too afraid to take a meaningful stand because they don’t want to alienate the vendors who helped fund their twice-yearly annual conference.”

This imbalance in power places libraries in a difficult position—they have to choose between keeping their patrons’ information out of the hands of vendors (who may not prioritize patron privacy), or tolerate the vendors’ collection behavior as the price for providing critical services.⁶

Libraries have limited financial resources. Participants describe a lack of (or uncertainty around) financial resources, hindering their ability to implement their IT goals. P11 has had their budget cut to zero:

P11: “My IT budget is zero. Like the line item literally says zero for it. So everything

⁶The ALA’s Library Privacy Guidelines for Vendors [11] provides some guidelines for how libraries should choose and manage third-party vendors, emphasizing privacy as a core value.

I have to do, I have to do either free, or I just ended up paying for it myself and not telling them.”

Some specifically noted challenges in purchasing essential privacy software like Deep Freeze.

P05: “More libraries in the last five years are getting over to [Deep Freeze], but again, it costs money. If you’re a small library, you know, it’s probably not as available as an option to you.”

In the face of limited resources, participants report having to sometimes compromise on their goals and values. For example, P06 discussed implementing web filtering—in their opinion, a form of undesirable censorship—in order to obtain federal grant money, which was needed to maintain high-quality internet service.

“When we brought fiber into the building, the only way for us to do that was to get federal grant funding, which then requires us to filter to some extent. That was something that I did not do for a very long time, because I have a philosophical issue with it. I don’t like it, I wish we didn’t have to do it, but the internet service that we [were previously] getting... we couldn’t afford it. It was a choice, like, do we take the grant funding and do some basic level of filtering so that we can provide internet for our community, which is so important now? Or do we keep it unfiltered and have terrible service, which, you know, does affect our community.”

The lack of resources, therefore, exacerbates the tension between availability and usefulness of computing services and other goals and values.

One participant explicitly associated budget cuts (and therefore a smaller IT staff) with transfer of IT responsibility from the library to the city, noting that the city's use of software-as-a-service and remote servers also reduced their workload:

Library IT teams are often small. We learned from our intake survey that half (6 out of 12) of our participants have teams of three or fewer people dedicated to IT at their libraries—including three participants who don't have *any* staffing support besides themselves. Two participants (one who works for a major U.S. city's library system) said they worked in teams of 15–20 dedicated to an entire library system, rather than one branch. In one surprising case, the participant we spoke to was both the primary IT provider for their library and the library's director.

P09 explicitly linked small teams—often a function of the limited resources previously discussed—to decreasing ability to protect patrons.

“Especially now that there's cutbacks, they just don't have the people and the resources and the know-how... And I do kind of worry that that means there's less focus on IT, and there's going to be fewer people, and things aren't going to be done, things aren't going to be updated. No one's going to be paying attention to the network and things are going to be less safe and secure at these places.”

Library IT staff often do not have formal security training. Several of our participants describe either themselves, their coworkers, or their supervisors as coming from primarily non-technical backgrounds. This hampers their ability to identify and implement technology that could mitigate surveillance, provide better privacy, enable ease of use, and/or reduce costs. P03, for example, started as a circulation assistant at their university, and eventually came to work there full time as an IT staff member, but did not pursue a degree in computer science or information

technology.

P04 indicates that they have little understanding of how the infrastructure works because they have not been trained:

P04: “To be honest, I don’t fully understand the tech infrastructure that we are utilizing. There’s no effort made to educate me on it.”

This is, in part, because the library administration does not necessarily prioritize technical qualifications when they employ IT staff, and because personnel with both library and IT qualifications can be hard to find:

P02: “The way they hire people is a little bit strange. Like the way they hire library technical assistants, for example, you don’t need any IT credential. They can have library technical assistants that are just like, library-centric, and then library technical assistants that are just IT-centric.”

Only one participant, P08, reported having formal education in both library science and IT, via a library degree with an IT concentration. However, P08 noted that this background is unusual, and that a lack of similarly trained professionals (which, according to them, seems to be the norm) hurts libraries’ ability to achieve their privacy goals:

“Individual libraries... don’t have the resources to have a privacy person... dedicated privacy trainers or data governance... I’m also seeing a whole number of libraries who are struggling with privacy issues, and there are a lot of gaps that are being filled by some people in the profession, but there’s not many people who are dedicated database privacy and security folks.”

Library IT staff can be limited by outside decision-makers. Participants note that sometimes those making decisions about technology that affect patron privacy and ease of access are not familiar with the library or deeply knowledgeable about how it operates. As a result, changes made for the sake of security are not always in line with patron interests and capabilities, and IT staff are then left to deal with the resultant frustration that patrons experience.

For example, P06 described the administrative bodies responsible for purchasing system upgrades' lack of engagement with patrons' day-to-day challenges as a barrier to availability and usefulness of computing services, especially when updates disrupt tasks and interfaces that patrons are accustomed to:

P06: "Most directors are fully administrative. And so I think sometimes directors or IT managers might not consider the same things I do, because I'm the one helping the patrons, so I know what a pain that's going to be if Windows decides that, what was the version where they decided to put the start menu in the upper right hand and some of the lower left? It was horrible, you had to explain it to every single person that used it. I think, when it comes to upgrading public stuff. . . I think very hard about that."

P01 notes frustration with administrative bodies comprised non-librarians and people without strong commitment to the library:

P01: "Most of the people who are in library administration are not librarians and [in] that mindset. A lot of them, they're like MBAs, and even worse, like the library board is usually a bunch of local real estate developers who just want to pad their resumé. They don't actually use the library, they don't really understand how it functions."

P02 says that decision-makers who are out of touch also don't know how to make staffing decisions that could better support their goals and resource constraints:

P02: "If they hired a real IT tech, they could just handle everything here for a much cheaper cost, but maybe they're just scared of that. A lot of the administration in the library doesn't really understand technology too well."

Library decisions and policies are also sometimes determined by city or county government and their centralized IT staff. P05 described how this can cause tensions in priorities and values:

P05: "If you're not a large library, your IT is probably the city's IT department, which makes it really challenging because they don't understand libraries. . . [County name] County's IT staff doesn't really have an understanding of how libraries work and you know, if the manager doesn't really care, the staff isn't going to really care to learn."

4.5 Discussion

In this section, we explore Sen's concept of capabilities and discuss larger themes and take-aways from these results, particularly around high-level goals and conflicts with other resources and stakeholders. We echo those in the Privacy Enhanced Technologies (PETs) community that have called for a systematically nuanced approach to understanding people's context, freedoms, values, motives, and abilities, advocating for a departure from an idealized usability approach to a capabilities approach [56].

4.5.1 Sen's capabilities applied to library IT

Sen stipulates that a community should define their own capabilities. In the context of our study, we agree with Sen that capabilities are most usefully defined by those who have insight into day-to-day library activities, including the nuances of what patrons want and need from library IT resources, and library IT staff are often well positioned to do so. We therefore allow our participants to define those capabilities, rather than using an existing framework like that provided by Nussbaum [154]. We find that library IT staff take a contextual or intersectional approach to supporting their patrons [133]: recognizing, for instance, that mitigating surveillance concerns can be highly valued capabilities because they support other sets of capabilities like freedom to access information for employment and social services, etc.

Like Sen, library IT staff think about capabilities both as negative rights (e.g., protection from police surveillance) and as positive rights (e.g., access to coding software that would allow them to perform their job) [193]. Library IT staff perceive that information freedom is tied to privacy guarantees; while those freedoms are articulated through different capabilities in different contexts (e.g., anonymous logins, non-city-dependent IT infrastructures, etc.), the relationship between the goal of information freedom achieved through those capabilities and economic, social, transparency and protective security is ultimately clear [196]. We point this out because privacy strategies are not simply ends in and of themselves, but part of a larger strategy toward achieving freedom.

Library IT staff also understand that while access to internet resources is a critical professional resource, access by itself is not enough. Privacy protocols can create a barrier that makes access less fluid or useful (e.g., security and privacy policies preventing patron usage of pre-

ferred software), but can also enable security for patrons to engage with internet resources (e.g., obtaining information about sensitive or stigmatized subjects without fear of retribution).

Throughout our discussion, we highlight the specific metrics that are relevant to understanding enhancement of capabilities of library IT patrons and to making recommendations and evaluating them.

4.5.2 Privacy needs and moral capabilities

Many of the participants were acutely aware of their role as defenders of values: allowing patrons to use internet technologies without being surveilled by the library, other patrons, or outside entities. This role is ingrained in library values, and, as one of our participants told us, these values are more critical (and possibly more difficult) to teach than IT skills.

Above all, library IT staff considered surveillance to infringe on patron's capabilities—such as freedom to find employment and information, or access other critical resources or social services, etc.—and were mostly willing to put aside library security concerns for this moral objective. Indeed, Sen's moral approach (which relies on the community to adjudicate value [198]) fits our context well because, as we find, library IT staff are most adamant about protecting their patrons from *law* enforcement. This view stands somewhat in opposition to the normative view of Rawls' distributive justice [172, 173] and Walzer's communitarianism [226], in that it challenges government and legal institutions and does not necessarily or entirely put faith in fair distribution of goods or the collective good. Rather, Sen's consequentialist conception of justice deeply considers social realities and aims to eliminate injustice [197], to consider value “not by the resources or primary goods the persons respectively hold, but by the freedoms they actually enjoy

to choose between different ways of living that they can have reason to value” [194].

Library IT staff are also attuned to the need to provide basic skills and privacy education to their patrons. Both of these services are critical to enhancing capabilities, which are about quality of access to achieve freedoms which patrons value, rather than simple availability of devices. That said, librarians face constraints in delivering on this promise because of external vendors, decision-makers, and limited financial resources, staff, and training.

However, security and privacy measures can sometimes be a deterrent or disappointment for patrons, especially those who use library resources in lieu of home computers. Despite library IT staff’s best efforts, these measures can hamper capabilities linked to fundamental economic freedoms, for example when data is erased after a session or when patron’s can’t install software they need for work or education. This paradoxically conflicts with library IT staff’s perception that they have nurtured capabilities in support of economic, social and social security. At other times, library IT staff prioritize providing easy access to PCs, as if they were patron’s own, even when this introduces privacy concerns, e.g., when patrons send their tax information to library email addresses to access the printer. Library IT staffers essentially must make nuanced assessments about which structural threats to privacy are (not) severe enough to warrant limiting other useful capabilities.

Recommendations. Computing researchers and software developers should consider how to create better affordances that can support security and privacy while still allowing people to have an experience on public computer that is more comparable with personal devices. For instance, libraries might offer the ability to run virtual environments on their PCs for patrons that want to use Zoom, Xcode or other unsupported software.

Library staff should be more explicit about their strategies and the tradeoffs they make when considering patrons' privacy. To better socialize the idea that devices are/should be used as impermanent tools, library IT staff should, at very least, put signs on the computers to indicate what data will be erased and why. Additionally, libraries might consider ways to offer printing that might better balance access and privacy. Building on Sen, we also recommend that library IT staff measure the success of these interventions based on how they are valued by patrons, once patrons are informed about the rationale [159]. For example, a metric for gauging the value of Deep Freeze might be to understand if more patrons use computers or use them for longer once staff provide context for Deep Freeze.

4.5.3 Values in conflict with resources, decision-makers, and other institutions

While avoiding monitoring and not requiring logins are essentially free, library IT staff must sometimes invest in software to ensure that patrons' data is never stored, and this can sometimes create barriers. For libraries with more funding, software enables them to do this quite easily—though, as some mention, they must remind patrons to sign out of computers to ensure that they do their part. Libraries with less funding may struggle since free, open-source alternatives are resource-intensive and possibly esoteric.

In contrast, protecting patrons from vendors that provide critical services like printing or ebook access may feel out of reach due to lack of leverage, and protecting against malicious threats like phishing or identity theft may feel beyond library IT staff's skill level. In other cases, library IT staff feel constrained by out-of-touch decision-makers (who sometimes consider only superficial access to IT resources, rather than more complex capabilities) from making nuanced

tradeoffs in support of their patrons' capabilities.

While the ALA advises libraries on how to navigate third-party vendor selection to choose vendors with policies that better secure patron privacy [11], library activists argue for providing encryption and ensuring that the library is the primary data steward [201]. The latter requires that libraries take on a more proactive role, that library IT staff be more sophisticated, and that vendors be more transparent about data policies. Managing these relationships as the ALA advises is therefore challenging, especially since libraries realistically have few options when selecting vendors.

Recommendations. When library staff are constrained by resources, third-party opacity, or skill level from better supporting patron capabilities—including security and privacy—there is an opportunity for the security and privacy community to provide support. Researchers could use measurement studies to provide some transparency into vendor-provided tools, develop open-source projects to better support library needs with lower cost and less surveillance, and provide more user-friendly interfaces as well as training to help library IT staff in their nuanced decision-making. Further, the security and privacy community has an obligation to develop new metrics that better align with enhancing capabilities: measuring whether and how security and privacy can enhance amount and quality of IT access when patrons feel protected.

4.5.4 Future work

This paper primarily discusses the perspectives of library IT staff members in the United States. However, while we did not include them in the final analysis, we did interview two participants that worked in libraries in non-U.S., albeit Western countries. While we believe that the

contexts of these two participants' libraries were not significantly different from those presented in this paper, we did exclude them from our findings given that they comprised a relatively small proportion of our participant pool. More work should be done to understand the security and privacy priorities and tradeoffs that library IT staff members make in other countries.

Recommendations. Researchers should seek to capture and convey the experiences of library IT staff members in countries around the world to better understand the international perspective of library IT staff's role in the greater security and privacy community.

4.5.5 Connecting to the at-risk framework

Our findings illustrate ways in which both library IT staff and library patrons can be *resource or time constrained*. Library IT staff have small teams and limited resources, leading to difficulty implementing secure setups for their public PCs and many patrons have low socioeconomic status, meaning they rely on the library for computing resources. Additionally, patrons might face other *contextual risk factors*, such as *marginalization*, as P05 described in Section 4.4.1 in relation to LGBTQ+ youth, or *legal or political* risk factors, as P08 described in the same section in relation to immigrants and undocumented people.

This highly-constrained environment requires cost-effective technical solutions to maintain the balance between freedom from surveillance and the capabilities of patrons to achieve their goals. Currently, there is a tradeoff between the two—future work should seek to reduce this tradeoff as much as possible.

Chapter 5: How Entertainment Journalists Manage Online Hate and Harassment

5.1 Introduction

As part of modern digital life, journalists often have public presences on the internet, both through direct publication and social media. However, as journalists report on things that engender a negative reaction in their audience, they may experience harassment as a result. This harassment is never justified, but often has discernible patterns — certain topics, characteristics of the journalist, or other factors — which might increase or reduce the likelihood that journalists experience harassment.

Significant prior work has discussed the digital-safety needs and practices of journalists when facing nation-state adversaries [63, 135–138]. However, harassment—defined by Citron as “a persistent and repeated course of conduct targeted at a specific person, that is designed to and that causes the person severe emotional distress, and often the fear of physical harm.” [48]—has not been studied as much in this specific context within the digital-safety research community. In this work, we seek to understand the experiences and practices of what we term “entertainment journalists”—those who write primarily about topics like movies, video games, and sports—when dealing with online harassment. Despite the less sensitive nature of their work, these jour-

nalists experience significant harassment online, which can lead to severe consequences both external and internal.

I and my collaborators chose to study these journalists due to their decade-long experience of dealing with this challenge, hypothesizing that these journalists can experience the intersection of the *prominence* and *marginalization* risk factors described in Chapter 3. Since online harassment can be pervasive and recurring, many entertainment journalists have *already* developed adaptive responses to mitigate harassment's harmful impact. By understanding their current *protective practices* (Chapter 3), we can better understand where technical and non-technical solutions may be most useful.

Additionally, this work serves as a case study of the potential interaction of *prominence* and *marginalization*—two of the *contextual risk factors* identified in Chapter 3. We hypothesize that the combination of a public profile that's used to perform and promote their work and the increased impact of harassment on journalists who experience *marginalization* makes harassment attacks particularly damaging, following prior scholarship on harassment and people who experience marginalization [40, 210, 234].

Accordingly, our research questions were as follows:

RQ1: How do entertainment journalists mitigate the negative effects of harassment, both proactively and reactively? How and where do these journalists learn these protective strategies? Are they effective?

RQ2: How is this harassment impacted by risk events (i.e. publishing something about #GamerGate, publishing an article about prejudice in a media franchise, tweeting about a current controversy, etc.)? What are the characteristics of these risk events?

RQ3: What *contextual risk factors* (Chapter 3) do these journalists experience? How do these impact their experience of online harassment? If multiple contextual factors are at play, how do they interact?

RQ4: How does institutional support play a role in these journalists' protective strategies?

5.2 Methods

We designed and conducted an interview study to answer our research questions, as described in this section.

5.2.1 Recruitment

Recruitment was conducted directly via email and Twitter direct message. We sent messages to candidates who followed the following criteria:

- Current or former professional journalist covering media and popular culture (film, television, video games, music, sports, etc.)
- Can work either independently (i.e. YouTube channel, blog, Substack, etc.) or for an outlet

The recruitment process began finding journalists' contact information through video game websites, given the specific context of #GamerGate (see Chapter 2.4). With a combination of prior knowledge and searching for news articles on popular video games, we made a systematic list of websites and platforms that publish journalism focused on the video game industry. We also invited interview participants to ask their colleagues if they would be interested in participating. This snowball sampling [151] was essential in getting more journalists to speak with us, given the

sensitive nature of the project. As a result, we also interviewed journalists who covered sports, as recommended by our participants who described sports journalists as additional common targets for harassment.

5.2.2 Protocol

We conducted nine interviews via video conference. All participants consented to have the interviews audio-recorded. As an extra step to protect participant privacy, we chose not use the audio transcription service built into our video conference software, as this requires sending the recording to a third party with no particular privacy guarantees. Instead, we used WhisperX [22], an open-source, speech recognition model which transcribes long-form audio while also identifying distinct speakers.¹ This software was run on the University of Maryland’s Nexus computing cluster—our data was only accessible by the researchers and system administrators at UMD.

Our interview protocol covered the following areas. The full protocol text can be found in Appendix C.

- **Consent procedures:** Participants were shown the consent form, and briefed especially about recording, transcription, and anonymity (described below)
- **Warmup and career overview:** We asked about the background of the participant: how long they had been in industry, areas of specialty, average readership, size of following on social media, etc.
- **Specific risk events:** We asked about the participant’s experience with harassment: what factors seemed to influence their experience of harassment, how common or uncommon

¹<https://github.com/m-bain/whisperX>

this was, and their strategies for managing it.

