How Experts Personalize Privacy & Security Advice for At-Risk Users

Wentao Guo
University of Maryland
wguo5@umd.edu

Alexander Yang
University of Maryland
ayang130@terpmail.umd.edu

Nathan Malkin
New Jersey Institute of Technology
nathan.malkin@njit.edu

Michelle L. Mazurek
University of Maryland
mmazurek@umd.edu

Abstract

Prioritizing privacy and security advice is a particular challenge for at-risk users, requiring difficult judgment calls with the risk of harmful consequences. We interviewed 18 experts who tailor privacy and security advice for at-risk users in personalized consultations and trainings, to understand their strategies and challenges. We identify five main objectives that experts balance as they differentiate the content and delivery of advice, and we explore how different types of context about clients are used to achieve those objectives. We also describe four methods used to ascertain this context, which have different trade-offs regarding priorities such as using time efficiently and managing the reliability of information. Through this, we especially focus on the challenges and rationales that motivate providers to follow particular practices. By surfacing choices that experts make about advice differentiation and by identifying areas for researchers and technologists to assist experts, our work can inform next steps in research, tool design, and ultimately better advice for at-risk users.

Keywords

at-risk users, advice, privacy, security, digital safety

1 Introduction

Many resources explain how to protect one's privacy and security, but people who wish to do so still face several problems. The advice[1] they find may be incorrect or outdated; it may be inapplicable or counterproductive for their circumstances; and they may struggle to prioritize among the huge diversity of advice [54]. These problems are compounded for *at-risk users*, defined as "individuals with risk factors that augment or amplify their chances of being digitally attacked and/or suffering disproportionate harms" [71]. Experts stress the importance of tailoring advice for at-risk users, noting that generic advice can be unproductive or even harmful [5]. Thus, a small number of advice providers[2] worldwide offer personalized consultations and trainings to, for example, journalists [1] and targets of tech abuse by intimate partners [36].

Tailoring advice for at-risk users can require hazardous judgment calls without obvious answers. Even reasonable advice sometimes risks dire consequences: e.g., deleting spyware installed by an abuser could give clients greater autonomy, but it could also lead to physical violence [30]. To complicate matters further, at-risk users are far from monolithic [67], and advice providers deal with many sources of uncertainty for each client—about the nature of threats, the effectiveness of mitigations, and their client's capabilities [69]. Different providers make different choices about how to approach these risks: as Cuomo et al. write in the Technology Abuse Clinic Toolkit, readers planning advice clinics should "adapt the content to their communities and locales as needed" [21]. While we know that providers do adapt to meet these challenges, there is little documented about *how* they do so on a day-to-day basis, including what information they seek about clients, how they learn that information, and how this affects the advice they give. Systematizing these practices could provide a foundation for those interested in empirically studying, improving, and adapting advice.

Researchers and technologists can also address challenges more directly. These efforts are ongoing: the Citizen Lab, the Clinic to End Tech Abuse (CETA), and other groups publish findings on the threat actors and technologies facing at-risk users [29, 22], while tools such as VeraCrypt [3] and iPhone's Safety Check [6] are sometimes recommended to or designed in part for at-risk users. However, we have limited knowledge of how providers feel these efforts affect the practical challenges they face in tailoring advice. Their perspectives could inform research and technology design on how to better support providers.

Finally, self-service tools that deliver tailored advice straight to users—in the vein of modular guides like Surveillance Self-Defense [27] and interactive questionnaire-based tools like Security Planner [19]—can complement personalized consultations and trainings by reaching users at scale who are not able or motivated to meet with providers. Though self-service tools introduce additional challenges (e.g., lacking real-time human guardrails) that may complicate or dissuade their use for high-risk cases, they typically still tailor advice based on risk factors to some extent. However, we have limited understanding of how they should account for risk to provide better advice without causing more harm than good. Based on their direct experience, providers may have useful insights.

To address these gaps, we interviewed 18 advice providers who work with a variety of at-risk populations, including human rights

[1]For brevity, we use *advice* to refer specifically to privacy, security, and digital safety advice unless otherwise specified.
[2]We refer to experts who give advice as *providers* and those they advise as *clients*.

Proceedings on Privacy Enhancing Technologies 2026(1), 321–337

https://doi.org/10.56553/popets-2026-0017

defenders (activists, journalists, etc.), targets of tech abuse, and people facing online harassment due to gender, racial, and other identities. Participants are privacy and security researchers engaging in long-term applied work and/or professionals affiliated with advocacy organizations. We ask the following research questions:

RQ1: What kinds of choices do providers make to differentiate advice, and why? We analyze five objectives that providers seek to balance by tailoring the content and delivery of advice. To do so, providers gather various kinds of context about clients, which we characterize. We also identify four types of methods for determining this context, and we explore trade-offs, particularly regarding reliability. Based on this framework, we suggest directions for future work on tailoring advice.

RQ2: How can research and technology design better support providers in tailoring advice? We identify three main areas for improvement: empirical research on threats and countermeasures, better infrastructure for sharing information among providers, and more accessible privacy and security tools. We dive into why providers find existing resources lacking in these areas, what compromises they have made to mitigate these challenges, and how they could be better supported in their critical work.

RQ3: How should self-service tools account for risk factors when tailoring advice? Many participants expressed reservations about high-risk users receiving individually tailored advice from self-service tools. We explore why, focusing on the capabilities they believe would be hard to scale up or automate. Based on their concerns and recommendations, we propose preliminary guidelines on how self-service tools should account for risk factors when providing advice.

2 Background

Our work builds on that of scholars and practitioners who have sought to understand and meet the needs of at-risk users.

Privacy and security for at-risk users. All online users face privacy and security threats, but these are more pronounced for some people due to who they are (e.g., transgender people [61]), what they do (e.g., dissident activists and journalists [44]), or the nature of their relationships (e.g., targets of tech abuse by intimate partners [36]). Prior works have studied the risks, practices, and needs of these and many other at-risk users, including refugees [64] and undocumented immigrants in the U.S. [35], political activists in Sudan [24], and women in South Asia [60]. This at-risk user research is systematized in a framework by Warford et al. [71], who find ten major contextual factors common to people who face elevated digital safety risks, such as marginalization, resource or time constraints, and access to a sensitive resource.

Some have attempted to improve privacy and security for at-risk users: e.g., informed by research on the needs of people with visual impairments [37], researchers designed new tools and interfaces to help them navigate email and Internet browsing more safely [75, 40]. Some universities have organized specialized clinics to provide advice and other digital safety support to targets of tech abuse [36, 30, 69, 68], staff at politically targeted organizations [13], and managers of critical public infrastructure [45]. Some nonprofit organizations also operate similar clinics [1, 57, 59].

Researchers have interviewed some of these advice providers, particularly in the tech abuse sphere. They emphasize the need to provide care, as understood by feminist and other critical theories, in addition to technical support [69, 66]; they also highlight the difficulties of providing support remotely [**tseng21**]. While we share overlapping research goals, our work differs in its focus on providers' practices and challenges in tailoring advice to specific contexts.

Some providers have published guides on setting up new clinics [21, 18] and training new providers [38, 46, 39]. We extend this knowledge by comparing and contrasting perspectives across a greater variety of organizations and contexts. By interviewing practitioners directly, we also avoid relying on published documents, which can represent idealized and abstracted processes.

Guides and advice. To reach more users, organizations have published privacy and security guides. Some are aimed at the general population [50], while others target specific at-risk groups, such as activists [32, 52], journalists [31, 49, 57], and migrant domestic workers [65]. In addition, some guides are interactive, incorporating user input to match advice to specific contexts and priorities; examples include Security Planner [19] (for general users), the Digital First Aid Kit [52] (for human rights defenders), and the SOAP policy generator [9] (for civil society organizations).

Scoping and prioritizing advice is difficult. Several studies find that security experts [55, 54] and advice writers [48] do not agree among themselves on advice prioritization for general users. This pattern holds true in narrower contexts as well: when interviewing 24 subject matter experts on mitigating online hate and harassment, Wei et al. found they agree only loosely on the most important threats for general Internet users to prioritize [72]. Although prior work often poses this as a problem, it may be that prioritizing advice without a particular individual's context is inherently underspecified and/or subjective—a challenge that interactive tools such as Security Planner attempt to address. By interviewing providers who necessarily give advice contextually, our findings contain larger lessons for tailoring advice to individual situations.

Though most privacy and security advice online is aimed at the general population, some targets at-risk users. For example, Boyd et al. examined advice given to Black Lives Matter protesters [11], Geeng et al. interviewed LGBTQ+ adults about their experiences with advice [33], and Schmüser et al. studied advice on Twitter for potential targets of cyberattacks during the 2022 Russian invasion of Ukraine [62].

At the same time, at-risk users do not necessarily obtain privacy and security advice online: Munyendo et al. found that cybercafe customers in Kenya rely on cafe staff for advice and support, though these staff often prioritize convenience over secure practices [47]. There may also be equity gaps in who benefits from online advice, as evidenced by surveys from Redmiles et al. [53] and Coopamootoo and Ng [20].

Social media is one oft-mentioned source of privacy and security advice. Platforms like Twitter host discussions on subjects such as passwords [26], and TikTok has featured problematic "anti-security"

advice detailing how to spy on partners and children, alongside defensive techniques [73]. Engagement with this content varies: Bhagavatula et al. found that users rarely interacted with constructive advice on Facebook and Twitter [10].

Some have imagined generative AI and large language models becoming sources of privacy and security advice, although Chen et al. found these models promoted misconceptions in almost one-third of scenarios [12].

These mixed findings about the evolving state of privacy and security advice highlight the value of creating multiple avenues for advice and simultaneously the harms of doing so without adequate care. Through our interviews with providers, we explore this tension as we try to find ways to build capacity to tailor advice for at-risk users while minimizing harmful outcomes.

3 Method and participants

We conducted 18 semi-structured interviews between March and December 2023 with experts who tailor privacy and security advice for at-risk users. As experienced providers are rare, we conducted interviews on Zoom. Participants received a $60 Tango card.[3]

3.1 Recruitment

We required participants to be 18 years or older, speak English, and have experience working directly with individuals or groups with elevated risks to provide digital privacy, security, or safety advice tailored to their needs. Example settings include

- one-on-one meetings to address tech abuse,
- long-term consultations with dissident media groups, and
- trainings for surveilled activists with time for questions.[4]

By default, we did not count creating advice guides (not direct), providing advice as part of a short-term project such as a single research paper (not enough experience), or providing advice mainly focused on other topics such as mental health (not digital privacy and security). However, none of these was used to definitively exclude participants: instead, we provided guidance for potential participants to gauge their own eligibility, and we invited them to participate if they felt their work fit the goals of our study.

We began recruitment by compiling a list of organizations and providers we believed to be qualified, screening them using online information about advice services. We assembled this list using resource hubs such as CiviCERT [14] and the Coalition Against Online Violence [16]; online searches such as "security advice consultation for journalists," across a variety of phrasings and risk factors; and our own pre-existing knowledge. We also asked participants and relevant researchers, advocates, and organizations to recommend other candidates.

Ultimately, we reached out directly to 48 organizations and providers. We interviewed eight participants from organizations we identified ourselves; ten more were referred. We aimed to recruit a diverse set of privacy and security experts, with cumulative experience across as many risk factors and contexts as possible. However, as Table 1 illustrates, some are more common than others: e.g., many participants have experience with human rights defenders, but we were unable to interview providers with extensive experience tailoring advice for older adults. Conversations with providers suggest this imbalance is due in part to allocation of funds by governments and nonprofits.

