User Perceptions of Security and Privacy for Group Chat

SEAN OESCH, The University of Tennessee
RUBA ABU-SALMA, King's College London
OUMAR DIALLO, The University of Tennessee
JULIANE KRÄMER, TU Darmstadt
JAMES SIMMONS, The University of Tennessee
JUSTIN WU, Brigham Young University
SCOTT RUOTI, The University of Tennessee

Secure messaging tools are an integral part of modern society. To understand users' security and privacy perceptions and requirements for secure group chat, we surveyed 996 respondents in the US and UK. Our results show that group chat presents important security and privacy challenges, some of which are not present in one-to-one chat. For example, users need to be able to manage and monitor group membership, establish trust for new group members, and filter content that they share in different chat contexts. We also find that respondents lack mechanisms for determining which tools are secure and instead rely on non-technical strategies for protecting their privacy—for example, self-filtering and carefully tracking group membership.

To better understand how these results relate to existing tools, we conduct cognitive walkthroughs (a form of expert usability review) for five popular group chat tools. Our results demonstrate that while existing tools address some items identified in our surveys, this support is partial and is insufficient in many cases. As such, there is a need for improved group chat tools that better align with user perceptions and requirements. Based on these findings, we provide recommendations on improving the security and usability of secure group chat.

CCS Concepts: • **Security and privacy → Usability in security and privacy**; **Domain-specific security and privacy architectures;**

Additional Key Words and Phrases: Group chat, user perceptions, security, privacy

ACM Reference format:
Sean Oesch, Ruba Abu-Salma, Oumar Diallo, Juliane Krämer, James Simmons, Justin Wu, and Scott Ruoti. 2022. User Perceptions of Security and Privacy for Group Chat. *Digit. Threat.: Res. Pract.* 3, 2, Article 15 (February 2022), 29 pages.
https://doi.org/10.1145/3491265

This work was in part funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—SFB 1119—236615297.
Authors' addresses: S. Oesch, O. Diallo, J. Simmons, and S. Ruoti, The University of Tennessee, Knoxville, Tennessee TN 37996; emails: {toesch1, osouleym, jsimmo58}@vols.utk.edu, ruoti@utk.edu; R. Abu-Salma, King's College London, London WC2R 2LS, UK; email: ruba.abusalma@kcl.ac.uk; J. Krämer, TU Darmstadt, Darmstadt 64289, Germany; email: juliane@qpc.tu-darmstadt.de; J. Wu, Brigham Young University, Provo, Utah UT 84602; email: justinwu@byu.edu.

2576-5337/2022/02-ART15 $15.00
https://doi.org/10.1145/3491265

1 INTRODUCTION

Secure messaging tools have become an integral part of modern society—for example, the top-3 messenger tools have more than billion users each.[1] Despite the popularity of these tools, researchers have identified significant problems with their security and usability [3, 16, 29, 46, 48]. In particular, when users misunderstand or misconfigure these tools, they risk exposing their sensitive data. For example, it was recently demonstrated that users' private chats could be found using Google search if links to those chat sessions had previously been shared online.[2]

There is extensive research into improving secure messaging, both from traditional security [43] and usable security [3, 8] perspectives. Still, most research has focused on improving the security and usability of one-to-one communication. As such, while there have been some attempts to develop secure group chat protocols [7, 35, 39], there remains a lack of research into users' expectations and requirements for secure group chat.

In this work, we attempt to address this gap via an exploratory survey of group chat users. The goals of this survey are threefold. First, we want to understand better users' security and privacy perceptions of and requirements for group chat. Second, we want to suggest improvements to existing tools to help users stay secure. Third, we want to understand how national differences might impact perceptions and requirements. Taken together, we believe this information will be invaluable to both tool designers and security researchers.

To achieve these goals, we designed a 44-question survey that examines respondents' attitudes toward and experiences of privacy and security in group chat. The survey was split into four major topic areas—tool usage, group dynamics, privacy, and security. This survey was administered using Prolific (previously known as Prolific Academic), with 996 individuals completing the study. Half of our survey respondents resided in the US and half in the UK, helping us conduct an initial exploration on whether national differences would impact the results.

Our results show that users share sensitive information in group settings and are concerned with their privacy. Still, instead of selecting tools that help protect their privacy, users select tools based on which tools their contacts use. As such, users manage their security and privacy by using various non-technical strategies, with the two most important strategies being (1) self-filtering the content they share and (2) carefully managing and monitoring chat group membership. This second strategy is crucial as users' most significant privacy concern is that individuals they do not know will see the information they share in group chat. Additionally, we find that users of group chat suffer from significant alert fatigue—driven mainly by the fact that they receive notifications for group messages that they have no interest in—making it unlikely that users see important security and privacy notifications (e.g., when a new user is added to the group).

To help understand how user perceptions and expectations relate to existing group chat tools, we also conduct cognitive walkthroughs for five popular secure group chat tools: WhatsApp, Telegram, Signal, Facebook Messenger, and iMessage. Our results demonstrate that while existing tools address some items identified in our surveys, this support is partial and is insufficient in many cases. As such, there is a need for improved group chat tools that better align with user perceptions and requirements. For example, users need better methods for managing and monitoring group membership, establishing trust for new group members, and filtering the content they share in different chat contexts. To this end, we conclude our article with recommendations for improvements to existing tools and areas that require future research.

2 BACKGROUND AND RELATED WORK

In this section, we cover important background material and related work.

[1] https://www.statista.com/statistics/258749/most-popular-global-mobile-messenger-apps/.
[2] https://www.independent.co.uk/life-style/gadgets-and-tech/news/whatsapp-google-group-chat-private-link-messages-search-public-a9354391.html.

2.1 Secure Communications

Secure communication tools became widely available with the release of PGP in 1991, which was followed by the creation of many PGP-based tools (e.g., Enigmail, GPGMail, GPG4WiN, Mailvelope, OpenPGP). Nowadays, the smartphone era has seen an explosion of new communication tools, typically called "instant messaging applications" or "messengers". Unlike PGP—which was designed for asynchronous, high-latency e-mail communications—instant messaging applications are fast and responsive.

Unger et al.'s work remains the best resource for understanding the technological underpinnings of secure messaging. Unger et al. [43] created a framework for evaluating the security, usability, and ease of adoption of secure messaging primitives. They found that establishing trust, conversation security, and transport privacy (hiding communication metadata) were key challenges for these tools, with transport privacy the most challenging property to maintain. The work of Unger et al. [43] also highlights the fact that no existing research examined users' expectations for secure group chat protocols and that usability research is "sorely needed". Several studies have found weaknesses in Signal's group communication protocol [7, 35] and proposed alternative solutions [7, 39].

Looking at this research and the features emphasized on the websites for various secure chat tools, we believe that existing researchers and tools are largely focused on the following security properties:

(1) Preventing or detecting man-in-the-middle attacks—i.e., using end-to-end encryption.
(2) Ensuring that the tool's owner or operator has not interfered with the trust negotiation process.
(3) Limiting the sensitive information stored on the tool owner or operator's server. This protects against both hacks by malicious adversaries as well as legal requests made by law enforcement organizations.
(4) Ensuring that conversations do not leak metadata to other users or the tool's owner or operator.
(5) Allowing users to use cryptographic keys to prove their identity.

2.2 Defining Privacy

There are many ways to define privacy. In this article, we use the privacy categories described by Finn et al. [13]—privacy of the person, privacy of behavior, privacy of communication, privacy of data and image, privacy of thoughts and feelings, privacy of location and space, and privacy of association. In our results, respondents mentioned privacy concerns falling into all these categories except privacy of the person (i.e., biometric information). In addition to considering the categories related to privacy, we also consider the dimensions of privacy described by Mulligan et al. [31]—the dimension of provision (how is privacy provided?), the dimension of harm (who is the threat?), and the dimension of scope (what are the social boundaries of privacy and what is its time span?).

