Account Security Interfaces: Important, Unintuitive, and Untrustworthy

Alaa Daffalla and Marina Bohuk, *Cornell University;* Nicola Dell, *Jacobs Institute Cornell Tech;* Rosanna Bellini, *Cornell University;* Thomas Ristenpart, *Cornell Tech*

https://www.usenix.org/conference/usenixsecurity23/presentation/daffalla

This paper is included in the Proceedings of the 32nd USENIX Security Symposium.

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the 32nd USENIX Security Symposium is sponsored by USENIX.

Account Security Interfaces: Important, Unintuitive, and Untrustworthy

Alaa Daffalla[‡], *Marina Bohuk*[‡], *Nicola Dell*[*], *Rosanna Bellini*[‡], *Thomas Ristenpart*[α]

[‡]*Cornell University,* [*]*Jacobs Institute Cornell Tech,* [α]*Cornell Tech*

Abstract

Online services increasingly rely on user-facing interfaces to communicate important security-related account information—for example, which devices are logged into a user's account and when recent logins occurred. These are used to assess the security status of an account, which is particularly critical for at-risk users likely to be under active attack. To date, however, there has been no investigation into whether these interfaces work well.

We begin to fill this gap by partnering with a clinic that supports survivors of intimate partner violence (IPV). We investigated hundreds of transcripts to identify ones capturing interactions between clinic consultants and survivors seeking to infer the security status of survivor accounts, and we performed a qualitative analysis of 28 transcripts involving 19 consultants and 22 survivors. Our findings confirm the importance of these interfaces for assessing a user's security, but we also find that these interfaces suffer from a number of limitations that cause confusion and reduce their utility.

We go on to experimentally investigate the lack of integrity of information contained in device lists and session activity logs for four major services. For all the services investigated, we show how an attacker can either hide accesses entirely or spoof access details to hide illicit logins from victims.

1 Introduction

Web authentication used to be relatively straightforward: just type in a valid username and password. Now account access, whether by such login or via account recovery mechanisms, is more complex, with wider deployment of multi-factor authentication, risk-based authentication, recovery through backup codes, and more. At the same time, users often have multiple devices from which they need ongoing access to accounts.

To help users make sense of this more complicated landscape, major services deploy various interfaces for configuring access challenges and obtaining information about historical or ongoing accesses made to the account. We call such interfaces account security interfaces. These are important for users to understand their account's security posture including determining whether compromise has occurred, and they are critically important for at-risk users [69] who face active and often complex attacks. Examples include survivors of intimate partner violence (IPV) [25, 46], journalists [47], activists [17], undocumented immigrants [29], and refugees [61].

Despite its importance, little work has investigated user understanding of account access. Prior work has looked at the complexity of configuring particular access mechanisms (c.f., [4, 41, 56, 67]) but not whether users can understand the current configuration of their account. In terms of access notifications, Markert et al. [45] performed an in-lab study of email login notifications, and Redmiles [54] interviewed users that had service-identified suspicious login incidents on Facebook. Both these studies consider a subset of account access notifications and focus on users who are not necessarily under attack. In short, none of the prior work focuses on how users interact with account security interfaces and whether the interfaces succeed in helping users assess security.

We therefore initiate a study of these interfaces, including how users interact with them, how they are used to assess security of accounts, and whether the interfaces themselves are secure. We first conduct a survey of modern account access challenges and related user interfaces (UIs) across four major online services—namely Google, Facebook, Apple, and WhatsApp. The survey provides a background on what kinds of account security interfaces are currently deployed.

We then perform a case study to understand the role of such interfaces for a particular at-risk population: IPV survivors. IPV survivors face a plethora of technological risks, and prior work has identified account takeover by the abuser as a frequent problem [11, 25, 26, 46, 64]—abusers can exploit their physical and social proximity to their partner to bypass authentication mechanisms. These concerns have led to deployment of clinical computer security approaches [16, 24, 31] in which trained consultants work directly with survivors to help them with digital abuse. Several of the authors are volunteers at such a tech clinic and have experienced how account security interfaces play a critical role in survivor safety.

Our anecdotal experiences suggested that clinical settings can be a rich source of data about user experiences with account security interfaces. We partnered with the Clinic to End Tech Abuse (CETA)[1]—an IPV clinic in New York City that handles referrals for hundreds of survivors each year. After obtaining IRB approval, we used a keyword search of 220 transcripts to identify a set of 28 transcripts for further analysis and performed a qualitative inductive content analysis [23].

Our findings confirm that account security interfaces play a key role in IPV survivor safety. They can help survivors understand their security posture and identify likely compromises by their abuser. We also identify a number of limitations to current interfaces, including difficulties finding and interpreting account access information; and we find that survivors often hesitate to make security improvements due to uncertainty about the impact of potential configuration changes.

Our findings highlight that survivors and consultants rely particularly heavily on the information about accesses available in these interfaces, such as the operating system, device type, access date, and location to help them gauge if an access is malicious. Given this important role, we analyze the integrity of these interfaces in the face of malicious adversaries. We introduce two types of attacks: access hiding attacks in which an adversary can arrange for their accesses to go unreported by a service and access spoofing attacks in which the adversary manipulates their access to appear like it comes from a different device (e.g., the victim's). Hiding or spoofing attacks can therefore prevent discovery of malicious monitoring of and full control over user accounts.

We show that all four major services are vulnerable to attacks by knowledgeable but technically unsophisticated adversaries. An example of a spoofing attack on Facebook's active sessions interface appears in Figure 1. Our results are similar to recent attacks on risk-based authentication (RBA) [40], which also undermine a service's ability to correctly identify the client device.

Summary. Our paper makes the following contributions:

- We initiate work on user understanding of account security interfaces, including their ability to assess security posture and compromise status. We survey four major services, providing a snapshot of the diversity of current designs.
- We perform a qualitative study of IPV technology abuse consultations that assess the security posture and compromise status of survivor accounts. Our findings highlight the importance and limitations of account security interfaces.
- We discover access hiding and spoofing attacks that work in some form against all four services studied. Our attacks are simple, exploiting the services' reliance on untrustworthy client-provided values to populate these interfaces. These weakness could put at-risk users in danger.

Despite the discovered limitations, we emphasize that these interfaces play an important role in user safety—the IPV consultations we analyzed would have struggled without them. We therefore believe account security interfaces deserve further attention to improve their usability and security. As such, we provide a discussion of various directions for future work.

[1] https://www.ceta.tech.cornell.edu/

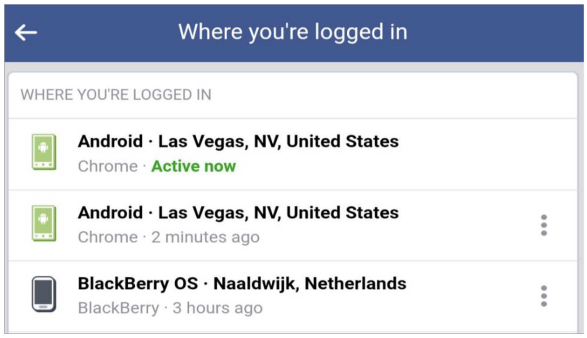

Figure 1: An example of a user's view of their Facebook account's recent logins interface, showing three logins (from top to bottom): (1) the legitimate user's; (2) a malicious login from the same city using an Apple Mac computer spoofed to look like the user's device; and (3) a malicious login from the same city and an Apple Mac computer spoofed to appear as coming from a Blackberry device in another country.

Ethics. Our research shows how relatively unsophisticated adversaries can perform attacks on account security interfaces. While there is a risk that adversarial users learn new strategies from our work, we posit that this risk is marginal and that benefits outweigh potential harms: making progress on improving safety requires frank discussion of security problems and abuse. That said, we avoid step-by-step instructions on how to perform spoofing attacks, and we are in the process of performing responsible disclosure to each service. We received Institutional Review Board (IRB) approval for our transcript analysis study and performed several rounds of quote reviews to assess deanonymization risk (e.g., references to potentially unique situations); quotes were modified where needed to mitigate risk while maintaining the voice of the participant as much as possible.

We performed responsible disclosure, contacting relevant teams at Google, Apple, Facebook, and WhatsApp about the discovered hiding and spoofing attacks. We met with teams from Facebook, WhatsApp, and Apple to answer questions and make suggestions about potential near- and longer-term mitigations (see Section 7). They also reproduced or otherwise confirmed our results.

2 Related Work

Usability of authentication mechanisms. A now long line of work has focused on usability of various authentication mechanisms. Early work focused primarily on password-based authentication (e.g., [5]) and has expanded to encompass two-factor authentication (2FA) (see below), biometric or

other passwordless authentication (e.g., [37, 41]), and studies that directly compare different approaches (e.g., [13, 58]).

Because we work in modern authentication systems using some form of multi-factor authentication, work on its usability is particularly relevant. There has been considerable work studying the usability of multifactor authentication [18, 19, 51] and two-factor authentication (2FA) more specifically [4, 20, 53, 55]. A subset of work studying 2FA focuses on simulated in-lab studies to evaluate users' understanding of setting up 2FA on their accounts. For example, Acemyan et al. [4] conducted a usability assessment of Google's 2FA methods and found that users struggle with completing the 2FA setup task and take awhile to do so. Furthermore, Petsas et al. [53] investigated the adoption of two-factor authentication for Google accounts and found that at the time only 6.4% of users used 2FA.

Risk-based authentication. Most 2FA approaches require an explicit secondary communication channel. Risk-based authentication (RBA) is, instead, an adaptive authentication measure deployed by several major online services to augment password-based login by taking into account additional login parameters to trigger a given challenge [75]. In their work investigating RBA across eight popular online services, Wiefling et al. [74] were able to determine the underlying feature sets governing RBA implementations across the different services. They also confirmed that only a limited set of client features are useful for practical deployments [73].

Another line of work explores the privacy of RBA systems [76] and usability of RBA. Wiefling at al. [71, 72] conducted a lab study with 65 participants and found that users considered RBA to be more secure than passwords and more usable compared to 2FA. These studies did not investigate how users understand the configuration of accounts, nor the ability to assess whether illicit accesses have occurred.