3.2 Participant information

Table 1 lists participants. All are privacy and security researchers engaging in long-term applied work and/or professionals affiliated with advocacy organizations recommended by trusted groups such as the Coalition Against Online Violence [14, 16, 17]. Participants provide advice covering a wide array of risk factors, with significant individual variation from client to client. For instance, even if a provider specializes in journalists of a specific ethnicity, they may also advise journalists of other ethnicities, laypeople of the same ethnicity, and other clients from a helpline for online hate and harassment. Clients include both individuals and representatives of small organizations (e.g., an activist or journalist group).[5]

Sixteen participants give advice at least sometimes as part of an organization (fifteen nonprofit or academic organizations and one private business), and two do so only as unaffiliated volunteers or consultants. As P15 and P16 are colleagues, there are fifteen organizations total, and the scale of advice operations varies across them: two have only 1 person giving advice; four have 2–4; six have 5–9; and three have 10+. Some participants are based in a specific region and give advice in that context: six in Africa, two in Asia, and one in Latin America. The other nine are based in the U.S. and Europe, though many advise clients internationally, often focusing on a few specific regions. In total, thirteen participants are female or use the pronouns she/they; five are male or use he/they. Nine are Black, African, or African American; five are White; three are Asian, Native Hawaiian, or Pacific Islander; and one is Hispanic or Latino.

3.3 Interview design

Our semi-structured interviews included five main parts:

- **Background:** Participants' context and experience giving advice, including the setting and goals.
- **Information sought:** What information participants seek about clients and how it affects the advice they give. To develop an in-depth understanding of how each participant tailors advice, we asked them to describe one or two specific instances of giving advice in as much detail as they were comfortable providing, and we asked if they had examples of advice that is good for some clients but bad for others. During this part especially, we asked follow-up questions when we were missing context: e.g., how a participant determined something about a client, or how a client reacted to a recommendation.
- **Advice in a high-risk context:** How participants structure their advice process to meet the needs of particular at-risk clients. To understand how participants' current practices were shaped by their experience as a provider, we asked about

[3] These are redeemable for a variety of gift cards; see https://www.tangocard.com/.
[4] Though we initially did not focus on group trainings, we broadened our scope as we found that most participants also conduct group trainings, in which they find ways to tailor advice following similar strategies.
[5] While differences between these settings would be interesting to explore in future work, this is not our focus, as we observed that providers largely rely on the same toolbox of methods and face similar challenges across both settings.

Table 1: This table lists the risk factors each participant described primarily addressing through advice, as well as their years of experience. In practice, they often advise a broad range of at-risk users, who may not fall cleanly into one category. For example, a provider who focuses on journalists of a specific ethnicity may also advise journalists of other ethnicities, laypeople of the same ethnicity, and a variety of other clients facing online hate and harassment.

	Clients' risk factors	Years
P1	HRD, tech abuse	10+
P2	HRD, gender, LGBT+	10+
P3	HRD, race/ethnicity	10+
P4	HRD, race/ethnicity	10+
P5	Journalists	10+
P6	HRD, tech abuse, race/ethnicity	5–9
P7	HRD, gender, LGBT+, religion	5–9
P8	HRD, race/ethnicity, gender, LGBT+, religion	5–9
P9	HRD, gender, LGBT+, rural	5–9
P10	HRD	5–9
P11	HRD	5–9
P12	Journalists	5–9
P13	HRD, gender, LGBT+	2–4
P14	HRD, gender	2–4
P15	Tech abuse	2–4
P16	Tech abuse	2–4
P17	HRD	0–1
P18	Gender, disability	0–1

HRD = human rights defenders broadly

changes they have made over time, strategies for mitigating potential harm, and reflections on their process.

- **Challenges:** Factors that make it difficult to tailor advice. We also asked them to brainstorm hypothetical strategies and resources to address these challenges.
- **Self-service tools:** Considerations for designing tools that automatically personalize advice. We generically described a tool like Security Planner and asked for considerations for both design and usage. Beyond the specific implications of such tools, these questions also prompted broader conversations about the kinds of tools and resources that providers find useful or counterproductive, as well as reflections on key elements and challenges of their work.

Our interview protocol is in Appendix A. Interviews lasted 57 minutes on average. Participants also completed a one-minute pre-questionnaire on their organization and demographics (Appendix B).

We conducted four pilots to test and refine our interview protocol. One was with a lab member who is an active member of a privacy and security advice clinic; the others were conducted, using a slightly modified protocol, with professionals at schools and nonprofits who give advice for a general audience. Based on these pilots, we focused our research questions specifically on advice for at-risk users.

3.4 Thematic analysis

We transcribed audio recordings of interviews using OpenAI's Whisper model locally, and we corrected transcripts manually. As we aimed to both describe advice providers' practices and explore interpretive themes to understand how and why they adopt those practices, we followed a template analysis approach, which is a form of thematic analysis that does not draw a clear line between descriptive and interpretive coding [41]. The first two authors developed a qualitative codebook as they coded the first three transcripts collaboratively. At that point, the codebook's high-level structure was stable, so the first two authors double-coded the remaining interviews separately, meeting once or twice per interview to resolve differences and discuss themes. Our codebook is linked in Appendix C.

Our analysis was deeply oriented around our research questions, especially RQ1: How do providers differentiate advice, and why? We aimed not only to capture individual answers to (1) what information providers seek about clients, (2) how they learn it, (3) how it impacts advice, and (4) what challenges they face; but also to understand how providers connect these pieces to develop nuanced practices and address specific problems. Thus, we organized codes into hierarchical groups that address different aspects of providers' practices, and the themes we developed usually relate multiple aspects together to explain *why* providers give advice in a certain way. These themes form the narrative of Section 4.

3.5 Ethics

This study was approved by the University of Maryland (UMD) Institutional Review Board. We obtained informed consent, including for automated transcription; data was only stored locally and by third parties that have contractual agreements with UMD to protect privacy. We emphasized to participants that we understood the confidential nature of their work with clients and that they should withhold details as needed. We have limited the reporting of personal information and quotes to reduce the risk of re-identifying participants or their organizations.

3.6 Limitations

As this is a qualitative study with 18 interviews, we refrain from making strong quantitative claims such as about the proportion of providers who follow any given practice. The process of thematic analysis is inherently subjective, requiring researchers to use their own experience and knowledge to develop nuanced interpretations of data; our analysis method with two coders promoted thoroughness and consistency.

Participants' descriptions may have left out aspects of their process that are subconscious or seem obvious, recounted idealized experiences due to social desirability bias, or omitted details deemed too sensitive to share. We followed interview best practices to mitigate these limitations, such as emphasizing that there are no right or wrong answers.

We note that recruitment is a challenge for this topic of research, which focuses on advice providers who are few and far between to begin with. As their work is often sensitive, potential participants may have declined interviews or kept a low profile that prevented us from inviting them. Though participants live and work in several regions of the world, our sample is limited because we conducted interviews in English. And although participants branch out to many different risk factors, most have at least some background in privacy and security for human rights defenders or targets of tech abuse (Table 1). Related works have addressed this challenge by taking an introspective turn and studying the experiences of a single organization [36] or by recruiting domain experts more broadly and

asking them to recommend strategies based on whatever relevant knowledge or experiences they have [72]. We chose a different but complementary approach—interviewing providers with hands-on experience across a wide array of organizations and contexts—so that we can study a variety of tested strategies for tailoring advice to at-risk users without limiting our scope to a narrow subset of risk factors and experiences. Nonetheless, providers in contexts under-represented in our study (e.g., older adults, teachers, sex workers) might make different choices to tailor advice and might experience other challenges.

As we did not ask about education level or technical knowledge, we cannot describe participants' level of technical expertise in a specific and consistent way. However, we believe our screening process (§3.1) establishes a strong-enough baseline to draw reliable insights from participants' rare and precious real-world experience tailoring advice for at-risk users.

4 Findings

We now present results of our thematic analysis. After background on advice format, we describe how participants differentiate advice and how they determine the necessary context. We then explore their challenges and their perceptions of self-service tools.

4.1 Advice format

Providers tailor advice in both one-on-one consultations and group trainings (mentioned, respectively, by 16 and 12 participants).

Personalization happens with individuals and groups. In *consultations*, providers typically meet with a single client—either one person or representatives of one organization. Consultations often include much client input, mostly through conversation, though providers may also send intake forms or conduct organizational audits to gather information beforehand. Sometimes, clients arrive at consultations through a helpline run by the provider's organization, or through a referral from another provider. Consultations often involve clients seeking help with an urgent crisis, though they do include more proactive cases, such as trying to secure devices before traveling to a high-risk region.

In contrast, during *trainings*, providers typically present information simultaneously to multiple clients, sometimes from different organizations. While trainings generally have less room to tailor advice to individuals, providers may still adjust based on information gathered beforehand from intake forms and audits, engage with clients during the training, and take individualized questions and follow-ups afterward.

A wide range of durations meet different client needs. Most of P6's consultations last under an hour, per clients' preferences for an answer to a specific problem rather than a full assessment. P2's consultations with targets of online harassment can last three to four hours to provide care for clients under acute stress. On the other end of the spectrum, P10 compared their process to a camp: after training human rights defenders from an organization for days, they often spend several more days observing them at work and helping to implement advice.

Follow-up strategies also vary widely. Some providers do not follow up at all; others share resources and guides; one participant has an active group chat with former clients; another mentors clients for multiple years.

These differences in engagement strategies highlight a recurring theme in our findings as a whole: determining what advice to give is only one of many tasks when meeting clients. As the rest of this section will explore, participants' priorities can also include teaching clients technical skills, persuading clients to follow advice, and providing emotional support.

4.2 Objectives for tailoring advice

To tailor advice, providers incorporate many different types of context about clients, leading to differences in not only the content of advice but also the pace and process for delivering it. The purposes for which providers use context can be organized under five main objectives: triaging threats, ruling out advice that violates constraints, ensuring that advice is acceptable, avoiding unnecessarily upsetting clients, and upholding providers' own principles. We begin by characterizing these objectives, focusing on why each is important and how context is used to achieve it. Figure 1 contains a conceptual diagram representing different types of context and objectives, along with methods of determining context, described in §4.3. For more detail, Appendix D lists examples of the types of context that providers seek.

Objective 1: Triage and assess clients' threats and risks. For many providers, assessing and communicating clients' threats and risks is the primary task when providing advice. This can be complicated for at-risk users: in addition to threat actors, their capabilities, and their motivations, threat models include other crucial components shaping the likelihood and nature of attacks, such as clients' accounts and devices, privacy and security hygiene, home and family situations, cultural and geographic contexts, and identities.

Triaging active threats and risks is often one of the first steps in providing advice, as urgent threats may alter the pacing and structure of the advice process. For example, P1 normally advises clients who appear overwhelmed to take a gradual, account-by-account approach to adopting password managers and multi-factor authentication (MFA)—but in the case of immediate danger, they make changes to all the relevant accounts together during the consultation. Even when threats and risks are not time-sensitive, understanding them early dramatically reshapes how providers approach the rest of a client engagement. If clients' communications are likely to be targeted, P4 tailors advice to specific smartphone models, working with clients through device settings to ensure data is not stored in the country of their nation-state adversary. In contrast, if clients have less sensitive communications, P4 gives higher-level advice such as recommending the use of secure communication apps and VPNs. In a similar vein, P17 usually focuses on general security hygiene such as password managers and MFA, but when threats escalate to physical acts such as stalking, they instead advise human rights defenders to seek protection from the police and make significant life changes such as quitting their job.