2.3 Differences between the US and Europe

The US and Europe are known to differ in their attitudes toward privacy in relation to society, legislation, and culture [41]. In Europe, privacy is viewed as a human right [18, 20, 23], whereas in the US, it is the responsibility of individuals to protect their own privacy [18, 20, 41]. Moreover, European nations have centralized privacy agencies and laws ensuring the privacy of their citizens [5, 18, 30]. In contrast, the US does not have much legislation protecting the privacy of its citizens [18, 23, 41].

Considering these known differences between privacy attitudes in the US and Europe, we believed that comparing these two groups would provide an ideal initial exploration into how national differences might impact perceptions and requirements regarding group chat. Instead of surveying Europe generally, we instead chose a single locality within Europe—the UK. We chose to use the UK as our representative locality because we had a researcher from the UK on our team, and experience has shown that studying users in a nation not represented among researchers can lead to difficulties in interpreting results because of cultural differences and language barriers. Future research could expand upon these results by doing similar studies of other nations, giving a

greater view of how national differences impacts perceptions and requirements for group chat; still, we believe that studying these two localities provides a helpful initial exploration of the topic.

2.4 Usability and Adoption

Poor usability has often been studied as one possible impediment to adoption [37]. While it is difficult to ascertain whether poor usability is the primary impediment to adoption [3, 11], research has found that users tend to prefer more usable but less secure tools over more secure, less usable tools [4]. Similarly, research has shown that users select tools based on peer influence (which is driven by usability) rather than provided security properties [8]. Research has also found that many users feel that secure chat tools are only useful for people that are either paranoid (perhaps rightfully so) or "up to no good" [38, 56]. Other factors impacting adoption of secure messaging tools include small and fragmented user groups [3].

2.5 Incorrect Mental Models

Even when users do adopt secure messaging applications, they do not always configure them correctly or take advantage of the security features [46]. For example, Telegram users were found to use the less secure default chat mode [1]. Such misuse is likely explained by the fact that users have incorrect mental models about how security works [1, 2, 28, 56]. This lack of understanding can also lead users to distrust tools making claims about security [9, 16].

Users especially struggle with the authentication ceremony in apps like Signal and WhatsApp due to incorrect mental models [47, 48, 55]. The most successful modification of the authentication ceremony was done by Wu et al. [55], who conducted a three-phase modification of the warning notifications surrounding this ceremony in Signal and found that it enhanced usability without weakening security.

2.6 Group Chat Dynamics

While no prior work analyzed the security and privacy of group chat, prior research has studied group chat usage more generally. Prior research has found that group chat is used to discuss a wide array of topics, including both personal and work-related topics [17, 19]. Ling et al. [26] found that users struggled to classify their chats into appropriate topics, which could impact users' ability to self-filter what information they share during simultaneous communications. Other research has found that in group chat, users heavily leverage non-textual group chat features (e.g., attaching images and videos) [42]. For tool selection, users focus on tools that provide control, enjoyment, reliability, speed, and ease of use [44], with many respondents indicating that what their peers use plays a vital role as well [26]. Research has also shown that individuals from different nations and cultures use group chat tools differently [22, 27].

3 METHODOLOGY

To understand users' perceptions and requirements regarding secure group chat, we conducted two surveys approved by our institutions' IRB and ethics board, respectively. We conducted the first survey from February 4 to February 11, 2019, and collected 500 responses from individuals in the US. We then conducted the second survey from March 11 to March 19, 2019, and collected 501 responses from individuals in the UK. Both surveys were administered using Prolific (formerly Prolific Academic), a platform comparable to MTurk, and had identical content. We used Prolific to select representative samples from the US and the UK before collecting responses.[3] The consent form is included in Appendix A and the full survey in Appendix B.

[3] https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-Samples-on-Prolific.

3.1 Survey Contents

Because no prior work has addressed security and privacy in group chat, we designed our survey to explore this space. We wanted to identify the security and privacy challenges facing users, understand how they navigate those challenges, and understand how the social dynamics of group chat impact security and privacy. We intentionally avoided defining security and privacy in our survey instead of seeking to understand users' perceptions regarding these terms. While this prevented us from exploring specific privacy issues, we believe it was necessary to understand what is foremost in users' minds. The survey was comprised of four sections: tool usage, group dynamics, privacy, and security.

Tool usage: First, we asked respondents which group chat tools they used and for which purpose they used them. Next, we asked them about what they liked and disliked about the tools they used. We also asked how they selected the group chat tools they wanted to use.

Group dynamics: We then asked respondents which groups respondents used group chat tools to communicate with (e.g., family, work colleagues, friends), including the average size of their chat groups. We asked if respondents had ever been removed from a group and, if so, why they had been removed and how it affected them. Next, we asked whether permission should be requested from the group before new users were added and why they felt that way. We also asked when, if ever, users had reviewed the member list of a group. Finally, we asked respondents how they felt about group respondents who rarely participated in the group and why they felt that way.

Privacy: We started by asking respondents whether they were concerned about others sharing screenshots of their chat logs. We then asked them what topics they felt uncomfortable discussing in chat. Next, we asked whether someone had ever shared something that put them in an awkward position and, if so, how it had affected them and how they had responded. Finally, we asked whether respondents had ever joined a chat group for a given topic (e.g., politics) and, if so, what those topics were and what privacy concerns they had for those groups.

Security: We asked respondents if they had ever shared sensitive information and, if so, what types of sensitive information they had shared. Next, we asked how they determined whether a group chat tool was secure and what steps they took to secure their group chat communication. We also asked whether they considered any existing tools to be secure and, if so, which ones. Finally, we asked if they had ever been concerned with impersonation in their group chats and, if so, how they handled those situations.

Demographics: We ended the survey by asking basic demographic questions: age, gender, ethnicity, education, and frequency of group chat tool usage.

3.2 Survey Development

After developing the initial questions of our survey, we conducted cognitive interviews with 10 demographically diverse respondents. The interviews were used to evaluate the survey and glean insights into how survey respondents might interpret and answer questions [34]. As respondents answered each survey question, we asked them to answer the following questions: (1) Was this question difficult to answer? (2) Was there an answer choice missing? (3) How did answering this question make you feel? We incorporated this feedback to improve our study.

After the fifth cognitive interview, we also asked five computer security and privacy researchers with survey expertise, as well as our institution's IRB consultant, to review our survey questionnaire and assess question wording, ordering, and bias. Expert review is a method used to complement cognitive interviews in identifying questions that require clarification or further revision [34]. Based on these reviews, we updated some of our survey questions and then conducted the remaining five cognitive interviews to ensure no further problems would emerge.

3.3 Quality Control

Coders examined respondents' responses to open-ended questions during coding. The coders discarded the responses if they left most questions blank or did not address the questions asked. In total, we excluded five respondents, four from the US and one from the UK, leaving 996 responses for analysis.

3.4 Data Analysis

Our survey included 19 open-ended questions in the survey. For each of these questions, we created a codebook and divided responses into categories. We then applied pair-coding, having two researchers work together to code responses. Pair-coding does not require calculating an inter-rater reliability metric as all disagreements are resolved as the items are coded. The coders also noted responses that they perceived as particularly interesting.

After coding, we met together as a research team to discuss the results and identify themes within the data. This included analyzing the data to find differences based on nationality, age, and gender. To search for meaningful groupings within the responses that we might not have considered, we conducted k-means and k-modes clustering of our data.

3.5 Limitations

Because we used a survey, it was not possible to identify areas where respondents might have misrepresented their behaviors. Prior work showed that users often claim to be more concerned about privacy than they are in practice [54]. Future work could use interviews, diary studies, or direct observation to determine if there is a gap between our results and actual user behavior.

Our respondent demographics were slightly skewed toward a young female population and primarily Caucasian. Future research could expand on these results by studying specific sub-populations in more detail.

Our work provides an initial exploration into how national differences impact perceptions and requirements related to group chat. While we used the UK as a representative locality in Europe, future research should look at Europe and the world more generally to see how other national differences could impact our results. Ideally, researchers who undertake this work should understand the cultural perceptions of respondents and know the language.

4 RESULTS

In this section, we highlight key themes from our results.