Lin et al. [40] present a practical phishing attack that undermines the use of browser fingerprinting in RBA by exploiting services that remember users' devices upon login. Our results in Section 6 exploit the same root cause problem with RBA mechanisms: the reliance on untrustworthy client-provided data. But this prior work did not investigate the impact on user understanding of account security status.

User understanding of account security. A key aspect of our study is to investigate whether users can understand the security of their account and in particular if others can have illicit access.

Several studies have investigated account compromise detection, as well as users' reactions to account compromise [8, 63, 68]. For example, Shay et al. [60] found that 30% of surveyed participants indicated that they have experienced an account compromise, 50% of which discovered the compromise because of suspicious activity originating from their account. Prior work has also looked at the efficacy of user security notifications. Security notifications have been studied in a variety of contexts such as browser warnings [6] and password reuse [28], but only a few studies have focused on notifications for account compromise detection.

Redmiles [54] interviewed 67 participants for which Facebook flagged login attempts to their accounts as suspicious. She found that the lack of sufficient information in the notifications led users to believe that these notifications were false positives and that no protective action was required. Our results will similarly highlight the importance of having detailed information about accesses to assess security.

Markert at al. [45] conducted measurements of existing services' email-based login notifications as well as user studies to understand user interactions with these notifications. They found that users want to be informed about suspicious logins but are often confused why an email login notification is triggered and can experience warning fatigue. In a follow-up study, they found that websites rarely provide advice to prevent unwanted access [44].

Our study complements and expands on those above, in that we look at whether and how users suffering active attacks can holistically make sense of their account security with the aid of both notifications and interfaces.

Integrity of access descriptions. Logs of access or system events have long been considered important for assessing security posture (c.f., [50, 62]), along with the integrity of such logs [9, 10, 42]. Our work similarly touches on the theme of integrity in logging, but in the previously unexplored context of modern user accounts for web services. The spoofing attacks we discuss in Section 6 have some similarity to prior attacks aimed at arranging for a UI that tricks users such as phishing (c.f., [32]) or clickjacking [33]. To the best of our knowledge, we are the first to investigate the integrity of login notifications or access identification interfaces.

3 Account Security Interfaces

As mentioned earlier, we refer to a user interface that allows a user to control or monitor access to an online account as an account security interface (ASI). Traditionally, ASIs include interfaces for configuration of authentication mechanisms such as passwords, second-factor authentication, and recovery information. These ASIs are relatively standard across the services we have explored.

In addition to configuration, many services now provide ASIs that help users make sense of current or historical access to their account. They can provide the user with information about active sessions, authorized devices, and any suspicious (atypical) account access activity. We have come across four distinct types of such ASIs: *device lists*, *session lists*, *activity logs*, and *access notifications*—although as we will see, there is often not a one-to-one mapping between a specific interface (i.e., a web page) and its type (e.g., because a single ASI includes both a device and session list). We explain each of these four ASI types in turn, using as representative running

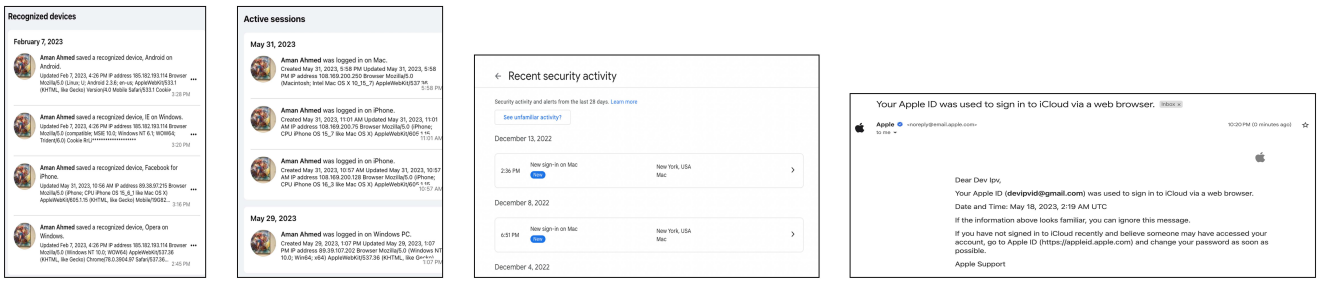

Figure 2: Screen captures of an example device list, session list, activity log, and access notification from left to right. The account information shown (including the name and email) is that of a fake test account created for the purpose of this work.

examples the four services that will be the focus of the rest of the paper: Google, Apple, WhatsApp, and Facebook. See Figure 2 for screen captures of examples of each interface.

Device lists. Device lists provide users with information on devices that are authorized to access their accounts. Authorized devices have been given a bearer token (e.g., [12,35,36]) that authenticates ongoing access to an account. The precise details included in a device list vary across services, but typically include information to help identify the device, including the device model, operating system (OS), and platform used (web browser or app). They may also include the IP address, geographic location, and the date and time of the first or most recent access by the device.

Google, Facebook, and WhatsApp all provide users with lists of currently or previously authorized devices. A device list on Apple services, on the other hand, is limited only to Apple devices from which a user has already authenticated using a password plus 2FA. They call these *trusted devices* (see [2]). Google refers to trusted devices instead as devices that can *bypass* 2FA, and although they do not provide a list of trusted devices, they provide an ASI in which the user can render all devices untrusted [1]. As a final example, Facebook provides a separate list of *recognized devices*: devices where 2FA has been enabled and which bypass 2FA for future logins. The terminology around authorized devices is therefore inconsistent across services, with nuanced semantics.

In Figure 3 we give examples of device list ASIs. We find that across the four services studied, device lists have common elements such as system information, platform information, and date and time identifiers. We find that Facebook and Google's device lists also include location information, while WhatsApp's and Apple's do not. On the other hand, Apple's device lists do not include date and time identifiers, but they do contain additional system information (OS version and serial number), IMEI, and associated phone number.

Most device list ASIs allow users to log out a suspicious device and prompt users worried about account compromise to take further measures to secure the account, such as changing their password and configuring 2FA. Device lists are considered a common access identification approach across other services like Twitter and Telegram, but there are exceptions to this trend, such as Amazon, where there are no interfaces that embed device lists.

Session lists. Session lists provide users with a list of currently active or terminated sessions. Some services, such as Google, display device lists and associated session lists in a single interface so that for each device listed, the user can expand to see the associated sessions on the device. Users can see a list of currently active or past sessions on Facebook using the *Active sessions* interface, which displays the device's user agent string, along with the time and date the session was initiated. Both Apple and WhatsApp do not provide users with a log of currently active or past sessions.

Activity logs. Activity logs interfaces show recent security activity on an account. Activity may encompass suspicious logins, changed passwords, and 2FA and recovery configurations, among other information. Some services, like Google, have a separate interface for activity logs containing all sign-ins on new devices, password and recovery information changes, and requests to download user data. On Google, this activity log is only available for 28 days. Similarly, Facebook has a *Logins and Logouts* interface that lists all login and logout activity on an account. This interface uses the IP address to determine from where the login or logout occurred.

Access notifications. Email and in-app notifications are used by services to communicate recent account activity to users. For example, users are notified via email when Google detects suspicious account activity, and Apple notifies users when non-trusted devices access `icloud.com`. Facebook allows users to opt into receiving both in-app and email notifications, while WhatsApp does not support access notifications at all.

Generally, email access notifications do not provide much information about login instances. The Google email notification only informs users of the device model they are using to log into an account, while Apple provides no device identifiers and only informs them of the time of login (see Figure 2).

In addition to the services reviewed in this study, other services also use access-based notifications to inform users of account access. Each Amazon login prompts an email notification, and each Twitter login appears as an in-app notification.

We find that most services provide access notifications, device lists, activity logs, and session lists across one or more interfaces. Navigation flows—the routes taken by clicks or

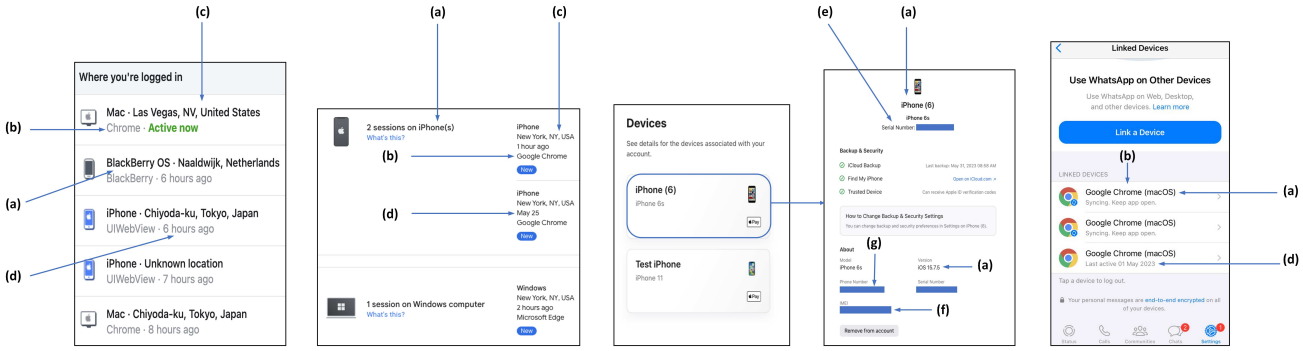

Figure 3: Screen captures of example device lists, in order from left to right: Facebook's device list, Google's combined device and session list, Apple's device list and expansion of a single device on that list, and finally WhatsApp's device list. We also label different identifiers across device lists: (a) system information (device model and/or OS), (b) browser information, (c) location information, (d) date and/or time information, (e) serial number, (f) IMEI information, and (g) phone number. For privacy reasons, we redact some information as seen in blue.

intermediary pages to reach a specific interface—differ between interfaces, however. Complex or multiple navigation flows can make it harder to reach these interfaces. For example, on Facebook there are five identification interfaces. To reach the set of *Where you're logged in and Authorized logins* interfaces, a different navigation flow is required than that for the set of *Active sessions, Recognized devices, and Logins and logouts* interfaces (see Figure 7). In contrast, on Apple and Google, access identification information is accessed through a single navigation flow.