Building a nuanced threat model is crucial to providing advice that protects clients instead of endangering them. P1 described

Figure 1: A conceptual diagram of key decisions providers make to tailor advice. Providers choose from among four main methods to learn different types of context about clients. Providers take this context into account as they attempt to balance five main objectives.

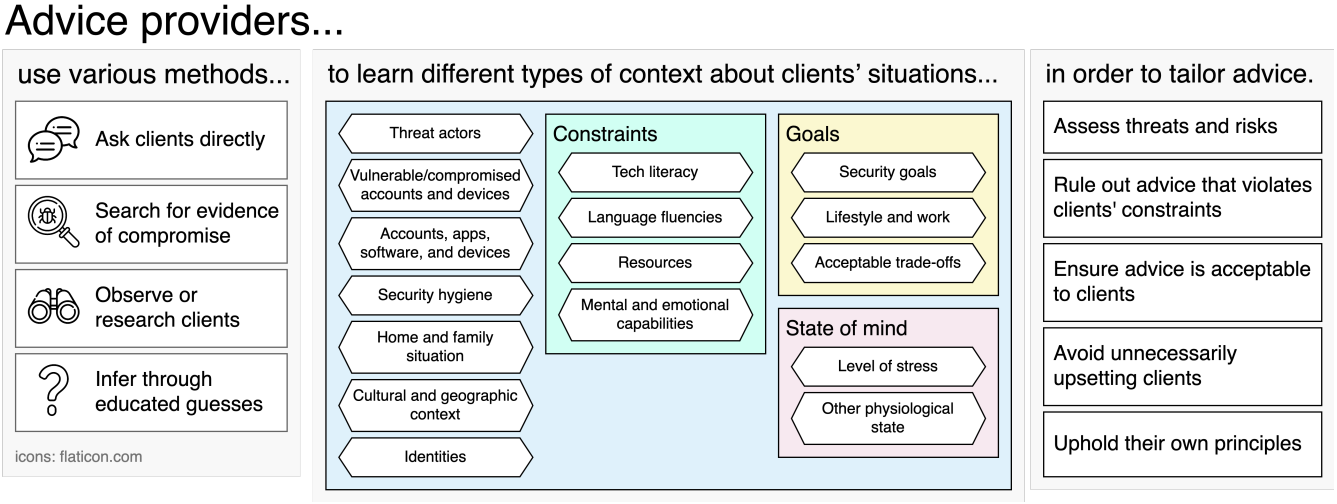

a case in which they did not ask enough questions about how and why a client needed to protect their web browsing, leading them to recommend a mitigation that did not protect the client's IP address against the right adversary: "I didn't ask him if he was optimizing for disguising his address from the website that he was going to, or if he was just optimizing for speed, or if he was optimizing for security against interception in transit. And if I had asked these things, I would have given different advice." Similarly, P12 said that while using multiple phones can be a useful risk management strategy for journalists, there are regions in which having more than one phone can raise suspicion. Participants gave other examples of assessing threats and risk to prevent harmful advice, including not recommending VPNs in areas where VPNs are uncommon, which can raise suspicion, and not reporting security issues to law enforcement in cases where officers themselves might commit violence against clients who are LGBTQ+ or of a persecuted religious minority.

Objective 2: Rule out advice that violates clients' constraints. Clients also face constraints on their capabilities and resources that make certain advice less accessible. While all users need advice that is accessible to them, assessing constraints is a higher priority with at-risk users. This is because broadly accessible advice is often sufficient for lower-risk users, whereas many at-risk users can benefit greatly from advice that is less accessible.

Tech literacy stood out as a key concern, driving some—but not all—providers to differentiate advice. Based on early experiences with clients who were uninterested in advice that seemed too difficult to adapt to, P7's organization developed a needs assessment process to categorize clients into three tiers based on their understanding of digital technologies, their current privacy and security hygiene, and their privacy and security goals. Some advice is given to all clients, such as using Signal, whereas some advice is only given to clients in higher tiers, such as using a VPN, Tor, or PGP.

On the other hand, even though P11 views tech literacy as an important factor, it does not generally change the advice they give: "If their technical competency is a little bit lower, then they still need to use the same tools. It's more that they need a little more guidance, and so the training will be a lot more in depth." Still, P11 said this occasionally leads to inadvertent differentiation in advice, as they may run out of time to cover more advanced tools like firewalls when spending more time on basics.

While tech literacy was a particular focus for many, participants also mentioned other factors limiting clients' ability to carry out advice. Clients' language fluencies, financial resources, and accounts and devices are useful for ruling out tools and resources that are inaccessible or irrelevant: for example, for clients who are freelance journalists with limited money, P12 considers recommending alternatives to paid password managers, such as writing passwords down somewhere secure in the home or even reusing passwords. These findings align with barriers to protective practices identified in prior work focusing on at-risk users [71].

Objective 3: Ensure that advice is acceptable to clients. For various reasons, a client may decide not to adopt advice even though it is technically possible; for this reason, providers often work to prioritize advice that will not be objectionable, and to persuade clients to follow it. Almost all participants emphasized the importance of getting on the same page with clients about their privacy and security goals, the work and other tasks they need to accomplish, and the trade-offs they are willing to make to accomplish these (sometimes conflicting) aims. Without this, providers cannot be sure that their advice will protect clients in ways that clients find acceptable. This is especially important for at-risk users because it can be more difficult for them to achieve a level of privacy and security that feels safe; understanding clients' goals enables providers to find the right balance between safety and acceptability.

Failing to define clients' goals can lead to inaction or unwanted results. P2 described a time they were so personally invested that

they filed reports on their client's behalf, without checking what the client actually wanted: "I basically run over all the red lights, connect them directly to the main source, which is the people on the social media platform. ... I gave my all, and [the client] decided not to do anything about it." While no clear harm came of this, P2 described it as a drain on their own limited capacity. P11 highlighted the consequences of recommending what many privacy and security experts consider to be "right ways of doing things" without regard for clients' existing lifestyles: "[If] people use a particular chat app and you tell people, 'No, that app is trash and surveilled, don't use it,' people will likely not listen to you—number one, because you come off hostile, and number two, because that app is so intertwined to their social life."

Sometimes, clients hold misconceptions that inhibit advice uptake: P1 recalled clients who insisted on using a VPN, even though that would not address their threats, and P10 recalled clients who worried a digital safety app would use more data than it actually does. Another barrier to advice uptake is downplaying privacy and security risks, especially when other risks and priorities seem more immediate: P4 noted that hacked emails don't "really compute as much" when colleagues have been imprisoned or killed.

In these cases, providers may try to understand clients' beliefs about technologies and risks in order to be persuasive. When clients downplay their own level of risk, P8 sometimes advises them to talk with a friend, who will hopefully remind them that their risk is not normal: "The thing that we get used to, it will be weird for others." Several elaborated on how to explain things, noting the need to prioritize information because technical detail can overwhelm clients. P12 said clients need to know the rationale behind advice in order to take it seriously, but overly technical explanations can be confusing and time-consuming: when explaining the difference between messaging apps like Signal, WhatsApp, and Telegram, for example, "they don't care about how encryption works or anything like that; they just want to know, which one should I use and why?" They may connect these rationales back to clients' goals unrelated to privacy and security: P12 points out that by preventing identity theft, MFA protects journalists' ability to make a living.

P11 described another strategy to avoid overwhelming clients with questions and caveats: they give advice based on "general security principles" and trust clients to speak up if their needs differ. For example, they often recommend deleting sensitive messages, and they expect clients to speak up if, for instance, authorities might inspect their phone and see that messages were deleted. Despite the potential to miss important context if clients do not volunteer it, P11 was optimistic about this approach, given their clients' already-high vigilance: "By the time they realize they need security training, they've already established that they have a threat model in mind."

Still, providers' strategies for boosting adoption are far from guarantees. For a client whose email had been hacked, P6 recommended a security key, giving real-world examples of how security keys had protected others, and demonstrating how to use one. Unfortunately, they said, the client did not follow the advice and was hacked again: "I know it's a bet, because it's quite a big change for someone who wasn't very technical. But it was like, 'Look, this has happened to you, and you're at risk of it happening again. This is the right thing if you want to avoid it.' And it didn't work."

Objective 4: Avoid unnecessarily upsetting clients. When at-risk users seek help, they are often already in crisis, prompting some providers to slow down. P15's clients often have an intense emotional reaction upon learning that their device is compromised, which may prompt a pause before taking further action: "It might not actually be the best time to walk them through what to do now, right? [The abuser] having that access for another day or two is not necessarily going to be life-changing, if they want to take time to call back tomorrow." Following the model of psychological first aid [28], P2's first question is to understand clients' state of mind; they explained that jumping straight to changing passwords can trigger panic attacks if clients are not emotionally stable.

Based on clients' state of mind, some providers refrain from giving advice that might prove detrimental to clients' well-being. P15 sometimes refrains from recommending privacy and security hygiene practices such as monitoring login history, if they believe clients will be unduly anxious about benign occurrences such as a login being recorded with an imprecise geographic location. Similarly, when targets of tech abuse are very stressed, P6 avoids talking about threat models and risks, in order to avoid adding to that stress.

Deciding how to account for clients' mental state requires judgment and tact. In many cases, P15 and P6 both will directly address clients' beliefs about threats that appear to be unrealistic. P15 said some targets of tech abuse find it helpful to learn that certain attacks are technically difficult to carry out, and that many abusers overstate their capabilities; P6 noted that although it can be difficult to tell a client, "Look, you're afraid by something, but this thing is not true," doing so has reassured clients in the past. However, when they perceive unrealistic beliefs as particularly strong, both avoid pushing back directly, instead diverting the conversation to general steps that clients can take to protect themselves. These findings echo prior work on providers' strategies for navigating clients' trauma in the tech abuse domain [51].

Objective 5: Uphold one's own principles. Providers' choices are necessarily mediated by subjective values and priorities belonging to them and/or their community. Two main core principles emerged from some participants: prioritizing clients' goals, and empowering clients.

Many believe it is a matter of respect to prioritize clients' privacy and security goals, even at the expense of giving advice that could provide stronger protection. Despite describing the CETA's systematic protocols [15] as "awesome," P6 eschews them out of respect for clients' priorities: "People can come to me because they are having an issue. And so if I put [an assessment-based] process first, instead of their concern, then it's going to be dismissive of their concern. ... They don't want full assessment; they want a solution to their problem." Similarly, P1 said that securing a single account or device is "often enough" for clients to meet their most essential needs.

Clients' goals also play an elevated role for some participants who prioritize *empowering* their clients. P1 is against advising clients to make significant sacrifices such as lowering their public profile, as "that's just blaming the victim." Others stated that it was unacceptable to let fear stop human rights defenders from carrying on their work: P4 summed up their goal as "How do

you make sure that people become empowered, rather than self-censor themselves?" while P8 said, "People should [reduce their technology usage] if they want to because it's good for their mental health or . . . because of a life choice, not fear. But the fact that people are doing it because of fear—that's not okay, and that's not empowering." Not all providers share this principle; P2 and P17 both mentioned that they sometimes advise clients to consider quitting their work due to the risks.

4.3 Methods of determining context

Learning useful and reliable context about clients is not always straightforward: it can take time, trigger distress, and still result in incomplete or inaccurate understanding. In this section we detail four different approaches to determining context about clients, each with its own trade-offs: asking directly, searching for evidence of compromise, observing and researching clients, and inferring information. We observe that managing the reliability of information is a key factor that providers consider when deciding which methods to use and how to apply them.