4.1 Demographics

Respondent demographics are shown in Table 1. The only recruitment requirement was the respondent's location—the US or the UK. There was slightly more female (58%) than male respondents, and over half of respondents were millennials (62%), which we define as those under 34. Nearly all survey respondents had completed high school, and over half had completed some level of higher education after high school (58%). Less than half a percent (.025%) of respondents never used group chat tools, while 15.7% used them rarely, indicating that tools need to address the security of infrequent users.

4.2 Tool Usage

Most respondents used a group chat tool at least 2–3 times per week (n = 727; 73.0%). Facebook Messenger (n = 798; 80.1%) and WhatsApp (n = 588; 59.0%) were the most common tools used, with WhatsApp being more popular in the UK than in the US (see Figure 1). Few respondents used tools commonly associated with a security mindset, such as Signal (n = 19; 1.9%), Telegram (n = 64; 6.4%), or Viber (n = 64; 6.4%).

While there were several strategies for selecting tools, the most common strategy (n = 696; 69.9%) was to use whatever application was popular among friends and colleagues (see Figure 2). Only a small number of

Table 1. Respondent Demographics

	UK		US	
	#	%	#	%
Gender				
Male	180	36.0%	233	47.0%
Female	318	63.6%	258	52.0%
Other	1	0.2%	3	0.6%
No answer	1	0.2%	2	0.4%
Age				
Under 21	78	15.6%	28	5.6%
21–34	250	50.0%	261	52.6%
35–44	94	18.8%	123	24.8%
45–54	51	10.2%	45	9.1%
55–64	19	3.8%	34	6.9%
65+	8	1.6%	4	0.8%
Education				
No diploma	12	2.4%	8	1.6%
High school	95	19.0%	54	10.9%
Some college	128	25.6%	129	26.0%
Associate's	24	4.8%	54	10.9%
Bachelor's	156	31.2%	161	32.5%
Master's	65	13.0%	72	14.5%
Doctoral	9	1.8%	7	1.4%
No answer	2	0.4%	2	0.4%
Ethnicity				
Black or African American	15	3.0%	34	6.9%
Asian	28	5.6%	16	3.2%
Mixed race	20	4.0%	25	5.0%
Pacific Islander	0	0.0%	1	0.2%
Caucasian	428	85.6%	409	82.5%
No answer	1	0.2%	3	0.6%
Usage Frequency				
Daily	289	57.8%	182	36.7%
4–6 times/week	55	11.0%	70	14.1%
2–3 times/week	56	15.0%	3	0.6%
Weekly	38	7.6%	50	10.1%
Rarely	53	10.6%	103	20.8%
Never	9	1.8%	16	3.2%

respondents (n = 12; 1.2%) indicated that security was a key factor in tool selection. This result is in line with prior work [8, 26, 29].

4.3 What Users Chat about

Most respondents (n = 756; 75.9%) used these tools for chatting, with a third using it for coordinating events (n = 350; 35.1%), and another third using it for work or school-related discussions (n = 277; 27.8%). Most respondents also used group chat to talk to friends (n = 873; 87.7%), immediate family members (n = 672; 67.5%),

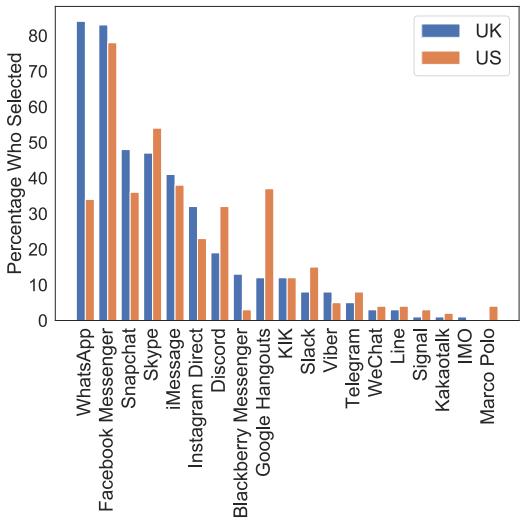

Fig. 1. Percent of respondents who used each tool.

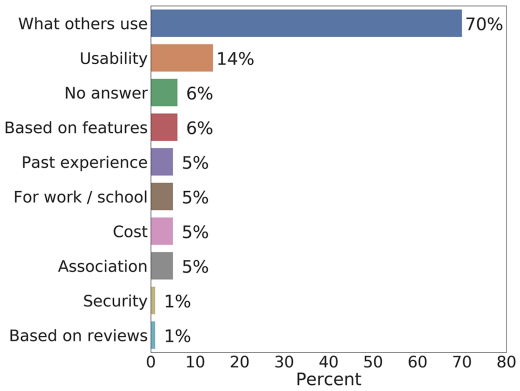

Fig. 2. Reasons respondents selected a group chat tool.

and work colleagues (n = 524; 52.6%), with a third using it to talk to extended family members (n = 387; 38.9%). The most common chat group size was 3–5 people (n = 480; 48.2%), followed by 6–10 people (n = 263; 26.4%). Some respondents (n = 142; 14.3%) also indicated that their group size varied widely. In explaining why group size varied, R529 shared:

> *Because I use discord, a server-based platform with a lot of members I do not know directly as well as group messaging which is private and only with people I know directly.*

4.3.1 Topic-Based Groups. A quarter of respondents said they had joined a group chat precisely because of the topic being discussed (n = 245; 24.6%). Figure 3 provides a breakdown of these topics. Games (n = 59/245; 24.1%) and hobbies (n = 39/245; 15.9%) were the most common answers. Interestingly, respondents from the US mentioned games and finance twice as often as respondents from the UK.

Figure 4 reports topics that made respondents uncomfortable, with politics (n = 270; 27.1%), religion (n = 227; 22.8%), and sexuality (n = 194; 19.5%) being the most common responses. US respondents were more likely than UK respondents to feel uncomfortable discussing politics, religion, and sexuality.

Of respondents who participated in topic-based conversations, just under a fifth (n = 44/245; 18.0%) had privacy concerns regarding the content being discussed (privacy of thoughts and feelings). The two most common concerns revolved around not knowing some group members (n = 17/44; 38.6%) and a fear that knowledge that they participated in that group could negatively impact how family or coworkers perceived the user (n = 9/44; 20.5%):

> R46: *"Sometimes I might be interested in a game or other topic that I wouldn't necessarily want to be associated with my professional profile, so I take extra precautions to alter my name/appearance in the group."*
>
> R193: *"As a freelance professional, I need to be aware that all my web presences are my "game face" professionally and that potential clients or coworkers may see what I post no matter where it is."*

4.3.2 Awkward Situations. While respondents did discuss sensitive topics, most respondents (n = 757; 76.0%) did not indicate that they had ever felt awkward due to what was shared in a group chat. Of those who had been placed in an awkward position (n = 140; 14.1%), most said it was either due to gossip (n = 88/140; 62.9%) or their sensitive information being shared without their permission (n = 19/140; 13.6%):

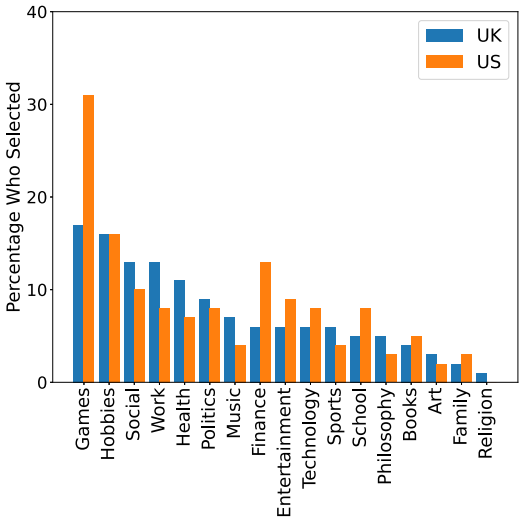

Fig. 3. Topics respondents joined groups to discuss.