4 Case Study: Intimate Partner Violence

To better understand the need for and efficacy of ASIs, we perform a case study in the context of a particular at-risk user population. Technology-facilitated abuse or 'tech abuse' is a common occurrence in many cases of intimate partner violence (IPV) [30, 70], as technology can be used by abusers to harass, stalk, threaten, or otherwise harm their victims [26]. Security experts have documented a range of different technical attacks that rob a survivor of their right to privacy, including GPS tracking [52, 52, 57], doxxing [26], harassment [65, 77], and surreptitiously monitoring a survivor's digital activity [11, 64].

Account compromise is a frequent abuse strategy. As many intimate partners have close physical and social proximity [25, 39], abusers can often both physically access a victim's accounts or devices and leverage intimate knowledge of a victim to bypass access challenges like passwords or knowledge-based questions. Abusers may also compel victims to disclose passwords or force the victim to respond to other types of access challenges [25]. Most often these attacks represent what Freed et al. [25] term an *UI-bound adversary*, meaning the abuser can accomplish their goals using only standard UIs and without the aid of any sophisticated tools or attack techniques (c.f., [11, 64]).

Computer security clinics. Given the complexity of technology abuse in IPV, there has been growth in providing direct, expert support to survivors in the form of technology clinics [16, 24, 31]. While clinics vary in terms of services and methods [16, 24, 31, 65, 66], common elements include working with individual survivors to help them navigate technology abuse, trauma-informed care approaches, and integration into community survivor advocacy organizations.

Four of the authors of this work are volunteers at CETA, which has handled more than 400 referrals to date. They work alongside 30+ technology experts (*consultants* hereafter) in privacy and security, IPV, and trauma-informed care to deliver tailored advice to survivors of tech abuse (*clients*). Clients experiencing technology abuse are referred to the clinic from a variety of IPV support organizations in the community. A referred client is assigned a consultant team to work with across a series of appointments that may last anywhere from ten minutes to a few hours. The number of appointments per client varies as well, with one to two appointments being the norm. While some appointments are in-person, most are remote via a conference call; we only analyze transcripts of the latter appointment type.

The clinic follows a high-level procedure introduced by Havron et al. [31] in which consultants work to understand a client's technology concerns, use tools to investigate potential digital threats to a client's well-being, and advise how a client may make changes to their security and privacy. To investigate a client's concerns about tech abuse, a consultant may ask to see a client's devices or ask the client to read aloud information shown on relevant ASIs. While aiding clients, most consultants use their own devices to view the same interfaces as the client (but for another account) to help them guide the client to the correct screens. We call this an *interface walkthrough*: an active back-and-forth discussion between a client and consultant about a client's device lists, activity logs, and session history. Either a consultant or a client might initiate conversations about these interfaces. However, we find that consultants most often brought them up in conversation due to the nature of the clinic setting where the client is at the

receiving end of technology services.

The clinic has an ongoing IRB-approved research protocol to better understand technology abuse in IPV under which our study falls. All volunteers and consultants in the clinic complete a human subjects research training equipped with information on trauma-informed care. Upon first working with a client, consultants ask if the client consents to contributing their case to ongoing research into tech abuse in IPV contexts. Clients receive the same quality of service irrespective of their response to participate. Consultants ask if clients additionally consent to audio recordings of appointments and whether they agree to allow consultant notes to be used for research. Consent is verbal to protect the privacy and safety of participants. Anonymized notes and recordings may then be used for research, monitoring, and ongoing case management.

Transcript selection. Based on the authors' experience in the clinic, we knew that appointments frequently involved detailed discussions between the client and consultant about ASIs during walkthroughs. To identify relevant transcripts, we (1) confirmed the presence of discussion of ASIs in clinic data, (2) devised a keyword-based search strategy to surface relevant transcripts, and (3) sampled these for analysis. We discuss each step in turn.

To assess the viability of clinic data, we performed an initial scoping search in the clinic CMS (content management system) of 220 transcripts of client consultations conducted between May 2019 and June 2022. To do this, we trialed a sample of phrases related to account security (*"known devices"*, *"last signed in"*, and *"recovery information"*), which returned several positive matches of consultant and client conversations on ASIs in clinic contexts. However, in many cases, ASIs were merely referenced passively and not fully discussed between client and consultant in a consultation.

Thus, we designed a search strategy based on a regular expression, keyword-based search on all transcripts in the clinic's CMS. Our keywords were selected after reviewing a small sample of ASIs across different online services, where an asterisk matches arbitrary suffixes: *"recognized device*"*, *"list of device*"*, and so on. We extended this keyword set to terms that related to device identification at a later stage, including *"two-factor"*, *"two-step"*, *"sign in history"*, and *"last active"*. A full list of keywords is in Appendix A: Figure 6.

We executed each keyword search in turn and manually reviewed search results to exclude transcripts that did not include ASI walkthroughs. This returned a smaller, relevant data set of 96 transcripts. Upon inspection of the 96 transcripts by the fourth author, we found six services that were discussed: Google, Facebook, Apple, WhatsApp, Instagram, and Microsoft—four of which we examined further for this study. Instagram and Microsoft were excluded from this coding effort due to a lack of a representative sample for each. We randomly sampled five transcripts for each of the four chosen services (Google, Facebook, Apple, WhatsApp) for further analysis to avoid skewing our findings towards one service. Eight transcripts (two per service) were added at a later stage to accommodate for 2FA functionality.

Our search and sampling approach resulted in a total dataset of 28 transcripts consisting of 9.7k words. This set of transcripts pertained to 22 clients (S1–S22) who were supported by 19 individual consultants (C1–C19) of varying ranges of expertise (1–3 years of volunteering at the clinic).

Qualitative coding. Motivated by the richness of the transcript segments, we chose to use a qualitative inductive content analysis [23] of the 28 transcript segments. Three authors performed three rounds of open coding following guidelines established by Saldana [59] across each transcript segment. The first round of coding 20 transcripts generated 35 codes in a shared codebook being careful to distinguish client and consultant unique codes and to label each interface using descriptive notation. The shared codebook was reconciled after two further rounds of open coding eight additional transcripts, resulting in a final codebook of 50 codes—a higher number to reflect the addition of new data from two-factor segments also being included in the dataset (see Appendix A).

The coding team met to generate high-level categories that accurately represented the use of ASIs. We present each of these categories in turn to explicate how interfaces are used in a clinical context and what barriers or limitations exist in their use to relay account access information. We address potential routes for improvement of these interfaces in Section 7.

5 Case Study Findings

Overall, we find that ASIs play a significant role in supporting survivors of IPV by allowing investigation of suspicious login activity and account compromise on their accounts. We present our findings in two sections: (1) why people use these interfaces, and (2) the limitations faced when using them.

Where relevant, we discuss how our findings relate to typical affordances in ASIs. However, the nature of our data–transcripts of verbal discussions—does not include ground truth on what ASI (if any) a client or consultant was looking at, nor do we know if these ASIs have changed between when the consultation occurred and our survey of contemporary ASIs. Nevertheless, we believe our analyses give valuable insight into the utility and efficacy of modern ASIs.

5.1 Functions of Account Security Interfaces

Here we explicate the main functions of ASIs in the context of clinical computer security.

To summarize account activity and security settings. As listing all owned digital devices from memory can be cognitively taxing, device lists acted as a prompt to assist clients in providing an overview of their digital footprint—the set of devices and accounts that they use—to the consultants. We

found that consultants frequently requested for clients to review these interfaces to jog their memory of potential devices that are currently logged in or have logged into their account in the past. This enables them to clarify potential security vulnerabilities in account access. In the following interaction, a consultant guides a client through Apple's *Devices* interface.

> **C2:** *"And when you scroll down right above sign out, it will show a list of devices that are logged into your Apple ID. So I will just ask you to take a moment to see if you recognize all these devices."*
>
> **S7:** *"Yeah. I mean, there's something here that says, [device]. We had- my husband and I have two [devices]. He currently has one; I have another. I'm not sure if this is mine or his."*

Also, both consultants and clients leverage activity logs to check recent account activity such as password changes and account recovery information. On Google, these interfaces include the *Security Checkup* and *Recent security activity* interfaces to which the consultant is referring in this interaction.

> **C7:** *"So right now, I just want to also ask, underneath devices, are there any recent security events?"*
>
> **S4:** *"It just says, 'You signed in on Windows three minutes ago [...]' And there's something about three apps has access to your data. This is very interesting. Email, Edison Mail has access to Gmail Google contacts. I don't know what's Edison mail, may be something he's associated with."*

As illustrated in this interaction, making sense of a client's digital footprint is a critical first step to assessing the security of their accounts. We found that consultants mostly rely on device lists and activity logs in this initial assessment. Unrecognized devices, linked accounts, or linked apps often end up as a potential concern, and the uncertainty about device identification and access exhibited in both quotes above was common in consultations. Both quotes highlight the insufficiency of device descriptions in current ASIs. While devices the client may not recognize may be under the control of an abuser, they might also be devices the client simply forgot or abandoned.

To prove suspicions of account compromise. The wide range of information on an ASI makes them valuable for investigating account compromise or other suspicious login activity. Our findings demonstrated that clients use these interfaces as evidence that would help in their abuse situation. Clients sometimes described how they take screenshots of unfamiliar devices or a recognized abuser's device that appeared on these device lists.

> **S1:** *"I am in the process of divorce. I wanted to have it as evidence. Every time I see that he is connected, I [take] a screenshot to have it as evidence."*

We also find that these interfaces can indicate account compromise through the presence of backdoors (i.e., when an abuser can access the account using recovery access methods irrespective of whether they can use a primary authentication method). For example, in this interaction between a client and consultant, a consultant recommends that the client go through Google's *Security Checkup* interface to check the security issues that might be present in the account.

> **C4:** *"Okay, perfect, so yeah, let's go through those security issues found. So it says take action."*
>
> **S7:** *"And then it says, go to password checkup, so that's one. And then there's sign-in and recovery, confirm your recovery phone, which let me check what the recovery... That's my husband's phone number, lovely."*

In this example, the husband's phone number may give access to the account via recovery workflows. For some clients, such evidence of vulnerability proved so compelling that clients reported confronting their abuser with screenshots of their ASIs.

To make decisions around digital safety. Information on such interfaces also assisted decisions about managing personal risk and safety, including whether a compromised device should be removed from a survivor's account to prevent further access from an abuser. As digital abuse is usually accompanied by other forms of abuse (e.g., physical, emotional, and psychological) [25], limiting an abuser's capacity to inflict harm through digital means may exacerbate abuse through other channels—this is known as *escalation*. Here a consultant talks to a client about the risks associated with signing out an abuser's device from the client's Google account.