Method 1: Asking directly is common but can provide unreliable information. Asking clients direct questions is a ubiquitous method of gathering information. Specific, focused questions can target key details, while open-ended questions can cause unexpected, important information to surface organically. Open-ended questions can also serve a therapeutic purpose: P1 said, "Often survivors really need somebody who will sit there and just patiently listen to them tell their story, and to believe them." However, open-ended questions come with a natural trade-off in time spent: P2 said, "The first question that you always ask is, 'How are you?' . . . And probably you'll be 40 minutes receiving the story."

Several participants often send an intake form to clients before meeting. Despite low response rates, these can help providers prepare for consultations. P10 even changes their questionnaire based on what information they think will be most relevant for a particular client: for clients in rural areas, they may add questions about Internet connectivity. On the other hand, P12 limits their questionnaire to asking about clients' privacy and security concerns, because they sense that clients dislike filling out forms.

Providers worry about the reliability of answers to direct questions. When describing threats, clients may lack details or jump to unlikely conclusions. Some participants attributed clients' belief in infeasible or unlikely attacks to stress or trauma: as P1 put it, "If you don't have a lot of technical knowledge and you have been abused, it's very easy to feel like your abuser is capable of anything and everything." On the other hand, clients may underestimate threats: P3 said, "The person you're talking to is often wrong, and human beings tend to rationalize down threats," and P11 said a woman of color who faces regular harassment might acclimate and downplay their level of risk. Self-reported tech literacy can also be particularly unreliable: P7 mentioned that clients who report being ready for the highest tier of protective measures sometimes lack the basics, such as strong and unique passwords.

Clients may even intentionally stretch or omit parts of the truth for various reasons. P10 said clients have exaggerated their technical knowledge on an intake form, out of a misguided fear that low-tech clients would be deprioritized for consultations. And P15 always revisits questions from their intake form at the start of a consultation, because some clients are uncomfortable documenting certain information.

Method 2: Searching for evidence of compromise can provide unequivocal results, but uncertainty causes issues. When providers need more information about threats, a natural option is to search for direct evidence of account or device compromise, via manual inspection or using tools such as iPhone's Safety Check [6]. For example, when a client attributed glitchy phone behavior to surveillance by an abuser, P6 checked all the apps on her phone for evidence of compromise, ultimately concluding that said behavior was probably due to ordinary technical problems. Evidence can also be non-technical: P15 described sometimes doing "detective work," asking clients to list things an abuser shouldn't have known about (e.g., an upcoming vacation) and where clients had discussed them (e.g., over email or text messages), in order to deduce potential compromises.

Evidence of compromise undeniably helps to mitigate threats. However, searching often provides unclear answers, wasting time and even causing distress. P16, for example, hesitates to check for spyware, which is often "not very provable" but can make clients "really freaked out." Remote consultations add complications, as providers cannot inspect clients' devices themselves. P7 said existing workarounds, such as instructing clients to run scans on their devices, are not good enough: "[Unless] we can technically look at the device, it's very difficult for us to determine what is actually going on."

Method 3: Observing and researching clients is time-consuming but provides information that providers trust. Some participants learn about clients by observing them at work. This is time-consuming: after giving initial advice, P10 often spends another day or two with at-risk organizations to understand daily life and provide support in implementing advice. However, this time investment can yield unique insights: P10 once realized only after visiting that clients were running cracked, unpatched anti-virus programs and operating systems.

Sometimes, an alternative is to create opportunities to observe clients during consultations. To assess tech literacy, P3 recounted having a client try out different encryption tools, starting by using VeraCrypt to create an encrypted hidden volume on a USB drive: "It's all the steps that a normal person should have problems with, and I'm banking that you will mess up." If they struggled, P3 would teach them a simpler alternative, such as the GUI-based Cryptomator; if they still struggled, P3 would advise them to acquire a password-protected hard drive as the simplest option.

Some participants mentioned researching clients' identities or background beforehand. This can provide context about needs and threats without taking clients' time or triggering additional distress, and it can unearth information clients do not know about their digital footprint. When a potential client reaches out, P2 may vet them with their network of providers. In addition to learning valuable context, this can also inform whether to take the case; P2 described turning away a potential client who had already sought help from other experts and had "filled the space with false claims for a very long time." Vetting clients also protects providers from malicious

clients. P2 once realized one of their clients was malicious only after consulting their provider network, saying, "It was very scary to talk to the perpetrator and see how the perpetrator was using the system also to get information [about providers]." Similarly, P13 tries to vet clients in case they are secretly with the government, as it is illegal to give advice such as using VPNs; however, they noted the dilemma posed by clients with urgent issues, which forces them to balance their own safety with a timely response.

Vetting clients creates a privacy trade-off, especially if it reveals information to third parties, and it can clash with some providers' attempts to minimize information they learn about clients. P5 establishes boundaries to protect clients' privacy, "Your duty of care as a practitioner is not to pry too much into exactly what they're doing unless they offer it—unless that is 100% germane to the type of advice you're trying to give. . . . I'm not interested to know what it is that you're working on, who you are working with."

P2, for whom it is a "basic feminist value" to show clients that they are believed, said checking clients' claims with the provider network helps address serious concerns about reliability without being skeptical directly to clients. This contrasts with others, such as P11, who try to actively dispell seemingly inaccurate beliefs: "I really try to soothe that kind of wild paranoia that people can sometimes spiral into." We note that although these approaches conflict, both are motivated by care for clients' emotional health. In many cases, providers try to strike a middle ground, showing empathy and understanding for clients' unrealistic beliefs but cautiously pushing back against them if clients seem open.

Method 4: Inference fills gaps but is error-prone. Finally, providers sometimes infer context about clients—making educated guesses and generalizations based on other characteristics. This can be unreliable, but it enables creative solutions to practical problems. For example, as it can be tricky to evaluate tech literacy directly, P3 sometimes estimates it by asking questions ranging from experience with programming and video games to attitudes toward technology, while P9 assumes that clients in rural areas of their country will have lower tech literacy and usage.

Inference can also help fill in context that is simply missing. P4 sometimes needs to learn about a client's contacts without interacting with those contacts, in order to recommend an appropriate secure communication setup. They do so by making assumptions about contacts' tech literacy and threats based on their education level and home region: "A lot of [regional residents] who are from that area with similar backgrounds will have some kind of similarities. . . . You can't gauge the technical skills, you can't gauge the exact risk sometimes—but you can get a bit of an approximation."

Note: Balancing consistency and flexibility is hard. Some providers prefer to follow well-defined protocols to ensure thoroughness and avoid mistakes. In high-risk cases, P3 uses a branching checklist to avoid relying on memory, "[even for] the things that I do literally almost every day—I still use a checklist to make sure we've gone through it." Similarly, P8 uses a protocol so "it's not an ad hoc approach." P1 generally forgoes a checklist with clients in the interest of saving time, but uses one when teaching others to give advice because "that's extremely important for replicating one's methodology." P12, who doesn't use a documented protocol, perceives the lack of standard processes and questions for receiving information about clients as "one of the issues with digital security."

At the same time, some participants emphasized the importance of flexible protocols. P5 stressed that providers should use their knowledge to go beyond their organization's documented protocols: "I want to avoid where people feel that they can just go to a checklist and feel like they've gotten it all." P12 worries a protocol "might sound a little bit robotic," undermining human connection.

Standard protocols can help providers collaborate efficiently and responsibly. P2's helpline uses a "very specific case filing format" to enable delegation of cases from one responder to another without asking clients to re-explain anything: "The idea is never to ask the person to tell the story twice, because that is re-victimizing." However, sharing cases between multiple providers does come with other downsides that may not be addressed by standardizing protocols: P12 once referred a client to another organization for advice, but the client found it "very disorienting" because "she just kept getting different people all the time on their helpline."

4.4 Opportunities for research and tool design

Previous sections have discussed how participants' advice processes are designed to address various challenges. Here we explore in greater detail the challenges that participants attribute to limitations in current research and technology design. We also discuss the consequences that result and the compromises that participants have developed to mitigate these challenges.

More empirical research on threats and countermeasures would empower risk assessment. Tailoring advice is challenging with incomplete information. Even beyond challenges to obtaining reliable context about individual clients (§4.3), a lack of empirical research can make it difficult to assess risks and countermeasures.

Participants expressed a need for better knowledge of how specific technologies work. This includes privacy and security properties and vulnerabilities of technologies their clients use: one participant noted the "completely different ecosystem" of Chinese devices and apps, raising questions such as how Chinese messaging apps surveil text and voice communications, and how different Chinese Android phones monitor side-loading of off-store apps. Inconclusive answers to these questions distort this participant's advice away from their ideal; they err on the side of caution, advising higher-risk clients to use iPhones even if that isn't their client's preference. Knowledge gaps also include the behavior of specific offensive technologies, especially new and fast-changing ones. P12 elaborated on how unknowns regarding spyware limits the advice they can give: "I always tell the journalists, we're not 100% sure that this [solution] will work, so you should just operate under the assumption that everything you do on your device could be compromised at some point." They noted that reports by groups like Citizen Lab can be helpful but do not solve the problem of uncertainty: "[We think] turning off and turning on your phone again could erase Pegasus from your device, but I don't think anyone's 100% sure of that."

Participants also expressed a need for better knowledge of the behavior of threat actors in practice. For example, regarding physical tracking devices such as AirTags, P1 wants more research on

"actual cases in which these things are being used for harm, the ways in which they're being used for harm, and whether or not the mitigations that these companies are implementing work." They see this as an impactful area for potential research, as more knowledge could not only lead to better advice but also help pressure companies and inform technological standards as physical trackers become more common. Better knowledge of threats in practice would help providers with the notedly difficult task of predicting future threats for the sake of prioritizing advice. As P6 said, their key task when gathering information from clients is to "identify scenarios of surveillance and try to estimate how likely they are"; while they can draw on studies showing that, for example, phone hacking is not common, they noted that "there are no clear numbers." P3 summed it up by explaining that their process involves "gambling" on future attacks: "We've got to hope that's the thing that hits you first, the one that you're protected against."

Privacy and security tools that match clients' needs more closely would give providers flexibility. Providers also struggle to provide advice when available mitigations are insufficient in some way. Some privacy and security tools pose usability barriers: P7 commonly recommends the password manager KeePassXC and the encryption software BitLocker, but they noted that many clients find the user interfaces difficult. In contrast, they noted that Signal is "very easy to use," and that encrypted email service Proton Mail is a usability improvement over their previous recommendation for clients to generate their own keys and use PGP. Lack of localization is a related barrier that can require both translation and accounting for cultural context. Translation and cultural context may go hand in hand: P4 explained that important privacy and security concepts such as encryption do not always have a clear translation in local languages. Other barriers mentioned include tools' exceeding clients' budgets, dependency on a stable Internet connection, and unavailability due to regional or organizational restrictions.

Privacy and security tools also sometimes lack desired capabilities. For example, participants called out the physical tracker Tile's anti-stalking mitigations as ineffective, email encryption tool Mailvelope as buggy, and VPNs as too distinctive in areas with low VPN usage. Similarly, services intended to provide clients recourse can be unreliable: several participants noted that social media platforms often fail to take down harmful content affecting their clients, and law enforcement agencies often fail to respond to technology-facilitated harms. One participant expressed significant frustration, arguing that social media companies such as Meta and Twitter overlook crucial cultural context when they refuse to take down, for example, photos or videos outing a transgender person in countries with particularly high levels of anti-trans violence.