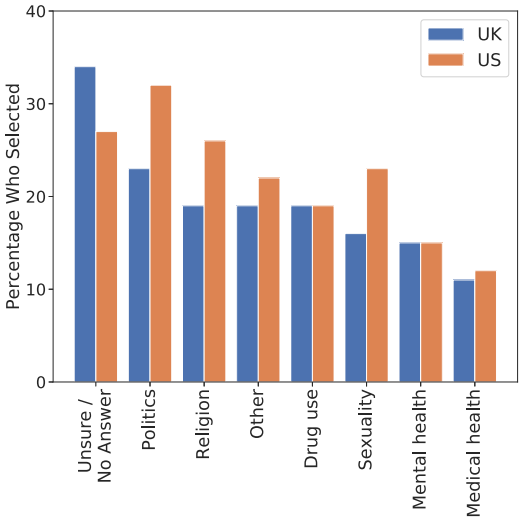

Fig. 4. Discussion topics that caused discomfort.

R390: *"Sometimes people get a bit too personal in what they share, e.g., I've had conversations where people start to share things about their married life, which I really don't need to be a part of."*

R400: *"I've accidentally said things to a group chat that were meant for my wife. I've also said things about people accidentally who I did not know were on the group."*

R371: *"Someone in my work group shared a schematic that wasn't approved for sharing and people who weren't authorized to have knowledge of the product saw it."*

R932: *"I messaged a committee chat about how one of the union representatives at our university wanted us to vote for a specific candidate, and one of the girls in the chat asked me to screenshot and forward the message so she could report it."*

When asked how they responded to being put in an awkward position, most said that they either reached out directly to the person that had caused the situation (n = 65/140; 46.4%) or that they did nothing about it (n = 56/140; 40.0%). Only a small minority of respondents (n = 11/140; 7.9%) left a group chat over an incident.

4.4 Managing Group Membership

When asked whether people should ask for permission before adding someone to join a group chat, most respondents indicated that they felt this was important: "yes" (n = 403; 40.5%), "it depends" (n = 399; 40.1%), "no" (n = 148; 14.9%). When asked why they felt it was necessary to ask for permission, most respondents (n = 627; 63.0%) focused on ensuring that the user was a good fit for the group. Protecting the privacy of information shared by group participants (privacy of communication) was also important to many respondents (n = 196; 19.7%):

R283: *"In some groups, the members share private and sensitive information. In groups with this sort of trust, it's important to get permission for the whole group before adding new members so the atmosphere of security can be maintained."*

Over half of respondents (n = 620; 62.2%) said that they checked group membership when they first joined a group. Many respondents (n = 399; 40.1%) also checked the member list occasionally to see if anything had changed. Some respondents (n = 187; 18.8%) also checked the member list whenever a new member was added to the group.

We also asked respondents whether they had ever been removed from a group chat, with only a small minority of respondents (n = 96; 9.6%) answering in the affirmative. In most cases, the removal from a group stemmed from the end of a relationship (n = 47/96; 49.0%) (personal or professional), misbehavior (n = 25/96; 26.0%), or a joke (n = 18/96; 18.8%):

> R46: *"I've been removed from group chats of gaming groups/guilds because of being inactive or someone having a personal issue with me."*

> R932: *"One of my friends removed me as a joke because I kept on sending the same link to the chat every second. I was added again later."*

Inactivity in a group had been observed by most respondents (n = 812; 81.5%) and was usually not seen as a concern (n = 665/812; 81.9%).

4.5 Fears Regarding Impersonation

Just over a 10th of respondents (n = 111; 11.1%) indicated that they had at some point been concerned that a member of a group chat was not whom they claimed to be. Most often, this concern came from a general fear about deception (n = 59/111; 53.2%), though it was also triggered by observing group members that were acting in a way not congruent with how they normally acted (n = 40/111; 36.0%):

> R218: *"I didn't know the person and so I felt uncomfortable that they could be anyone - and I wasn't quite sure how they were invited to the group chat in the first place."*

To validate the identity of a group member, most respondents (n = 28/111; 25.2%) relied on personal knowledge of that member, including how they normally talked. Respondents would also cross-reference the group members account with a social media account (n = 26/111; 23.4%), have the group member send a picture of themselves (n = 11/111; 9.9%), or video chat with the person (n = 8/111; 7.2%). Other respondents (n = 16/108; 14.8%) did not feel there was a way to verify a group member's actual identity:

> R625: *"If you know the person in real life then you could ask them questions that only that person would know. Also if you have another way of contacting the person, you could contact them and ask them to verify that it is in fact them."*

4.6 Perceptions of Tool Security

When asked what it meant to them that an instant messaging tool was secure for group communication, nearly half (n = 432; 43.4%) of respondents said a group chat tool would be secure if non-group members could not read messages:

> R821: *"No public is allowed to let themselves in. Messages are encrypted and only those who are in the group can view them."*

Some respondents also expected that the list of members in a group should be confidential:

> R535: *"If it is secure there is no way that an outside can gain access the messages being sent or the list of group members without the permission of a group admin."*

Another security concern was centered around strict control of who could enter a group (n = 42; 4.2%):

> R271: *"This means that my conversation thread doesn't get hacked, non-members can't join or have access to my conversation thread without permission and that I am able to permanently delete my conversation and not have it saved by the IM tool in some universal back-up file for other organizations to have access to."*

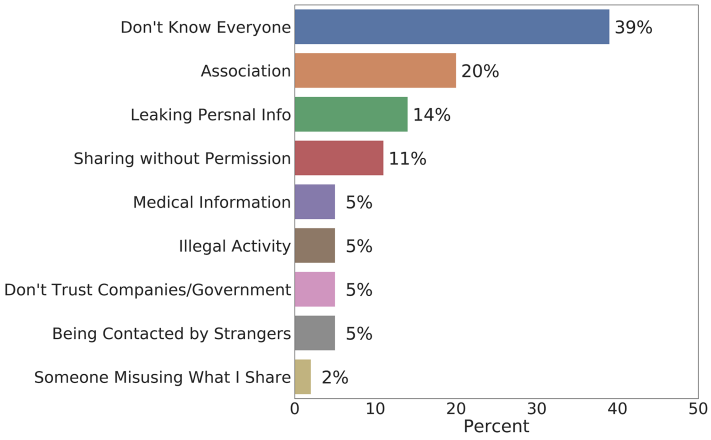

Fig. 5. Privacy concerns when joining a group chat.

Aside from these specific concerns, another third of respondents (n = 375; 37.7%) indicated that security was important to them but did not specify what security features they were most interested in.

There was a wide range of strategies respondents indicated using to validate the security of a group chat tool: reading reviews (n = 158; 15.9%), whether the tool supported encryption (n = 122; 12.2%), tool popularity (n = 96; 9.6%) or reputation (n = 57; 5.7%), recommendations from friends or family (n = 50; 5.0%), tool documentation (n = 37; 3.7%), and personal experience (n = 40; 4.0%). As with tool selection, many of these strategies focused on popular opinions:

> R719: *"If there are a lot of positive reviews from users of the instant messaging tool, it leads me to believe that it is secure for group communications. In addition, I look for what my friends and family use because I trust their opinions in what is safe. However, I also understand that full security is unlikely and I must make efforts to maintain security in group chats."*
>
> R831: *"I don't, but if my friends are on it I am pretty sure its safe."*
>
> R224: *"If the IM tool is well known and widely used I just blindly trust that it will be secure. If it was some shady app or software that had very little reputation or reviews behind it, I would probably not use it."*

Additionally, over a third of respondents (n = 384; 38.6%) admitted that they would not check the security of a tool. When asked if there were instant messaging tools that they believed they were secure for group communications, responses were split: "yes" (n = 412; 41.4%), "uncertain" (n = 370; 37.1%), and "no" (n = 214; 21.5%).