> **C16:** *"You recognize the devices under this list, right?"*
>
> **S12:** *"One of the devices I didn't recognize it. I just signed out from it."*
>
> **C16:** *"For safety reasons, and also based on our experience, it is best to log all the devices out that you don't recognize. [We] also encourage people to do safety planning, because in certain cases people do not feel comfortable signing a device out, because they say that the person they're concerned about is going to be more aggressive with them."*

In IPV settings, safety planning is an important step following a tech consultation and requires a social worker or case manager who is knowledgeable about the survivor's abuse history and present situation to create a personalized strategy for a survivor that maximizes safety and minimizes risk. In their work introducing clinical computer security in IPV, Havron et al. [31] argue that tech support services in isolation are not sufficient to address tech abuse, in large part because of the complicated risks associated with escalation.

5.2 Challenges Using Account Security Interfaces

Clients and consultants rely on ASIs to understand digital footprints and potential compromise. In specific, we find that referencing device lists is particularly common in clinic settings by both consultants and clients. However, our findings also surface a variety of challenges faced when attempting to do so. Both clients and consultants had difficulty navigating to and within ASIs, and there was often confusion about terminology and information presented on ASIs. Often, device lists, session lists, access notifications, and activity logs proved insufficient to assess account security, including whether illicit accesses are presently occurring or had occurred in the past.

Difficulties in navigation and pageflow. We discover that consultants and clients spend a significant portion of the consult navigating through different pages to find relevant ASIs. Often, clients are surprised to find out that these interfaces exist in the first place. As we see in the following quote, while the client may have had suspicions that their abuser compromised their account, they lacked the knowledge to confirm. Even though Facebook provides an activity log of in-app actions, this log does not inform users whether messages were read and by whom.

> **S15:** *"The entire time he was actually logged into my Facebook account. He was reading my messages. I just didn't know. I'm aware now that I could have gone into the security part of Facebook and checked to see what devices were logged in, but I didn't even think about it"*.

This is also consistent with our survey of ASIs (Section 3), in which we observed that Facebook's and Google's device list interfaces require navigating through three distinct pages after login, as well as two separate navigation flows on Facebook. This complexity was clearly reflected in our data, as most investigations during consults start with consultants providing clients with step-by-step navigation instructions to get to a specific interface or access device-specific information.

> **C4:** *"Can we try one thing—can we go to settings again? And then the general tab. And then if you scroll down, you will see a profile and device management button. And click on it."*

Consultants sometimes tried to rely on a client's prior knowledge or understanding of one service to give navigation instructions. For example, in one case a consultant (**C11**) knew that the client was familiar with ASIs on Google and provided instructions that attempted to leverage similarities between different services: *"Click on that ... it's similar to Google."* However, the significant differences across services (as shown in Section 3) makes this less effective than it otherwise could be and may cause confusion at a later stage.

ASIs are not only confusing to clients: consultants can also be confused either because they are unfamiliar with the navigation flow or because updates to the interfaces render them less recognizable. In one consultation, a consultant had to pause midway through a security walkthrough:

> **C19:** *"And then go to, let me see. Give me one second. Sorry, the new UI, I don't know what's going on ... Actually, just click on settings. Settings and privacy and then settings."*

These findings are in line with Tseng et al. [65], where they analyzed remote clinical computer security for IPV during the COVID-19 pandemic and found similar challenges with remote device and account investigations—many also rooted in a lack of familiarity with such interfaces. The data here suggests that friction in usability is not due solely to conducting appointments via audio conferencing: both clients and consultants had difficulty aligning on interface terminology and locating specific menus and features within those interfaces.

Confusion around identifying devices. Our data consistently shows that clients and consultants use ASIs (particularly device lists) to infer whether unauthorized account access has occurred. Device lists aid identification by displaying information like the OS, platform, location, and time of access (see Figure 2). In some cases, this works as intended, allowing a client to confirm that an entry on a device list is associated to the abuser given the device model or login location:

> **S10:** *"There were two phones on my thing. He's from [country], and one of them said it was in [country] when you look at the location."*

But more often clients and consultants struggled with device identification, as the information provided within ASIs proved ambiguous and insufficient (as shown with S4 and S12 in Section 5.1). This confusion can sometimes arise from the misunderstanding of the permanence or duration of a session or device log on these interfaces. Here a client is confused about a second iPhone on the device list which they used previously to log into the account but have since turned off.

> **C19:** *"So it's apparently your Mac that you're currently using [...] and the two iPhones. [...] So one iPhone is probably your phone. [...] The other iPhone, can you see the time?"*
>
> **S3:** *"It says February [date and time redacted]."*
>
> **C19:** *"Alright. February [date and time redacted]. So your phone was turned off for a long time, is that correct?"*
>
> **S3:** *"Yeah, I still haven't turned it on. I have it in a cabinet, in a box [...] If the phone isn't turned on, is it still somehow logged in?"*

For most services, there is no formal documentation on the lifetime of the logs. Google is the only service that specifies that device, session, and activity logs are available for 28 days.

Both clients and consultants can sometimes confuse sessions from a single device for multiple devices in an ASI that combines both sessions and devices. In this interaction, a client (S10) is confused about a second Mac device appearing on Google's *Your devices* interface.

> **S10:** *"I don't know why there's a second one there."*
>
> **C16:** *"That can mean a couple of things. [...] Sometimes you can get different browsers showing as the same computer. For example, if you signed in on Safari or if you signed in on Chrome, then sometimes they can show as separate devices."*

Finally, we also discovered that clients cannot assume ownership based on where their device is physically located according to the ASI. In one case, a client is confused about a refurbished device they bought online which they think might have appeared in the device list on Facebook as an iPad that is logging in from an unfamiliar location.

> **S21:** *"You know, I'm worried because it says iPad, [redacted city]. I don't know why. But it's this phone that I bought refurbished from [online retailer]. That's the iPad. I don't know what's this iPad from [redacted city]. Maybe previously it was in [redacted-city]."*

As the interface was unclear, the client had to make jumps in reasoning such as presuming the physical location of the device rather than easily interpreting its past location.

Device- and phone number-to-account mapping. We find that it can be confusing for clients to determine how devices and phone numbers are connected to an account, as well as what an abuser has access to at what time. For example, here a client expresses concern following a physical compromise that an abuser knowing her phone number might enable them to monitor their messages:

> **S14:** *"He followed me to the store because I tried to change my number. [...] He bought the same phone, and the guy said my phone number out loud, and he has my phone number. And so there were a few times after that [...] where he has grabbed my phone and not given it back. [...] He has my phone number—is he getting my messages, you know? It's just very confusing."*

Confusion over what devices and phone numbers are associated to an account can also render some security tasks such as recovering account access more challenging. For example, some of the clients struggled to make sense of Apple's recovery process because of their limited understanding of Apple's trusted devices and phone numbers (see Section 3).

In the following interaction, a client (S11) expresses their frustration with an inflexible Apple account recovery process at a time where they no longer have access to a trusted phone number: *"[Apple] told me that it was sending a message to the phone number; and I don't have the phone number anymore, so it won't let me go to the next step"*. The consultant then goes on to explain to the client that they have to be logged in on a trusted Apple device to recover access to their account. This suggests the need for services to delineate between different terminologies mentioned on the interfaces in a way that a user may understand.

> **C7**: *"Because your phone is a new phone and you have not accessed the iCloud account from this phone, it does not recognize the device, so it might be helpful if you try to recover your account from the device that you think you already logged into the iCloud [...] so it's a device that Apple recognizes once you try to recover your account. Do you have access to such a device?"*

Account access discussions could motivate a consultant to provide further account security advice, such as changing passwords, configuring or changing recovery information, and turning on 2FA. While prior studies show that configuring 2FA and account recovery is challenging for most users [4, 53], we found that clients also struggled with configuring 2FA or account recovery when they were unable to anticipate how such changes could potentially alert an abuser. For IPV survivors, who may frequently change devices and numbers seeking to avoid their abuser, this makes it challenging to even remember what accounts and devices can be trusted.

> **S11:** *"[2FA] is already turned on. [...] I never used it before. [...] [It] gives you [a] one time code, right?"*
>
> **C7:** *"I see. In that case, you probably want to turn off the two-factor authentication so that somebody else is not getting a notification every time you log into Facebook."*
>
> **S11:** *"So should I put that I want to receive it as a text message?"*
>
> **C7:** *"Sure, text message as long as it is your phone number and not somebody else's phone number."*
>
> **S11:** *"Well, I mean I'm hoping it is my phone number."*

6 Integrity of Access Identification

Our research has shown that in IPV contexts, ASIs are critical resources for assessing account security. Clients and consultants rely particularly heavily on the device and access details presented in device lists, session lists, and activity logs to diagnose whether illicit accesses have occurred. Thus it is critical that these details are trustworthy.

Service	Account Security Interface (ASI)	Spoofability					
		Device Model	Operating System	Browser	Location	Date	Time
Google	Recent Security Activity	●	n/a	n/a	×	○	○
	Your Devices	●	●	●	×	○	○
	Security Checkup	●	●	●	×	○	○
	Email login notifications	●	n/a	n/a	n/a	n/a	n/a
	Find Your Phone	●	n/a	n/a	×	○	○
Facebook	Where you're logged in	n/a	●	●	●	○	○
	Authorized Logins	●	●	●	●	○	○
	Recognized Devices	●	●	●	●	○	◐
	Logins and Logouts	n/a	n/a	n/a	●	○	○
	Active Sessions	●	●	●	●	○	◐
	Login Alerts (email & in-app)	n/a	●	●	●	○	○
Apple	Devices	○	○	n/a	n/a	n/a	n/a
	Email Notifications	n/a	n/a	n/a	n/a	○	○
	Login Push Notifications	n/a	n/a	n/a	●	n/a	n/a
WhatsApp	Linked Devices	n/a	●	●	n/a	○	○

Figure 4: Spoofability of device identification fields shown on the specified ASI or notification mechanism. A ● symbol indicates that the referenced field is fully spoofable; a ◐ symbol indicates that the field is partially spoofable; a ○ symbol indicates that we could not spoof the field; a × symbol indicates the field cannot be spoofed but can be suppressed (location is hidden); and *n/a* means that the interface does not display the corresponding field.