Providers' strategies to address these challenges include providing translations, offering vouchers to help pay for privacy and security tools, and facilitating communication with companies and organizations. However, it is often infeasible to overcome the current limitations of tools and resources. In drastic cases—for example, when the risk of physical violence is high—some providers suggest dramatic lifestyle changes, such as quitting a job with a public profile. As P2 said, "There's always the possibility of turning the device off and walking away. People need to know that they're in control." On the other hand, providers may consider it a success to provide validation and reassurance, even if they cannot solve the underlying problem. As P6 summarized, "Sometimes that's even I would say the key part of the help; the advice is even less important than them having a space to talk about the concerns." P2 elaborated, "99.9% of the cases, going to the police is never going to work. But for some people, they need it because they just have to do something. . . . [They go] into the tribunal or into the police station, and we are with them on the phone. Sometimes people need to feel that they are doing everything they can."

Providing non-solution-based support takes time, and not all providers agree about its value. P15 said they struggle to justify providing "comfort as a service," given limited resources: "We'll have clients that are waiting weeks, and they have really urgent concerns. Is that the right time to walk someone through how to set up their router more safely, when there's no sign that's actually going to improve their material safety, but it might provide them a lot of benefit or comfort?"

Better information infrastructure would help providers keep up with changing threats and circumstances. While participants have their areas of expertise, each client brings their own configuration of circumstances and risk factors, some of which may be less familiar. This is compounded by the shifting nature of technologies and threat landscapes, as providers must take into account the potential for defenses and knowledge to become outdated.

Participants described a constant, patchwork process to fill in personal knowledge gaps. P11 constantly scours different sources: "Lots of it is just kind of staying abreast of the latest research that comes out, talking with friends who are also doing this kind of work. . . . Mostly it's kind of informal gathering." As an example, they explained, "I know lots of folks—we chat online and in person—and if they know that Company X actually just got a contract with Spyware Firm Y, then that's a big data point." P5 stressed the importance of learning from peers in spaces such as journalism and technology conferences, which was disrupted by the COVID-19 pandemic: "Giving us the space to exchange ideas and to create opportunities to connect is so vital for this work to continue." When dealing with an unfamiliar region, P3 works with the client to build a "dossier" on threats, asking questions such as "What's the best and worst thing that you remember or know historically happening to people like you in this place?" This can be helpful, in a limited scope: "People who are dealing with the worst of things, they have a lot of knowledge on how to keep themselves safe. . . . They maybe don't know the technical tool exactly, but they have some solutions."

Participants noted downsides to their current strategies for staying informed, often related to the amount of work required. In the area of gender-based violence, P6 said CETA is one of few organizations that publish any details about observed digital surveillance. P6 feels that the lack of wider information sharing impedes their ability to diagnose clients' issues, forcing them to fall back to trial and error: "I have to find my own idea of what's most common. . . . In some cases, I couldn't really find a scenario that matched everything, and things were not clear—but I said, 'Okay, let's do this first step, like securing your emails, and then see if that initial evidence that there is digital surveillance is continuing.'"

As a potential solution, several participants expressed interest in some form of central resource or repository for sharing information with and among providers. For example, P13 wishes there were a unified "pool of information," because they feel that every organization and every provider independently comes up with and updates their own materials and protocols. Similarly, P8 "dreams" of a database listing resources such as privacy and security tools and services, filterable by criteria such as country, cost, and use. This could improve on P8's current system for keeping track of tools and services, which depends on their memory and on documentation regularly maintained at significant effort by their organization.

Participants also expressed that increasing funding for advice providers is crucial to solve this and other capacity-related challenges. In some cases, such as P8's, providers do not have any funding at all to do this work and are involved out of a sense of ethical duty: "We do it because we really cannot say no to people, and people are in distress. . . . We don't actually have funding for it. And it is emotionally taxing, and it's really rough to do this work." P1 added that in many organizations that offer clinics, "There's like one person doing the actual work with clients, and they are burned out and poorly paid. So, step one is fundraising, because this is very difficult and exhausting work."

4.5 Perspectives on self-service tools

We prompted participants to consider self-service tools that personalize advice for at-risk users. While expressing significant concerns about the potential for harm in high-risk cases, participants also provided insights on how self-service tools could account for risk factors in appropriate contexts.

Providers are concerned that critical aspects of their advice process are fundamentally difficult to scale up or automate. Several are skeptical that self-service tools would be as thorough as humans, potentially leading to unsuitable advice. As P12 pointed out, deciding whether to recommend a password manager to journalists can depend on many nuances, including the potential for spyware on their device, risk of arrest, and travel patterns. In greater detail, P4 explained that consultations with human rights defenders can start with basic questions about devices and threat actors, but questions need to then build on previous responses: for example, if an American client is being targeted by the U.S. government, a provider might then ask which telecommunication companies they use, as different companies may provide different protections against the specific adversary. While a tool could handle the basic questions well, P4 said, the follow-up questions would quickly become highly individualized and less feasible: "There are certain questions that are beyond a survey, beyond a structure." P4 added that human providers are not perfect at being thorough, but they felt more comfortable with that responsibility when dealing with clients one by one, rather than deploying a tool at scale.

Some providers are concerned that self-service tools would grow out of date with the rapidly changing landscape of privacy and security advice. P11 noted that over the span of five years, as a result of breaches, LastPass went from the most widely recommended password manager to being perceived as the worst. While a human provider keeping up with the latest would know to update their recommendations, tools need to be proactively updated, which is not always supported by the typical funding life cycle: "Oftentimes when these projects start, they get this nice box of funding in the beginning and everything is great. . . . Over the years, they basically rot. They start breaking down, and there's no longer that funding available to maintain them, to bring them up to date."

Finally, self-service tools lack the same ability to tailor advice to clients' emotional state and provide reassurance. P6 noted the human connection as a crucial distinction between self-service tools and human providers: "If you work with survivors of domestic violence, really talking to people I think is a key part of it. And sometime I feel I give more help by listening to their concerns than by really providing them meaningful advice. . . . I don't see how you can do that in a tool."

Providers have recommendations for how self-service tools should be designed and used. Above all, some providers urge an attitude of humility: P3 said, "I think designers [of self-service tools for at-risk users] need to know that it's a fool's errand, and it won't work well, but they should still try to make it, right? So, if they start thinking, 'We can actually properly solve this,' they're already the wrong people. They need to be like, if this just helps one person and harms another person the least, then it's okay." Similarly, some emphasize the need to make users aware of tool limitations. As self-service tools may grow out of date, P11 suggested they should not only list when content was last updated but also actively warn users when appropriate: "This material was last updated in 2017; that's now more than five years out of date. You should probably seek updated advice." Since the provided advice may be wrong to begin with, P3 recommended having "a disclaimer on every single page that says, 'If you're extremely high risk, we do not recommend you use this, because it could be wrong.'"

To avoid harms, several providers suggested alternative use cases for self-service tools. P4 suggested that Security Planner would be a great way for some of their clients—members of a government-oppressed minority—to get general security advice, but using such tools to escape censorship or targeted surveillance would be too risky. Some believe advice personalization tools are better suited for providers than clients. As P1 explained, "If you are [a client] at high risk, then you are essentially engaging in a crapshoot. Those might be the appropriate advice, but you don't necessarily know, and the consequences of getting it wrong are high. . . . That is why you need the interpretive layer [of a provider]." To that end, P1 envisioned a provider walking a client through the recommended advice, while adding their own filtering and nuanced interpretation.

5 Discussion

In this section, we discuss the two main contributions of our work. In §4.2 and 4.3, we characterized approaches to personalizing advice for at-risk users employed by experienced privacy and security experts, alongside the underlying rationales; here in §5.1, we propose directions for future work on tailoring advice. In §4.4 and 4.5, we explored existing challenges to tailoring advice and potential solutions; here in §5.2, we develop recommendations and guidelines for researchers and technologists based on providers' perspectives.

5.1 A framework for tailoring advice

Our work provides a framework for understanding how privacy and security advice is tailored for at-risk users. We characterize the information that providers seek about clients, the methods they use to obtain that information and manage its reliability, and the objectives they seek to balance by tailoring advice based on context. Our framework captures how providers tailor not only what advice to give but also how to give it, echoing prior work that emphasizes the value of providing care to clients [69, 66]. While advice delivery is arguably important for all users, participants illustrated how it takes on greater urgency with at-risk users, who are often under stress or in crisis, need more or more complex advice, or face such a strong threat that they downplay the potential benefit of advice. We find that providers think deeply and strategically about how to increase the likelihood of successful advice adoption, including varying the pacing of advice, the amount of explanation, the arguments for adoption, and more.

We believe this framework provides a level of abstraction that could enable future work to describe and compare different providers' advice processes. For example, providers may seek different types of context, in different orders, and using different methods to obtain that information. Our study focused on characterizing providers' most important strategies for tailoring advice; other work could, for example, ask providers to rank different types of context by priority, or conduct case studies in the field to document advice processes in practice. Understanding differences between providers' practices under this framework could help improve existing processes as well as adapt tailored advice to new contexts. This framework could also be used to compare self-service tools' workflows to those of other tools or human providers.

Future work can also investigate how providers navigate inherent tensions and trade-offs as they attempt to balance multiple competing objectives. For example, participants expressed conflicting perspectives on whether or not it is appropriate to suggest dramatic changes to clients' lifestyle and work in order to lower their risk, revealing a tension between client safety (Objective 1: triaging threats) and empowerment (Objective 5: upholding providers' principles). Contradictory approaches may drive at the same goal: providers who proactively address perceived misconceptions and providers who believe clients' accounts of events as a matter of principle (again, a tension between Objectives 1 and 5) both do so ultimately to care for clients' emotional well-being. Future work could build on ours by investigating not only how providers navigate specific tensions between objectives but also outcomes and clients' perspectives, in order to develop guidelines for when different approaches may be appropriate.

These tensions arise in part because tailoring advice remains an unsolved problem: we still have no good solutions for many threats, and only very limited empirical information on the relative effectiveness of different tools and techniques in different circumstances, meaning that making recommendations to at-risk users remains murky at best. Even beyond these limitations, however, we find that tailoring advice fundamentally lacks one-size-fits-all solutions. Providers bring their personal priorities and values to what is in large part care work with an inherently human dimension. Perhaps this finding—that advice prioritization is based not only on technical questions but also on deeply subjective inputs such as providers' values and clients' definitions of what privacy, security, or safety mean to them—sheds light on why experts do not agree on advice prioritization even outside the at-risk context [54].

5.2 Challenges and solutions

There are opportunities for a variety of stakeholders to help address providers' challenges in tailoring privacy and security advice.

Researchers can help providers understand evolving threats and technologies. While providers face many challenges, limited knowledge was the one for which the most participants expressed a desire for external help. Those seeking to fund or conduct research to fill knowledge gaps should look to reports by groups such as Citizen Lab [29] and Amnesty Tech [4], both of which participants mentioned, as a potential model.