4.7 Privacy Concerns

Respondents expressed a variety of privacy concerns regarding group chat, with the majority being concerned about sharing information in groups with individuals they did not know (see Figure 5). The concerns shared by respondents generally fell into six of the seven privacy categories previously identified by Finn et al. [13]:

- **Privacy of behavior:** *"If we are talking about, like say, something you smoke that might or might not be legal, I do not feel comfortable." (P11)*
- **Privacy of communication:** *"A member of a discord was screenshotting sensitive discussions and sharing them on social media." (P735)*

— **Privacy of data and image:** *"People got personal and we're finding images of people and taking cheap shots at them as they could not seem to broaden their mindset." (P767)*
— **Privacy of thoughts and feelings:** *"Things i might have said to my friends (secrets) regarding home issues" (P584)*
— **Privacy of association:** *"I have been involved in anime fanfiction circles, and was fearful of my writing being associated with my real name/identity." (P803)*
— **Privacy of location and space:** *"Some applications allow others to see certain information by default that is somewhat intrusive such as your location, these settings sometimes have to be disabled or may not be obviously enabled." (P866)*

About a quarter of users (n = 249; 25.0%) indicated that they had shared sensitive information in group chat, with the most common types of information being PII (n = 85/249; 34.1%), personal feelings (n = 51/249; 20.5%), medical (n = 44/249; 17.7%), mental health (n = 34/249; 13.7%), family issues (n = 32/249; 12.9%), sexuality (n = 30/249; 12.0%), and romantic relationships (n = 26/249; 10.4%).

4.8 Protection Strategies

Only a small number (n = 62; 6.2%) of respondents indicated that they relied on the tool to protect their privacy. Instead, respondents employed a variety of strategies for protecting their confidentiality and privacy. Most commonly (n = 303; 30.4%), respondents indicated that they self-filtered their messages, being very careful with what they shared. This is in line with previous research on how users cope with web security challenges [38]. For example,

> R271: *"I realize that I can't trust that my communication is 100% secure at this point so I am just careful in what I say. Especially with the way the creators of internet apps, social media platforms etc. are constantly breaching or violating the privacy of users and selling information."*

> R155: *"I never say anything I couldn't say in front of my grandmother."*

The second strategy reported by many respondents (n = 272; 27.3%) was to carefully monitor group membership:

> R148: *"I check to make sure I know who all the people in the group are."*

> R288: *"I don't feel that I can control them but I do take into account who is a member of the group and what I feel comfortable sharing with them."*

> R802: *"Dependent on the group members and my trust of them rather than technology."*

While self-filtering and group membership maintenance were by far the two most common strategies, there was a range of other strategies listed by respondents: ensuring the device is up-to-date and properly configured (n = 65; 6.5%), ensuring the messaging tool is up-to-date and properly configured (n = 56; 5.6%), using a password (n = 47; 4.7%), and setting messages to expire (n = 16; 1.6%). Additionally, a quarter of respondents (n = 250; 25.1%) indicated that they took no proactive steps to protect their privacy when using group chat tools.

These behaviors can also be understood using the dimensions of privacy previously identified by Mulligan et al. [31]:

— **Dimension of provision:** Respondents relied on themselves, not their tools, to protect their privacy.
— **Dimension of harm:** Harm depended on the category of privacy [13] respondents were concerned about. For example, with privacy of association, users were concerned about friends or coworkers knowing about their activities, whereas with privacy of behavior, respondents were afraid of central authorities. Most prevalent, respondents were concerned about other users contacting them without their permission or

misusing private information to shame or manipulate them. Only a few respondents were concerned about threats like government surveillance.

— **Dimension of scope:** Respondents did not want their private information to be visible to anyone outside of the group with whom it was shared at any time in the present or the future.

4.9 Likes and Dislikes

When we asked respondents about what they liked about group chat tools, most answers focused on its ease of use (n = 414; 41.6%), speed (n = 324; 32.5%), ability to share media content (n = 81; 8.1%), and ability to easily lookup old messages (n = 42; 4.2%). When asked what they disliked about group chat, nearly half of respondents (n = 398; 40.0%) indicated they were overwhelmed with messages and notifications:

> R944: *"When you have lots of people in one chat everyone talking at once can send me lots of notifications and that can get annoying."*

> R609: *"Sometimes there can be too many people messaging at once and certain messages could be ignored that may be important."*

A 10th of respondents (n = 82; 8.2%) indicated that they experienced negative personal or social effects because of the "always-on" and impersonal nature of group chat tools. These effects included constant pressure to be available to respond to messages in group chat, which resulted in an inability to find rest and solitude, and frustration trying to have meaningful discussions through textual communications when face-to-face may be more appropriate. Ironically, these negative effects were the direct result of what many respondents (n = 324; 32.5%) indicated was what they liked about group chat–that it was instantaneous:

> R348: *"Sometimes it can be hard to disconnect – I always feel like I am within reach and cannot take time away from work or social interactions."*

> R389: *"Not always the best for certain topics – though I guess this is more of an issue with text as a form of medium since you're missing out on body language and all the other things we use to communicate."*

> R530: *"If you've been in an important meeting or been busy in any way shape or form and you come back to a chat gone bonkers - with sooooo much information that's come in while you've been gone…that can get rather annoying as you then have a loooot of catching up to do, even though you were perhaps taking some valuable and much needed time for yourself in the meantime."*

5 NATIONALITY, GENDER, AND AGE

In addition to calculating overall statistics for our results, one of the goals of our study was to understand whether national differences (e.g., culture) between the US and the UK would impact responses. Also, we examined how gender and age impacted responses. To account for multiple tests, we calculated the Bonferroni correction for each category tested and included the α value with our results.

First, we found that there were significantly different group chat tools used by US and UK respondents. For example, most UK respondents used WhatsApp (n = 421/500; 84.2%), whereas a much smaller number of US respondents did (n = 167/496; 33.7%). A χ^2 analysis found this difference to be significant ($\chi^2(16) = 256, p < .0001$, $\alpha = .00625$). We also found that gender ($\chi^2(18) = 92.4, p < 5.26\text{e-}12, \alpha = .0056$) and age ($\chi^2(18) = 143, p < .0001$, $\alpha = .00625$) had a statistically significant effect on tool usage. For example, Millennials were more likely to use Snapchat, Instagram Direct, or Discord.

UK respondents were also more likely to say they used group chat daily (n = 289/500; 57.8%) than US respondents (n = 182/496; 36.7%), with the result being statistically significant ($\chi^2(5) = 48.5, p < .0001, \alpha = .00625$). Similarly, younger respondents were more likely to use group chat frequently ($\chi^2(5) = 96.7, p < .0001, \alpha = .00625$).

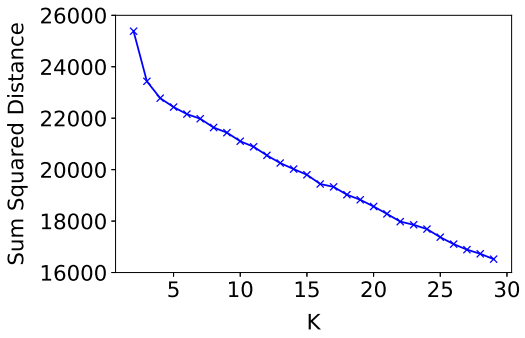

Fig. 6. K-Means elbow.

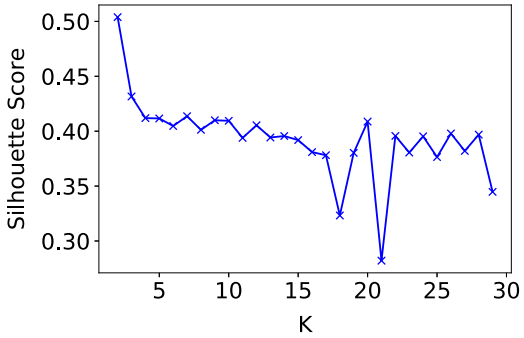

Fig. 7. K-Means silhouette score.

We did not find a statistically significant difference for frequency of usage based on gender ($\chi^2(5) = 4.57, p = .47, \alpha = .0056$).

We found no significant difference between the topics discussed by respondents (country—($\chi^2(15) = 15.5, p = .41, \alpha = .00625$), gender—($\chi^2(15) = 11.7, p = .70, \alpha = .0056$), age—($\chi^2(15) = 14.2, p = .51, \alpha = .00625$)). Interestingly, we did find that respondents from the US were more uncomfortable discussing politics, religion, and sexuality, though this result was not statistically significant after correction ($\chi^2(7) = 18.9, p = .0083, \alpha = .00625$). We failed to find a similar effect for gender ($\chi^2(7) = 7.62, p = .37, \alpha = .0056$) or age ($\chi^2(7) = 13.7, p = .057, \alpha = .00625$).