- **Advice, given and received:** We asked about where and how participants sought (or gave) advice about dealing with harassment in their networks.
- **Debrief:** We asked participants to describe any relevant personal characteristics that may have impacted their experience, like race, gender, or sexual orientation. We decided *not* to collect these systematically, in order to prevent participant deanonymization.

Online hate and harassment can be intensely emotionally challenging and traumatic. Our interview protocol was developed using a trauma-informed lens, as defined in Chapter 2.1.3. We used Chen et al.'s trauma-informed computing framework for computing [46] and Wong's guidelines for trauma-informed care in qualitative research [241], which both provide some recommendations for conducting user-centered research. We used these principles, following the below guidelines:

- We conducted interviews online so that participants could talk to us wherever they were most comfortable.
- We conducted warmup and debrief sessions to establish rapport and check on participant well-being, as well as address any content they may have wished us to not include in our research.
- We were transparent about the goals of this work.
- We prefaced sections of the interview which discuss difficult topics to give participants a chance to prepare their response.

Anonymity. Harassment events are highly public, so we took extra care in reporting on these events by not just removing names but other identifying details. After each interview, we asked participants if there were any details they would like us to remove or elide to protect their privacy and avoid inciting future harassment. For the same reason, we choose *not* to provide a summary of participant demographics in aggregate to help avoid reidentification. We have additionally chosen to use “they/them” singular pronouns to refer to our participants, regardless of gender identity.

5.2.3 Limitations

We do not claim that our results are representative of a general population, or even of all entertainment journalists, following appropriate guidelines for reflexive thematic analysis [37]. Instead, the goal of this study was to identify key themes and ideas related to our research questions that are richly situated in our participants’ contexts and experiences. Accordingly, we did not seek *data saturation*, as that implies an approach contrary to the methods and goals of reflexive thematic analysis [34].

Additionally, since harassment is a highly-sensitive and often-traumatic subject, some potential participants with valuable perspectives may have declined to speak with us to avoid further psychological harm. In particular, most of our participants were cisgender men, or present as such online in a way that shields them from gender-based harassment. Many of our participants told us that women, people of color, and trans people experienced even greater levels of harassment than themselves (see Section 5.4.3), and that perspective is largely missing from this work due to this challenge.

5.3 Analysis

To analyze our results, we used Braun and Clarke’s process of reflexive thematic analysis [33], described broadly in Section 2.5. Coding and theme development was conducted collaboratively by two researchers, guided by a combination of a *deductive* approach—using the framework developed in Chapter 3 to understand the data—and an *inductive* approach—using the data itself as a way to understand our participants’ experiences. I discuss this work’s relationship to the at-risk framework (Chapter 3) in Section 5.6.

Reflexive thematic analysis also exists on a spectrum between *semantic* meaning—focusing on the explicit content of the data—and *latent* meaning—focusing on the implied or subtextual meaning of the data. On this spectrum, our analysis process tended towards the semantic, focusing on concrete practices and attitudes. However, some latent meanings were analyzed in relationship to how our participants understood the general experience of online hate and harassment.

Participant quotes were lightly edited for clarity, primarily by removing some (but not all) filler words or repeated phrases. We seek to capture the tone and style of how our participants speak while still retaining clarity.

5.4 The context of harassment

First, we describe what it’s like to experience harassment while being an entertainment journalist across three core themes.

5.4.1 The inevitable price of admission

Social media—most often Twitter, but not exclusively—was commonly described by our participants as essential for their jobs. They use it to contact sources, advertise their work, and receive tips, which are all critical elements of their profession. However, the need for a public presence simultaneously exposes them to targeted harassment that is impossible to completely avoid.

Many participants, particularly the most experienced ones, had a resigned attitude toward this problem, characterizing it as an inevitable feature of online life. Less experienced participants had also had to reckon with this—only two years ago, P09 was warned about the difference between a private and public persona on social media when starting their current job as a games journalist. While individual participants all characterized harassment as happening at different frequencies—from rarely to occasionally to daily—it was a familiar experience for each of them.

Confusion over the effectiveness of moderation tools, shifts in Twitter after Elon Musk's purchase, and the general distortion of reality these participants experience on social media were commonly referenced as platform features that made harassment particularly harrowing. Some participants described looking for other platforms, like replacements to Twitter such as Mastodon or Blue Sky, or moving to more direct-to-subscriber business models—such as a paid Substack newsletter, a Patreon account, or Aftermath, the recent subscriber-supported, worker-owned project from former Kotaku writers [165]—that reduced their need to be active on public social media.

Harassment is ephemeral. Although every participant was familiar with harassment, many char-

acterized it as ephemeral, temporary, or impersonal. A common sentiment was that harassers will inevitably find a new target, particularly if you do not engage (further details on non-engagement as a strategy can be found in Section [5.5.1](#)). P07 told a story about tweeting that it was important for more women to be involved in sports writing, and described the outcome as follows:

Some right wing troll account that doesn't cover sports amplified it for whatever reason. And so for like three days I got people on Twitter yelling at me that I was anti-man or that I supported mediocrity in sports writing or four-letter epithets and everything. And just because I wasn't used to it, it really hurt my feelings. And to the point where like my colleagues were like, just give me your phone. And like, don't look at it. Or we're taking away your privileges here for a little bit. And then, you know, in three days it goes away. Because the people yelling at me don't know who I am. They don't know what my beat is. And it goes away.

Here, P07 also describes a common strategy for managing harassment—having a colleague or trusted other take away your device so you can step away from the online world. Further detail on this strategy can be found in Section [5.5.3](#).

Where some participants emphasized a “just a name on a screen” (P07) mentality, focusing on social media’s inherent disconnection from reality, some instead emphasized that taking action was indeed important. P02, for example, specifically contradicted the common advice to ignore harassment:

Like you have to take steps. You might not want to engage, but you do have to take a step back or make some kind of statement. But just ignoring it does not help and will often make it significantly worse. Once that fire is going, you have to at least

start digging trenches around it so it doesn't spread. So when people start to say, you can just let go and it will flare down, no. You have to do something. Even if it's just getting away from it and taking that step for your own safety. The don't poke the bears thing, yeah, I don't agree with that either. [...] I don't think there's a way to ignore the trolls and then they'll go away. We've proven that's not the case.

While P02 still mentions not engaging, here referring to arguing with or making fun of the harasser, they make clear that *some* kind of action is still required. Shortly thereafter, P02 later vividly compared harassers to toddlers:

It feels like dealing with a toddler who is trying to get their parents' attention any way they can. And it might, you know, start really innocent and sweet, and then if they're still completely ignored, they might start pushing over lamps. And this is the version of pushing over lamps. And then once it gets to that point, the people who just want to push over lamps, like, yeah, f***ing game on.

Harassment is dangerous. Some participants experienced dramatic escalations of harassment beyond social media comments and hateful emails. P06 called to mind the E3 leak—an incident where the Entertainment Software Association leaked the addresses and contact information of hundreds of journalists attending the video game industry's largest trade show [150]—which exposed their home address. Describing the impact of this experience, P06 said:

I've had people like paste, like in an email or a DM, just my address. And like, that's it. Like, they don't say anything. [...] They're just able to post that they like, "Hey, we know where you live." Um, that's just information that's out there. [...] It's warm

out. It's summer. I don't close every window every night. It's nice to have air going, but you can't help sometimes but wonder. It's like, all right, well, you know, people out there have my address. All it takes is one person to [...] have the wrong idea. And it's like, oh, because I didn't lock the screen door last night, someone can just come into the house.

Having one's address exposed was a common fear, especially for participants with children. As P06 describes above, having an exposed address meant that there was potential for direct, physical danger due to this leak. Receiving strange and upsetting mail or being swatted² were other escalations that were top of mind for our participants.

Dangerous outcomes of harassment may also be more likely or more severe for women, people of color, and trans people. Although none of our participants described an escalation that they perceived as related to their identity, the underlying sentiment that harassment is worse for people who experience marginalization was very common.

Harassment has a persistent emotional toll. All participants described harassment as having at minimum a moderate emotional toll and impact on their mental and emotional well-being. For some, it was a constant reminder that there was a group of people who would seek to tear down them and their work, at times leading to insecurity, as P05 describes:

There is sort of like a seed planted of just like insecurity. Just like if I voice my opinion online, will people get so mad at me that they're just gonna cuss out my entire existence?

²Swatting is a common colloquial term for a *false reporting* [210] attack where an attacker will call in a spurious police report in order to get a SWAT team to descend on the home of their target.

For others, it escalated into complex PTSD or severe anxiety based on their experiences. P02 describes the impact of the #GamerGate controversy on their life as having this type of severe impact:

I have complex PTSD³. I'm hypervigilant because it was a legitimate, long-time, traumatizing event. So in a really sad way, that hypervigilance is super helpful to anyone who is doing any of this stuff.

Here, P02 also points to the adaptive elements of this condition, discussed further in Section 5.5.2.

Tensions between characterizations of harassment. The three characteristics of harassment outlined above demonstrate a core tension: how can harassment be ephemeral in one sense, but dangerous and emotionally draining on the other? These can all simultaneously be true. A specific harassment event, like an insulting tweet or a threatening email, may indeed be a one-time event, and social media furor will eventually pass. Despite this, the *pattern* of these individually-ephemeral events leads to genuine fear about physical harm and attendant psychological consequences.

5.4.2 Hurt people hurting people

Based on their experiences, our participants held very strong beliefs about who, exactly, was doing the harassment. Harassers were characterized variously as “angry”, “lonely”, “cowards”, and in general, poorly-adjusted. Some participants pointed out that they believed harassers might be seeking some sort of emotional need, often as simple as getting attention or having

³Post-traumatic stress disorder

their voice heard, which occasionally engendered some compassion. P05 said, “There are a lot of instances where like I want to engage with these accounts and just say like “Are you okay?” Because these responses are not like coming from a person who’s in the right mindset.” It was not clear how these participants reconciled this more sympathetic framing with the real harms that can result from harassment.

Some participants made reference to particular intellectual properties or groups that would incite harassment. P05, who covers college sports, described how making a small mistake about a team or incorrectly predicting the outcome of a game could lead to harassment far out of proportion to the perceived slight. P03 described the Star Wars and Game of Thrones franchises as attracting particular harassment, especially in relation to “spoiler-phobic culture.”

5.4.3 White cisgender maleness is a shield

Our participants who presented as White, cisgender and male on the internet made it abundantly clear that this presentation protected them from worse harassment. They frequently mentioned that women, people of color, and trans people got harassed much more often and with greater severity, in some cases to the point of being chased out of the industry entirely.

Presentation is a key element of what these participants described. P07 mentioned that despite having a name that sounded stereotypically White and having lighter skin, they were Latino, which changed how people spoke to them:

I’m half Brazilian, like literally whole Brazilian citizenship—mother was from São Paulo, have a Brazilian flag in my account. And so occasionally Latino things come up or about Brazil, I’ll talk about those things, but because of my name and because

a lot of Americans don't realize that Brazilians can be white too, or white-ish, [they] will feel very comfortable saying something pretty anti-Mexican to me before realizing, oh, he's an immigrant kid too.

P03 mentioned the same phenomenon in terms of their gender identity—despite identifying as genderqueer, they were often perceived as male and used both “he” and “they” pronouns. As they say:

A lot of the time, by being White and presenting as a White guy, I think a lot of time people will take me a little more seriously or be a little less cruel to me than they might otherwise.

Neither P03 nor P07 described intentionally presenting as White or male in order to avoid harassment. Instead, they referred to this phenomenon as passively protective—harassers did not use their identity as a way to attack them, since they were not part of a commonly-marginalized group at first impression.

Our participants who presented as people of color described this having a particularly negative impact on their experience. P05 described needing to preemptively mute negative words relating to their identity as a Mexican-American and child of immigrants, which required them to think through the worst names someone could call them.

I can think of any hateful words to describe people, so like the way they antagonize me and like my people, I will add those to the list of muted words so that I don't see them randomly when I'm tweeting about sports or tweeting about like current events.

This task, even though it was protective from future bad actors, still required a significant amount of emotional resilience and energy, as they go on to discuss:

It requires me to actually sit down. And at least for a little bit, when I'm in a good, like mental space to come up with as many ways to, you know, to insult me. And, you know, that's no fun either but at least I'm in like a good place to come up with all those as opposed to when I'm feeling down or when I'm currently being harassed.

P09, who is Black, mentioned that commiseration with other Black writers in the industry was essential for both emotional support and seeking advice—having other folks who shared the same particular experience as them was helpful in navigating the particular challenges they faced due to racism on the internet.

It just kind of became a thing where, like, if I ever had, like, a more specific kind of question about harassment and things that I face, I'd feel more comfortable to ask someone in this space who [...] would be Black and would probably have the same kind of, like, avenue of harassment that I'd face. So I kind of would reach out to them as like point people to be like, “so what can I expect getting into this role?” And I'd like kind of talk to [them] on and off kind of just about like specific harassment stuff. So like we'd kind of, like, be like go-to points for each other for the most part for stuff.

P08 expressed distress at this phenomenon and described how harassment has driven people without this shield out of their industry entirely:

The thing that bothers me the most, though, is that, like, I am a privileged cis white male and I have a lot more tools at my disposal that makes it easier for me to, when needed, kind of stand in front of this wind that blows occasionally. And a lot of

women and a lot of people of color and a lot of folks less advantaged than me have not been able to do that. And they have had to, pick up and leave their lives and give real aid to those closest to them.

And some of those voices are the voices that brought me to this career path. There are literally people that brought me into this field who are not here because they weren't able to kind of weather this storm. And it hurts to know that their voices have been silenced and that they are no longer doing this work because of what happened to them. And that hurts at a very intrinsic level because I value this work so much and I miss them so much. So it lessens my own work not to have their work here next to me being read as well.

5.5 Strategies for dealing with harassment

Now, we talk about how our participants mitigate the effects of harassment, and why they chose these methods.

5.5.1 “Just ignore it”: the best worst option

Many of the techniques our participants used to protect themselves from harassment involved some variation on simply ignoring it. As described above, harassment is ephemeral, yet simultaneously dangerous and emotionally taxing. Nevertheless, most of our participants believed that rather than engage with harassers, it was better to try to ignore it and move on. Many participants stated that any kind of engagement, especially trying to call out harassers' bad behavior, would encourage rather than chastise the harassers—in P06's words, “The moment you

start talking about your harassment, you are going to get harassed. [...] That's just fuel for people because they're noticing like, oh, it's getting to them." P03 also describes this in terms of their perception of the harasser's goal:

These are people just trying to get a rise out of me. And if I give them that, then they get what they're seeking for. And if I ignore them, then they don't.

Participants often described this approach with considerable ambivalence. Although it may be the best response most of the time, several participants did not want to diminish the impact of harassment on others, especially women, trans people, and people of color. P06 experienced particular conflict between ignoring harassment and retaliating, particularly when observing others' behavior:

I have trouble squaring [ignoring harassment] [...] I don't consider that to mean I'm endorsing, like just letting harassment happen. But I do sometimes see amongst people like, they're getting harassed, and their response to it is to like, quote tweet a harasser, and be aggressive with them. Harassment tends to beget more harassment. And so unfortunately, it's like you have a person who is hurting and being harassed. And then, of course, what they're going to want to do is punch back because the platforms are not built in a way to handle this. That is the form of recourse that some people have is just to get angry in response and that frequently seems to just agitate, you know, that's what [harassers] are looking for. And so I don't like that my solution is essentially to just ignore it.

P08 also expressed the impossibility of ignoring if harassment reached a certain level of severity, calling to mind the earlier theme *harassment is dangerous*:

At the same time that you can't engage, you do kind of need to keep your head on a swivel so that you're aware of what's going to be outside your door when you open it.

Blocking vs. muting. A key technical affordance of Twitter⁴ is providing ways to block or mute users. Blocking an account means they cannot follow you, see what you are saying, or tag you in their own tweets, as well as preventing you from seeing the account. Muting, on the other hand, simply prevents them from showing up on your feed—they can still reply and see your account, but if they tag you, it will simply not show up.

The primary difference between these two approaches, which have fairly similar outcomes from the perspective of the target of harassment, is that blocking is observable by the blocked user, whereas muting is silent. Being blocked was perceived as potentially a badge of honor to the harasser—a sign that they had successfully gotten an emotional reaction from you, and thus achieved their goal. Muting's silence prevented the harasser from getting what they wanted; the target would not have given them the satisfaction of a strong response. P09 describes this rationale:

“When people see that they get like, blocked, they see it as like some like, badge of honor, like, look, we, like, defeated this person in some verbal spat, [...] some debate or whatever and [...] they] see it as like a badge. So in a weird way, I'd like rather not give you that weird, like, victory in your head that you're right because I blocked you or whatever.”

Participants were familiar with block lists—automated tools that could block huge num-

⁴In this context, we focus on Twitter because our participants did. Other platforms have various blocking affordances that may differ from what is described here.

bers of accounts that had been collated by others—but did not use them, due to false positives. Although they might get rid of large amounts of potential harassers at once, the tradeoff of accidentally blocking a colleague or friend was seen as not worth it.

Stepping away from the screen. Despite the potential physical danger of harassment described above, our participants often expressed that most harassment has little-to-no relationship with the offline world. When asked about what advice they would give other colleagues, P07 made a point to emphasize telling others that “it’s important to remember that what you see on TweetDeck is not real life.” Accordingly, simply stepping away, closing the laptop or phone, and connecting with friends and family in the real world was often a crucial strategy. P05 described the following when advising other members of their team, echoing their own strategy to step away when harassment got bad:

You can scroll through Twitter or Instagram for so long, and then your brain just gets fried. So go outside, go take a break. It might seem counterintuitive for the manager of a social media team to tell you, but like, it’s really important to just like log off, tune out and sort of refresh your brain in that way.

This distinction, however, remains difficult in the context of real physical threats. P06, who had experienced a variety of high-harassment events, still characterized harassment as being “99.9% online”, despite also purchasing physical security after their address was leaked to the public. This suggests a pattern of low-probability, high-harm events when harassment escalates past insulting comments online.