One useful topic for research is better understanding privacy and security for specific categories of tools and platforms, such as Chinese Android phones or mitigations against physical trackers and spyware. Participants also felt they lacked knowledge on the behavior of threat actors more broadly, such as whether tech abuse plays out similarly in different countries, and on more journalistic topics such as which companies have business relations with spyware firms. Prior work has analyzed privacy and security across various contexts, such as encryption protocols for Android apps [70], backup mechanisms for MFA apps [34], and women's health apps' data policies [2]. While this is a fantastic start, the authors of a recent literature review on mobile app security in developing countries were surprised by how little research exists in these areas; they note that some app categories are under-studied, and that more research should conduct dynamic (not only static) analyses and develop proof-of-concept attacks [25]. We also urge more longitudinal studies that can keep providers updated as they progress; otherwise, providers may have no way to know whether privacy and security issues uncovered by researchers have been resolved. Examples of existing studies in this vein include longitudinal analyses of app privacy labels [7] and data access requests [42].

Empirical studies to build a better understanding of proper prioritization—i.e., what advice is most likely to benefit a client the most under given circumstances—could help providers who presently feel that they are guessing at what top priorities should be. Significant, causal results based on data might be prohibitively difficult to obtain, especially given ever-shifting threat landscapes, but research could at least help inform reasonable heuristics for prioritizing advice, such as P16's approach of checking for common forms of compromise (email forwarding) before less common forms (installed spyware). Providing a potential starting point, some in other domains have empirically evaluated the outcomes of personalized legal and health advice using administrative data [8], surveys and interviews [43], and other methods [74].

Organizations could collaborate on a central repository to facilitate sharing knowledge and protocols. Not only does the privacy and security community lack knowledge that would help tailor advice to at-risk users, but also providers need to expend significant effort keeping up to date with the knowledge that does

exist. This is especially true when switching contexts, as knowledge about threats to journalists in one part of the world does not fully translate to activists in another region, let alone targets of tech abuse or older adults.

Several participants suggested creating a centralized repository to share knowledge, as well as updated configuration guides, walkthroughs, and advice protocols (§4.4). This repository should be filterable by context such as specific threats and constraints, and it could support advice personalization more directly by guiding providers through more complex workflows (with the expectation that providers would use their own judgment rather than blindly relay the recommended advice to clients). Providers, researchers, and journalists could all contribute to the repository, and technology developers could save providers stress and effort by notifying them of updates to settings and interfaces.

The design—and, importantly, maintenance—of any such repository should be informed by providers' input, and deployment should be undertaken with caution for how it could provide useful information to not only providers but also adversaries. One crucial consideration is regulating the quality and reliability of the content, as a repository runs the risk of becoming clogged up with unreliable, less relevant, or outdated resources. In a worst-case scenario, adversaries could even add misinformation. The repository could include social mechanisms for gauging trustworthiness—e.g., associating content with the real-world identity of its creator, or allowing identified users to vouch for the quality of other content—but we note that some providers may want to use or contribute to the repository anonymously, to protect the safety of themselves and their clients.

Technologists could empower providers and clients by making tools more accessible. Providers sometimes struggle to obtain complex technical information from clients quickly and reliably, especially in remote consultations. While some diagnostic tools such as ISDi (for identifying spyware on mobile devices) exist [36], they are arguably better suited for in-person use. Building diagnostic tools natively into devices to detect spyware, cracked software, and other red flags—like iPhone's Safety Check, which summarizes users' interpersonal data sharing [6]—is one way to make these tools more accessible.

Privacy and security tools for clients also could be improved. Despite increasing awareness and consideration of usability, advice providers still do not find such tools usable enough, often differentiating advice for clients with lower tech literacy. Encryption tools in particular remain a usability challenge; participants specifically called out VeraCrypt and Cryptomator (encrypted storage), as well as Mailvelope (encrypted email), as tools they sometimes avoid recommending to clients with lower tech literacy. Many tools and resources also have limited utility due to being available only in English or a small number of other languages; translation would significantly broaden advice options for many providers working outside of the English-speaking world. Researchers have studied the usability of some relevant tools for at-risk clients, ranging from Mailvelope [58] and Signal [63] to account security interfaces [23]. This work has generated recommendations for improving tool accessibility; as improvements are made, more usability studies should be conducted on these and other tools. We note that identifying usability issues and even implementing fixes is not always enough, though, as these efforts need to grapple with how to approach and communicate the problem of usability in a way that is salient to other stakeholders: in a project that aimed to improve VeraCrypt, researchers developed usability improvements that were ultimately not adopted due to pushback and suspicion from the open-source community [56].

We propose guidelines for how self-service tools should account for risk factors. Self-service tools may not be a good fit for users facing the highest risks, given high stakes, dynamic situations, and an inability to adapt and fill in knowledge gaps on the fly. However, these tools do and should take risk into account to some extent when tailoring advice for lower-risk users. We propose three preliminary guidelines based on participants' concerns about self-service tools. First, providers are concerned about the capacity of tools to assess context thoroughly when small nuances can be the difference between useful and harmful advice; therefore, if users indicate they have serious risk factors, self-service tools should accompany advice where relevant with strong caveats and examples of situations where the advice could be counterproductive. Next, providers are concerned about tools reflecting outdated information; therefore, self-service tools should not only clearly display when they were last updated or reviewed (and how substantial that update or review was) but also proactively notify users that content may be outdated if a certain amount of time has passed. Finally, providers are concerned that tools miss a crucial human connection; therefore, if users indicate they are distressed, self-service tools should recommend talking to a supportive human, suggesting expert providers if at all possible.

6 Conclusion

Through 18 interviews with providers of privacy and security advice for at-risk users, we characterized processes for tailoring advice, as summarized in Figure 1. In §4.2, we describe how providers use different types of context to achieve key objectives. In §4.3, we examine different methods for determining this context and their trade-offs. In §4.4, we describe opportunities for researchers and technologists to address challenges that providers face. Finally, in §4.5, we explore providers' concerns and recommendations for tools that automatically personalize advice.

We hope our characterization of advice processes offers practical insights for providers and researchers who carry out and study personalization of privacy and security advice, as well as for experts creating self-service resources. There are many paths forward to support the tailoring of privacy, security, and digital safety advice; by understanding current challenges and lessons learned in practice, we are laying the groundwork to provide more users facing diverse sets of risk factors with the help they need.

Acknowledgments

This work was supported in part by DARPA grant HR00112010011 and in part by a gift from Google. We thank our participants, the individuals who helped with recruitment, and our reviewers (particularly our revision editor).

References

[1] Access Now. [n. d.] Digital Security Helpline. https://www.accessnow.org/help/. ().

[2] Najd Alfawzan, Markus Christen, Giovanni Spitale, and Nikola Biller-Andorno. 2022. Privacy, data sharing, and data security policies of women's mHealth apps: scoping review and content analysis. *JMIR mHealth and uHealth*, 10, 5, (May 2022). https://doi.org/10.2196/33735. Retrieved Aug. 20, 2025 from.

[3] AM Crypto. [n. d.] VeraCrypt. https://veracrypt.io/en/Home.html. (). Retrieved Aug. 15, 2025 from.

[4] Amnesty Tech. 2021. Forensic methodology report: how to catch NSO Group's Pegasus. https://www.amnesty.org/en/documents/doc10/4487/2021/en/. (July 2021). Retrieved Feb. 8, 2024 from.

[5] Sarah Aoun. 2023. Working on the frontlines: privacy and security with vulnerable populations. https://www.usenix.org/conference/enigma2023/presentation/aoun. (2023). Retrieved Jan. 17, 2024 from.

[6] Apple. [n. d.] How Safety Check on iPhone works to keep you safe. https://support.apple.com/guide/personal-safety/how-safety-check-works-ips2aad835e1/web. ().

[7] David G. Balash, Mir Masood Ali, Monica Kodwani, Xiaoyuan Wu, Chris Kanich, and Adam J. Aviv. 2025. Longitudinal analysis of privacy labels in the Apple App Store. https://doi.org/10.48550/arXiv.2206.02658. (Apr. 2025). Retrieved Aug. 20, 2025 from.

[8] Nigel J. Balmer, Marisol Smith, Catrina Denvir, and Ash Patel. 2012. Just a phone call away: is telephone advice enough? *Journal of Social Welfare and Family Law*, 34, 1. https://doi.org/10.1080/09649069.2012.675465. Retrieved Jan. 22, 2025 from.

[9] Gem Barrett. 2020. SOAP: Securing Organizations with Automated Policymaking. https://usesoap.app/. (2020).

[10] Sruti Bhagavatula, Lujo Bauer, and Apu Kapadia. 2022. "Adulthood is trying each of the same six passwords that you use for everything": the scarcity and ambiguity of security advice on social media. *Proceedings of the ACM on Human-Computer Interaction*, 6, CSCW2, (Nov. 2022). https://doi.org/10.1145/3555154.

[11] Maia J. Boyd, Jamar L. Sullivan Jr., Marshini Chetty, and Blase Ur. 2021. Understanding the security and privacy advice given to Black Lives Matter protesters. In *CHI '21*. https://doi.org/10.1145/3411764.3445061. (May 2021). ISBN: 978-1-4503-8096-6. Retrieved Aug. 15, 2023 from.

[12] Yufan Chen, Arjun Arunasalam, and Z. Berkay Celik. 2023. Can large language models provide security & privacy advice? Measuring the ability of LLMs to refute misconceptions. In *Computer Science Applications Conference*. https://doi.org/10.1145/3627106.3627196. ISBN: 979-8-4007-0886-2.

[13] Citizen Clinic. [n. d.] https://www.citizenclinic.io/. ().

[14] CiviCERT. [n. d.] Computer Incident Response Center for Civil Society. https://www.civicert.org/. ().

[15] Clinic to End Tech Abuse. [n. d.] Step-by-step how-to guides. https://www.ceta.tech.cornell.edu/resources. ().

[16] Coalition Against Online Violence. [n. d.] About the Coalition Against Online Violence. https://onlineviolenceresponsehub.org/about-the-coalition-against-online-violence. (). Retrieved Aug. 11, 2025 from.

[17] Coalition Against Stalkerware. [n. d.] Resources. https://stopstalkerware.org/resources/. ().

[18] Consortium of Cybersecurity Clinics. 2023. Clinic Development Toolkit. https://cybersecurityclinics.org/wp-content/uploads/2023/06/CCDS-Clinic-Development-Toolkit-2023.pdf. (June 2023).

[19] Consumer Reports. [n. d.] Security Planner. https://securityplanner.consumerreports.org/. ().

[20] Kovila P.L. Coopamootoo and Magdalene Ng. 2023. "Un-equal online safety?" A gender analysis of security and privacy protection advice and behaviour patterns. In *USENIX Security '23*. https://www.usenix.org/conference/usenixsecurity23/presentation/coopamootoo. (Aug. 2023). ISBN: 978-1-939133-37-3.

[21] Dana Cuomo, Nicola Dell, Alana Ramjit, and Thomas Ristenpart. 2023. The Technology Abuse Clinic Toolkit. https://www.techabuseclinics.org/the-toolkit. (Apr. 2023).

[22] Alaa Daffalla, Arkaprabha Bhattacharya, Rahul Chatterjee, Nicola Dell, Rosanna Bellini, and Thomas Ristenpart. 2025. A framework for abusability analysis: the case of passkeys in interpersonal threat models. In *USENIX Security '25*. https://www.usenix.org/conference/usenixsecurity25/presentation/daffalla. (Aug. 2025).

[23] Alaa Daffalla, Marina Bohuk, Nicola Dell, Rosanna Bellini, and Thomas Ristenpart. 2023. Account security interfaces: important, unintuitive, and untrustworthy. In *USENIX Security '23*. https://www.usenix.org/conference/usenixsecurity23/presentation/daffalla. (Aug. 2023). ISBN: 978-1-939133-37-3. Retrieved Aug. 20, 2025 from.