Compared to the US (n = 30/496; 6.0%), twice as many respondents from the UK (n = 72/500; 14.4%) had been removed from an instant messaging group without their permission, with the difference being statistically significant ($\chi^2(2) = 19.3, p < .0001, \alpha = .00625$). For both nations, younger respondents were more likely to have been removed from groups ($\chi^2(2) = 19.1, p < .0001, \alpha = .00625$). We did not find any similar effect for gender ($\chi^2(2) = 4.19, p = .12, \alpha = .0056$).

Other than these few differences, we did note a couple of other small differences, but they were too small to be practically significant.

5.1 Other Factors

We were also interested in whether there were other meaningful breakdowns of the data besides nationality, gender, and age. To explore this issue, we attempted to apply machine learning-based clustering algorithms to our dataset.

First, we explored our data using k-means [52]. We built models using values of k from 2 to 30. We evaluated our data using the elbow method [51] (see Figure 6) but did not find an elbow. We also examined each model's silhouette score [53] (see Figure 7), but all scores were below the 0.5 threshold, indicating no good clusters. Taken together, these results strongly suggest there is no robust clustering within our data.

We then attempted to cluster our data using k-modes [12]. At $k = 2$, responses were clustered by nationality. At $k = 4$, responses were split by gender, two of the splits that we had already considered. Otherwise, there were no meaningful splits in our data, providing some level of confidence in the completeness of our analysis.

6 ANALYZING SECURE CHAT TOOLS

Our survey helps reveal user perceptions regarding security and privacy in group chat settings. In this section, we describe a usability inspection [32] of five popular group chat tools—WhatsApp, Telegram, Signal, Facebook Messenger, and iMessage—investigating how well these tools map to user perceptions and what gaps may exist.

Usability inspections are a family of methods for evaluating the user interface of a computer system [32]. They primarily aim at finding usability problems in the user interface's design and evaluate the overall usability of an entire system. Unlike empirical user studies, which rely on recruiting representative participants, users inspections instead leverage expert evaluators. While usability inspections do not replace the need for user studies [10, 21, 25, 32], they provide a means to rapidly study a system and understand key usability features and issues [21].

More specifically, in this article, we leverage the cognitive walkthrough methodology. A cognitive walkthrough is a usability inspection method used to evaluate a user interface design for its *exploratory learning*, an essential aspect of usability testing [24, 33, 40, 50]. First-time users of a system may prefer to learn how to use it by exploring it, rather than investing time in comprehensive formal training or reading long tutorials [6, 14]. Cognitive walkthrough identifies problems that users could have as they approach an interface for the first time. It also identifies mismatches between how users and designers conceptualize a task and how designers make assumptions about users' knowledge of a specific task (which could, for example, impact the labeling of buttons and icons). The cognitive walkthrough process follows a structured series of questions derived from the theory of exploratory learning to evaluate each step (or action) in the workflow. A detailed overview of the cognitive walkthrough process is described by Wharton et al. [50].

In the remainder of this section, we first provide more details regarding our methodology for conducting the cognitive walkthroughs. We then finish by sharing observations and lessons learned from the cognitive walkthroughs.

6.1 Methodology

Our cognitive walkthroughs examined five tools: WhatsApp (v2.21.81.1), Signal (v5.11.1), Telegram (v7.7), Facebook Messenger (v310.0), and iMessage (v14.4.2). These tools were tested using iOS 14.5 on an iPhone. These tools were selected to include the most popular group chat tools (WhatsApp and Facebook Messenger), security-focused tools (Signal and Telegram), and the most popular texting-based tool (iMessage).

For each of these tools, we evaluated eight tasks:

(1) **Creating a new chat group.** As part of this, the evaluator explored available security and privacy options.
(2) **Inviting a member to join the group,** including exploring who can perform this operation as well as investigating how to actually complete the task.
(3) **Joining a group.** In particular, this involved looking for notifications or other signals that would allow the user joining the group to know what security and privacy options were supported by the group chat and who was already a member of the group.
(4) **Watching a member join a group.** In regard to the survey results, this involved monitoring group membership. We considered two variants of this task, both where the user was active when the member was added and when they were inactive for several hours (and thus may have missed any initial notifications).
(5) **Removing a member from a group,** and what steps or permissions were required.
(6) **Being removed from a group,** with particular attention to what capabilities were still available to the user after they were removed from the group.
(7) **Watching a member be removed from a group.** As with watching a user get added, we were primarily concerned with the ability of group participants to continue monitoring group membership.
(8) **Managing notifications.** This included setting what notifications should be received, how often they should be received, and whether a do not disturb option was set. We also considered how turning off notifications might impact the results from the previous tasks (e.g., if notifications are off, do you ever learn when members join the group).

The cognitive walkthroughs were completed by a single member of our research team, who role-played a mildly technical user (roughly what could be expected of a college graduate). While this role is far from

representative of all users, we believe it is a good medium and roughly represents the demographic being targeted by most group chat tools. The reviewer was selected based on their experience conducting cognitive walkthroughs and their familiarity with group chat tools. During the cognitive walkthroughs, the evaluator took careful notes of their findings, and these were used to generate the observations lessons learned reported below.

6.2 Observations

Below we discuss observations from each of the five tools we studied: WhatsApp, Telegram, Signal, Facebook Messenger, and iMessage.

6.2.1 WhatsApp. When creating a group in WhatsApp, the user creating that group is set as the admin. For the duration of the group's existence, any user added to the group can see who created the group and when.

Only admins can add users to a group by directly adding them or creating an invite link. This latter functionality allows other members to invite other users to the group. If an admin adds a user to a group directly, that users is immediately added to the group, without any need to approve the operation by the added user.

Upon joining a group, the added user is shown a message in the chat history telling them that messages and calls are end-to-end encrypted (see Figure 8). This is helpful for new users, allowing them to know what content they can share within the channel. Still, this message can be missed as new chat messages will displace it from the visible part of the chat history. For very active chat groups, new members may never see this message unless they scroll back through the chat history. New users are only able to see chat messages that occur after they have been added.

When users are added to the group, all users in the group receive a message in the chat window informing them of the group membership change. This also happens when users are removed (see Figure 11). However, these notifications are only visible within the chat history and can easily be missed if ongoing chat messages displace them from the currently visible portion of chat history.

As with adding a user, only an admin can remove a user from the group. This is done by selecting the user and choosing the "Remove From Group" option (see Figure 12) and then confirming the operation on a second screen. After being removed, the former group member can still see the chat history from when they were members of the group. Interestingly, admins can remove themselves from groups. If this happens and there is only a single admin, then a group user is chosen at random to be the new admin (without requesting permission from that user). Additionally, while removed owners can no longer participate, they can still delete the group.

Users have coarse-grain abilities to control what notifications they receive—they can entirely disable notifications or temporarily mute them. Unfortunately, this means that users still must choose between a potential fire hose of notifications or risk entirely missing notifications for important events (e.g., group membership changes). For example, after a group is unmuted, users are not informed of who was added or removed while the group was muted, even though our survey results indicate this is critical information to users. While users can manually scan the chat log, we do not believe this is likely to happen in practice or effective for active/large chat groups.

6.2.2 Telegram. As with WhatsApp, the user creating the group is set as the admin, and users can look up who the admins are. By default, new groups are private (not discoverable in Telegram's group search function), though they can be made public by the admin to allow anyone to find the group and join (up to 200,000 users).

Unlike WhatsApp, all users in a (private) group can add members to the group, though only users in their contact list (unless changed by a group admin). Users who are added in this manner automatically join the group. Admins can generate invite links to add users not in any group member's contact list, and they can revoke these invites at any time.