Look for good-faith actors. Occasionally, participants would describe being willing to look for good-faith actors amongst their social media replies. If someone was unskillfully but honestly

engaging with the participant's argument, rather than attacking their identity or character, they might engage with that person rather than block them. This did require energy and effort that other participants were not willing to spend. P04 describes doing this with a surprising sense of playfulness:

If someone's a little mean but comes in sort of wanting to have a conversation, sometimes I'll send one or two Twitter replies. [...] If someone takes the time to find my email address and sends, like, a mean email, I will sometimes get a little cheeky and be like, "thanks for reading" or, you know, "glad you liked it." Or one time I said, "so does that mean you won't be RSVPing to my birthday party?" You know, stuff like that, depending on the kind of mood that I'm in. If someone took the time to email me.

5.5.2 Constant vigilance

Potential harassment clearly had a persistent impact on how our participants chose to use social media and how they approached publishing their work. Participants drew an explicit connection to writing and publishing about sensitive topics, particularly when speaking with a politically left orientation. In some cases, this led to a chilling effect; deliberate self-censorship, even about issues important to the participants, was sometimes seen as necessary to protect themselves. P06 describes this effect with a tone of resignation, framing this decision in terms of protecting their family:

[I have] not really weighed into certain topics to the degree that I might've done in the past, because [...] I have a, broader obligation to think about, which is my family.

And there are younger people with more energy than me to sling those arrows and to take them these days.

When describing the impact of harassment on their behavior, P02 described themselves as “hypervigilant,” (Section 5.4.1) implying a watchfulness that went over and above what was needed to protect themselves due to their prior experiences with severe harassment. They go on to describe their strategy around *any* social media activity as follows:

It’s high stakes, low risk. It’s very rare for a tweet or any message to blow up in a negative way. But if it does, everything’s in play on all aspects of your life or the company [which hired you to run their social media]. So when you think about it that way, low risk, high stakes, it’s like you kind of do need to bring that care to every single time.

Writing about issues of particular sensitivity would often lead to preparation for potentially being the target harassment. P01 described in detail their outlet’s decision to cover or not cover Hogwarts Legacy in early 2023, due to J.K. Rowling’s persistent promotion of anti-trans myths [179]. While ultimately deciding to cover it, the outlet decided to turn off comments in advance and make sure that any journalists who covered the game would be prepared to receive harassment after an article was published. P01’s strategy was the most concrete—many participants made reference to simply being mentally prepared for harassment when covering a sensitive topic.

Other sensitive topics included the war in Ukraine, violence in video games, anti-racism, feminism, and trans rights, succinctly summarized by P04 as “any issue that’s in the culture war at any given time.” Participants took varying levels of care when reporting on these topics, de-

pending on their perception of the likelihood and severity of potential harassment. P02 describes a strategy commonly seen amongst colleagues:

You pay attention to what the big topics are. Right now, for whatever reason, trans athletes, we know that's going to be a big one. Anything having to do with Pride, we know that's going to be a big one. You just kind of keep a list of like things that everyone is talking about in a positive way and things that might be controversial. And most people I know who do this kind of work keep that list in their head. So do I. you just kind of get a sense for it.

Participants with children were particularly mindful of what they said or shared on social media. In addition to the above self-censorship, P08 took active care to delete images of their children off of Instagram when Twitter, in their mind, started to decline:

Preemptively, as Twitter began to melt down, I went to my Instagram and I deleted every image of my children out of my Instagram to kind of sanitize that and make that a place where I could land professionally, if need be.

5.5.3 External support is critical

Participants relied strongly on their social networks and employers in order to support them in mitigating harassment. This was seen as crucial—P08 described that in comparison to having going independent (as opposed to working for a larger organization) would be terrifying, as not having “DC lawyers” on his side would make facing harassment vastly more challenging. As they described, “That’s a good feeling to go to bed with at night, the next morning, no matter

what happened at work that day. to know that there are some angry, smart motherf***ers with law degrees in my corner.”

Concrete support. Our participants often relied on their news organizations as a primary vector for concrete solutions. These companies provided a variety of resources, often footing the bill for particular tools and services. For more severe incidents, these organizations also provided legal support or paid for cameras to protect their journalists’ homes. Not every organization always had these policies—P08 described how their publication did not have a concrete harassment mitigation policy until the #GamerGate controversy targeted them and their colleagues. However, this led to improvements for *other* publications in the same parent news organization who learned from P08 and their colleagues’ experiences.

Participants made frequent mention of data deletion services like DeleteMe⁵—companies that, as a service, will trawl the internet for one’s personal details and get them deleted or taken down—which were sometimes paid for by their employer, and other times paid for by the participants themselves.

Colleagues and management also offered support mechanisms, both formal and informal. P01 and P03 both described a Slack channel where their colleagues could share stories, information, and seek support when facing harassment. Several participants made reference to giving colleagues their phones to perform triage in the face of severe harassment—this meant that the person who was suffering harassment did not have to deal with blocking, muting, and generally managing the incident and could instead take some time for themselves to recover and let the incident pass.

Local police were mentioned by two participants, although in contradictory terms. P08 de-

⁵<https://joindeleteme.com/>

scribed educating their local police department about swatting (defined in Section 5.4.1). They remained in regular contact with their police department to ensure the threat was accurately understood. P06, on the other hand, described their police department as completely oblivious to this threat, despite repeated attempts to educate them. Police were not mentioned by any other participants as helping to handle or track down threats.

Emotional support. Even more so than the above concrete support strategies, commiserating with and getting emotional support from colleagues was essential for managing harassment. P03 describes their colleagues setting an example for how to respond to harassment:

I remember when I came up on the college football team, 'cause that was, like, the work environment that really cemented a lot of how I approach things in me. You know, it was my first full-time job in journalism and [...] I was the youngest person on the team, you know?

So it's like one of those situations where everyone else there, you're kind of looking up to them to set an example. And a lot of those guys were, really funny, smart, like the smartest, funniest people I knew. They were all kind of brash Southern guys who like were aligned with me on like morals and stuff, but also were like really into college football.

And I say this stuff about them being the funniest, smartest guys I know, because then they would still just get so many stupid people saying just, like, heinous things to them on social media in the comments and the way they responded was mostly to laugh.

Our participants often described how simply sharing their experience and getting validation

and shared recognition from their colleagues was valuable for keeping themselves healthy in the face of extreme harassment. P09 also referenced specifically seeking out other Black writers in the space, because their particular experiences around harassment and race were more specific and useful.

Some participants experienced severe mental health consequences, and sought external support from therapists. Sometimes this was helpful, but P04 describes a therapist dismissing their concerns after a period of particularly intense harassment:

I even went to therapy because I was bothered so much by [harassment]. Therapist saw me once and said he wouldn't see me again because this wasn't actually a mental illness. I was like, "yeah, that's fair." But I was looking for any tips on how to deal with it because I was letting it occupy too much of my brain space.

Support from family and friends was also crucial, especially when stepping away from the screen and spending time away from the source of the harassment.

5.5.4 IT hygiene

A certain amount of basic "IT hygiene" (P04) was a part of the toolkit of our participants. Multi-factor authentication was mentioned as a tool to protect them, particularly to prevent a harasser from taking over their account and causing severe reputational harm. P05 describes their approach:

The one thing that I'm sort of more focused on sort of in conjunction with harassment is just getting my accounts hacked by people who sort of want to troll me or just want to take it a step further. And because of that, I've just gotten into just having my

accounts as secure as possible, whether it's like setting up two-factor authentication or physical security keys, so that even if somehow my accounts get hacked, they still won't be able to access stuff and post as me, impersonating as me, so that they ruin my professional and maybe personal life because I've seen instances of that.

P05 goes on to describe being careful to watch for malicious links and limiting which devices were logged into their accounts. Perhaps surprisingly to the digital-safety research community, only P04 and P05 made reference to specific, traditional security advice. For targets of harassment, security advice may need to be more contextual to be useful, as prior work has demonstrated [177,234].

5.6 Discussion

In this chapter, we show that entertainment journalists experience severe harassment—while any *individual* insult or attack may be ephemeral and temporary, the *pattern* of harassment exacts a severe emotional toll and can sometimes escalate to real danger. In response, our participants largely adopted an “ignore it and move on” attitude, choosing to disengage in order to protect themselves. They rely on external support to help manage this, both practically and emotionally. We also investigate the intersection of *prominence* and *marginalization* (Chapter 3)—harassment directed at entertainment journalists who experience marginalization may target their identity, leading to more severe outcomes. These findings echo the taxonomy of harassment presented in Thomas et al. [210]—our participants experienced *toxic content*, *content leakage* and *overloading*. They were also worried about the possibility of *false reporting* in the form of swatting. These findings suggest a need to account for and support *non-technical* means

of dealing with harassment. As harassment is currently facilitated by affordances of modern platforms, community-based aid may be effective for those who suffer severe harassment.

5.6.1 Contextualizing with prior scholarship on harassment

Thomas et al. [210] categorize harassment attacks based on several features described in Section 2.4. In this work, we see several of the attacks from that framework described by our participants. *Toxic content* was the most salient concern for our participants — our participants faced toxic content ranging from insulting their intelligence and credentials to racial slurs and death threats. This worked hand-in-hand with *overloading*, where the volume and severity of these attacks would sometimes make social media platforms unusable, needing triage. In some cases, our participants also experienced *content leakage*, with home addresses or phone numbers being leaked online with very little recourse for removing them. Our participants also worried about *false reporting*, particularly swatting, even if they had not experienced it themselves.

In work interviewing content creators (defined as “social media personalities with large audiences on platforms like Instagram, TikTok, and YouTube”), Thomas et al. [211] identified many similar patterns of responses to harassment. For content creators, harassment is also characterized as unavoidable and largely takes the form of *toxic content* and *overloading*. Like our participants, many creators opted to censor themselves or stop creating content entirely in response to this, particularly in relationship with marginalized identities. One key difference, however, is that these creators often don’t have the institutional support described by our participants in this study; they must instead create or find community. This again puts personal responsibility on these users to manage harassment.

The “inescapable Gamergate experience” of Twitter described by Trice and Potts [215] is echoed by our theme “The inevitable price of admission.” Our participants described a sense of being stuck on the platform—if they wanted to participate on Twitter (which was viewed as required for their work), they had to endure harassment. Paradigm repair, as described by Perreault and Vos [163], was less salient. Our participants did not seem interested in engaging with harassers, either directly on Twitter or indirectly through posts on their websites. The latter strategy was a core element of the repair observed by Perreault and Vos, but our participants preferred a strategy of non-engagement. Perreault and Vos’s work, however, was conducted closer to the height of #GamerGate; our participants may now see their reputation as more secure and thus needing less direct repair.

5.6.2 Prominence and marginalization: Relating to the at-risk framework

The at-risk framework provided a key lens to understand the contextual needs of these participants. First, the intersection of *prominence* and *marginalization*, observed first in other work as described in Chapter 3, directly led us to working with this particular population; understanding these factors also led us to developing an interview protocol which focused on these factors and their intersection. Prior understanding of the *prominence* and *marginalization* contextual risk factors also informed the thematic development around how these factors impacted the experiences of these journalists. *Prominent* individuals often experience highly focused targeting due to their fame or notoriety, which is a key factor of *overloading* attacks [210]. This was a key feature of our participants’ experiences, all of whom wrote for large public audiences. In addition to our participants’ general public presence, *prominence* could also occur temporarily. An account

might choose to quote tweet a post from a journalist in order to leverage their audience in an *overloading* attack, amplifying that journalist’s words in front of a new audience.

This work also continues to show the impact of *marginalization* on users’ digital-safety needs and strategies. Our participants who were people of color needed to take more care than other participants to avoid hate speech and slurs. In addition, our participants pointed to the outsized impact of harassment on their colleagues who experienced harassment based on marginalized identities; as P06 described in Section 5.4.3, “their voices have been silenced.”

5.6.3 An unequal distribution of labor

Many protective strategies described by our participants emphasized individual action in the face of harassment. From muting hateful words that targeted a participant’s identity (Section 5.4.3) to blocking users who send harassing comments (Section 5.5.1), our participants’ strategies required taking personal responsibility for managing online harassment. Since the attackers in this scenario are anonymous online mobs, this is inherently unbalanced; this echoes findings by Wei et al. [234] where security experts describe the burdens of individuals managing harassment. As this scholarship describes, mandating individual responsibility for dealing with harassment can actually lead to perpetuating and entrenching these patterns by dismissing the societal factors that lead to *marginalization*. Institutional support, both concrete and emotional, can lessen this individual burden, making it a key part of dealing with the problem of harassment.

Distancing behaviors were a particularly common response to intense harassment, aligning with our findings in Chapter 3. Our participants characterized harassment as inevitable and unavoidable; the best option was to simply ignore this broad category of attacks. For security

researchers, this should be troubling. Existing security solutions do not serve the needs of our participants; when prompted about tools and technical solutions, participants were unable to call to mind even *ineffective* technical supports other than blocklists. *Technical solutions* to better support these *distancing behaviors* could also be beneficial; tools that allow users to step away from their accounts and have these accounts triaged on their behalf could be of great use to our participants and other users who face extreme harassment.

Our results also suggest that *collective* action may have great benefit. Support from colleagues and employers made a dramatic difference in our participants' ability to deal with harassment, both practically and emotionally. Further work that enables this support, like HeartMob [29], could be of great benefit. Support organizations like Tall Poppy⁶ and PEN America⁷ exist to address these concerns, but our participants did not mention them in the context of support strategies. Leveraging a community of support could be much more effective than stemming the tide of harassment on one's own.

This work also suggests a need for advocacy on the part of the digital-safety research community. As a parallel example, the Clinic to End Tech Abuse not only provides clinical support to survivors of intimate partner violence [68, 87], but advocates for policy solutions to address the problem of technology abuse from intimate partners.⁸ Social media platforms can enable harassment through their design [127], so public pressure and advocacy should aim to change these affordances so that individuals bear less personal responsibility for this systemic problem.

⁶<https://www.tallpoppy.com/>

⁷<https://pen.org/>

⁸<https://www.ceta.tech.cornell.edu/advocacy>

Chapter 6: Discussion

In this thesis, I (with my collaborators) developed a framework for unifying at-risk user research so that researchers and technology designers need not reinvent the wheel when aiming to address the unique challenges that at-risk users face. I then applied this framework to two populations—library IT staff and entertainment journalists—in order to more effectively understand their concerns and demonstrate the utility of this framework.

The framework presented in Chapter 3 is an essential step toward generalizing the needs and concerns of at-risk users. By using this framework, researchers and technology creators can focus their efforts on addressing *contextual risk factors*, rather than the disparate needs of dozens of at-risk populations.

Using this framework, the results in Chapter 4 show that the *resource and time constraints* that govern both library IT staff and their patrons mean that privacy and utility in fact end up in conflict with one another. In order to protect the privacy of patrons, who themselves may be at-risk, library IT staff take protective actions that diminish the capabilities of their patrons to do important digital tasks. In Chapter 5, we show that *prominence* exposes entertainment journalists to often-severe risks and harms, particularly when they also experience *marginalization*. These journalists often find that in the face of severe harassment, it is generally better to attempt to ignore it and move on than to try to stop it from happening—it is the price of admission into

their industry. This echoes the prevalence of *distancing behaviors* found in our at-risk framework (Chapter 3).

6.1 Reflecting on the at-risk framework

The at-risk framework provided a key analytical lens for the work presented in Chapters 4 and 5. For both studies, we started the analysis already having an understanding of potential *contextual risk factors*—*legal or political* and *resource or time constraints* for library IT staff and *prominence* and *marginalization* for entertainment journalists. This framework provided a deductive frame for the reflexive thematic analysis of both studies by providing a starting point for considering potential challenges and constraints faced by the participants in both studies. Without this framework, we would have needed to rediscover all of these factors from the ground up; instead, our analysis began with rich context that allowed us to focus on the specifics of the participants' experiences.

In addition, the work in Chapter 5 was explicitly inspired by the framework—in Chapter 3, we describe how *prominence* and *marginalization* might interact, but point to the lack of prior research on this interaction. Our study on entertainment journalists starts to address this gap, which was a key motivation of this work; we show that these two factors do interact to create uniquely challenging experiences for users facing harassment online. The at-risk framework is robust, but also temporally bound—other researchers should use the framework and test its limits, as it is based on scholarship from a particular time period. Future work might surface new or unknown *contextual risk factors*, *protective practices*, and *barriers* to digital safety. This can be done, however, without needing to reinvent the wheel when studying new populations; unstudied

at-risk users’ digital-safety needs might be met by tools and techniques already surfaced in work that studies the relevant *contextual risk factors*.

My work in Chapters 4 and 5, however, also exposed an important gap in the framework: *marginalization* is a broad category that still requires specific cultural context to understand. For example, detection of racial slurs might be accomplished through robust dictionaries of existing hate speech and zero-shot hate speech detection, but misgendering transgender people could be *harder* to detect, since this manifests through people simply using incorrect (but not obviously hateful) pronouns. Automatic hate speech detection is an ongoing area of research [67, 147, 166], but current scholarship points at gaps with hate speech that is more subtle [57] or hate speech that is not written in English [153]. Further work should explore the different ways *marginalization* manifests and what shared patterns and factors might be generalized so that scalable solutions for addressing the particular harms of *marginalization* can be developed. Splitting *marginalization* into dozens of separate subcategories would defeat the purpose of the framework, but future work could attempt to explore if different experiences of *marginalization* share any unifying characteristics that could be useful to digital-safety researchers. In addition, future work should broadly consider what tools and techniques might be useful for each *contextual risk factor*; such a mapping could be greatly useful for future technology developers to both create and apply state-of-the-art techniques to the existing needs of at-risk users.

This shows the at-risk framework’s utility in general. It can point researchers to understudied populations and *contextual risk factors*, driving further understanding of current gaps in the at-risk digital safety literature. When conducting these studies, researchers can more comprehensively ask questions that capture *contextual risk factors*, *protective practices*, and *barriers to protective practices*, with confidence that they are not missing an important element of their

participants' experience. Finally, this framework prevents new work from having to start from nothing when understanding at-risk users' experiences. By situating their work in this framework, researchers can bring important context to the specific circumstances of populations they study, and move to the deeper, more specific elements of these at-risk populations' experiences.