[24] Alaa Daffalla, Lucy Simko, Tadayoshi Kohno, and Alexandru G. Bardas. 2021. Defensive technology use by political activists during the Sudanese revolution. In *IEEE S&P '21*. https://doi.org/10.1109/SP40001.2021.00055. (May 2021).

[25] Alioune Diallo, Jordan Samhi, Tegawendé F. Bissyandé, and Jacques Klein. 2025. (In)Security of mobile apps in developing countries: a systematic literature review. *Empirical Software Engineering*, 30. https://doi.org/10.1007/s10664-025-10689-z. Retrieved Aug. 20, 2025 from.

[26] Paul Dunphy, Vasilis Vlachokyriakos, Anja Thieme, James Nicholson, John McCarthy, and Patrick Olivier. 2015. Social media as a resource for understanding security experiences: a qualitative analysis of #password tweets. In *SOUPS '15*. https://www.usenix.org/conference/soups2015/proceedings/presentation/dunphy. ISBN: 978-1-931971-24-9.

[27] Electronic Frontier Foundation. [n. d.] Surveillance Self-Defense. https://ssd.eff.org/. (). Retrieved Aug. 10, 2025 from.

[28] George S. Everly Jr. and Jeffrey M. Lating. 2022. *The Johns Hopkins Guide to Psychological First Aid*. Johns Hopkins. ISBN: 978-1-4214-4399-7.

[29] Alberto Fittarelli and Lokman Tsui. 2023. Beautiful Bauhinia: "HKLeaks" – the use of covert and overt online harassment tactics to repress 2019 Hong Kong protests. https://citizenlab.ca/2023/07/hkleaks-covert-and-overt-online-harassment-tactics-to-repress-the-2019-hong-kong-protests/. (July 2023). Retrieved Feb. 8, 2024 from.

[30] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. 2019. "Is my phone hacked?" Analyzing clinical computer security interventions with survivors of intimate partner violence. *Proceedings of the ACM on Human-Computer Interaction*, 3, CSCW, (Nov. 2019). https://doi.org/10.1145/3359304.

[31] Freedom of the Press Foundation. [n. d.] Security guides & training. https://freedom.press/training/. ().

[32] Front Line Defenders. [n. d.] Security-in-a-Box. https://securityinabox.org/. ().

[33] Christine Geeng, Mike Harris, Elissa Redmiles, and Franziska Roesner. 2022. "Like lesbians walking the perimeter": experiences of U.S. LGBTQ+ folks with online security, safety, and privacy advice. In *USENIX Security '22*. https://www.usenix.org/conference/usenixsecurity22/presentation/geeng. (Aug. 2022). ISBN: 978-1-939133-31-1. Retrieved Sept. 11, 2022 from.

[34] Conor Gilsenan, Fuzail Shakir, Noura Alomar, and Serge Egelman. 2023. Security and privacy failures in popular 2FA apps. In *USENIX Security '23*. https://www.usenix.org/conference/usenixsecurity23/presentation/gilsenan. (Aug. 2023). ISBN: 978-1-939133-37-3. Retrieved Aug. 20, 2025 from.

[35] Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H. Mhaidli, Kentaro Toyama, and Florian Schaub. 2018. Keeping a low profile? Technology, risk and privacy among undocumented immigrants. In *CHI '18*. https://doi.org/10.1145/3173574.3173688. (Apr. 2018). ISBN: 978-1-4503-5620-6.

[36] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. 2019. Clinical computer security for victims of intimate partner violence. In *USENIX Security '19*. https://www.usenix.org/conference/usenixsecurity19/presentation/havron. (Aug. 2019). ISBN: 978-1-939133-06-9.

[37] Jordan Hayes, Smirity Kaushik, Charlotte Emily Price, and Yang Wang. 2019. Cooperative privacy and security: learning from people with visual impairments and their allies. In *SOUPS '19*. https://www.usenix.org/conference/soups2019/presentation/hayes. (Aug. 2019). ISBN: 978-1-939133-05-2.

[38] Internews. 2023. Field guide to incident response for civil society and media. https://internews.org/resource/field-guide-to-incident-response-for-civil-society-and-media/. (Nov. 2023).

[39] Internews. [n. d.] Security Auditing Framework and Evaluation Template for Advocacy Groups (SAFETAG). https://safetag.org/. (). Retrieved Feb. 9, 2024 from.

[40] Smirity Kaushik, Natã M. Barbosa, Yaman Yu, Tanusree Sharma, Zachary Kilhoffer, JooYoung Seo, Sauvik Das, and Yang Wang. 2023. GuardLens: supporting safer online browsing for people with visual impairments. In *SOUPS '23*. https://www.usenix.org/conference/soups2023/presentation/kaushik. (Aug. 2023). ISBN: 978-1-939133-36-6.

[41] Nigel King and Joanna M. Brooks. 2017. *Template Analysis for Business and Management Students*. SAGE. ISBN: 978-1-4739-5288-1.

[42] Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. How do app vendors respond to subject access requests? A longitudinal privacy study on iOS and Android apps. In *International Conference on Availability, Reliability and Security*. https://doi.org/10.1145/3407023.3407057. (Aug. 2020). ISBN: 978-1-4503-8833-7. Retrieved Aug. 19, 2025 from.

[43] Gemma Lasseter et al. 2022. Exploring the impact of shielding advice on the wellbeing of individuals identified as clinically extremely vulnerable amid the COVID-19 pandemic: a mixed-methods evaluation. *BMC Public Health*, 22. https://doi.org/10.1186/s12889-022-14368-2. Retrieved Jan. 22, 2025 from.

[44] William R. Marczak and Vern Paxson. 2017. Social engineering attacks on government opponents: target perspectives. *PoPETs*, 2017, 2. https://doi.org/10.1515/popets-2017-0022. Retrieved Feb. 8, 2024 from.

[45] MIT Cybersecurity Clinic. [n. d.] Cybersecurity: a social engineering approach at MIT. https://urbancyberdefense.mit.edu/CybersecurityClinic. ().

[46] Natasha Msonza. 2019. Digital Safety Trainer's Assistant. https://safesisters.org/wp-content/uploads/2019/09/Digital-Safety-Trainers-Assistant-smaller.pdf. (Sept. 2019).

[47] Collins W. Munyendo, Yasemin Acar, and Adam J. Aviv. 2023. "In eighty percent of the cases, I select the password for them": security and privacy challenges,

advice, and opportunities at cybercafes in Kenya. In *IEEE S&P '23*. https://doi.org/10.1109/SP46215.2023.10179410. (May 2023).

[48] Lorenzo Neil, Harshini Sri Ramulu, Yasemin Acar, and Bradley Reaves. 2023. Who comes up with this stuff? Interviewing authors to understand how they produce security advice. In *SOUPS '23*. https://www.usenix.org/conference/soups2023/presentation/neil. (Aug. 2023). ISBN: 978-1-939133-36-6.

[49] PEN America. [n. d.] Online Harassment Field Manual. https://onlineharassmentfieldmanual.pen.org/. ().

[50] Privacy Guides. [n. d.] https://www.privacyguides.org/. ().

[51] Lana Ramjit, Natalie Dolci, Francesca Rossi, Ryan Garcia, Thomas Ristenpart, and Dana Cuomo. 2024. Navigating traumatic stress reactions during computer security interventions. In *USENIX Security '24*. https://www.usenix.org/conference/usenixsecurity24/presentation/ramjit, 2011–2028. ISBN: 978-1-939133-44-1. Retrieved Sept. 16, 2025 from.

[52] Rapid Response Network. [n. d.] Digital First Aid Kit. https://digitalfirstaid.org/. ().

[53] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2016. How I learned to be secure: a census-representative survey of security advice sources and behavior. In *ACM CCS '16*. https://doi.org/10.1145/2976749.2978307. ISBN: 978-1-4503-4139-4.

[54] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. 2020. A comprehensive quality evaluation of security and privacy advice on the web. In *USENIX Security '20*. https://www.usenix.org/conference/usenixsecurity20/presentation/redmiles. (Aug. 2020). ISBN: 978-1-939133-17-5. Retrieved Oct. 10, 2022 from.

[55] Robert W. Reeder, Iulia Ion, and Sunny Consolvo. 2017. 152 simple steps to stay safe online: security advice for non-tech-savvy users. *IEEE Security & Privacy*, 15, 5. https://doi.org/10.1109/MSP.2017.3681050.

[56] Felix Reichmann, Annalina Buckmann, Konstantin Fischer, M. Angela Sasse, and Alena Naiakshina. 2025. Bridging the gap between usable security research and open-source practice - lessons from a long-term engagement with VeraCrypt. In *CHI '25*. https://doi.org/10.1145/3706598.3713983. ISBN: 979-8-4007-1394-1. Retrieved Aug. 20, 2025 from.

[57] Reporters Without Borders. [n. d.] Digital security for journalists. https://helpdesk.rsf.org/. ().

[58] Scott Ruoti, Jeff Andersen, Daniel Zappala, and Kent Seamons. 2016. Why Johnny still, still can't encrypt: evaluating the usability of a modern PGP client. https://doi.org/10.48550/arXiv.1510.08555. (Jan. 2016). Retrieved Aug. 20, 2025 from.

[59] Safe Sisters. [n. d.] https://safesisters.org/. (). Retrieved Feb. 9, 2024 from.

[60] Nithya Sambasivan et al. 2019. "They don't leave us alone anywhere we go": gender and digital abuse in South Asia. In *CHI '19*. https://doi.org/10.1145/3290605.3300232. ISBN: 978-1-4503-5970-2.

[61] Morgan Klaus Scheuerman, Stacy M. Branham, and Foad Hamidi. 2018. Safe spaces and safe places: unpacking technology-mediated experiences of safety and harm with transgender people. *Proceedings of the ACM on Human-Computer Interaction*, 2, CSCW, (Nov. 2018). https://doi.org/10.1145/3274424.

[62] Juliane Schmüser, Noah Wöhler, Harshini Sri Ramulu, Christian Stransky, Dominik Wermke, Sascha Fahl, and Yasemin Acar. 2022. "Please help share!": Security and privacy advice on twitter during the 2022 Russian invasion of Ukraine. https://doi.org/10.48550/arXiv.2208.11581. (Aug. 2022). Retrieved Feb. 8, 2024 from.

[63] Svenja Schröder, Markus Huber, David Wind, and Christoph Rottermanner. 2016. When SIGNAL hits the fan: on the usability and security of state-of-the-art secure mobile messaging. In *EuroUSEC '16*. https://doi.org/10.14722/eurousec.2016.23012. (July 2016). ISBN: 978-1-891562-45-7. Retrieved Aug. 20, 2025 from.

[64] Lucy Simko, Ada Lerner, Samia Ibtasam, Franziska Roesner, and Tadayoshi Kohno. 2018. Computer security and privacy for refugees in the United States. In *IEEE S&P '18*. https://doi.org/10.1109/SP.2018.00023. (May 2018).

[65] Julia Slupska, Marissa Begonia, and Nayana Prakash. 2021. Digital Privacy & Security Guide for migrant domestic workers. https://domesticworkerprivacy.github.io/. (Sept. 2021).

[66] Julia Słupska and Angelika Strohmayer. 2022. Networks of care: tech abuse advocates' digital security practices. In *USENIX Security '22*. https://www.usenix.org/conference/usenixsecurity22/presentation/slupska-networks. (Aug. 2022). ISBN: 978-1-939133-31-1. Retrieved Sept. 11, 2022 from.