Upon joining a group, users are not shown any indication regarding the security of their chat messages (see Figure 9). This may leave some users unclear regarding how safe it is to discuss specific topics within the chat

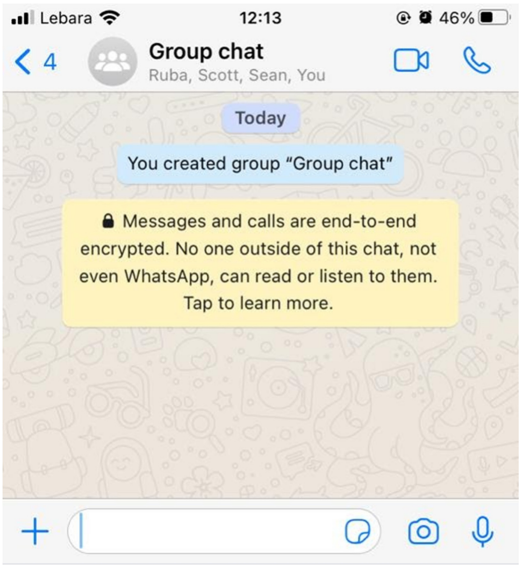

Fig. 8. WhatsApp: new group screen.

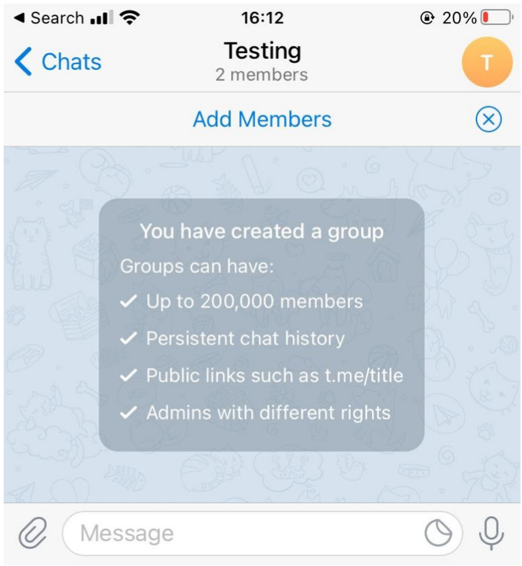

Fig. 9. Telegram: new group screen.

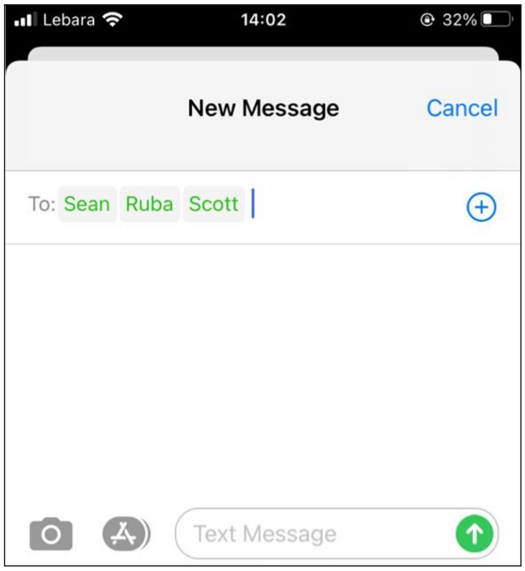

Fig. 10. iMessage: new group screen.

group. By default, new users do not have access to the past chat logs for the group, though the admins can change this.

Similar to WhatsApp, when users are added or removed from the group, they are shown a message indicating the group membership change in the chat history. As before, these messages can easily be missed if new chat messages displace them from the currently visible portion of the chat history.

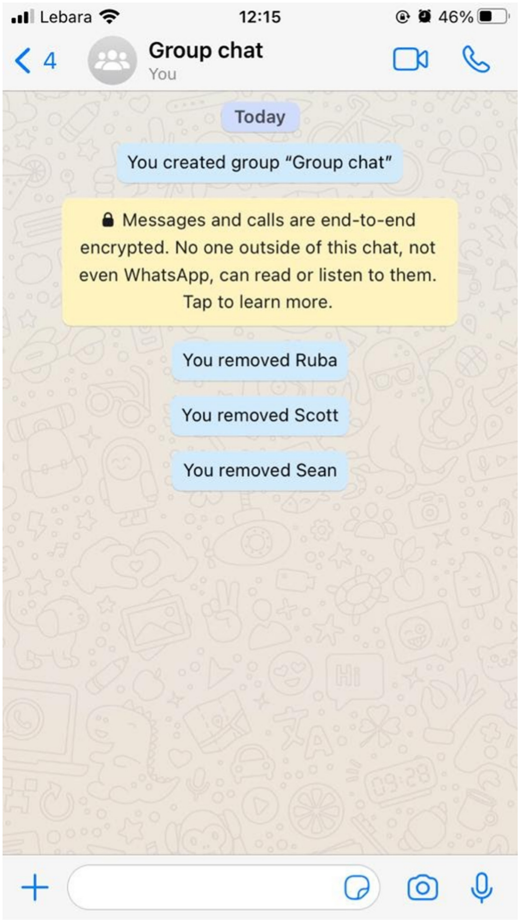

Fig. 11. WhatsApp: removing a user.

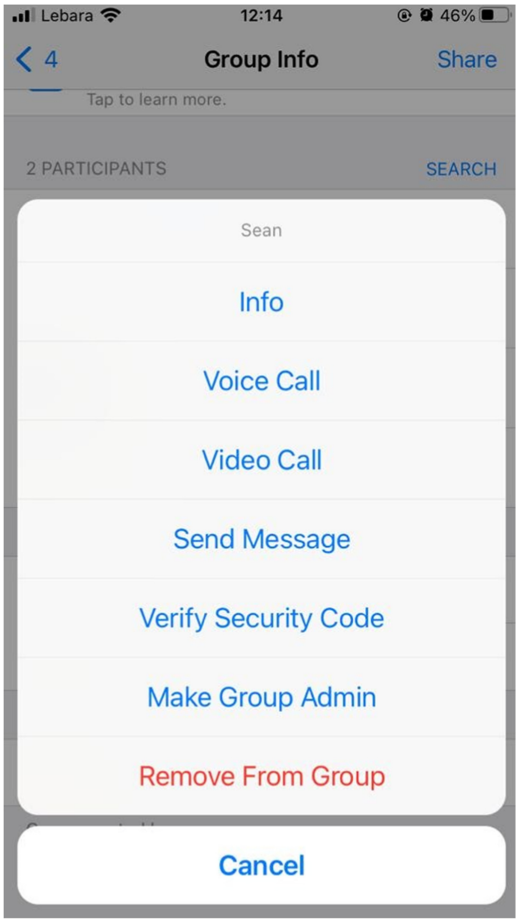

Fig. 12. WhatsApp: user options.

Only an admin can remove group members, and like WhatsApp, removed users can still see the chat history for the group up to the moment they were removed from the group. If an admin is removed, they lose all rights to the group (they cannot delete it), and no users are set as an admin to replace them, even if they are the only admin.

Like WhatsApp, users only have coarse-grained options for handling notification frequency, along with all the limitations this entails.

6.2.3 Signal. Like Telegram, any user can invite other members to join a group, though the admin can disable this. Unlike either WhatsApp or Telegram, users are only invited to the group, not automatically added—i.e., they receive an invitation and must accept it before they start receiving group chat messages. Admins can create links for joining the group, enabling easier public participation.

Similar to Telegram, upon joining a group, users are not shown a message informing them of the security properties provided by the group. Additionally, users only have access to chats sent after they are added to the group. Notifications regarding added and removed users also work as they do in WhatsApp and Telegram.

Only admins can remove group members, but unlike Telegram and WhatsApp, after a user is removed, they lose access to past chats from the group within the Signal app. If the admin removes themselves and they are the last admin, they are prompted to select a new admin, after which they will lose all access to the group. The new admin is notified within the chat window, though, like all notifications, this message might be missed by the new admin if other chat messages displace the notification in the chat history.

Like WhatsApp and Telegram, users only have coarse-grained options for handling notification frequency, along with all the limitations this entails.

6.2.4 Facebook Messenger. Users create groups by adding multiple users in a chat, with the creating user being the admin. Groups can have multiple admins.