6.2 How can technology creators and security researchers use the at-risk framework?

The *contextual risk factors* developed in Chapter 3 can point technology creators and security researchers toward maximally impactful technical solutions. For technology creators, this framework can be used to help justify why it is worth focusing on a particular type of solution, since they can be scaled up. As an example, it might be difficult to justify to a large company why the company should spend its resources helping survivors of intimate partner violence, a relatively small demographic group, if the company wants to focus on helping as many people as possible with its resources. This can be frustrating for advocates for survivors, given their deep experience and understanding that these users face extremely severe challenges and harms. Instead, with this framework, one might propose focusing resources on helping protect people with a *relationship with their attacker*, which could encompass not just survivors of IPV, but also women in South Asia, foster teenagers with abusive parents, and survivors of human trafficking, as well as future populations with the same *contextual risk factor*.

The *contextual risk factors* may also point to solutions that can be developed by technologists and researchers. Artificial intelligence and machine learning researchers could continue to refine zero-shot approaches for automatic detection of hate speech to reduce the burden of filter-

ing described by our *marginalized* participants in Chapter 5. Similarly, zero-shot detection of a harassment event could be especially valuable for *prominent* users, or even those who *become* prominent due to amplification by an attacker with a large audience. This is a real challenge—due to the affordances of social media, it could be quite difficult to distinguish between a social media post that goes viral for positive reasons or for negative ones. However, as this sudden rise to *prominence* exposes users to additional attacks and harms, quick detection would be extraordinarily valuable for quick triaging and management of harassment events, especially with regards to reducing individual burden on the user for dealing with these events.

Resource and time constraints also play an enormous role in how users address their digital-safety needs. The dataset in Chapter 3 began to point at this challenge, and this finding was reinforced by our study with library IT staff in Chapter 4. Technology creators might consider how to better address secure, private methods for low-resourced users to do persistent work in public environments like libraries.

The *protective practices* described in Chapter 3 can also direct research and technology development towards at-risk users' priorities. Currently, at-risk users rely heavily on *distancing behaviors* and *social strategies*, in addition to *technical solutions*. Future work should support all three pillars of at-risk users' responses to the particularly digital-safety problems they face. As described in Chapter 5, enabling users to step away from platforms where they experience harassment (a *distancing behavior*) may in fact be quite helpful—technology should enable this through providing automatic triage for these accounts while users take the time that they need. Current social media platforms are usually designed to keep users on them as much as possible, but providing structured ways to step away from the platform, perhaps for a certain amount of time, could be of use to people experiencing harassment. This would be particularly useful when

paired with automatic detection of high harassment events; one could step away until the platform notifies you that the wave of harassment is over and return to repair and reset their social media experience. In addition, formal ways to allow trusted colleagues and friends to manage one's social media accounts would support the existing triage strategy described in Chapter 5. Instead of needing to provide a password and access to all past account history, platforms might be able to provide a means of temporary, limited authentication that limits the ability of the helper to snoop on the target, even accidentally.

6.3 Understanding context is crucial

This thesis clearly demonstrates that an understanding of the contexts that at-risk users experience is crucial to begin to address their digital-safety challenges. The framework presented in Chapter 3 reduces the task from “consider and understand all at-risk populations” to “consider and understand ten *contextual risk factors*.” Further work will help strengthen this framework, by comparing it to current and future at-risk populations to ensure that everyone's needs are considered when developing new technologies and systems.

In practice, this can be seen in the results from Chapters 4 and 5. Chapter 4 shows that the particular constraints and values of library IT staff force them to choose between privacy—a strong library value—and utility due to the *legal and political* environment in which they operate, favoring strong privacy guarantees over other priorities. These decisions are not made for no reason, but rather to deal with strong external threats. Future work could focus on developing tools to increase the utility of library resources for patrons—like permitting persistent work—while maintaining the strong privacy needed by our participants.

This thesis shows the importance of accounting for the specific threat models at-risk users face, especially in relation to assumptions normally made about the “average user.” Entertainment journalists who experience harassment do not have a profit-motivated adversary, but rather face a diffuse group of loosely-organized attackers with the aim of causing emotional, psychological, or even physical harm. Harassment attacks do not require access to an account or privileged information, but are made possible by the affordances of modern social media platforms. Future work should investigate ways to help users *prevent* harassment proactively, rather than responding to it after the fact.

Solutions in this space must not be burdensome—at-risk users already face the difficult challenge of dealing with elevated risk. By understanding their context, we can develop tools and support that can be successfully woven into how at-risk users already operate online.

6.4 Non-technical support is just as important as new tools

A key finding in Chapter 3 was that at-risk users often employ non-technical solutions to solve digital-safety problems, spanning a range of *social strategies* and *distancing behaviors*. These solutions were chosen because they were perceived as effective—at least, as effective as the users could manage with their available resources. While a prevalence of *distancing behaviors* is troubling, *social strategies* and other non-technical solutions are a promising avenue for support for at-risk users.

Library IT staff can be supported by their library systems and library organizations like the ALA. The resources and legal support from these organizations should work in concert with the digital-safety goals of these staff members. Library IT staff need resources like money, time,

training, and staff to do their jobs effectively, and library systems should support them in getting those resources. The work presented in Chapter 4 provides a first step in making clear why devoting resources to digital safety in libraries is crucial. Continued advocacy is needed to ensure that library IT staff are supported in their goals.

In Chapter 5, we show that community support is one of the most important elements of entertainment journalists' strategies for dealing with the persistent impact of online harassment. Emotional support from colleagues was critical to managing harassment from media journals; this is the primary way our participants dealt with the *emotional* toll of suffering these attacks. This finding suggests that community-based solutions should continue to be explored in the context of digital safety. These solutions might take the form of online support networks which allow people to share experiences and strategies, like HeartMob [29]. The simple fact of shared experience, noted by our participants, allows users to process and understand what's happening to them, as well as share workable strategies for dealing with their particular contextual needs. Further work should consider the power of community connections to address digital safety—rather than have researchers hop in and assume expertise, allowing communities of at-risk users to work together to find solutions could be both empowering and effective.

Organizations could also provide formal, practical support, as has been attempted in the past by projects like the now-defunct Crash Override.¹ For investigative journalists, organizations like PEN America² can already provide similar support, and some companies like Tall Poppy³ provide harassment prevention as a service. However, none of our participants in Chapter 5 mentioned organizations beyond their own employers, suggesting a lack of reach or accessibility

¹<http://www.crashoverridenetwork.com/>

²<https://pen.org/>

³<https://www.tallpoppy.com/>

for these services. Further development of *free, open-access* support is crucial, so as to not add financial burden on top of an already-challenging threat, particularly for at-risk users who are not already supported by a larger organization. Community support *in parallel* with continued development of privacy-enhancing technologies will be the most effective way to address the multifaceted challenges at-risk users face.

6.5 Digital safety cannot be isolated from social conditions

A central theme of this thesis is that many of the factors that impact at-risk users are larger than the digital-safety research community. Library IT staff face constraints due to external pressures from the government and large corporations; these challenges are often larger than technical solutions. As our participants in Chapter 4 pointed out, the control that third-party vendors have over their resources is completely outside of their influence and control. Open-access alternatives which can provide ebooks and academic journal articles would be very valuable, but this is a problem of marketplaces and copyright, rather than technical tools. The marginalization that some entertainment journalists experience cannot be solved by a new digital-safety tool; marginalization is a societal problem that has broad-reaching impact, *including* digital safety.

However, this thesis does not argue that these problems are impossible to solve. Security researchers can and should act as advocates in addition to performing their research. As an example, the Clinic to End Tech Abuse regularly engages in advocacy⁴ that is informed by their research and clinical practice, advocating for policy to help survivors of intimate partner violence. This approach, which combines high-quality research with effective advocacy, should be applied to the social conditions faced by other at-risk users. Existing technical solutions also need

⁴<https://www.ceta.tech.cornell.edu/advocacy>

to be practically incorporated in the environments in which at-risk users operate. Further research on zero-shot hate speech detection, for example, will be of no benefit unless platforms like Facebook or Twitter actually employ it; if this does not happen, people like our participants will still have to manually input hateful words into their list of muted terms. Digital-safety researchers should additionally investigate the barriers to implementing these solutions, both on a technical and human-centered level; in order to implement them, we must understand who and what is currently preventing tools like these from being implemented in the places where users need them most. By combining research and advocacy, using each to inform the other, we can create a safer and more equitable world for all users.

Appendix A: At-risk framework paper list

Title	Venue	Year	Authors
Privacy concerns and behaviors of people with visual impairments	CHI	2015	Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia
Security practices for households bank customers in the Kingdom of Saudi Arabia	SOUPS	2015	Deena Alghamdi, Ivan Flechais, and Marina Jirotko
Investigating the computer security practices and needs of journalists	USENIX	2015	Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner
How I learned to be secure: A census-representative survey of security advice sources and behavior	CCS	2016	Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek
Dear diary: teens reflect on their weekly online risk experiences	CHI	2016	Pamela Wisniewski, Heng Xu, Mary Beth Rosson, Daniel F. Perkins, and John M. Carroll
ICT use by prominent activists in Republika Srpska	CHI	2016	Borislav Tadic, Markus Rohde, Volker Wulf, and David Randall
LGBT parents and social media: Advocacy, privacy, and disclosure during shifting social movements	CHI	2016	Lindsay Blackwell, Jean Hardy, Tawfiq Ammari, Tiffany Veinot, Cliff Lampe, and Sarita Schoenebeck
Mediating the undercurrents: Using social media to sustain a social movement	CHI	2016	Yong Ming Kow, Yubo Kou, Bryan Seaman, and Waikuen Cheng
Understanding social media disclosures of sexual abuse through the lenses of support seeking and anonymity	CHI	2016	Nazanin Andalibi, Oliver L. Haimson, Munmun De Choudhury, and Andrea Forte

Vulnerability, sharing, and privacy: Analyzing art therapy for older adults with dementia	CSCW	2016	Raymundo Cornejo, Robin Brewer, Caroline Edasis, and Anne Marie Piper
Individual versus organizational computer security and privacy concerns in journalism	PoPETS	2016	Susan E. McGregor, Franziska Roesner, and Kelly Caine
Straight talk: New Yorkers on mobile messaging and implications for privacy	Simply Secure Report	2016	Ame Elliot
Addressing physical safety, security, and privacy for people with visual impairments	SOUPS	2016	Tousif Ahmed, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia
Intuitions, analytics, and killing ants: Inference literacy of high school-educated adults in the us	SOUPS	2016	Jeffrey Warshaw, Nina Taft, and Allison Woodruff
Navigating relationships and boundaries: Concerns around ICT-uptake for elderly people	CHI	2017	Dominik Hornung, Claudia Müller, Irina Shklovski, Timo Jakobi, and Volker Wulf
Parents' and children's preferences about parents sharing about children on social media	CHI	2017	Carol Moser, Tianying Chen, and Sarita Y. Schoenebeck
Privacy considerations when designing social network systems to support successful ageing	CHI	2017	Andrew R. McNeill, Lynne Coventry, Jake Pywell, and Pam Briggs
Privacy, security, and surveillance in the Global South: A study of biometric mobile SIM registration in Bangladesh	CHI	2017	Syed Ishtiaque Ahmed, Md. Romael Haque, Shion Guha, Md. Rashidujjaman Rifat, and Nicola Dell
Stories from survivors: Privacy & security practices when coping with intimate partner abuse	CHI	2017	Tara Matthews, Kathleen O'Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F. Churchill, and Sunny Consolvo
Toys that listen: A study of parents, children, and internet-connected toys	CHI	2017	Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner
Where is the digital divide?: A survey of security, privacy, and socioeconomics	CHI	2017	Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek

“Our privacy needs to be protected at all costs”: Crowd workers’ privacy experiences on amazon mechanical turk	CSCW	2017	Huichuan Xia, Yang Wang, Yun Huang, and Anuj Shah
Classification and its consequences for online harassment: Design insights from HeartMob	CSCW	2017	Lindsay Blackwell, Jill Dimond, Sarita Schoenebeck, and Cliff Lampe
Digital privacy challenges with shared mobile phone use in Bangladesh	CSCW	2017	Syed Ishtiaque Ahmed, Md. Romael Haque, Jay Chen, and Nicola Dell
Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders	CSCW	2017	Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell
Filtered out: Disability disclosure practices in online dating communities	CSCW	2017	John R. Porter, Kiley Sobel, Sarah E. Fox, Cynthia L. Bennett, and Julie A. Kientz
Interpretations of online anonymity in Alcoholics Anonymous and Narcotics Anonymous	CSCW	2017	Sabirat Rubya and Svetlana Yarosh
Most teens bounce back: Using diary methods to examine how quickly teens recover from episodic online risk exposure	CSCW	2017	Bridget Christine McHugh, Pamela J. Wisniewski, Mary Beth Rosson, Heng Xu, and John M. Carroll
“No telling passcodes out because they’re private”: Understanding children’s mental models of privacy and security online	CSCW	2017	Priya Kumar, Shalmali Milind Naik, Utkarsha Ramesh Devkar, Marshini Chetty, Tamara L. Clegg, and Jessica Vitak
On making data actionable: How activists use imperfect data to foster social change for human rights violations in mexico	CSCW	2017	Adriana Alvarado Garcia, Alyson L. Young, and Lynn Dombrowski
Parental control vs. teen self-regulation: Is there a middle ground for mobile online safety?	CSCW	2017	Pamela Wisniewski, Arup Kumar Ghosh, Heng Xu, Mary Beth Rosson, and John M. Carroll
Parents just don’t understand: Why teens don’t talk to parents about their online risk experiences	CSCW	2017	Pamela Wisniewski, Heng Xu, Mary Beth Rosson, and John M. Carroll
Privacy leakage in event-based social networks: A Meetup case study	CSCW	2017	Taejoong Chung, Jinyoung Han, Daejin Choi, Ted Taekyoung Kwon, Jong-Youn Rha, and Hyunchul Kim

Would you Slack that?: The impact of security and privacy on cooperative newsroom work	CSCW	2017	Susan E. McGregor, Elizabeth Anne Watkins, and Kelly Caine
Social engineering attacks on government opponents: Target perspectives	PoPETS	2017	William R. Marczak and Vern Paxson
How effective is anti-phishing training for children?	SOUPS	2017	Elmer Lastdrager, Inés Carvajal Gallardo, Pieter Hartel, and Marianne Junger
Learning system-assigned passwords: a preliminary study on the people with learning disabilities	SOUPS	2017	Sonali Tukaram Marne, Mahdi Nasrullah Al-Ameen, and Matthew Wright
When the weakest link is strong: Secure collaboration in the case of the Panama Papers	USENIX	2017	Susan E. McGregor, Elizabeth Anne Watkins, Mahdi Nasrullah Al-Ameen, Kelly Caine, and Franziska Roesner
“A stalker’s paradise”: How intimate partner abusers exploit technology	CHI	2018	Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell
A matter of control or safety?: Examining parental use of technical monitoring apps on teens’ mobile devices	CHI	2018	Arup Kumar Ghosh, Karla Badillo-Urquiola, Mary Beth Rosson, Heng Xu, John M. Carroll, and Pamela J. Wisniewski
Co-designing mobile online safety applications with children	CHI	2018	Brenna McNally, Priya Kumar, Chelsea Hordatt, Matthew Louis Mauriello, Shalmali Naik, Leyla Norooz, Alazandra Shorter, Evan Golub, and Allison Druin
Gender recognition or gender reductionism?: The social implications of embedded gender recognition systems	CHI	2018	Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham
Keeping a low profile?: Technology, risk and privacy among undocumented immigrants	CHI	2018	Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H. Mhaidli, Kentaro Toyama, and Florian Schaub
Safety vs. surveillance: What children have to say about mobile apps for parental control	CHI	2018	Arup Kumar Ghosh, Karla Badillo-Urquiola, Shion Guha, Joseph J. LaViola Jr, and Pamela J. Wisniewski
The use of private mobile phones at war: Accounts from the Donbas Conflict	CHI	2018	Irina Shklovski and Volker Wulf

Examining security and privacy research in developing regions	COMPASS	2018	Aditya Vashistha, Richard Anderson, and Shrirang Mare
“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	CSCW	2018	Jessica Vitak, Yuting Liao, Mega Subramaniam, and Priya Kumar
Participatory design of technologies to support recovery from substance use disorders	CSCW	2018	Zachary Schmitt and Svetlana Yarosh
Privacy unraveling around explicit HIV status disclosure fields in the on-line geosocial hookup app Grindr	CSCW	2018	Mark Warner, Andreas Gutmann, M. Angela Sasse, and Ann Blandford
Safe spaces and safe places: Unpacking technology-mediated experiences of safety and harm with transgender people	CSCW	2018	Morgan Klaus Scheuerman, Stacy M. Branham, and Foad Hamidi
Computer security and privacy for refugees in the united states	IEEE	2018	Lucy Simko, Ada Lerner, Samia Ibtasam, Franziska Roesner, and Tadayoshi Kohno
On enforcing the digital immunity of a large humanitarian organization	IEEE	2018	Stevens Le Blond, Alejandro Cuevas, Juan Ramón Troncoso-Pastoriza, Philipp Jovanovic, Bryan Ford, and Jean-Pierre Hubaux
The spyware used in intimate partner violence	IEEE	2018	Rahul Chatterjee, Periwinkle Doerfler, Hadas Orgad, Sam Havron, Jackeline Palmer, Diana Freed, Karen Levy, Nicola Dell, Damon McCoy, and Thomas Ristenpart
“Privacy is not for me, it’s for those rich women”: Performative privacy practices on mobile phones by women in South Asia	SOUPS	2018	Nithya Sambasivan, Garen Checkley, Amna Batool, Nova Ahmed, David Nemer, Laura Sanely Gaytán-Lugo, Tara Matthews, Sunny Consolvo, and Elizabeth Churchill
Cameras on beds: The ethics of surveillance in nursing home rooms	AJOB Empirical Bioethics Volume 10	2019	Clara Berridge, Jodi Halpern, and Karen Levy