There is ample reason for concern, as previous research suggests that risk-based authentication (RBA) mechanisms that rely on similar information (such as user agents and IP addresses) can be defeated by special-purpose tools [75]. In this section, we therefore investigate the following question: *can abusers easily undermine the integrity of access identification information provided on user-facing ASIs?* The answer, unfortunately, is yes, which means abusers can conceal ongoing monitoring and full control over victim accounts.

Threat model. We assume the adversary has the ability to log into a victim's account, but is doing so from a distinct device. Note that for WhatsApp, this translates to having temporary physical access to the unlocked device of the victim—an assumption that holds in various settings including some IPV situations. A subset of our results also make sense in the context of using the same device (a situation that arises in IPV and other domestic abuse scenarios); but we mostly focus on the distinct device setting since it is harder for the adversary.

The adversary's goal is to log into the victim's account, while ensuring that the device and associated login session either (1) do not appear on any of the account's ASIs or (2) do appear, but the information displayed makes it appear to be a benign login from the victim's device. In the first case, we say that the adversary has hidden their login, and in the second case we say that they have spoofed access identification information.

In terms of capabilities, we focus on less sophisticated attacks that may be within reach of a broader class of adversaries. In IPV and other abuse scenarios, most abusers are not employing technically sophisticated approaches and, in fact, are considered UI-bound adversaries [25]. Such adversaries rely only on readily available software and only operate within the confines of that software's features as provided by the standard UI. As we will see, our attacks will fall somewhere between standard UI-bound adversaries and the traditional worst-case adversaries assumed in computer security: they will sometimes use widely available but arguably arcane existing UIs and tools.

How access identification works. As discussed in Section 3, ASIs present to the user information about what devices have accessed an account. The information about such accesses varies across platforms but often includes the device type (web or mobile), operating system (OS), OS version, browser, browser version, location of the device, and finally the date and time of the most recent access.

How specific services infer this information is not documented in detail, but the modern web architecture means that this information must be inferred from HTTP requests which includes the user agent, the HTTP date header, and network-level information such as the IP address. A typical user agent consists of a number of different identifiers: a general Mozilla compatibility token that signifies a browser's compliance with Mozilla web standards, a platform identifier that identifies the native platform that the browser is running on, a layout or browser engine string, and a browser version. While there are plans to improve privacy by deprecating the user agent in the future and replace it with a feature called *client hints* [3], the status quo is that user agents remain in wide use.

There exist other RBA mechanisms such as browser and device fingerprinting libraries [7, 22, 27, 73]. We hypothesized and our experiments indicate that even if these are used by a service, they do not impact hiding or spoofing attacks.

Methodology. We investigate access hiding and spoofing

attacks for our running example set of four major services. To do so, we set up test accounts (with fake user information) on each of the four platforms for experiments; these accounts play the role of the victim. We then use an Apple Macbook Pro laptop running Mac OS 12 (Monterey) to simulate the role of an adversary's device and log into all services through the browser (including WhatsApp, for which we used WhatsApp Web). We also ran experiments using an Android phone and a Mac OS virtual machine for the adversarial system, but there were no changes in process or results; for simplicity, we report solely on the results obtained on the Macbook.

For each service, we experimented with various modifications to local settings, including modifying the local clock, changing the user agent, and configuring use of VPNs. We then simulated an adversarial session by logging into the target service using an incognito browsing window. We tested both explicitly logging out of the adversarial session (using service UIs) and not doing so. Finally, we log in from a separate device that plays the role of the victim's device and inspect access interfaces to determine what is relayed to the victim about the adversarial session. We experimented with using both a mobile platform (Android phone) and desktop (Macbook using the Chrome browser) for the simulated victim device. Unless otherwise indicated, the victim's views were consistent across the different devices. The explicit interfaces investigated are listed in Figure 4.

Experiments were conducted between August 2022 and February 2023; we didn't observe any changes in the ASIs under study in that timeframe. However, we experimented with Facebook's email login alerts in May 2023 after Facebook introduced the *Accounts center* feature that surfaced the option to opt into email-based login alerts and a *recent emails* interface.

Access hiding attacks. We start by investigating the ability for an adversary to hide access completely from user-visible interfaces. In this case, we perform login from the adversarial device and then perform explicit logout.

First, we note that Apple is trivially vulnerable to hiding attacks, as the only access identification interface supported is the "Devices" list, and this only ever includes Apple devices where the OS handled login. In more detail, any logins through a browser to `appleid.apple.com` fail to trigger changes in any user-facing interface: this gives an attacker full control over a victim's account (including the ability to change Apple account password, reset account recovery configurations, and more). However, logins through a browser to `icloud.com` do trigger an email login notification that contains a date and time identifier. They also may trigger a 2FA push challenge with location information to previously authorized devices. Given that notifications have been shown to be confusing and often ignored [45, 54], we view the fact that no ASI allows determining that an access occurred as a serious deficiency.

Other services are also vulnerable to hiding attacks. In WhatsApp, a user can just log out their device, removing any trace of the login. Thus an abuser who has temporary access to a victim's unlocked device can set up another abuser-controlled device to receive all WhatsApp messages and impersonate the victim. Should they later log out, all trace of the access is gone from the WhatsApp ASI.

On Facebook's *Where you're logged in* and *Active sessions* interfaces, we find that adversarial sessions that are correctly terminated are hidden. The *Logins and logouts* interface remains the only interface on Facebook that contains information about a logged out session. Furthermore, this interface is the one with the least identifiers—only the IP address and time/date information is shown. We hypothesize that this inconsistency across the three different interfaces is a likely further source of confusion for victims.

Similarly, we find that on Google's *Your Devices* interface an adversarial session can be hidden entirely if it does not trigger a *New sign-in* activity alert on the *Recent security activity* interface (i.e., the session emanates from a device that is already listed on the interfaces).

Overall, we find it concerning that it is so easy to remove any trace of a login session from these various ASIs.

Access spoofing attacks. We then conducted experiments to assess how easily we can spoof active, ongoing sessions to appear as the victim's device and hinder identification of these illicit accesses. For these experiments, we assume that the simulated adversary did not explicitly log out after logging in. Instead, we investigated the extent to which the adversary can control the details of what is displayed on ASIs to the victim just by changing the adversary's own local settings. See Figure 4 for a summary of our experimental results across ASIs for the four services when the adversary attempts to spoof the device model, OS, browser, location, date, and time displayed to the victim.

This spoofing was simple to do, as the services in many cases appear to rely completely on untrusted client-chosen data. Most modern browsers—including Google Chrome, Safari, and Firefox—allow easily overriding the user agent sent by the client via in-browser developer tool features. We found that by using these developer tools and modifying the user agent field, the adversary can easily change the device model, OS, and browser displayed to the victim for the interfaces with these fields on Google, Facebook, and WhatsApp. Additionally, we find that external email-based interfaces and notifications are also spoofable for both Facebook and Google. We did not find this to be the case for Apple.

To modify the time of access, we experimented with changing the local time (hours and/or minutes) on the adversary's machine. We also disabled automatic time synchronization. After the adversary had successfully logged into the victim's account, we then used the victim's device (where the time was set correctly and automatic time synchronization was enabled) to log into the account and record the time of entries associated with the adversary's login. We find that on

two of Facebook's interfaces when the simulated victim accesses from a desktop, the time is not spoofed. But when the victim uses the mobile version of Facebook's website,[2] the time shown to the victim is the adversarially-specified one, meaning spoofing succeeded. We note that the adversary can even make the session appear as if it happened in the future. This suggests that this particular interface is pulling the time from the local client, which is untrustworthy. On WhatsApp and Google, however, date and time were not spoofable, indicating that the service does not pull the date and time from the client. Apple's email notifications contain a date and time that we could not spoof.

Finally, the adversary can easily spoof locations using virtual private network (VPN) tools. We confirmed this using the basic, free version of Proton VPN, which allows us to select the location of the VPN's IP address at the granularity of a country. We also confirmed this using Tor, setting the exit node to the spoofed location. This worked for Apple and Facebook—WhatsApp does not display location information on their *Linked Devices* interface. For Apple, the only interface that displays a location identifier is the login push notifications sent to a trusted device upon login to an Apple account. We confirmed that an adversary can spoof the location on the push notifications that are sent to all trusted devices associated with an Apple account. On Google, spoofing the location leads to hiding the location identifier from the interface rather than displaying the spoofed location; in such instances we say that the location is suppressed.

7 Discussion

The results of the past few sections show how both experts and non-experts respond to ASIs in suspected or actual attack situations, complementing prior studies on logins [54] and login notifications [45]. Our findings in particular highlight the importance and limitations of these interfaces in assisting both experts (consultants) and non-experts (survivors) in diagnosing security posture. Our findings speak to broader issues that affect a wider range of users, but with specific lessons for at-risk users. We discuss the need for future work towards improvements and in so doing, highlight key tensions that make solutions difficult.

Making security interfaces easier to find. The challenges surfaced by our studies partially stem from confusion over how a user navigates to relevant ASIs. Some services have multiple navigation workflows to access a feature or perform a given task. Clients in our study who were purportedly less competent in their level of technical expertise found these navigation routes complicated (Section 5.2).

Cognitive walkthroughs for the web [43] (a host of task-based usability-inspection methods) and visual customer flows [15] (design tools to assist user navigation by removing obstacles) are well established tools that can be used in response to the usability problems that we encountered. Universal metrics such as the *predicted mean total clicks* (how many clicks a user needs to use before accomplishing a task) or the *event count* (the number of concrete steps required to achieve a goal) of a funnel analysis can help to demonstrate the severity of an issue. Usability metrics can be one helpful guide to redesigning such interfaces. The use of visual cues or 'signifiers' such as open-text boxes to type usernames and passwords are now ubiquitous indicators of access and authentication [48]. We suggest that there could be efforts to find how device and account access signifiers could also breed familiarity and user trust.

One tension facing improved usability is that in some threat models (like IPV), authenticated attackers may also benefit from improvements—for example, it might be easier for attackers to find activity logs for covert surveillance or configuration interfaces to lock the legitimate user from the account. Whether it is even possible to add friction to abuse use cases while easing friction for legitimate use cases is an interesting open question.