[67] Xinru Tang, Gabriel Lima, Li Jiang, Lucy Simko, and Yixin Zou. 2025. Beyond "vulnerable populations": a unified understanding of vulnerability from a socio-ecological perspective. *Proceedings of the ACM on Human-Computer Interaction*, 9, 2, (May 2025). https://doi.org/10.1145/3710935. Retrieved Aug. 20, 2025 from.

[68] Emily Tseng, Rosanna Bellini, Yeuk-Yu Lee, Alana Ramjit, Thomas Ristenpart, and Nicola Dell. 2024. Data stewardship in clinical computer security: balancing benefit and burden in participatory systems. *Proceedings of the ACM on Human-Computer Interaction*, 8, CSCW1, (Apr. 2024). https://doi.org/10.1145/3637316.

[69] Emily Tseng, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. 2022. Care infrastructures for digital security in intimate partner violence. In *CHI '22*. https://doi.org/10.1145/3491102.3502038. (Apr. 2022). ISBN: 978-1-4503-9157-3. Retrieved Apr. 26, 2024 from.

[70] Mona Wang, Jeffrey Knockel, Zoë Reichert, Prateek Mittal, and Jonathan Mayer. 2025. WireWatch: measuring the security of proprietary network encryption in the global Android ecosystem. In *IEEE S&P '25*. https://doi.org/10.1109/SP61157.2025.00224. (May 2025). ISBN: 979-8-3315-2236-0. Retrieved Aug. 20, 2025 from.

[71] Noel Warford et al. 2022. SoK: a framework for unifying at-risk user research. In *IEEE S&P '22*. https://doi.org/10.1109/SP46214.2022.9833643. (May 2022). ISBN: 978-1-6654-1316-9. Retrieved Aug. 19, 2022 from.

[72] Miranda Wei, Sunny Consolvo, Patrick Gage Kelley, Tadayoshi Kohno, Franziska Roesner, and Kurt Thomas. 2023. "There's so much responsibility on users right now": expert advice for staying safer from hate and harassment. In *CHI '23*. https://doi.org/10.1145/3544548.3581229. ISBN: 978-1-4503-9421-5.

[73] Miranda Wei, Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2022. Anti-privacy and anti-security advice on TikTok: case studies of technology-enabled surveillance and control in intimate partner and parent-child relationships. In *SOUPS '22*. https://www.usenix.org/conference/soups2022/presentation/wei. (Aug. 2022). ISBN: 978-1-939133-30-4.

[74] Charlotte Woodhead, Cerisse Gunasinghe, James Kenrick, and Dame Hazel Genn. 2022. Social welfare advice and health among young people: a scoping review. *Journal of Social Welfare and Family Law*, 44, 1. https://doi.org/10.1080/09649069.2022.2028409. Retrieved Jan. 22, 2025 from.

[75] Yaman Yu, Saidivya Ashok, Smirity Kaushik, Yang Wang, and Gang Wang. 2023. Design and evaluation of inclusive email security indicators for people with visual impairments. In *IEEE S&P '23*. https://doi.org/10.1109/SP46215.2023.10179407. (May 2023).

A Interview protocol

Warm-up

This interview is about how you decide what security and privacy advice to give folks who face elevated risks in order to help them stay safe going forwards. We're especially interested in the kind of things that you might learn from giving advice over time, that might not be intuitive without experience—for example, what might lead you to give different advice to two people who seem to be in similar situations on the surface.

Before we get into the details of how you give advice, I want to establish some background about who you work with and how.

- In broad strokes, who do you primarily give advice to?
- Could you please describe the context for giving advice? Typically, what is the setting, and what is the goal?
- How long have you been doing this kind of work?

Past experiences giving advice

Please think back to a time you gave advice about security and privacy, a time that you remember well.

- What information did you try to learn about your client(s) before giving advice?
- How did you go about learning this information? For example, did you ask a set series of questions, or have your client(s) describe their situation to you open-endedly, or some other approach?
- Did this information affect what advice you gave, and if so, how?

Differentiating advice

Next, we have some questions about whether and how you decide to give different security and privacy advice to different clients.

- To what extent do you give different advice to different clients about how to protect themselves going forward?

- Based on your own experience giving advice, is there any advice you give that you think is good for some people and bad or counterproductive for others? (Why?)
- Imagine that you're advising a client and trying to determine whether to give this advice and how to prioritize it. How do you make that decision? (Why?)

Advice process
Now I'd like to take a step back and talk about your overall process for advice.

- Have you or your organization made choices or changes over time, in order to make the process of giving advice work better specifically for [participant's at-risk population]?
- Do you ever follow up with clients, and if so, have you ever learned anything that changed how you give advice?
- People at high risk are often urged to seek security and privacy advice from experts who know their context really well, since following the wrong advice can have serious consequences. As someone who has knowledge and experience with a high-risk context, is there anything you do on a daily basis to mitigate the potential for harmful advice?
- Can you think of an aspect of your current process for giving advice that works particularly well and an aspect that you would change if you could?

Challenges
I want to talk for a bit about challenges you may face while giving security and privacy advice.

- Let's set aside some basic logistical concerns, like not having enough time to help as many people as you'd like. Once you're meeting with a client, are there significant challenges that make it harder for you to decide what advice to give or to give advice effectively?
- Is there anything you think is missing that would help make it easier for you to give good advice? This could be anything, from technical tools you want people to develop to studies you want scientists to do.
- Let's say you wanted to increase the number of clients that you or your organization are able to provide tailored advice to—say, 10 times or 100 times more than your current rate. How would you do so?

Building capacity
Some advocacy groups are working on tools that automatically personalize security and privacy advice—for example, a questionnaire that asks about the user's devices and other info, and then provides a list of advice prioritized according to their needs. We're interested in your thoughts on how these tools should be designed and used, if at all.

- Is there anything you think would be particularly important to take into account when designing an automated tool that personalizes security and privacy advice, especially for at-risk populations like [participant's at-risk population]?
- In an ideal world, what role do you envision an automated tool like this playing in the process of giving advice to at-risk populations?

Wrap-up

- Before we finish up, I want to ask, is there anything else you'd like us to know about how you tailor security and privacy advice to [participant's at-risk population]? Anything you wanted to say but didn't come up in our conversation?

B Pre-interview questionnaire

(1) Do you give security and privacy advice as part of an organization?
   - Yes | No

   Questions #2 and #3 are skipped if participants answer No to #1.
(2) Which of the following best describes this organization?
   - Nonprofit | Private company | Government agency | Educational institution (e.g., a university research group running advice clinics)
(3) About how many people are currently involved in giving security and privacy advice at your organization? _____
(4) What is your gender?
   - Male | Female | Non-binary | Prefer to self-describe _____| Prefer not to state
(5) Which of the following best describe your race/ethnicity? Select all that apply.
   - American Indian or Alaska Native | Asian | Black or African American | Hispanic or Latino | Native Hawaiian or Pacific Islander | White | Prefer to self-describe _____ | Prefer not to state

C Interview codebook

Our codebook is in the supplementary material: https://osf.io/pqh84/

D Types and examples of information sought about clients

Table 2: Types of information sought about clients, alongside illustrative examples from interviews.

Info type	Examples
Threat actors	**Type of threat actor**: "If the implementer of the attack is just a friend, you can report to the police, and the police will come into action. But if it's the government, even the police will decide to withdraw from you." (P17)
	Type of attack: "[A client] she was undergoing cyber harassment and had gone offline, and people had started following her. . . . If something has escalated to the physical, then I advise that person to seek protection from the police." (P17)
	Red flags: "If the person knows the harasser . . . or they know that they have a history of escalation and erratic behavior, this is another red flag. So there are specific red flags that I ask people about." (P8)
Vulnerable/ compromised accounts and devices	**Attack surfaces**: "I always do some research also about individuals and organizations to try and identify how much information is out there about them on the Internet that they might not even know about." (P14)
	Evidence of account compromise: "We go to the security page, we take a look at what's been logging into this account. Are these familiar IP addresses, are these familiar geographical regions? Is someone using an operating system or a device that is that is completely new or unknown to you, or that you know belongs to your abuser?" (P1)
	Evidence of physical tracking: "If the person has received a message about an AirTag following them, I will usually spend some time thinking about AirTags. But if their location seems to be being tracked but the contents of their communications are not, then I start looking for the other types of physical trackers, which do not have any anti-stalking mitigations." (P1)
Privacy and security hygience	**Password hygiene**: "Something that comes up loads with clients is . . . they either they use the same password for everything, or they just use easily guessable passwords, like kids or names or birthdays or whatever." (P5)
	Previous advice: "I'll ask them, have you gotten advice from anyone? Do they share your passport and identity? If you get advice from an American who's straight, and you're a gay Ugandan, that's bad advice." (P3)
Home and family situation	**Relationship with an abuser**: "I sometimes need to know at least the basics, like are you still with this [abuser] or have you left? Because you have that complete change in the threat model." (P6)
	Family members: "You have to talk about the child's devices, or what the child knows, or where the child is going, or the objects that they are taking back and forth between between the different homes." (P1)
Cultural and geographic context	**Local norms**: "Is he living in a remote area? Is he living in the city base? If you're thinking about VPN usage, like, what is the percentage of people using VPNs in that area? If it's very low, then it kind of sticks out, right?" (P4)
Identities	**Gender and sexual orientation**: "If it's a person from the transgender community or the LGBTQI community, we don't advise them to go to the law enforcement agencies." (P7)
Privacy and security goals	**Minimum goals**: "Where do you need to be able to go? What devices do you absolutely need to be able to trust? . . . If you have a particularly bad situation, [I] may start with just like a minimum viable product: what is the smallest bubble of privacy and security in which you can continue to do your work or potentially get away from your abuser?" (P1)
Lifestyle and work	**Occupational goals**: "In terms of what they are trying to achieve, is it more about just having conversations? Is it more about trying to get more information out of [a region]? Is it trying to send information back inside?" (P4)
Acceptable trade-offs	**Acceptable amount of inconvenience**: "The most important last question [of a threat model] is how much trouble are you willing to go through in order to prevent the consequences of this sort of compromise?" (P1)
Tech literacy	**Level of comfort with digital technologies**: "There might be people who are, even though they know that they're supposed to use a password manager, you look at them and you say like, well, you know, given their relationship to technology in general, and general aversion to it, we might find safer ways to incorporate a password manager, but also to give them less reliance on something that could for a variety of reasons just disappear from them." (P5)
Language fluencies	**Language**: "If that person is conversant in [official language], that is really helpful in terms of technology, because that allows you to share materials that are already available." (P4)
Resources	**Financial resources**: "If I'm talking to a freelancer that cannot really afford to pay the annual subscription of [a data deletion tool] and stuff like that, I will be more realistic about the advice I give." (P8)
Accounts, apps, software, and devices	**Devices**: "Generally they don't explain the setup they have, so that's a question I have to ask: to be like, okay, what's your device, what phone you use?" (P6)
Mental and emotional capabilities	**Composure**: "What's your comfort level doing complicated tasks? How do you respond to anxiety and stress?" (P3)
Level of stress	**Current level of stress**: "You have to ask and ascertain the state of mind of the person and their emotional stability. So, the first question that you always ask is, how are you?" (P2)
Other physiological state	**Influence of substances**: "Lately, we've been having a lot of people who are under the influence. And that is also a big no-no. Have you consumed or are you not at your capacity? Maybe call tomorrow." (P2)