“Disadvantaged in the American-dominated internet”: Sex, work, and technology	CHI	2019	Catherine Barwulor, Allison McDonald, Eszter Hargittai, and Elissa M. Redmiles
EarTouch: Facilitating smartphone use for visually impaired people in mobile and public scenarios	CHI	2019	Ruolin Wang, Chun Yu, Xing-Dong Yang, Weijie He, and Yuanchun Shi
“Everyone has some personal stuff”: Designing to support digital privacy with shared mobile phone use in bangladesh	CHI	2019	Syed Ishtiaque Ahmed, Md. Romael Haque, Irtaza Haider, Jay Chen, and Nicola Dell
“I make up a silly name”: Understanding children’s perception of privacy risks online	CHI	2019	Jun Zhao, Ge Wang, Carys Dally, Petr Slovak, Julian Edbrooke-Childs, Max Van Kleek, and Nigel Shadbolt
“If it’s important it will be a headline”: Cybersecurity information seeking in older adults	CHI	2019	James Nicholson, Lynne Coventry, and Pamela Briggs
Privacy and security considerations for digital technology use in elementary schools	CHI	2019	Priya C. Kumar, Marshini Chetty, Tamara L. Clegg, and Jessica Vitak
Risk vs. restriction: The tension between providing a sense of normalcy and keeping foster teens safe online	CHI	2019	Karla Badillo-Urquiola, Xinru Page, and Pamela J. Wisniewski
Spaces and traces: Implications of smart technology in public housing	CHI	2019	Sandjar Kozubaev, Fernando Rochaix, Carl DiSalvo, and Christopher A. Le Dantec
Technologies for social justice: Lessons from sex workers on the front lines	CHI	2019	Angelika Strohmayer, Jenn Clamen, and Mary Laing
“They don’t leave us alone anywhere we go”: Gender and digital abuse in south asia	CHI	2019	Nithya Sambasivan, Amna Batool, Nova Ahmed, Tara Matthews, Kurt Thomas, Laura Sanely Gaytán-Lugo, David Nemer, Elie Bursztein, Elizabeth Churchill, and Sunny Consolvo
Threats, abuses, flirting, and blackmail: Gender inequity in social media voice forums	CHI	2019	Aditya Vashistha, Abhinav Garg, Richard Anderson, and Agha Ali Raza
Tough times at transitional homeless shelters: Considering the impact of financial insecurity on digital security and privacy	CHI	2019	Manya Sleeper, Tara Matthews, Kathleen O’Leary, Anna Turner, Jill Palzkill Woelfer, Martin Shelton, Andrew Oplinger, Andreas Schou, and Sunny Consolvo

Upside and downside risk in online security for older adults with mild cognitive impairment	CHI	2019	Helena Mentis, Galina Madjaroff, and Aaron Massey
“Is my phone hacked?”: Analyzing clinical computer security interventions with survivors of intimate partner violence	CSCW	2019	Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell
“Privacy is not a concept, but a way of dealing with life”: Localization of transnational technology platforms and liminal privacy practices in Cambodia	CSCW	2019	Margaret C. Jack, Pang Sovannaroth, and Nicola Dell
Towards digitization of collaborative savings among low-income groups	CSCW	2019	Hamid Mehmood, Tallal Ahmad, Lubna Razaq, Shrirang Mare, Maryem Zafar Usmani, Richard Anderson, and Agha Ali Raza
Cooperative privacy and security: Learning from people with visual impairments and their allies	SOUPS	2019	Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang
Privacy and security threat models and mitigation strategies of older adults	SOUPS	2019	Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman
Clinical computer security for victims of intimate partner violence	USENIX	2019	Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart
Computer security and privacy in the interactions between victim service providers and human trafficking survivors	USENIX	2019	Christine Chen, Nicola Dell, and Franziska Roesner
Being (in)visible: Privacy, transparency, and disclosure in the self-management of bipolar disorder	CHI	2020	Justin Petelka, Lucy Van Kleunen, Liam Albright, Elizabeth Murnane, Stephen Volda, and Jaime Snyder
Circle of trust: A new approach to mobile online safety for families	CHI	2020	Arup Kumar Ghosh, Charles E. Hughes, and Pamela J. Wisniewski
Civic empowerment through digitalisation: The case of Greenlandic women	CHI	2020	Nicola Wendt, Rikke Bjerg Jensen, and Lizzie Coles-Kemp
Evaluating ‘prefer not to say’ around sensitive disclosures	CHI	2020	Mark Warner, Agnieszka Kitkowska, Jo Gibbs, Juan F. Maestre, and Ann Blandford

Fragile masculinity: men, gender, and online harassment	CHI	2020	Jennifer D. Rubin, Lindsay Blackwell, and Terri D. Conley
‘Lime’, ‘open lock’, and ‘blocked’: children’s perception of colors, symbols, and words in cybersecurity warnings	CHI	2020	Rebecca Jeong and Sonia Chiasson
Privacy and activism in the transgender community	CHI	2020	Ada Lerner, Helen Yuxun He, Anna Kawakami, Silvia Catherine Zeamer, and Roberto Hoyle
Under surveillance: technology practices of those monitored by the state	CHI	2020	Pedro Sanches, Vasiliki Tsaknaki, Asreen Rostami, and Barry Brown
Understanding cybersecurity practices in emergency departments	CHI	2020	Elizabeth Stobert, David Barrera, Valérie Homier, and Daniel Kollek
“They just don’t get it”: Towards social technologies for coping with interpersonal racism	CSCW	2020	Alexandra To, Wenxia Sweeney, Jessica Hammer, and Geoff Kaufman
Privacy threats in intimate relationships	Journal of Cybersecurity	2020	Karen Levy and Bruce Schneier
Blind and human: Exploring more usable audio captcha designs	SOUPS	2020	Valerie Fanelle, Sepideh Karimi, Aditi Shah, Bharath Subramanian, and Sauvik Das
Lessons learnt from comparing WhatsApp privacy concerns across Saudi and Indian populations	SOUPS	2020	Jayati Dev, Pablo Moriano, L. Jean Camp
Realizing choice: Online safeguards for couples adapting to cognitive challenges	SOUPS	2020	Nora McDonald, Alison Larsen, Allison Battisti, Galina Madjaroff, Aaron Massey, and Helena Mentis
Towards understanding privacy and trust in online reporting of sexual assault	SOUPS	2020	Borke Obada-Obieh, Lucrezia Spagnolo, and Konstantin Beznosov
DatashareNetwork: A decentralized privacy-preserving search engine for investigative journalists	USENIX	2020	Kasra Edalatnejad, Wouter Lueks, Julien Pierre Martin, Soline Ledésert, Anne L’Hôte, Bruno Thomas, Laurent Girod, and Carmela Troncoso
“I am uncomfortable sharing what I can’t see”: Privacy concerns of the visually impaired with camera based assistive applications	USENIX	2020	Taslina Akter, Bryan Dosono, Tousif Ahmed, Apu Kapadia, and Bryan Seaman

“I have too much respect for my elders”: Understanding South African mobile users’ perceptions of privacy and current behaviors on Facebook and WhatsApp	USENIX	2020	Jake Reichel, Fleming Peck, Mikako Inaba, Bisrat Moges, Brahmnoor Singh Chawla, and Marshini Chetty
Defensive technology use by political activists during the Sudanese Revolution.	IEEE S&P	2021	Alaa Daffalla, Lucy Simko, Tadayoshi Kohno, and Alexandru G. Bardas
“Why wouldn’t someone think of democracy as a target?”: Security practices & challenges of people involved with U.S. political campaigns	USENIX	2021	Sunny Consolvo, Patrick Gage Kelley, Tara Matthews, Kurt Thomas, Lee Dunn, and Elie Bursztein

Table A.1: A list of the 95 papers included in the dataset for the analysis in Chapter 3

Appendix B: Library IT staff protocols

B.1 Intake survey

1. What is your name?
2. What is your email address?
3. What is the name of the library where you are currently employed?
4. How many branches are there?
5. Are you responsible for IT at one or multiple branches?
6. Roughly how many patrons does the library system serve per month?

- (a) Do you collect any information about public terminal usage? If so, what usage statistics do you have? (Free text response)

7. If applicable, roughly how many patrons does your branch serve per month?

(a) Do you collect any information about public terminal usage? If so, what usage statistics do you have? (Free text response)

8. How long have you been doing IT?

9. How long have you been doing IT in libraries, specifically?

10. How long have you been doing IT at this particular library?

11. Have you had other roles at this library besides IT?

12. How many people are employed by the library system specifically to work on the library's IT system?

13. If applicable, how many people are employed by this particular branch specifically to work on the library's IT system?

14. How many of each type of computing device does your library oversee? Please specify for the branch and for the overall library system, and indicate for any that you don't know.

Devices	# at your branch (If applicable)	# at your system
Workstation computers or kiosk	_____ [don't know]	_____ [don't know]
Laptops	_____ [don't know]	_____ [don't know]
Tablets	_____ [don't know]	_____ [don't know]

15. Do you contract your IT administration to an outside company?

(a) If so, how long have you been using a contractor?

16. What operating systems and versions does your library use for its public computers?
17. Have you heard of session management tools like DeepFreeze, Windows SteadyState or some other session management tool? **Yes/No**
18. Does your library branch currently use a session management tool? **Yes/No**
19. We are interested in hearing about some of the vulnerable populations that you serve. Please indicate below what types of patron populations come to your specific library **branch**? There may be overlaps among communities or categories of identity. *Participants checked all that applied.*

- Undocumented immigrants
- Individuals without permanent residence
- Individuals with criminal history
- Low income communities
- Low income, single mothers
- LGBTQ2S+ (Lesbian, Gay, Bisexual, Transgender, Queer or Questioning and Two-Spirit)
- Racial/ethnic minorities (including Latinx, Black/African American, Native American, mixed, non-white)
- Journalists
- Political activists
- Older adults

- People with disabilities (e.g., learning, motor, vision, etc)

B.2 Interview protocol

Questions very similar, if not identical, to those listed below were asked of participants. Some follow-up questions, not listed, were asked depending on the conversation with the participants. Not all participants were necessarily asked all questions.

Introduction. Hello. My name is [INSERT NAME]. I am part of a multi-institutional group of researchers studying privacy and security IT administration in libraries. During this interview, we would like to discuss your library's policies and practices relating to the protection of public computers used by library patrons. Your identity, your library's identity, and the fact of your participation in this study will be held in confidence. If the results of this research are published, we will not include any identifying details in the publications. Your participation is voluntary, and you may decline to answer any questions or stop the interview at any time.

I'll call your attention back to the consent form, which you previously electronically agreed to. Do you have any questions about it, before we begin?

Finally, we would like to record these interviews for our analysis. Results of our research will be reported in aggregate with other libraries. Is it okay if I record this interview?

[If they agree, the interview was recorded. If not, the interviewer took notes.]

Do you have any final questions before we start?

Background Information. Thank you for filling out the pre-screener survey. We'll use your

responses to that to guide our discussion.

1. Can you briefly describe your role at [Insert Library Name]?
 - (a) What do you think about your role?
 - (b) How do you feel about what you're doing?
 - (c) (If applicable) You mentioned you had a previous role at this library. Could you tell me about that?
2. Can you briefly describe the structure of your library staffing? To whom do you report?
3. (If they responded that they use a contractor) I saw that on your survey you indicated that your library contracts out some or all of your IT services to a contractor.
 - (a) Can you detail which aspects you use a contractor for?
 - (b) Why does your library use a contractor for these services?
 - (c) What is the division of responsibility between your role and the contractor's?
4. (If applicable) For the remainder of the interview, I would like you to delineate between what the library does and what the contractor does.
5. Can you describe how your library has been impacted by COVID-19?

For the rest of this interview, I'll be asking you about how you operate and make decisions under normal operating times.
6. Can you describe your patron population and how they use/interact with your library computers? (Probe about populations not mentioned but checked in survey)

7. Can you walk me through the steps that a patron coming into the library needs to take in order to access one of your public computers?

(Listen for these questions)

- (a) How do you manage which patrons use which terminals and when?
- (b) Are the time or frequency limits set on the use of computers?
- (c) Do patrons log in and out of the terminals?
- (d) What happens upon login and logout?
- (e) What, if anything, gets wiped?
- (f) What happens upon reboot?
- (g) Do they have to agree to some terms (either the first time they use a terminal, or each time they log on)?
- (h) A user agreement of some sort?
- (i) Do these steps differ for patrons when using tablets or laptops? Explain.
- (j) Do these steps differ for any of the patron populations you mentioned earlier? (If so) How?

8. How would a patron access the WiFi?

(Listen for these questions)

- (a) Is the WiFi secured? How often does the WiFi password change?
- (b) Do you contract the WiFi administration/security out to a company?

9. What sorts of concerns do your patrons have regarding their **security** when they are engaging with library IT infrastructure?

(a) What types of patrons have this concern? (Listen for pop's mentioned earlier/in survey)

(b) What do you think of this concern?

(c) How do you address this concern?

(d) Can you think of any other security concerns that they have expressed to you?

10. What concerns do you hear from your patrons about their **privacy**?

(a) What types of patrons have this concern? (Listen for pop's mentioned earlier/in survey)

(b) What do you think of this concern?

(c) How do you address this concern?

(d) Can you think of any other privacy concerns that they have expressed to you?

11. Do you have any programs or initiatives, either formal or informal, that provide some education to patrons about their cyber security or online privacy?

(a) Could you describe them?

(b) What types of patrons use these programs/initiatives?

(c) Are there any issues with regard to privacy and security you feel you can't address in this format and wish you were able to address?

12. In general, are there any privacy or security concerns expressed by your patrons that you are unable to address? Explain.

Challenges and Mitigations.

1. As an IT professional in a library setting, what do you think are the different security and privacy challenges or threats you must overcome? Which of these are most pressing?

(Focus on just 5 for remaining questions in this section)

For each of these threats ask:

(a) What are the different ways that your library tries to address this challenge?

(b) Do you feel like you're able to fully address this challenge?

i. What additional things do you wish you could do?

(c) Are these threats specific to your branch or to the library system broadly? Explain.

(d) Are there certain patrons who may be more impacted by these threats? Explain.

(Talk with the participant about these challenges. Note if the participant brings up any of the topics in the Miscellaneous section and ask follow ups as appropriate.)

Policies. Now I was hoping to learn more about your specific library/system and how policies and practices are decided upon.

1. Can you explain to me what policies your library has concerning computers?
2. Does your specific library have a policy relating to the administration and oversight of the public computers in your library?

3. Does it have sections that pertain to the **protection of user data** while using them?
4. Does it have sections that pertain to the **physical security** of the computer?
5. Specifically, can you detail how the policies were developed for the public computers?
 - (a) Who was involved in the development of these policies?
 - (b) (If not already clear) Were you involved in the development of these policies? What was your role?
 - (c) Can you describe a recent policy decision that was made regarding the public computers? (Probe: What happened? What was the process? Who was involved?)
 - (d) Is that typical?
6. (If multiple branches) How do the policies that you have described apply to other branches in your system?

Other Follow-Ups.

1. Could you describe to me the layout of the public computers in your library?
2. What choices went into the decision of the physical location of the computers in the library space?

(Don't necessarily prompt them about the concerns below, but ask follow ups if they do)

 - (a) Are computers in a separate room/space? why?
 - (b) Are there separate computers for children and adults? Are they treated differently? why?

- (c) Are computer screens visible from a librarian's desk? Why?
 - (d) Are computer screens visible to someone walking by on the street or a public facing window? Why?
 - (e) Are computer screens visible to security cameras? Why?
 - (f) Are there dividers between computers? Why?
3. Are the terminals monitored for hardware abuse?
- (a) What type of abuse are you worried about?
 - (b) What do you worry about the most?
 - (c) Which are you able to address? Why? How?
 - (d) What are you not able to address? Why?
 - (e) Do you perform physical sweeps of the computers? Why/why not? How do you do them?
4. Are the terminals monitored for software abuse? If so: How frequently do sweeps occur?
- (a) How do you do them?
 - (b) What types of software abuse are you looking for?
 - (c) Can you describe to me an instance when you have discovered abuse? What happened? That's a great example, any others?
 - (d) (Reminder to clarify if IT or contractor administers for this and others as appropriate.)
5. Does the library filter or firewall internet content on any of its computers? If so:

- (a) What types of content?
 - (b) Where are the logs stored? Why? Who has access to these logs?
6. I see that you use (operating system). Why do you use this?
- (a) What were the factors in your library's decision to adopt this operating system?
 - (b) How long have you had it? (If recent:) Why prompted your library to change?
7. What software is on your computers?
- (a) Why do you use this?
 - (b) What were the factors in your library's decision to adopt this software/OS?
8. How long have you been using it? (If recent:) Why prompted your library to change?
9. How do you configure the software?
- (a) Is there a set of guidelines or best practices you follow? Why those guidelines/best practices?
 - (b) Do you have any documents or websites that open on application start (e.g., Firefox homepage)?
 - (c) What do you do with Microsoft Office? Why?
 - (d) What do you do with Adobe Acrobat? Why?
 - (e) Are there any other software that have special configurations? Explain.
10. What is your default browser and homepage?
- (a) Is it over HTTP or HTTPS?

(b) Does the browser have any plugins or addons installed?

(c) Does the browser use Tor, or have the ability to?

11. How do you update software, and how often?

(a) Why do they update at those intervals? What is the perceived benefit?

(b) Do you distinguish in your process between security updates and other updates? Is there a fast track?

12. I see that you (use/don't use) a session management software like SteadyState or Deep-Freeze. Why have you made that decision?

(a) Which software do you use?

(b) How do you handle protecting session data?

13. How do you ensure multiple computers all get the same configuration?

14. How often do print or scan jobs stay on your printing or scanning devices?

(a) Do you have access control mechanisms set up on these devices?

15. Are there any people or populations you feel like you are unable to serve with your IT services?

16. In light of COVID-19, are you or your library planning on making any changes going forward? How has covid changed any of your practices or protocols?

Wrap Up. Thank you for answering my questions/taking time to speak with me today.

1. Is there anything else you would like to add to the conversation?
2. Is there anyone else in IT that you think it would be important for me to speak with, either at your library or at another branch?

Thank you so much for your time by taking this interview. We would like to offer you \$35 as a thank you for your time and participation. We can provide this to you as an online gift card. Sometimes institutions have rules around receiving gifts. If that applies to you, we can also make a donation in your name. Would you like the gift card or the donation?

Appendix C: Entertainment journalist protocol

Introduction. Hi, my name is [researcher name], thanks for agreeing to participate in this research - we really appreciate your time.

First, let's quickly go over how this study will work. I will be interviewing you and [I/my colleague] will take notes. I expect the study to take approximately one hour. . One thing I'd like to mention is that the research interview process is somewhat different than the journalistic interview process - rather than seeking pull quotes or particularly interesting stories, we are instead looking for common themes between an entire corpus of interviews, even if those themes seem at first to be mundane.

[Describe everything on the consent form.]

We may cover some sensitive topics during this interview, so if you become uncomfortable at any time during the study and wish to withdraw, please let me know. You are also welcome to

skip any questions you do not wish to answer, and you only need to provide as much information as you're comfortable with. Do you have any questions so far?

[Give the participant the link to the consent form.]