Standardization of interface design. Our study surfaced difficulties users had with interpreting device lists and activity logs. One contributing factor could be that, across the services studied, different terminologies are used on the interfaces to refer to common authentication concepts and device identification features. As a quick recap, Facebook referred to a trusted device (a device that skips 2FA) as a *recognized* device and to a list of trusted devices as *authorized logins*—a differentiation that is not universal even across its own service. Such inconsistencies are carried into information architecture, whereby Facebook's list of trusted devices is located in *Security and Login* and recognized devices are located via *Logged Actions*—while Google does not provide a list of trusted devices, and Apple does not provide a list of authorized devices at all. When all these factors are taken into account, it is unsurprising that experts and non-experts alike are confused, and this unfortunately leads to elevated risks of recommending the wrong form of privacy and security fixes at critical moments.

We cannot comment on the justification for the considerations that went into current designs. However, we suggest that they be reconsidered because the use of inconsistent terminology alone goes against good design principles. We anticipate that future work might draw from Nielsen's usability heuristics for interface design [48], such as designing them to *match between system and the real world* (resisting introducing new words or concepts) and promote the *visibility of system status* (permit users to follow system activity).

We emphasize that consistency is really an industry-wide issue: in our study, clients could sometimes comprehend how the interface on one service worked, but we had numerous examples where these mental models were incompatible with another service.

[2] https://m.facebook.com

Improving device lists and activity logs. Our results indicate that users were unable to make judgments about account compromise or other security decisions. This is due to a lack of information about the devices logged into their account and the ambiguity of device identifiers.

All services could include a "Recent Logins" flow that provides a list of accesses with best-effort device identification. Ideally, device identifiers would be static (unlike IP addresses) and easy for the legitimate owner to associate to a given device. For example, serial or IMEI numbers are static and at least allow comparing with devices to which the legitimate owner has access. As the use of mobile devices is particularly widepsread (in IPV [21, 25] and beyond), associating a device's phone number to an access may be helpful to some users. Apple already supports this for relevant iOS devices. (see Figure 3). But phone numbers can be changed, some devices have multiple phone numbers associated to them (e.g., due to multiple SIMs), some devices (laptops, tablets, embedded devices) do not have phone numbers, and in some cases client software does not have the privilege to obtain the phone number from the OS.

For more advanced users or in clinical settings where an expert consultant is assisting a user, we can also imagine augmenting access lists with the ability to click through to obtain more detailed information about accesses. The more advanced interface could render a view closer to the one seen by the service—i.e., accesses should be based off the requested headers and session cookies that the service actually uses to identify sessions.

Designers might also consider including more detailed information about what happened during particular sessions—for example, a session activity summary to help users assess whether sessions were malicious or not. This activity summary could log actions on the account like reading messages and adding or removing 2FA and other security configurations among other things. Facebook currently provides users with a *Logged Actions* interface that keeps a record of in-app search history and other activity but other services do not.

Tensions with privacy. In contemplating such enriched ASIs, a key tension that emerges is between forensic benefit and privacy. First, any detailed logging that is user-visible can also be employed by an adversary that successfully logs into the account. One partial solution would be to adopt more broadly the pattern of forcing additional authentication challenges when accessing these pages, but in some threat models this will not prevent access. Another possibility would be allowing detailed activity logs to be opt-in (or opt-out): once turned on, it should not be possible to turn off without clear, permanent notice (such as a banner indicating when the feature was last enabled and last disabled).

Second, we must protect user privacy against services and adversaries that can access them (e.g., via system compromise, subpeona, etc.). Services may want to limit the duration of data they keep on user behavior as a matter of policy, such as Google's 28 day limit on past sessions. Again, allowing users the option of whether to set these limits and the types of information stored may be beneficial.

Practitioners and researchers have spent decades trying to make it difficult for web services to precisely track individual devices (for a small subset of recent work, see [14, 34, 38, 49]) because this could be abused by companies to track users. Thus the types of identifiers mentioned earlier (serial number, IMEI, phone number) may not be available to client applications as a matter of policy by the OS and giving access would allow malicious apps and services to track users.

We point out that this tension between device tracking and device identification does not seem to necessarily be fundamental. For the latter, we are concerned with the *legitimate user's* ability to track what devices are used with their account, rather than the *service's*. This observation suggests a provocative possibility—that we might improve device identification for users while avoiding service-based tracking by rearchitecting clients and web services to communicate just to authenticated users what devices have accessed an account. Done right, this could improve the integrity of device identification in access interfaces without enabling new, invasive tracking of users by services. But doing so would seem to require service-blind, persistent device identifiers, which are not provided by current OS and HTTP protocol designs. Thus future work is needed.

8 Conclusion

We explored the landscape of how services tell users about their security status. We are the first to look at how users understand and interact with security interfaces that describe devices and activity associated with an account. We performed a case study to understand how these interfaces are used to assess the security posture of IPV survivors' online accounts under imminent threat from an intimate partner abuser. Overall, we find that account security interfaces play a major role in detecting account compromise but that they need much improvement in their security and usability.

Acknowledgements

The authors would like to thank the CETA consultants and clients who made this work possible, as well as Emily Tseng for feedback on earlier drafts of the paper. This work was funded in part by NSF grants CNS-1916096 and CNS-2120651.

References

[1] Add or Remove Trusted Computers - Android - Google Account Help. https://support.google.com/accounts/answer/2544838?hl=en&co=GENIE.Platform%3DAndroid. (Accessed on 05/31/2023). 4

[2] Add or Remove Trusted Devices on Mac - Apple Support. https://support.apple.com/guide/mac-help/add-or-remove-trusted-devices-mchl2310b175/mac. (Accessed on 05/25/2023). 4

[3] Improving User Privacy and Developer Experience with User-Agent Client Hints - Chrome Developers. https://developer.chrome.com/articles/user-agent-client-hints/?utm_source=devtools. (Accessed on 01/21/2023). 10

[4] Claudia Ziegler Acemyan, Philip Kortum, Jeffrey Xiong, and Dan S Wallach. 2FA Might be Secure, But it's not Usable: A Summative Usability Assessment of Google's Two-factor Authentication (2FA) Methods. In *Human Factors and Ergonomics Society Annual Meeting*, 2018. 1, 3, 9

[5] Anne Adams and Martina Angela Sasse. Users Are Not The Enemy. *Communications of the ACM*, 1999. 2

[6] Devdatta Akhawe and Adrienne Porter Felt. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *USENIX Security*, 2013. 3

[7] Furkan Alaca and Paul C Van Oorschot. Device Fingerprinting for Augmenting Web Authentication: Classification and Analysis of Methods. In *ACSAC*, 2016. 10

[8] Julio Angulo and Martin Ortlieb. "WTH..!?!" Experiences, Reactions, and Expectations Related to Online Privacy Panic Situations. In *SOUPS*, 2015. 3

[9] Mihir Bellare and Bennet Yee. Forward Integrity for Secure Audit Logs. Technical report, 1997. Preprint. 3

[10] Mihir Bellare and Bennet Yee. Forward-security in Private-key Cryptography. In *Topics in Cryptology—CT-RSA*, 2003. 3

[11] Rosanna Bellini, Emily Tseng, Nora McDonald, Rachel Greenstadt, Damon McCoy, Thomas Ristenpart, and Nicola Dell. "So-called Privacy Breeds Evil": Narrative Justifications for Intimate Partner Surveillance in Online Forums. *ACM on Human Computer Interaction*, 2021. 1, 5

[12] Arnar Birgisson, Joe Gibbs Politz, Ulfar Erlingsson, Ankur Taly, Michael Vrable, and Mark Lentczner. Macaroons: Cookies with Contextual Caveats for Decentralized Authorization in the Cloud. 2014. 4

[13] Joseph Bonneau, Cormac Herley, Paul C Van Oorschot, and Frank Stajano. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *IEEE Security & Privacy*, 2012. 3

[14] Justin Brookman, Phoebe Rouge, Aaron Alva, and Christina Yeung. Cross-Device Tracking: Measurement and Disclosures. *PETS*, 2017. 13

[15] Laura M. Castro, David Cabrero, and Rüdiger Heimgärtner. *Software Usability*. BoD – Books on Demand, 2022. 12

[16] Dana Cuomo and Natalie Dolci. New Tools, Old Abuse: Technology-Enabled Coercive Control (TECC). *Geoforum*, 2021. 1, 5

[17] Alaa Daffalla, Lucy Simko, Tadayoshi Kohno, and Alexandru G Bardas. Defensive Technology Use by Political Activists During the Sudanese Revolution. In *IEEE Security & Privacy*, 2021. 1

[18] Sanchari Das, Bingxing Wang, and L Jean Camp. MFA is a Waste of Time! Understanding Negative Connotation Towards MFA Applications via User Generated Content. *arXiv preprint arXiv:1908.05902*, 2019. 3

[19] Sanchari Das, Bingxing Wang, Zachary Tingle, and L Jean Camp. Evaluating User Perception of Multi-factor Authentication: A Systematic Review. *arXiv preprint arXiv:1908.05901*, 2019. 3

[20] Emiliano De Cristofaro, Honglu Du, Julien Freudiger, and Greg Norcie. A Comparative Usability Study of Two-factor Authentication. *arXiv preprint arXiv:1309.5344*, 2013. 3

[21] Jill P Dimond, Casey Fiesler, and Amy S Bruckman. Domestic Violence and Information Communication Technologies. *Interacting with Computers*, 2011. 13

[22] Peter Eckersley. How Unique is Your Web Browser? In *PETS*, 2010. 10

[23] Christen Erlingsson and Petra Brysiewicz. A Hands-on Guide to Doing Content Analysis. *African Journal of Emergency Medicine*, 2017. 2, 6

[24] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. "Is My Phone Hacked?" Analyzing Clinical Computer Security Interventions With Survivors of Intimate Partner Violence. *ACM on Human Computer Interaction*, 2019. 1, 5

[25] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. "A Stalker's Paradise" How Intimate Partner Abusers Exploit Technology. In *CHI*, 2018. 1, 5, 7, 10, 13

[26] Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. Digital Technologies and Intimate Partner Violence: A Qualitative Analysis with Multiple Stakeholders. *ACM on Human Computer Interaction*, 2017. 1, 5