This consent form tells you who to contact if you have any problems or want to report any objections. Please print a copy of this consent form for your records.

[point out places the subject needs to mark checkboxes]

We would like to record the audio of this interview with your permission in order to properly represent your statements and point of view. However, recording is optional - if, either now or after the fact, you would like us to not use or delete the recording of this interview, please let me know. I'm also happy to answer questions about how we store and use these recordings. We will also be taking written notes during the interview. Do you give permission for us to audio-record this interview? (If yes, the interview)

[If they agree, the interview was recorded. If not, the interviewer took notes.]

Warmup/Career summary.

1. Can you please describe your career in media journalism?
2. How did you get started?
3. When did you get started?
4. What outlets have you worked at over the course of your career?
5. Do you have a particular specialty, like esports, a particular media property, interviews with creators, opinion pieces, or something else?
6. Can you tell me about something interesting you've been working on recently?

7. Could you tell us a little about your online engagement with readers/viewers?
8. What social media platforms do you use today? About how many followers on each platform did you have? How many readers or viewers do you usually reach?

Questions about Specific Risk Events. In the next part of the study, we are going to ask you about your experiences with harassment. We emphasize that we do not view any harassment as justified, but we also acknowledge that sometimes the amount of harassment one experiences might vary at different times or when writing about certain subjects. For the purposes of this study, we are going to define a “high harassment event” as a short period of time with an unusually high quantity or intensity of harassment.

1. Please describe your experience of online harassment on a day-to-day basis.
2. In the past two years, have you experienced any high harassment events? Please describe it in as much detail as you feel comfortable.
3. What patterns, if any, have you noticed in how and when high harassment events occur?
[below prompts if necessary]
 - (a) External events - either related to your industry or not
 - (b) Publishing articles or social media posts about a certain topic or issue (For example, when I tweet about X, I get a ton of angry DMs)
 - (c) Publishing any kind of article or social media post
4. How do you typically respond when you experience online harassment? [the following questions may be asked as needed for each protective strategy]

- (a) Did you take this protective action because you anticipated an increase in harassment, or after the increase in harassment started?
 - (b) How effective did you feel [this strategy] was? Why do you feel it is effective/ineffective?
 - (c) Do any colleagues or friends or people you know employ [this strategy] Is it effective for them? Why or why not?
 - (d) How did you learn about [this strategy]?
5. There are all kinds of strategies people use in situations like this - certain strategies work for some people, but not for others. Are there any harassment protection strategies you have heard about / considered but did not take? [the following questions may be asked as needed for each protective strategy]
- (a) Why or why not? [Prompts follow if the participant has difficulty answering]
 - (b) Does [strategy] not work generally? Why?
 - (c) Is [strategy] not applicable to your particular situation? Why?
 - (d) Does [strategy] have costs or downsides that make it difficult/unrealistic/undesirable to implement? What are those costs or downsides?
 - (e) Did you ever employ [strategy] in the past? Why did you stop employing it?
6. Are there any tools or technologies you use to respond to harassment? This can include affordances of various platforms (like blocking an individual on social media) or external tools (like blocklists that can be shared amongst users).
7. Just like strategies, different people use different tools and technology to deal with harassment for different reasons. Are there any tools or technologies for responding to harassment

that you know about but do not use?

(a) Why or why not? [Prompts follow if the participant has difficulty answering]

(b) Does [tool] not work generally? Why?

(c) Is [tool] not applicable to your particular situation? Why?

(d) Does [tool] have costs or downsides that make it difficult/unrealistic/undesirable to implement? What are those costs or downsides?

(e) Did you ever use [tool] in the past? Why did you stop using it?

8. In as much or as little detail as you like, could you please describe the impact of this experience on your life? You're welcome to discuss either specific events or your general experience of harassment.

(a) How does it impact you, emotionally?

(b) How does it impact your career?

(c) How does it impact your relationships with others?

Advice, Given and Received. Now, I'm going to ask some questions about specific strategies and advice you may have heard of or used when managing harassment.

1. What security advice have you received in the past that's relevant to your experience? [for each item, ask the following if needed]

(a) If unsure what I mean by security advice: Some examples of security advice might include "use a password manager", "don't respond to harassers", or "log off from social media for a while".

- (b) Did you follow this advice? Why or why not?
 - (c) How did you hear about this advice? If no response: could prompt for “on the Internet”, “from colleagues”, etc.
 - (d) In your opinion, how effective is this advice in relation to achieving your security goals?
 - (e) How difficult is this advice to follow?
 - (f) How time-consuming is it to implement this advice?
 - (g) How confident are you that you could implement this advice?
 - (h) How disruptive would it be to implement this advice?
2. Do you have any trusted people in your network you can turn to for advice, either on security specifically or in general with regards to responding to harassment?
3. Have you ever given advice to someone in a similar situation as yours? What advice did you give?
- (a) Did they follow this advice, to your knowledge? Why or why not?
 - (b) Is [advice] not applicable to their particular situation? Why?
 - (c) Does [advice] have costs or downsides that make it difficult/unrealistic/undesirable to implement? What are those costs or downsides?

Closing.

1. For this study, we are choosing not to systematically collect common demographic data, like race, gender, age and ethnicity, in order to help protect participant privacy. However,

we acknowledge these factors can have an impact on one's experience of harassment, so I'd like to give you the opportunity now to share any identifiers you are okay with us reporting as part of our analysis. We will, of course, not use your name, but we can also mask or share other factors.

2. Please share any comments, suggestions, or feedback you have about our study.

Bibliography

- [1] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40–46, 1999.
- [2] Sarah A. Aghazadeh, Alison Burns, Jun Chu, Hazel Feigenblatt, Elizabeth Laribee, Lucy Maynard, Amy L. M. Meyers, Jessica L. O’Brien, and Leah Rufus. *GamerGate: A case study in online harassment*, pages 179–207. Springer International Publishing, Cham, 2018.
- [3] Syed Ishtiaque Ahmed, Md. Romael Haque, Jay Chen, and Nicola Dell. Digital privacy challenges with shared mobile phone use in Bangladesh. *PACM HCI*, 1(CSCW):1–20, 2017.
- [4] Syed Ishtiaque Ahmed, Md. Romael Haque, Shion Guha, Md. Rashidujjaman Rifat, and Nicola Dell. Privacy, security, and surveillance in the global south: A study of biometric mobile SIM registration in Bangladesh. In *Proc. CHI*, 2017.
- [5] Syed Ishtiaque Ahmed, Md. Romael Haque, Irtaza Haider, Jay Chen, and Nicola Dell. “Everyone has some personal stuff”: Designing to support digital privacy with shared mobile phone use in Bangladesh. In *Proc. CHI*, 2019.
- [6] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. Privacy concerns and behaviors of people with visual impairments. In *Proc. CHI*, 2015.
- [7] Tousif Ahmed, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia. Addressing physical safety, security, and privacy for people with visual impairments. In *Proc. SOUPS*, 2016.
- [8] Taslima Akter, Bryan Dosono, Tousif Ahmed, Apu Kapadia, and Bryan Semaan. “I am uncomfortable sharing what I can’t see”: Privacy concerns of the visually impaired with camera based assistive applications. In *Proc. USENIX Security*, 2020.
- [9] Deena Alghamdi, Ivan Flechais, and Marina Jirotko. Security practices for households bank customers in the Kingdom of Saudi Arabia. In *Proc. SOUPS*, 2015.

- [10] Adriana Alvarado Garcia, Alyson L Young, and Lynn Dombrowski. On making data actionable: How activists use imperfect data to foster social change for human rights violations in Mexico. *PACM HCI*, 1(CSCW):1–19, 2017.
- [11] American Library Association. Library privacy guidelines for vendors, 2015. <https://www.ala.org/advocacy/privacy/guidelines/vendors>.
- [12] American Library Association. Professional ethics, 2017. <https://www.ala.org/tools/ethics>.
- [13] American Library Association. Access to library resources and services, 2021. <https://www.ala.org/advocacy/intfreedom/access>.
- [14] American Library Association. Resolution on the misuse of behavioral data surveillance in libraries, 2021. <https://www.ala.org/advocacy/intfreedom/datasurveillanceresolution>.
- [15] Nazanin Andalibi, Oliver L. Haimson, Munmun De Choudhury, and Andrea Forte. Understanding social media disclosures of sexual abuse through the lenses of support seeking and anonymity. In *Proc. CHI*, 2016.
- [16] Julia Angwin. First library to support anonymous internet browsing effort stops after DHS email. *ProPublica*, 2015. <https://www.propublica.org/article/library-support-anonymous-internet-browsing-effort-stops-after-dhs-email>.
- [17] Alissa N. Antle. The ethics of doing research with vulnerable populations. *Interactions*, 24(6):74–77, October 2017. <https://dl.acm.org/doi/10.1145/3137107>.
- [18] Paul M. Aoki and Allison Woodruff. Making space for stories: ambiguity in the design of personal communication systems. In *Proc. CHI*, 2005.
- [19] Entertainment Software Association. 2023 essential facts about the U.S. video game industry. <https://www.theesa.com/2023-essential-facts/>, 2023.
- [20] Australian Library and Information Association. Libraries and privacy guidelines, 2005. <https://www.alia.org.au/Web/Web/Research-and-Publications/Guidelines/Libraries-and-Privacy-Guidelines.aspx>.
- [21] Karla Badillo-Urquiola, Xinru Page, and Pamela J. Wisniewski. Risk vs. restriction: The tension between providing a sense of normalcy and keeping foster teens safe online. In *Proc. CHI*, 2019.
- [22] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. WhisperX: Time-accurate speech transcription of long-form audio. In *INTERSPEECH*. arXiv, July 2023. <http://arxiv.org/abs/2303.00747>.
- [23] Catherine Barwulor, Allison McDonald, Eszter Hargittai, and Elissa M Redmiles. “Disadvantaged in the american-dominated internet”: Sex, work, and technology. In *Proc. CHI*, 2021.

- [24] Elizabeth Behm-Morawitz and Dana Mastro. The effects of the sexualization of female video game characters on gender stereotyping and female self-concept. *Sex Roles*, 61(11-12):808–823, December 2009. <http://link.springer.com/10.1007/s11199-009-9683-8>.
- [25] Rosanna Bellini, Emily Tseng, Noel Warford, Alaa Daffalla, Tara Matthews, Sunny Consolvo, Jill Palzkill Woelfer, Patrick Gage Kelley, Michelle L. Mazurek, Dana Cuomo, Nicola Dell, and Thomas Ristenpart. SoK: Safer digital-safety research involving at-risk users. <http://arxiv.org/abs/2309.00735>, September 2023.
- [26] Clara Berridge, Jodi Halpern, and Karen Levy. Cameras on beds: The ethics of surveillance in nursing home rooms. *AJOB Empirical Bioethics*, 10(1):55–62, 2019.
- [27] Kristin MS Bezio. Ctrl-Alt-Del: GamerGate as a precursor to the rise of the alt-right. *Leadership*, 14(5):556–566, October 2018. <https://doi.org/10.1177/1742715018793744>.
- [28] Rasika Bhalerao, Vaughn Hamilton, Allison McDonald, Elissa M. Redmiles, and Angelika Strohmayr. Ethical practices for security research with at-risk populations. In *Proc. Euro S&P*, pages 546–553, Genoa, Italy, June 2022. IEEE. <https://ieeexplore.ieee.org/document/9799425/>.
- [29] Lindsay Blackwell, Jill Dimond, Sarita Schoenebeck, and Cliff Lampe. Classification and its consequences for online harassment: Design insights from HeartMob. In *PACM HCI*, 2017.
- [30] Lindsay Blackwell, Jean Hardy, Tawfiq Ammari, Tiffany Veinot, Cliff Lampe, and Sarita Schoenebeck. LGBT parents and social media: Advocacy, privacy, and disclosure during shifting social movements. In *Proc. CHI*, 2016.
- [31] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [32] Virginia Braun and Victoria Clarke. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4), 2019.
- [33] Virginia Braun and Victoria Clarke. *Thematic Analysis: A Practical Guide*. SAGE Publications, 1 edition, 2021.
- [34] Virginia Braun and Victoria Clarke. To saturate or not to saturate? Questioning data saturation as a useful concept for thematic analysis and sample-size rationales. *Qualitative Research in Sport, Exercise and Health*, 13(2):201–216, March 2021. <https://www.tandfonline.com/doi/full/10.1080/2159676X.2019.1704846>.
- [35] Sarah Brayne. Enter the dragnet. In *Predict and Surveil: Data, Discretion, and the Future of Policing*. Oxford University Press, 2020. <https://logicmag.io/commons/enter-the-dragnet/>.

- [36] Sarah Brayne. *Predict and Surveil: Data, Discretion, and the Future of Policing*. Oxford University Press, 2020.
- [37] Kelly Caine. Local standards for sample size at CHI. In *Proc. CHI*, 2016.
- [38] Matthew Cantor. US library defunded after refusing to censor LGBTQ authors: ‘We will not ban the books’. <https://www.theguardian.com/books/2022/aug/05/michigan-library-book-bans-lgbtq-authors>, 2022.
- [39] Will Caverly. Ransomware attacks at libraries: How they happen, what to do. <https://publiclibrariesonline.org/2021/05/ransomware-attacks-at-libraries-how-they-happen-what-to-do/>, 2021.
- [40] Kalyani Chadha, Linda Steiner, Jessica Vitak, and Zahra Ashktorab. Women’s responses to online harassment. *International Journal of Communication*, 14, 2020.
- [41] Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. You can’t stay here: The efficacy of Reddit’s 2015 ban examined through hate speech. *PACM HCI*, 1(CSCW), 2017.
- [42] Rahul Chatterjee, Periwinkle Doerfler, Hadas Orgad, Sam Havron, Jackeline Palmer, Diana Freed, Karen Levy, Nicola Dell, Damon McCoy, and Thomas Ristenpart. The spyware used in intimate partner violence. In *Proc. IEEE S&P*, 2018.
- [43] Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Athena Vakali. Hate is not binary: Studying abusive behavior of #Gamergate on Twitter. In *ACM Hypertext Conference*, 2017.
- [44] Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Athena Vakali. Measuring #Gamergate: A tale of hate, sexism, and bullying. In *International Conference on World Wide Web Companion*, 2017.
- [45] Christine Chen, Nicola Dell, and Franziska Roesner. Computer security and privacy in the interactions between victim service providers and human trafficking survivors. In *Proc. USENIX Security*, 2019.
- [46] Janet X. Chen, Allison McDonald, Yixin Zou, Emily Tseng, Kevin Roundy, Acar Tamer-soy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. Trauma-informed computing: Towards safer technology experiences for all. *Proc. CHI*, 2022.
- [47] Taejoong Chung, Jinyoung Han, Daejin Choi, Ted Taekyoung Kwon, Jong-Youn Rha, and Hyunchul Kim. Privacy leakage in event-based social networks: A Meetup case study. *PACM HCI*, 1(CSCW):1–22, 2017.
- [48] Danielle Keats Citron. Addressing cyber harassment: An overview of hate crimes in cyberspace. *Journal of Law, Technology & the Internet*, 2014.
- [49] Patricia Hill Collins. *Intersectionality as Critical Social Theory*. Duke University Press, 2019.

- [50] Sunny Consolvo, Patrick Gage Kelley, Tara Matthews, Kurt Thomas, Lee Dunn, and Elie Bursztein. “Why wouldn’t someone think of democracy as a target?”: Security practices & challenges of people involved with U.S. political campaigns. In *Proc. USENIX Security*, 2021.
- [51] Raymundo Cornejo, Robin Brewer, Caroline Edasis, and Anne Marie Piper. Vulnerability, sharing, and privacy: Analyzing art therapy for older adults with dementia. In *Proc. CSCW, 2016*, 2016.
- [52] Edward M. Corrado. Libraries and protecting patron privacy. *Technical Services Quarterly*, 37(1), 2020.
- [53] Sasha Costanza-Chock. Design justice: Towards an intersectional feminist framework for design theory and practice. *Proceedings of the Design Research Society*, 2018.
- [54] Kimberlé Crenshaw. Demarginalizing the intersection of race and sex: A Black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum*, pages 139–168, 1989.
- [55] Alaa Daffalla, Lucy Simko, Tadayoshi Kohno, and Alexandru G Bardas. Defensive technology use by political activists during the Sudanese revolution. In *Proc. IEEE S&P*, 2021.
- [56] Partha Das Chowdhury, Andrés Domínguez Hernández, Marvin Ramokapane, and Awais Rashid. From utility to capability: A new paradigm to conceptualize and develop inclusive PETs. *New Security Paradigms Workshop*, 2022. Publisher: Association for Computing Machinery (ACM).
- [57] Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *AAAI International Conference On Web and Social Media*, 2017.
- [58] Keith E. Davis and Irene H. Frieze. Research on stalking: What do we know and where do we go? *Violence and Victims*, 15(4):473–487, 2000.
- [59] Steve Dent. Security firm details how hackers stole \$1.3 million in wire transfers. <https://www.engadget.com/hackers-steal-1-3-million-wire-transfer-100039219.html>, 2020.
- [60] Jayati Dev, Pablo Moriano, and L Jean Camp. Lessons learnt from comparing WhatsApp privacy concerns across Saudi and Indian populations. In *Proc. SOUPS*, 2020.
- [61] Tracy L Dietz. An examination of violence and gender role portrayals in video games: Implications for gender socialization and aggressive behavior. *Sex Roles*, 38(5/6), 1998.
- [62] Paul Dunphy, John Vines, Lizzie Coles-Kemp, Rachel Clarke, Vasilis Vlachokyriakos, Peter Wright, John McCarthy, and Patrick Olivier. Understanding the experience-centeredness of privacy and security technologies. In *Proc. NSPW*, 2014.