[27] David Freeman, Sakshi Jain, Markus Dürmuth, Battista Biggio, and Giorgio Giacinto. Who Are You? A Statistical Approach to Measuring User Authenticity. In *NDSS*, 2016. 10

[28] Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. " What was that site doing with my Facebook password?" Designing Password-Reuse Notifications. In *CCS*, 2018. 3

[29] Tamy Guberek, Allison McDonald, Sylvia Simioni, Abraham H Mhaidli, Kentaro Toyama, and Florian Schaub. Keeping a Low Profile? Technology, Risk and Privacy Among Undocumented Immigrants. In *CHI*, 2018. 1

[30] Bridget Harris. Technology, Domestic and Family Violence: Perpetration, Experiences and Responses. 2020. 5

[31] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. Clinical Computer Security for Victims of Intimate Partner Violence. In *USENIX Security*, 2019. 1, 5, 7

[32] Jason Hong. The State of Phishing Attacks. *Communications of the ACM*, 2012. 3

[33] Lin-Shung Huang, Alexander Moshchuk, Helen J Wang, Stuart Schecter, and Collin Jackson. Clickjacking: Attacks and Defenses. In *USENIX Security*, 2012. 3

[34] Thomas Hupperich, Davide Maiorca, Marc Kührer, Thorsten Holz, and Giorgio Giacinto. On the Robustness of Mobile Device Fingerprinting: Can Mobile Users Escape Modern Web-Tracking Mechanisms? In *ACSAC*, 2015. 13

[35] Michael Jones, John Bradley, and Nat Sakimura. Json Web Token (JWT). Technical report, 2015. 4

[36] David M Kristol. HTTP Cookies: Standards, Privacy, and Politics. *ACM Transactions on Internet Technology*, 2001. 4

[37] Leona Lassak, Annika Hildebrandt, Maximilian Golla, and Blase Ur. "It's Stored, Hopefully, on an Encrypted Server": Mitigating Users' Misconceptions About FIDO2 Biometric WebAuthn. In *USENIX Security*, 2021. 3

[38] Adam Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016. In *USENIX Security*, 2016. 13

[39] Karen Levy and Bruce Schneier. Privacy Threats in Intimate Relationships, 2020. 5

[40] Xu Lin, Panagiotis Ilia, Saumya Solanki, and Jason Polakis. Phish in Sheep's Clothing: Exploring the Authentication Pitfalls of Browser Fingerprinting. In *USENIX Security*, 2022. 2, 3

[41] Sanam Ghorbani Lyastani, Michael Schilling, Michaela Neumayr, Michael Backes, and Sven Bugiel. Is FIDO2 the Kingslayer of User Authentication? A Comparative Usability Study of FIDO2 Passwordless Authentication. In *IEEE Security & Privacy*, 2020. 1, 3

[42] Di Ma and Gene Tsudik. A New Approach to Secure Logging. *ACM Transactions on Storage*, 2009. 3

[43] Thomas Mahatody, Mouldi Sagar, and Christophe Kolski. State of the Art on the Cognitive Walkthrough Method, Its Variants and Evolutions. *ACM on Human Computer Interaction*, 2010. 12

[44] Philipp Markert, Andrick Adhikari, and Sanchari Das. A Transcontinental Analysis of Account Remediation Protocols of Popular Websites. *arXiv preprint arXiv:2302.01401*, 2023. 3

[45] Philipp Markert, Leona Lassak, Maximilian Golla, and Markus Dürmuth. "I Knew It Was Me": Understanding Users' Interaction with Login Notifications. *arXiv preprint arXiv:2212.07316*, 2022. 1, 3, 11, 12

[46] Tara Matthews, Kathleen O'Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. Stories From Survivors: Privacy & Security Practices When Coping With Intimate Partner Abuse. In *CHI*, 2017. 1

[47] Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner. Investigating the Computer Security Practices and needs of Journalists. In *USENIX Security*, 2015. 1

[48] Jakob Nielsen. Enhancing the Explanatory Power of Usability Heuristics. In *CHI*, 1994. 12

[49] Nick Nikiforakis, Alexandros Kapravelos, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. Cookieless monster: Exploring the Ecosystem of Web-Based Device Fingerprinting. In *IEEE Security & Privacy*, 2013. 13

[50] United States. Department of Defense. *Department of Defense Trusted Computer System Evaluation Criteria*. Department of Defense, 1987. 3

[51] Aleksandr Ometov, Sergey Bezzateev, Niko Mäkitalo, Sergey Andreev, Tommi Mikkonen, and Yevgeni Koucheryavy. Multi-Factor Authentication: A Survey. *Cryptography*, 2018. 3

[52] Christopher Parsons, Adam Molnar, Jakub Dalek, Jeffrey Knockel, Miles Kenyon, Bennett Haselton, Cynthia Khoo, and Ron Deibert. The Predator in Your Pocket: A Multidisciplinary Assessment of the Stalkerware Application Industry. Technical report, University of Toronto, 2019. 5

[53] Thanasis Petsas, Giorgos Tsirantonakis, Elias Athanasopoulos, and Sotiris Ioannidis. Two-factor Authentication: Is the World Ready? Quantifying 2FA Adoption. In *EuroSEC*, 2015. 3, 9

[54] Elissa M Redmiles. "Should I Worry?" A Cross-Cultural Examination of Account Security Incident Response. In *IEEE Security & Privacy*, 2019. 1, 3, 11, 12

[55] Ken Reese, Trevor Smith, Jonathan Dutson, Jonathan Armknecht, Jacob Cameron, and Kent Seamons. A Usability Study of Five Two-Factor Authentication Methods. In *SOUPS*, 2019. 3

[56] Joshua Reynolds, Trevor Smith, Ken Reese, Luke Dickinson, Scott Ruoti, and Kent Seamons. A tale of Two Studies: The Best and Worst of Yubikey Usability. In *IEEE Security & Privacy*, 2018. 1

[57] Kevin A Roundy, Paula Barmaimon Mendelberg, Nicola Dell, Damon McCoy, Daniel Nissani, Thomas Ristenpart, and Acar Tamersoy. The Many Kinds of Creepware Used for Interpersonal Attacks. In *IEEE Security & Privacy*, 2020. 5

[58] Scott Ruoti, Brent Roberts, and Kent Seamons. Authentication Melee: A Usability Analysis of Seven Web Authentication Systems. In *WWW*, 2015. 3

[59] Johnny Saldana. *The Coding Manual for Qualitative Researchers*. SAGE Publications, 3rd edition, 2015. 6

[60] Richard Shay, Iulia Ion, Robert W Reeder, and Sunny Consolvo. "My religious aunt asked why I was trying to sell her viagra" Experiences With Account Hijacking. In *CHI*, 2014. 3

[61] Lucy Simko, Ada Lerner, Samia Ibtasam, Franziska Roesner, and Tadayoshi Kohno. Computer Security and Privacy for Refugees in the United States. In *IEEE Security & Privacy*, 2018. 1

[62] Marianne Swanson and Barbara Guttman. *Generally Accepted Principles and Practices for Securing Information Technology Systems*. NIST, 1996. 3

[63] Kurt Thomas, Frank Li, Chris Grier, and Vern Paxson. Consequences of Connectivity: Characterizing Account Hijacking on Twitter. In *CCS*, 2014. 3

[64] Emily Tseng, Rosanna Bellini, Nora McDonald, Matan Danos, Rachel Greenstadt, Damon McCoy, Nicola Dell, and Thomas Ristenpart. The Tools and Tactics Used in Intimate Partner Surveillance: An Analysis of Online Infidelity Forums. In *USENIX Security*, 2020. 1, 5

[65] Emily Tseng, Diana Freed, Kristen Engel, Thomas Ristenpart, and Nicola Dell. A Digital Safety Dilemma: Analysis of

Computer-Mediated Computer Security Interventions for Intimate Partner Violence During COVID-19. In *CHI*, 2021. 5, 8

[66] Emily Tseng, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. Care Infrastructures for Digital Security in Intimate Partner Violence. In *CHI*, 2022. 5

[67] Blase Ur, Fumiko Noma, Jonathan Bees, Sean M Segreti, Richard Shay, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. "I Added '!' at the End to Make It Secure": Observing Password Creation in the Lab. In *SOUPS*, 2015. 1

[68] Courtland VanDam, Jiliang Tang, and Pang-Ning Tan. Understanding Compromised Accounts on Twitter. In *International Conference on Web Intelligence*, 2017. 3

[69] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L Mazurek, Manya Sleeper, and Kurt Thomas. Sok: A framework for Unifying At-risk User Research. In *IEEE Security & Privacy*, 2022. 1

[70] Nicole Westmarland, Mariann Hardey, Hannah Bows, Dawn Branley, Mehzeb Chowdhury, Katie Wheatley, and Richard Wistow. Protecting Women's Safety? The Use of Smartphone 'Apps' in Relation to Domestic and Sexual Violence. SASS Research Briefing no. 12, University of Durham, 2013. 5

[71] Stephan Wiefling, Markus Dürmuth, and Luigi Lo Iacono. Verify It's You: How Users Perceive Risk-based Authentication. *IEEE Security & Privacy*, 2021. 3

[72] Stephan Wiefling, Markus Dürmuth, and Luigi Lo Iacono. More than Just Good Passwords? A Study on Usability and Security Perceptions of Risk-based Authentication. In *ACSAC*, 2020. 3

[73] Stephan Wiefling, Markus Dürmuth, and Luigi Lo Iacono. What's in Score for Website Users: A Data-driven Long-term Study on Risk-based Authentication Characteristics. In *Financial Cryptography and Data Security*, 2021. 3, 10

[74] Stephan Wiefling, Luigi Lo Iacono, and Markus Dürmuth. Is this Really You? An Empirical Study on Risk-based Authentication Applied in the Wild. In *IFIP International Conference*, 2019. 3

[75] Stephan Wiefling, Tanvi Patil, Markus Dürmuth, and Luigi Lo Iacono. Evaluation of Risk-based Re-authentication Methods. In *IFIP International Conference*, 2020. 3, 10

[76] Stephan Wiefling, Jan Tolsdorf, and Luigi Lo Iacono. Privacy Considerations for Risk-based Authentication Systems. In *IEEE EuroSEC Workshops*, 2021. 3

[77] Delanie Woodlock. The Abuse of Technology in Domestic Violence and Stalking. *Violence against women*, 23(5):584–602, 2017. 5

A Appendix

In this section, we present data that supports our work and findings.