- [63] Kasra EdalatNejad, Wouter Lueks, Julien Pierre Martin, Soline Ledésert, Anne L’Hôte, Bruno Thomas, Laurent Girod, and Carmela Troncoso. DatashareNetwork: A decentralized privacy-preserving search engine for investigative journalists. In *Proc. USENIX Security*, 2020.
- [64] Ame Elliott and Sara Brody. Straight talk: New Yorkers on mobile messaging and implications for privacy. Technical report, Simply Secure, 2015.
- [65] Valerie Fanelle, Sepideh Karimi, Aditi Shah, Bharath Subramanian, and Sauvik Das. Blind and human: Exploring more usable audio CAPTCHA designs. In *Proc. SOUPS*, 2020.
- [66] Andrew Guthrie Ferguson. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. NYU Press, 2017.
- [67] Paula Fortuna and Sérgio Nunes. A survey on automatic detection of hate speech in text. *ACM Computing Surveys*, 51(4):85:1–85:30, July 2018. <https://dl.acm.org/doi/10.1145/3232676>.
- [68] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. “Is my phone hacked?” analyzing clinical computer security interventions with survivors of intimate partner violence. *PACM HCI*, 3(CSCW):1–24, 2019.
- [69] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders. *PACM HCI*, 1(CSCW):1–22, 2017.
- [70] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. “A Stalker’s Paradise”: How intimate partner abusers exploit technology. In *Proc. CHI*, 2018.
- [71] Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman. Privacy and security threat models and mitigation strategies of older adults. In *Proc. SOUPS*, 2019.
- [72] Martin Garnar and Trina Magi. *Intellectual Freedom Manual*. ALA Editions, 10 edition, 2020.
- [73] Christine Geeng, Mike Harris, Elissa Redmiles, and Franziska Roesner. “Like lesbians walking the perimeter”: Experiences of U.S. LGBTQ+ folks with online security, safety, and privacy advice. In *Proc. USENIX Security*, 2022.
- [74] Arup Kumar Ghosh, Karla Badillo-Urquiola, Shion Guha, Joseph J. LaViola Jr, and Pamela J. Wisniewski. Safety vs. surveillance: What children have to say about mobile apps for parental control. In *Proc. CHI*, 2018.
- [75] Arup Kumar Ghosh, Karla Badillo-Urquiola, Mary Beth Rosson, Heng Xu, John M Carroll, and Pamela J. Wisniewski. Matter of control or safety? Examining parental use of technical monitoring apps on teens’ mobile devices. In *Proc. CHI*, 2018.

- [76] Arup Kumar Ghosh, Charles E. Hughes, and Pamela J. Wisniewski. Circle of trust: A new approach to mobile online safety for families. In *Proc. CHI*, 2020.
- [77] April Glaser. Long before Snowden, librarians were anti-surveillance heroes. *Slate*, June 2015. <https://slate.com/technology/2015/06/usa-freedom-act-before-snowden-librarians-were-the-anti-surveillance-heroes.html>.
- [78] Armanda Gonzalez, Elaine Gomez, Rebeca Orozco, and Samuel Jacobs. Entering the boys’ club: An analysis of female representation in game industry, culture, and design. In *Proc. iConference*. iSchools, March 2014. <https://www.ideals.illinois.edu/handle/2142/47355>.
- [79] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4), 2002.
- [80] Kishonna L. Gray, Bertan Buyukozturk, and Zachary G. Hill. Blurring the boundaries: Using GamerGate to examine “real” and symbolic violence against women in contemporary gaming culture. *Sociology Compass*, 11(3):e12458, 2017. <https://onlinelibrary.wiley.com/doi/abs/10.1111/soc4.12458>.
- [81] Kirsten Grind, Ben Fritz, and Sarah E. Needleman. Activision CEO Bobby Kotick knew for years about sexual-misconduct allegations at videogame giant. *Wall Street Journal*, November 2021. <https://www.wsj.com/articles/activision-videogames-bobby-kotick-sexual-misconduct-allegations-11637075680>.
- [82] Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H. Mhaidli, Kentaro Toyama, and Florian Schaub. Keeping a low profile?: Technology, risk and privacy among undocumented immigrants. In *Proc. CHI*, 2018.
- [83] Greg Guest, Arwen Bunce, and Laura Johnson. How many interviews are enough?: An experiment with data saturation and variability. *Field Methods*, 18(1), 2006.
- [84] Diana L. Gustafson and Fern Brunger. Ethics, “Vulnerability,” and feminist participatory action research with a disability community. *Qualitative Health Research*, 24(7):997–1005, July 2014. <https://doi.org/10.1177/1049732314538122>.
- [85] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proc. CHI*, 2018.
- [86] Eszter Hargittai. Second-level digital divide: Differences in people’s online skills. *First Monday*, 7(4), 2002. <http://journals.uic.edu/ojs/index.php/fm/article/view/942>.
- [87] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. Clinical computer security for victims of intimate partner violence. In *Proc. USENIX Security*, 2019.

- [88] Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang. Cooperative privacy and security: Learning from people with visual impairments and their allies. In *Proc. SOUPS*, 2019.
- [89] Cormac Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proc. NSPW*, 2009.
- [90] Dominik Hornung, Claudia Müller, Irina Shklovski, Timo Jakobi, and Volker Wulf. Navigating relationships and boundaries: Concerns around ICT-uptake for elderly people. In *Proc. CHI*, 2017.
- [91] John B. Horrigan. Libraries at the crossroads. Technical report, Pew Research Center, 2015. <https://policycommons.net/artifacts/618882/libraries-at-the-crossroads/1599903/>.
- [92] John B. Horrigan. Libraries 2016. Technical report, Pew Research Center, September 2016.
- [93] Margaret C Jack, Pang Sovannaroeth, and Nicola Dell. “Privacy is not a concept, but a way of dealing with life”: Localization of transnational technology platforms and liminal privacy practices in Cambodia. *PACM HCI*, 3(CSCW):1–19, 2019.
- [94] Rebecca Jeong and Sonia Chiasson. ‘Lime’, ‘open lock’, and ‘blocked’: Children’s perception of colors, symbols, and words in cybersecurity warnings. In *Proc. CHI*, 2020.
- [95] Anil Kalhan. Immigration Surveillance. *Maryland Law Review*, 74(1), 2014.
- [96] Chris Kanich, Nicholas Weaver, Damon McCoy, Tristan Halvorson, Christian Kreibich, Kirill Levchenko, Vern Paxson, Geoffrey M Voelker, and Stefan Savage. Show me the money: Characterizing spam-advertised revenue. In *Proc. USENIX Security*, 2011.
- [97] Os Keyes, Josephine Hoy, and Margaret Drouhard. Human-computer insurrection: Notes on an anarchist HCI. In *Proc. CHI*, 2019.
- [98] Dorothea Kleine. ICT4What? - using the Choice Framework to operationalise the capability approach to development. In *Proc. ICTD*, 2009.
- [99] Patrick Klepek. Nintendo employee ‘terminated’ after smear campaign over censorship, company denies harassment was factor. <https://kotaku.com/nintendo-employee-terminated-after-smear-campaign-over-1768100368>, March 2016.
- [100] Patrick Klepek. How restrictive contracts stifle and control creativity in the video game industry. <https://www.vice.com/en/article/g5va43/noncompete-contracts-video-game-industry>, February 2023.
- [101] Yong Ming Kow, Yubo Kou, Bryan Semaan, and Waikuen Cheng. Mediating the undercurrents: Using social media to sustain a social movement. In *Proc. CHI*, 2016.

- [102] Sandjar Kozubaev, Fernando Rochaix, Carl DiSalvo, and Christopher A. Le Dantec. Spaces and traces: Implications of smart technology in public housing. In *Proc. CHI*, 2019.
- [103] Klaus Krippendorff. Testing the reliability of content analysis data: What is involved and why. In Klaus Krippendorff and Mary Angela Bock, editors, *The Content Analysis Reader*, chapter 6.2, pages 350–357. SAGE, 2009.
- [104] Priya Kumar, Shalmali Milind Naik, Utkarsha Ramesh Devkar, Marshini Chetty, Tamara L Clegg, and Jessica Vitak. 'No telling passcodes out because they're private' understanding children's mental models of privacy and security online. *PACM HCI*, 1(CSCW):1–21, 2017.
- [105] Priya C. Kumar, Marshini Chetty, Tamara L. Clegg, and Jessica Vitak. Privacy and security considerations for digital technology use in elementary schools. In *Proc. CHI*, 2019.
- [106] David Kushner. 4chan's overlord reveals why he walked away, 6 2018.
- [107] April D. Lambert, Michelle Parker, and Masooda Bashir. Library patron privacy in jeopardy: An analysis of the privacy policies of digital content vendors. In *Proc. ASIST*, volume 52, 2015.
- [108] Sarah Lamdan. Librarianship at the crossroads of ICE surveillance. In *The Library With The Lead Pipe*, 2019. <https://www.inthelibrarywiththeleadpipe.org/2019/ice-surveillance/>.
- [109] Sarah Shik Lamdan. Social media privacy: A rallying cry to librarians. *The Library Quarterly: Information, Community, Policy*, 85(3), 2015.
- [110] Paul Lashmar. No more sources?: The impact of Snowden's revelations on journalists and their confidential sources. *Journalism Practice*, 11(6):665–688, July 2017. <https://www.tandfonline.com/doi/full/10.1080/17512786.2016.1179587>.
- [111] Elmer Lastdrager, Inés Carvajal Gallardo, Pieter Hartel, and Marianne Junger. How effective is anti-phishing training for children? In *Proc. SOUPS*, 2017.
- [112] Stevens Le Blond, Alejandro Cuevas, Juan Ramón Troncoso-Pastoriza, Philipp Jovanovic, Bryan Ford, and Jean-Pierre Hubaux. On enforcing the digital immunity of a large humanitarian organization. In *Proc. IEEE S&P*, 2018.
- [113] Amy E. Lerman and Vesla Weaver. Staying out of sight? Concentrated policing and local political action. *The ANNALS of the American Academy of Political and Social Science*, 651(1), 2013.
- [114] Ada Lerner, Helen Yuxun He, Anna Kawakami, Silvia Catherine Zeamer, and Roberto Hoyle. Privacy and activism in the transgender community. In *Proc. CHI*, 2020.

- [115] Kirill Levchenko, Andreas Pitsillidis, Neha Chachra, Brandon Enright, Márk Félegyházi, Chris Grier, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, et al. Click trajectories: End-to-end analysis of the spam value chain. In *Proc. IEEE S&P*. IEEE, 2011.
- [116] Karen Levy and Bruce Schneier. Privacy threats in intimate relationships. *Journal of Cybersecurity*, 6(1):tyaa006, 2020.
- [117] Calvin A. Liang, Sean A. Munson, and Julie A. Kientz. Embracing four tensions in human-computer interaction research with marginalized people. *ACM TOCHI*, 28(2):1–47, 2021.
- [118] Alan F. Luo, Rachel Greenstadt, Noel Warford, Michelle L. Mazurek, Samuel Dooley, and Nora McDonald. How library IT staff navigate privacy and security challenges and responsibilities. In *Proc. USENIX Security*, 2023.
- [119] Monica Maceli. Librarians’ Mental Models and Use of Privacy-Protection Technologies. *Journal of Intellectual Freedom & Privacy*, 4(1), 2019.
- [120] Mary Madden. Privacy, security, and digital inequality. Report, Data & Society Research Institute, September 2017. <https://apo.org.au/node/184986>.
- [121] Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. Privacy, poverty, and big data: A matrix of vulnerabilities for poor americans. *Washington University Law Review*, 95(1):53–126, 2017. <https://heinonline.org/HOL/P?h=hein.journals/walq95&i=59>.
- [122] Herbert Marcuse. *One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society*. Routledge, London, 1964.
- [123] William R. Marczak and Vern Paxson. Social engineering attacks on government opponents: Target perspectives. In *Proc. PETS*, 2017.
- [124] William R. Marczak, John Scott-Railton, Morgan Marquis-Boire, and Vern Paxson. When governments hack opponents: A look at actors and technology. In *Proc. USENIX Security*, 2014.
- [125] Sonali Tukaram Marne, Mahdi Nasrullah Al-Ameen, and Matthew K. Wright. Learning system-assigned passwords: A preliminary study on the people with learning disabilities. In *Proc. SOUPS*, 2017.
- [126] Alice Marwick, Claire Fontaine, and danah boyd. ”Nobody sees it, nobody gets mad”: Social media, privacy, and personal responsibility among low-SES youth. *Social Media + Society*, 3(2):2056305117710455, April 2017. <https://doi.org/10.1177/2056305117710455>.
- [127] Adrienne Massanari. #GamerGate and The Fappening: How Reddit’s algorithm, governance, and culture support toxic technocultures. *New Media & Society*, 19(3):329–346, March 2017. <https://doi.org/10.1177/1461444815608807>.

- [128] Tara Matthews, Kerwell Liao, Anna Turner, Marianne Berkovich, Robert Reeder, and Sunny Consolvo. “She’ll just grab any device that’s closer”: A study of everyday device & account sharing in households. In *Proc. CHI*, 2016.
- [129] Tara Matthews, Kathleen O’Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *Proc. CHI*, 2017.
- [130] Leslie McCall. The complexity of intersectionality. *Signs: Journal of Women in Culture and Society*, 30(3):1771—1800, 2005.
- [131] Nora McDonald, Karla Badillo-Urquiola, Morgan G Ames, Nicola Dell, Elizabeth Kenneski, Manya Sleeper, and Pamela J. Wisniewski. Privacy and power: Acknowledging the importance of privacy research and design for vulnerable populations. In *Proc. CHI Extended Abstracts*, 2020.
- [132] Nora McDonald and Andrea Forte. The politics of privacy theories: Moving from norms to vulnerabilities. In *Proc. CHI*, 2020.
- [133] Nora McDonald, Rachel Greenstadt, and Andrea Forte. Intersectional thinking about PETs: A study of library privacy. *PoPETS*, 2023(2):480–495, April 2023.
- [134] Nora McDonald, Alison Larsen, Allison Battisti, Galina Madjaroff, Aaron Massey, and Helena Mentis. Realizing choice: Online safeguards for couples adapting to cognitive challenges. In *Proc. SOUPS*, 2020.
- [135] Susan E. McGregor. Digital security and source protection for journalists. Technical report, Tow Center for Digital Journalism, 2014.
- [136] Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner. Investigating the computer security practices and needs of journalists. In *Proc. USENIX Security*, 2015.
- [137] Susan E. McGregor, Franziska Roesner, and Kelly Caine. Individual versus organizational computer security and privacy concerns in journalism. In *Proc. SOUPS*, 2016.
- [138] Susan E. McGregor, Elizabeth Anne Watkins, Mahdi Nasrullah Al-Ameen, Kelly Caine, and Franziska Roesner. When the weakest link is strong: Secure collaboration in the case of the Panama Papers. In *Proc. USENIX Security*, 2017.
- [139] Susan E. McGregor, Elizabeth Anne Watkins, and Kelly Caine. Would you Slack that? The impact of security and privacy on cooperative newsroom work. *PACM HCI*, 1(CSCW):1–22, 2017.
- [140] Bridget Christine McHugh, Pamela J Wisniewski, Mary Beth Rosson, Heng Xu, and John M Carroll. Most teens bounce back: Using diary methods to examine how quickly teens recover from episodic online risk exposure. *PACM HCI*, 1(CSCW):1–19, 2017.

- [141] Brenna McNally, Priya Kumar, Chelsea Hordatt, Matthew Louis Mauriello, Shalmali Naik, Leyla Norooz, Alazandra Shorter, Evan Golub, and Allison Druin. Co-designing mobile online safety applications with children. In *Proc. CHI*, 2018.
- [142] Andrew R. McNeill, Lynne Coventry, Jake Pywell, and Pam Briggs. Privacy considerations when designing social network systems to support successful ageing. In *Proc. CHI*, 2017.
- [143] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. Toys that listen: A study of parents, children, and internet-connected toys. In *Proc. CHI*, 2017.
- [144] Hamid Mehmood, Tallal Ahmad, Lubna Razaq, Shrirang Mare, Maryem Zafar Usmani, Richard Anderson, and Agha Ali Raza. Towards digitization of collaborative savings among low-income groups. *PACM HCI*, 3(CSCW):1–30, 2019.
- [145] Gautama Mehta. Proposal to install spyware in university libraries to protect copyrights shocks academics. *Coda Story*, 2020. <https://www.codastory.com/authoritarian-tech/spyware-in-libraries/>.
- [146] Helena M. Mentis, Galina Madjaroff, and Aaron K Massey. Upside and downside risk in online security for older adults with mild cognitive impairment. In *Proc. CHI*, 2019.
- [147] Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. Tackling online abuse: A survey of automated abuse detection methods. <http://arxiv.org/abs/1908.06024>, September 2020.
- [148] Torill Elvira Mortensen. Anger, fear, and games: The long event of #GamerGate. *Games and Culture*, 13(8):787–806, December 2018. <https://doi.org/10.1177/1555412016640408>.
- [149] Carol Moser, Tianying Chen, and Sarita Y. Schoenebeck. Parents’ and children’s preferences about parents sharing about children on social media. In *Proc. CHI*, 2017.
- [150] Maddy Myers. E3 expo leaks the personal information of over 2,000 journalists. <https://kotaku.com/e3-expo-leaks-the-personal-information-of-over-2-000-jo-1836936908>, August 2019.
- [151] Mahin Naderifar, Hamideh Goli, and Fereshteh Ghaljaie. Snowball sampling: A purposeful method of sampling in qualitative research. *Strides in Development of Medical Education*, 14(3), September 2017. <http://sdmejournal.com/en/articles/67670.html>.
- [152] James Nicholson, Lynne Coventry, and Pamela Briggs. “If it’s important it will be a headline”: Cybersecurity information seeking in older adults. In *Proc. CHI*, 2019.
- [153] Debora Nozza. Exposing the limits of zero-shot cross-lingual hate speech detection. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proc. ACL-IJCNLP*, pages 907–914, Online, August 2021. Association for Computational Linguistics. <https://aclanthology.org/2021.acl-short.114>.