Figure 5 shows the navigation flow for users to access ASIs on Google—specifically, the *Security Checkup, Your devices, and Recent security activity* interfaces. Similarly, in Figure 7, we present the UI paths on Facebook for a user to access the five access identification interfaces on the service.

Figure 6 and Figure 8 give additional context as to how we conducted our qualitative analysis; we show the list of keywords used in the transcript search and the complete codebook for our transcripts. In Figure 9, we provide the URLs for all the interfaces that we discuss in this work, including those in Figure 4. Figure 10 shows a screenshot of a user's view of the *Your devices* interface on Google. The interface shows both the victim's legitimate session on an Android Pixel phone and the adversary's spoofed session.

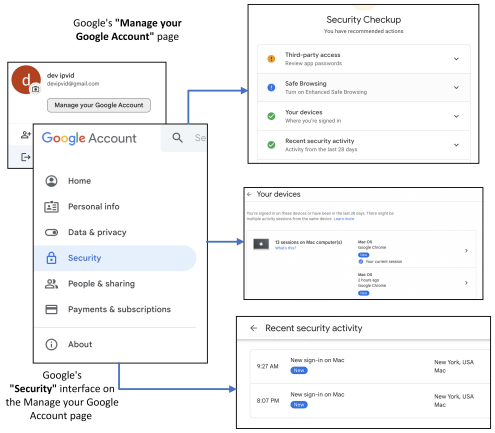

Figure 5: On Google's *Security* interface, account security interfaces include (from top to bottom) *Security Checkup*, *Your devices*, and *Recent security activity*. The account information shown (including the name and email) is that of a fake test account created for the purpose of this work.

Search category	Keyword or phrase
Scope	Known devices Last signed in Recovery information
Device identification	Known devices List of device* Recognize* Your device* "Where you're [you are] signed in"
Account access	Last active *Sign-in history Two-factor* Two-step* Recovery [information, devices]

Figure 6: A comprehensive list of keyword search terms used to find transcripts for our qualitative analysis in the clinic CMS. We include the initial keywords used to confirm the presence of consultant-client conversations in the clinic CMS. For our first search, we identified words related to device identification. We then searched for information about account logs and recovery.

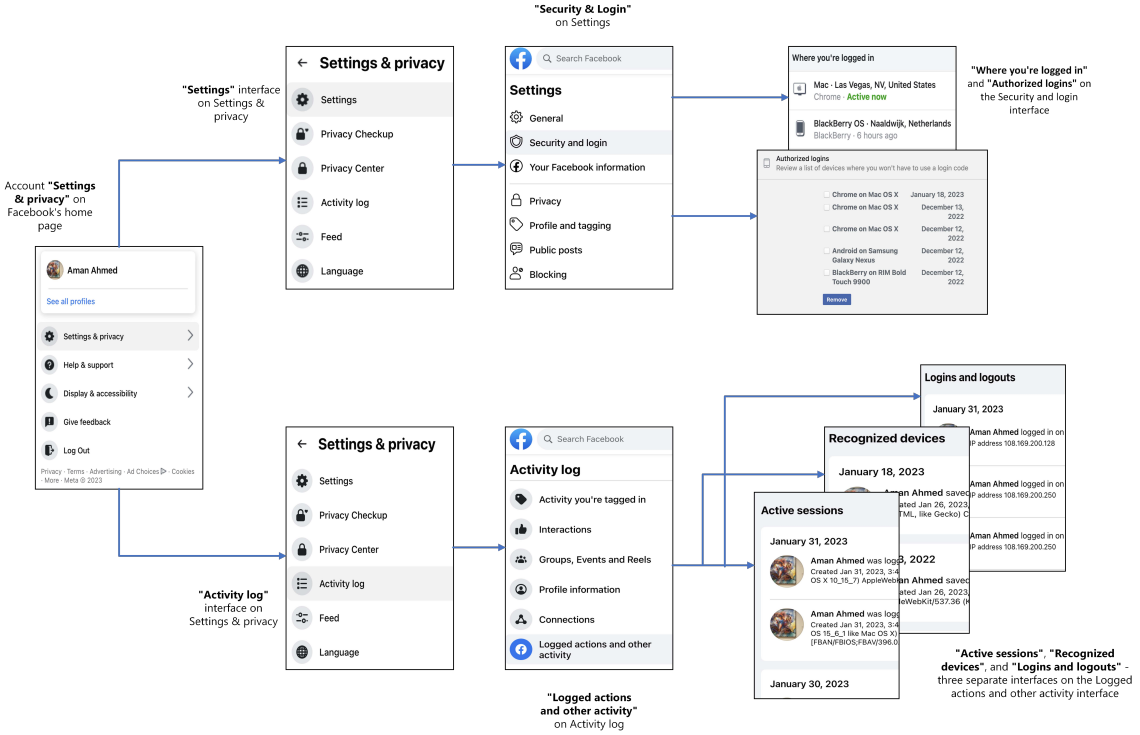

Figure 7: Two separate navigation flows for users to understand account access on Facebook. As shown, a total of five different interfaces show account access information: *Where you're logged in*, *Authorized logins*, *Active sessions*, *Recognized devices*, and *Logins and logouts*. The account information shown (including the name, location, and other identifiers) is that of a fake test account created for the purpose of this work.

account access identifiers/variables (AAIV)	client checks settings
client confirms device ownership to consultant	client confirms page navigation
client confirms they do not recognize a device	client confirms they recognize a device
client describes their motivations for use of DII	client describes how threat motivated reaching out to support services
client explains background to listed devices to consultant	client expresses confusion around information displayed on DII
client has device listed they do not use	client identifies POC's contact information in recovery fields
client is not sure about device ownership from DII	client shares device information from DII
client shares changes to their device/account privacy and security	client shares device list from DII
client shares past experience with suspicious device via DII	client states device model(s)
client uses evidence from DII to confront abuser	client/consultant theorizes how POC is abusing them or their device(s)
consultant asks client to share their concerns with technology	consultant asks client confirm ownership of listed device
client asks consultant for guidance on account recovery process	client confirms account recovery information
client/consultant navigate challenges with DII error message	consultant explains 2FA/account recovery to client
consultant asks client for information from DII	consultant decides to ignore possibly suspicious device
consultant explains differences/similarities in platform DII to client	consultant explains DII to client
consultant explains how client information synced across accounts	consultant explains risk of account compromise to client
consultant explains safety risks of changes to DII for client	consultant explains to client why checking device list is valuable
consultant expresses confusion around where and how information is displayed on DII	consultant gives navigation instructions
consultant identifies the limitations of DII to client	consultant instructs client to examine device list on DII
consultant places client in control of security decisions	consultant reassures client based on DII information
consultant suggests practices for client to better secure account	consultant recommends client sign out of suspicious device(s)
consultant shares past experience	client/consultant uses one or more DIIs to judge account/device compromise
consultant uses example to guide client	consultants ask client to take screenshot of DII for further investigation
consultants asks client if the devices are familiar	status of physical devices
client expresses confusion about the device to account mapping	consultant asks client about recovery information

Figure 8: The codebook (consisting of 50 codes) we used in our qualitative analysis

Service	Interface	URL
Google	Recent Security Events	https://myaccount.google.com/notifications
	Your Devices	https://myaccount.google.com/device-activity
	Security Checkup	https://myaccount.google.com/security-checkup
	Signing into Google Password	https://myaccount.google.com/signinoptions/password
	Personal info Password	https://myaccount.google.com/signinoptions/password
	App Passwords	https://myaccount.google.com/apppasswords
	2-Step Verification	https://myaccount.google.com/signinoptions/two-step-verification
	Recovery Email	https://myaccount.google.com/recovery/email
	Recovery Phone	https://myaccount.google.com/signinoptions/rescuephone
Facebook	Security & Login	https://www.facebook.com/settings?tab=security
	Mobile Settings	https://www.facebook.com/settings?tab=mobile
	General Account Settings	https://www.facebook.com/settings?tab=account
	Where you're logged in	https://www.facebook.com/settings?tab=security
	Authorized Logins	https://www.facebook.com/settings?tab=security
	Recognized Devices	https://www.facebook.com/100085069751845/allactivity?activity_history=false&category_key=RECOGNIZEDDEVICES&manage_mode=false&should_load_landing_page=false
	Logins and Logouts	https://www.facebook.com/100085069751845/allactivity/?activity_history=false&category_key=LOGINSLOGOUTS&manage_mode=false&should_load_landing_page=false
	Active Sessions	https://www.facebook.com/100085069751845/allactivity?activity_history=false&category_key=ACTIVESESSIONS&manage_mode=false&should_load_landing_page=false
	Login Alerts	https://www.facebook.com/login_alerts
Apple	Sign-In and Security	https://appleid.apple.com/account/manage/section/security
	Account Security	https://appleid.apple.com/account/manage/section/security
	Account Recovery	https://appleid.apple.com/account/manage/section/security
	Devices	https://appleid.apple.com/account/manage/section/devices
	App-Specific Passwords	https://appleid.apple.com/account/manage
WhatsApp	Linked Devices	https://web.whatsapp.com/
	Two-Step Verification	https://web.whatsapp.com/

Figure 9: URLs for ASIs across the four services. Accessing the interface might require signing into a service's account. All URLs were accessed between August 2022 and February 2023.

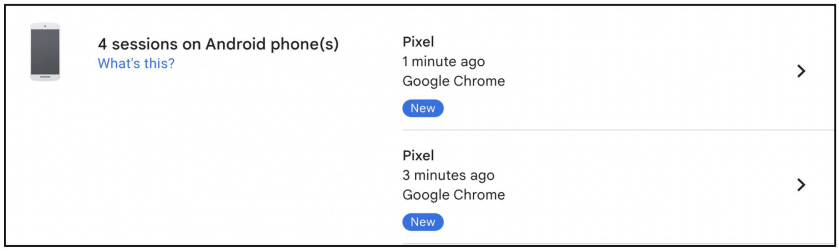

Figure 10: Google's *Your devices* interface showing a victim's legitimate session on an Android Pixel phone and an adversary's spoofed session on a Mac computer (the session was spoofed to look exactly like that of the victim's).