- [154] Martha C. Nussbaum. *Creating Capabilities*. Harvard University Press, 2011.
- [155] Martha C. Nussbaum. *Creating Capabilities: The Human Development Approach*. Harvard University Press, 2011.
- [156] Borke Obada-Obieh, Lucrezia Spagnolo, and Konstantin Beznosov. Towards understanding privacy and trust in online reporting of sexual assault. In *Proc. SOUPS*, 2020.
- [157] National Institute of Standards and Technology. Digital identity guidelines. Technical Report National Institute of Standards and Technology Special Publication 800-63-3, U.S. Department of Commerce, Washington, D.C., 2020.
- [158] Ilse Oosterlaken and Jeroen van den Hoven. Editorial: ICT and the capability approach. *Ethics and Information Technology*, 13(2), 2011.
- [159] Ilse Oosterlaken and Jeroen van den Hoven. Editorial: ICT and the capability approach. *Ethics and Information Technology*, 13(2):65–67, March 2011.
- [160] Zizi Papacharissi. Privacy as a luxury commodity. *First Monday*, 15(8), 2010.
- [161] Jessica Pater, Casey Fiesler, and Michael Zimmer. No humans here: Ethical speculation on public data, unintended consequences, and the limits of institutional review. *Proceedings of the ACM on Human-Computer Interaction*, 6(GROUP):1–13, January 2022. <https://dl.acm.org/doi/10.1145/3492857>.
- [162] Ruth Pearce. A methodology for the marginalised: Surviving oppression and traumatic fieldwork in the neoliberal academy. *Sociology*, 54(4), 2020.
- [163] Gregory P Perreault and Tim P Vos. The GamerGate controversy and journalistic paradigm maintenance. *Journalism*, 19(4):553–569, April 2018. <https://doi.org/10.1177/1464884916670932>.
- [164] Justin Petelka, Lucy Van Kleunen, Liam Albright, Elizabeth Murnane, Stephen Volda, and Jaime Snyder. Being (in) visible: Privacy, transparency, and disclosure in the self-management of bipolar disorder. In *Proc. CHI*, 2020.
- [165] Jay Peters. Former kotaku writers are launching a new video game site — and they own it this time. <https://www.theverge.com/2023/11/7/23949269/aftermath-video-games-kotaku-defector>, November 2023.
- [166] Flor Miriam Plaza-del-arco, Debora Nozza, and Dirk Hovy. Respectful or toxic? Using zero-shot learning with language models to detect hate speech. In Yi-ling Chung, Paul R\textbackslashottger, Debora Nozza, Zeerak Talat, and Aida Mostafazadeh Davani, editors, *Proc. WOAHI*, pages 60–68, Toronto, Canada, July 2023. Association for Computational Linguistics. <https://aclanthology.org/2023.woah-1.6>.
- [167] John R. Porter, Kiley Sobel, Sarah E Fox, Cynthia L Bennett, and Julie A Kientz. Filtered out: Disability disclosure practices in online dating communities. *PACM HCI*, 1(CSCW):1–13, 2017.

- [168] Julie Posetti. *Protecting Journalism Sources in the Digital Age*. UNESCO Publishing, May 2017.
- [169] Susannah Quick, Gillian Prior, Ben Toombs, Luke Taylor, and Rosanna Currenti. Cross-European survey to measure users’ perceptions of the benefits of ICT in public libraries. Technical report, Bill and Melinda Gates Foundation, 2013.
- [170] Lee Rainie. The information needs of citizens: Where libraries fit in. Technical report, Pew Research Center, April 2018. <https://policycommons.net/artifacts/617402/the-information-needs-of-citizens/1598222/>.
- [171] Justin M Rao and David H Reiley. The economics of spam. *Journal of Economic Perspectives*, 26(3):87–110, 2012.
- [172] John Rawls. *A Theory of Justice: Revised Edition*. Belknap Press, 1999.
- [173] John Rawls. *Justice as Fairness: A Restatement*. Belknap Press, 2001.
- [174] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. How I learned to be secure: a census-representative survey of security advice sources and behavior. In *Proc. ACM CCS*, 2016.
- [175] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. Where is the digital divide? a survey of security, privacy, and socioeconomics. In *Proc. CHI*, 2017.
- [176] Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek. I think they’re trying to tell me something: Advice sources and selection for digital security. In *2016 Symposium on Security and Privacy (SP)*, pages 272–288, 2016.
- [177] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. A comprehensive quality evaluation of security and privacy advice on the web. In *Proc. USENIX Security*, 2020.
- [178] Jake Reichel, Fleming Peck, Mikako Inaba, Bisrat Moges, Brahmnoor Singh Chawla, and Marshini Chetty. ‘I have too much respect for my elders’: Understanding South African mobile users’ perceptions of privacy and current behaviors on Facebook and WhatsApp. In *Proc. USENIX Security*, 2020.
- [179] Aja Romano. Is J.K. Rowling transphobic? Let’s let her speak for herself. <https://www.vox.com/culture/23622610/jk-rowling-transphobic-statements-timeline-history-controversy>, March 2023.
- [180] Pam Royse, Joon Lee, Baasanjav Undrahbuyan, Mark Hopson, and Mia Consalvo. Women and games: Technologies of the gendered self. *New Media & Society*, 9(4):555–576, August 2007. <https://doi.org/10.1177/1461444807080322>.
- [181] Sabirat Rubya and Svetlana Yarosh. Interpretations of online anonymity in Alcoholics Anonymous and Narcotics Anonymous. *PACM HCI*, 1(CSCW):1–22, 2017.

- [182] Michael Salter. From geek masculinity to Gamergate: The technological rationality of online abuse. *Crime, Media, Culture*, 14(2):247–264, August 2018. <https://doi.org/10.1177/1741659017690893>.
- [183] Nithya Sambasivan, Amna Batool, Nova Ahmed, Tara Matthews, Kurt Thomas, Laura Sanely Gaytán-Lugo, David Nemer, Elie Bursztein, Elizabeth Churchill, and Sunny Consolvo. “They don’t leave us alone anywhere we go”: Gender and digital abuse in South Asia. In *Proc. CHI*, 2019.
- [184] Nithya Sambasivan, Garen Checkley, Amna Batool, Nova Ahmed, David Nemer, Laura Sanely Gaytán-Lugo, Tara Matthews, Sunny Consolvo, and Elizabeth Churchill. “Privacy is not for me, it’s for those rich women”: Performative privacy practices on mobile phones by women in South Asia. In *Proc. SOUPS*, 2018.
- [185] Pedro Sanches, Vasiliki Tsaknaki, Asreen Rostami, and Barry Brown. Under surveillance: Technology practices of those monitored by the state. In *Proc. CHI*, 2020.
- [186] Shruti Sannon and Andrea Forte. Privacy research with marginalized groups: What we know, what’s needed, and what’s next. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–33, November 2022. <https://dl.acm.org/doi/10.1145/3555556>.
- [187] Charlie Savage and Leslie Kaufman. Phone records of journalists seized by U.S. *The New York Times*, May 2013. <https://www.nytimes.com/2013/05/14/us/phone-records-of-journalists-of-the-associated-press-seized-by-us.html>.
- [188] Morgan Klaus Scheuerman, Stacy M. Branham, and Foad Hamidi. Safe spaces and safe places: Unpacking technology-mediated experiences of safety and harm with transgender people. *PACM HCI*, 2(CSCW):1–27, 2018.
- [189] Ari Schlesinger, W. Keith Edwards, and Rebecca E Grinter. Intersectional HCI: Engaging identity through gender, race, and class. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 5412–5427, 2017.
- [190] Zachary Schmitt and Svetlana Yarosh. Participatory design of technologies to support recovery from substance use disorders. *PACM HCI*, 2(CSCW):1–27, 2018.
- [191] Sarita Schoenebeck, Amna Batool, Giang Do, Sylvia Darling, Gabriel Grill, Daricia Wilkinson, Mehtab Khan, Kentaro Toyama, and Louise Ashwell. Online harassment in majority contexts: Examining harms and remedies across countries. In *Proc. CHI*, pages 1–16, Hamburg Germany, April 2023. ACM. <https://dl.acm.org/doi/10.1145/3544548.3581020>.
- [192] Jason Schreier. The horrible world of video game crunch. <https://kotaku.com/crunch-time-why-game-developers-work-such-insane-hours-1704744577>, September 2016.

- [193] Amartya Sen. Rights and agency. *Philosophy & Public Affairs*, 11(1):3–39, 1982. Publisher: Wiley.
- [194] Amartya Sen. Justice: Means versus freedoms. *Philosophy & Public Affairs*, 19(2):111–121, 1990.
- [195] Amartya Sen. Capability and well-being. In *The Quality of Life*. Clarendon Press, 1993.
- [196] Amartya Sen. *Development as Freedom*. Alfred Knopf, 1999.
- [197] Amartya Sen. *Identity and Violence: The Illusion of Destiny*. W. W. Norton & Company, 2007.
- [198] Amartya Sen. *The Idea of Justice*. Belknap Press, 2009.
- [199] Irina Shklovski and Volker Wulf. The use of private mobile phones at war: Accounts from the Donbas conflict. In *Proc. CHI*, 2018.
- [200] Lucy Simko, Ada Lerner, Samia Ibtasam, Franziska Roesner, and Tadayoshi Kohno. Computer security and privacy for refugees in the United States. In *Proc. IEEE S&P*, 2018.
- [201] Sanhita SinhaRoy. Defenders of patron privacy. *American Libraries: The Magazine of the American Library Association*, 2021.
- [202] Manya Sleeper, Tara Matthews, Kathleen O’Leary, Anna Turner, Jill Palzkill Woelfer, Martin Shelton, Andrew Oplinger, Andreas Schou, and Sunny Consolvo. Tough times at transitional homeless shelters: Considering the impact of financial insecurity on digital security and privacy. In *Proc. CHI*, 2019.
- [203] Brian Stanton, Mary Theofanos, Sandra Spickard Prettyman, and Susanne Furman. Security fatigue. *IT Professional*, 18(5):26–32, 2016.
- [204] Joan Starr. Libraries and national security. *First Monday*, 9(12), 2004.
- [205] Elizabeth Stobert, David Barrera, Valérie Homier, and Daniel Kollek. Understanding cybersecurity practices in emergency departments. In *Proc. CHI*, 2020.
- [206] Angelika Strohmayer, Jenn Clamen, and Mary Laing. Technologies for social justice: Lessons from sex workers on the front lines. In *Proc. CHI*, 2019.
- [207] Substance Abuse and Mental Health Services Administration. Samhsa’s concept of trauma and guidance for a trauma-informed approach, 2014.
- [208] Borislav Tadic, Markus Rohde, Volker Wulf, and David Randall. ICT use by prominent activists in Republika Srpska. In *Proc. CHI*, 2016.
- [209] David R. Thomas. A general inductive approach for analyzing qualitative evaluation data. *American Journal of Evaluation*, 27(2):237–246, 2006.

- [210] Kurt Thomas, Devdatta Akhawe, Michael Bailey, Dan Boneh, Elie Bursztein, Sunny Consolvo, Nicola Dell, Zakir Durumeric, Patrick Gage Kelley, Deepak Kumar, et al. SoK: Hate, harassment, and the changing landscape of online abuse. In *Proc. IEEE S&P*, 2021.
- [211] Kurt Thomas, Patrick Gage Kelley, Sunny Consolvo, Patrawat Samermit, and Elie Bursztein. “It’s common and a part of being a content creator”: Understanding how creators experience and cope with hate and harassment online. In *Proc. CHI*, pages 1–15, New Orleans LA USA, April 2022. ACM. <https://dl.acm.org/doi/10.1145/3491102.3501879>.
- [212] Alexandra To, Wenxia Sweeney, Jessica Hammer, and Geoff Kaufman. “They just don’t get it”: Towards social technologies for coping with interpersonal racism. *PACM HCI*, 4(CSCW1):1–29, 2020.
- [213] Cherie Todd. Commentary: GamerGate and resistance to the diversification of gaming culture. *Women’s Studies Journal*, 29(1):64–67, August 2015. <https://www.ideals.illinois.edu/handle/2142/47355>.
- [214] Michael Trice. Putting GamerGate in context: How group documentation informs social media activity. In *Proc. SIGDOC*, pages 1–5, Limerick, Ireland, July 2015. ACM. <https://dl.acm.org/doi/10.1145/2775441.2775471>.
- [215] Michael Trice and Liza Potts. Building dark patterns into platforms: How GamerGate perturbed Twitter’s user experience. *Present tense*, 6, January 2018.
- [216] Emily Tseng, Diana Freed, Kristen Engel, Thomas Ristenpart, and Nicola Dell. A digital safety dilemma: Analysis of computer-mediated computer security interventions for intimate partner violence during COVID-19. In *Proc. CHI*, 2021.
- [217] Emily Tseng, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. Care infrastructures for digital security in intimate partner violence. In *Proc. CHI*, pages 1–20, New Orleans LA USA, April 2022. ACM. <https://dl.acm.org/doi/10.1145/3491102.3502038>.
- [218] U.S. Federal Bureau of Investigation. Business e-mail compromise the 12 billion dollar scam. <https://www.ic3.gov/Media/Y2018/PSA180712>, 2018.
- [219] U.S. Federal Bureau of Investigation. FBI Tech Tuesday: Strong passphrases and account protection. <https://www.fbi.gov/contact-us/field-offices/phoenix/news/press-releases/fbi-tech-tuesday-strong-passphrases-and-account-protection>, 2021.
- [220] Jan Van Dijk. *The Digital Divide*. Polity Press, 2020.
- [221] Aditya Vashistha, Richard Anderson, and Shirang Mare. Examining security and privacy research in developing regions. In *Proc. COMPASS*, 2018.
- [222] Aditya Vashistha, Abhinav Garg, Richard Anderson, and Agha Ali Raza. Threats, abuses, flirting, and blackmail: Gender inequity in social media voice forums. In *Proc. CHI*, 2019.

- [223] Jessica Vitak, Yuting Liao, Priya Kumar, and Mega Subramaniam. Librarians as information intermediaries: Navigating tensions between being helpful and being liable. In Gobinda Chowdhury, Julie McLeod, Val Gillet, and Peter Willett, editors, *Transforming Digital Worlds*, pages 693–702. Springer International Publishing, Cham, 2018. http://link.springer.com/10.1007/978-3-319-78105-1_80.
- [224] Jessica Vitak, Yuting Liao, Mega Subramaniam, and Priya Kumar. “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. *PACM HCI*, 2(CSCW):1–25, 2018.
- [225] Ashley Marie Walker, Yaxing Yao, Christine Geeng, Roberto Hoyle, and Pamela Wisniewski. Moving beyond ‘one size fits all’ research considerations for working with vulnerable populations. *Interactions*, 26(6):34–39, 2019.
- [226] Michael Walzer. *Spheres Of Justice: A Defense Of Pluralism And Equality*. Basic Books, 1984.
- [227] Nina Wang, Daniel Bateyko Allison McDonald, and Emily Tucker. American dragnet: Data-driven deportation in the 21st century. Technical report, Georgetown Law Center on Privacy & Technology, 2022. <https://americandragnet.org/>.
- [228] Ruolin Wang, Chun Yu, Xing-Dong Yang, Weijie He, and Yuanchun Shi. EarTouch: Facilitating smartphone use for visually impaired people in mobile and public scenarios. In *Proc. CHI*, 2019.
- [229] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L. Mazurek, Manya Sleeper, and Kurt Thomas. SoK: A framework for unifying at-risk user research. In *Proc. IEEE S&P*, 2022.
- [230] Mark Warner, Andreas Gutmann, M Angela Sasse, and Ann Blandford. Privacy unraveling around explicit HIV status disclosure fields in the online geosocial hookup app Grindr. *PACM HCI*, 2(CSCW):1–22, 2018.
- [231] Mark Warner, Agnieszka Kitkowska, Jo Gibbs, Juan F Maestre, and Ann Blandford. Evaluating ‘prefer not to say’ around sensitive disclosures. In *Proc. CHI*, 2020.
- [232] Jeffrey Warshaw, Nina Taft, and Allison Woodruff. Intuitions, analytics, and killing ants: Inference literacy of high school-educated adults in the US. In *Proc. SOUPS*, 2016.
- [233] Rick Wash and Emilee Rader. Too much knowledge? Security beliefs and protective behaviors among United States internet users. In *Proc. SOUPS*, 2015.
- [234] Miranda Wei, Sunny Consolvo, Patrick Gage Kelley, Tadayoshi Kohno, Franziska Roesner, and Kurt Thomas. “There’s so much responsibility on users right now:” expert advice for staying safer from hate and harassment. In *Proc. CHI*, pages 1–17, Hamburg Germany, April 2023. ACM. <https://dl.acm.org/doi/10.1145/3544548.3581229>.

- [235] Nicola Wendt, Rikke Bjerg Jensen, and Lizzie Coles-Kemp. Civic empowerment through digitalisation: The case of Greenlandic women. In *Proc. CHI*, 2020.
- [236] Dmitri Williams, Mia Consalvo, Scott Caplan, and Nick Yee. Looking for gender: Gender roles and behaviors among online gamers. *Journal of Communication*, 59(4):700–725, December 2009. <https://academic.oup.com/joc/article/59/4/700-725/4098413>.
- [237] Pamela Wisniewski, Arup Kumar Ghosh, Heng Xu, Mary Beth Rosson, and John M Carroll. Parental control vs. teen self-regulation: Is there a middle ground for mobile online safety? In *Proc. CSCW, 2017*, 2017.
- [238] Pamela Wisniewski, Heng Xu, Mary Beth Rosson, and John M Carroll. Parents just don’t understand: Why teens don’t talk to parents about their online risk experiences. In *Proc. CSCW, 2017*, 2017.
- [239] Pamela Wisniewski, Heng Xu, Mary Beth Rosson, Daniel F Perkins, and John M Carroll. Dear diary: Teens reflect on their weekly online risk experiences. In *Proc. CHI*, 2016.
- [240] Brittany Wong. Wait, what the heck is a ‘parasocial relationship’? https://www.huffpost.com/entry/parasocial-relationships-with-celebrities_l_60a56a18e4b0d45b75248115, 2021.
- [241] Rebecca Wong. Guidelines to incorporate trauma-informed care strategies in qualitative research. <https://www.urban.org/urban-wire/guidelines-incorporate-trauma-informed-care-strategies-qualitative-research/>, 2021.
- [242] William Wresch. Progress on the global digital divide: An ethical perspective based on Amartya Sen’s capabilities model. *Ethics and Information Technology*, 11(4), 2009.
- [243] Huichuan Xia, Yang Wang, Yun Huang, and Anuj Shah. “Our privacy needs to be protected at all costs”: Crowd workers’ privacy experiences on Amazon Mechanical Turk. *PACM HCI*, 1(CSCW):1–22, 2017.
- [244] Fujiko Robledo Yamamoto, Amy Volda, and Stephen Volda. From therapy to teletherapy: Relocating mental health services online. *PACM HCI*, 5(CSCW2), 2021.
- [245] Jun Zhao, Ge Wang, Carys Dally, Petr Slovak, Julian Edbrooke-Childs, Max Van Kleek, and Nigel Shadbolt. ‘I make up a silly name’: Understanding children’s perception of privacy risks online. In *Proc. CHI*, 2019.
- [246] Kathryn Zickuhr, Lee Rainie, and Kristen Purcell. Library services in the digital age. Technical report, Pew Research Center, 2013.
- [247] Shoshana Zuboff. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019.