Within Facebook Messenger, any group member can add other members to the group, though this can be disabled by the group admin. As with Telegram and Signal, users are not shown any details regarding the security or privacy features provided by the group chat. Unlike the other tools, messages predated by the user being added to the group will be visible by the newly added user in the chat history. Notifications regarding added and removed users also work as they do in WhatsApp, Telegram, and Signal.

Only admins can remove group users. Admins can remove themselves, even if they are the last admin, and if this happens, there will be no admins left in the group.

Like the other tools, users only have coarse-grained options for handling notification frequency, along with all the limitations this entails.

6.2.5 iMessage. Groups in iMessage are created by sending messages to multiple users. There are no admins for the groups. When creating groups, the security of the group relies on all members using iMessage. This is indicated by the color of the user's name in the participant list (see Figure 10)—green names indicate iMessage users, blue names indicate non-iMessage users. Only when all names are green will end-to-end encryption be used; otherwise, communication is conducted over MMS. This is very problematic as it is unlikely that most users understand the implication of having a blue name in the list, nor are its implications ever stated by the app.

Any member can add new members to the group, though this forks the group from the old group and requires users to continue using the new group. By default, new members do not have access to past group chat messages. As with other tools, notifications regarding added and removed users are shown in the chat history.

Anyone can also remove a user from a group, though, this forks the group and requires everyone to use the new group instead of the old one. Accidentally using the old (similarly named) group will allow the removed person to receive messages.

Like all the tools we tested, users only have coarse-grained options for handling notification frequency, along with all the limitations this entails.

6.3 Lessons Learned

All the tools handled group membership notifications in the same way—including those messages in the chat history. This makes it possible for those messages to be easily missed, especially in active groups where users may not review the entirety of chat history for messages that occurred since they last checked the group. If a user has notifications turned off or muted, it further increases the chance they will miss these notifications. As our survey shows, this is one of the key means by which users manage their privacy in group chat settings, and this is a significant problem.

Another common issue with the tools we studied is that except for WhatsApp, these tools do a poor job of communicating to users what security and privacy guarantees are provided by the tool. As such, users have incomplete information about the tools and may bring their own incorrect mental models for how the tool works. This is especially true in the case of iMessage, which requires users to understand how iMessage works

behind the scenes to understand the need for all group members to use iMessage for there to be any end-to-end guarantees. We identify two possible negative outcomes when users misunderstand the security of their group chat: (1) users overestimate the security and privacy of a given group, causing them to share information they would not be comfortable sharing if they had a correct understanding of the group's properties or (2) users underestimate the security and privacy of a group preventing them from sharing content they want to. In both cases, the severity of these problems will depend on the topics discussed in the group. However, they could lead to especially problematic issues if the topic is important—e.g., both oversharing and under sharing could be problematic if there is a significant medical issue that needs to be discussed.

We also find that users have very poor control over the notifications they receive. They need to either choose to receive all notifications or no notifications, though sometimes this can be per group. While this may address some needs, it is still overly coarse-grained. We expect users to be better served by notifications systems that allow them to mute unimportant chats while still highlighting security-critical notifications (e.g., group management). Similarly, once muted notifications are resumed, users need mechanisms that inform them about security-critical actions that occurred while their notifications were muted, but current tools do not provide this.

Regarding adding users to groups, we were surprised to find that in all tools besides Signal users did not need to agree before being added to a group. Depending on the group they are added to, this may result in users being suddenly inundated with notifications. This may act as a type of denial-of-service for the user's attention, and more work is needed to explore this potential issue.

There is support for auto-deleting messages in WhatsApp, Telegram, and Signal, but this is disabled by default, making it unlikely that users are aware of it. Tools supporting this functionality will need to do work to provide awareness of this functionality. Additionally, we believe that Facebook Messenger should add this functionality to address any concerns that new group members can see past group chat messages.

7 DISCUSSION AND FUTURE WORK

The purpose of our study was threefold: (1) to understand user perceptions of and requirements for the privacy and security of group chat tools, (2) to suggest improvements to existing tools to help users stay secure, and (3) to understand differences in attitudes between the US and the UK. In this section, we suggest fruitful directions for future work concerning the privacy and security of group chat tools.

7.1 Understanding Users

Our results shed light on the privacy and security concerns and requirements of users regarding group chat.

First, users share sensitive information in group chat settings and are concerned about their data privacy. Unfortunately, users do not select tools based on their ability to protect this data; instead, they choose tools widely used by their contacts. This leaves user data open to possible compromise.

Second, instead of relying on the security mechanisms of their group chat tools, most users leverage various non-technical strategies for maintaining their privacy and security in group chat settings. The most important of these strategies is to self-filter the content they share. This strategy ties closely into their second most important strategy, i.e., carefully managing and monitoring who is added and removed from their groups.

Third, users' biggest privacy concern has group members whom they do not know or trust read their messages. Relatedly, users struggle to know how to verify the identity of contacts other than through observing the actions of a contact and verifying that these actions match those of the contact in real life.

Fourth, at the same time, users want to ensure that some group chats remain pseudo-anonymous to prevent their messages from impacting how work colleagues, friends, or family may perceive them.

Fifth, despite their security and privacy concerns, users are largely desensitized to notifications. Users suffer a significant amount of alert fatigue, exacerbated by the fact that many group chat messages are irrelevant to the user receiving them. This fatigue could limit users' ability to pay attention to essential security notifications.

Finally, independent of their privacy and security concerns, users like the speed and features offered by group chat but may experience negative personal or social effects due to the "always on" and impersonal nature of online chat.

7.2 Understanding the Gap between Group Chat Tool Security and Users

As discussed in Background and Related Work (Section 2.1), researchers and tools are primarily concerned with the security of the underlying protocols. This focus on protocol security is reflected in the way that secure messaging tools advertise themselves to users. For example, on their homepage, Signal advertises that they use "state-of-the-art end-to-end encryption". Likewise, WhatsApp boasts end-to-end encryption as its key security feature.

Unfortunately, this focus on protocol security fails to address many of the users' self-identified needs. For example, users feel the need to self-filter their conversations. Yet, secure messaging tools lack functionality that could help simplify self-filtering (e.g., warning users before sharing sensitive-looking information with a large group, allowing users to have a moment or two to recall their message). Similarly, while existing tools ensure that only group members can read messages, they largely fail to help users monitor and manage group membership.

This is not to say that efforts to ensure protocol security are unimportant or misplaced. In truth, these protocol-level security properties are critical and address many issues that users are largely ignorant of [38, 49, 55, 56]. Instead, we believe our results indicate that *in addition* to a focus on protocol-level security, researchers and tool makers need to broaden the scope of their efforts to address users' stated security needs more fully. In part, this will include improving the usability and discoverability of existing features that might address these needs but are currently not used effectively (e.g., key verification).

7.3 Improving Group Chat Tools

Our research suggests several ways that group chat tools could be improved to provide better privacy and security to users.

Ensuring parity of functionality with one-to-one messaging. While one-to-one chat and group chat have many of the same security and privacy requirements, there is often a disconnect in how settings for these two use cases are handled. For example, in Facebook Messenger, WhatsApp, or GroupMe, if a user blocks another user, the user will no longer see one-to-one messages from the blocked user. However, he will continue to see messages from the blocked user if they are sent in a group chat. Similarly, some tools (e.g., Snapchat) default to end-to-end encryption for one-to-one chats, but not group chats. This behavior is especially undesirable as users are unlikely to notice this difference between the two modes.

Another related issue is that while it is feasible for a user to verify the private key fingerprint of the users whom they communicate with using one-to-one messaging, and this paradigm does not scale well to group chat. For example, when a new user joins a group, they would need to verify fingerprints for every other member of the group, which is not feasible for any moderately-sized group. This is only compounded by the difficulty of conducting key verification in existing tools [47, 48].

Adding additional group management options. Group management functionality across group chat applications is not consistent. In some cases, the feature set is simple, allowing anyone in the group to add or remove members of the group. Other tools allow for complex rules regarding who can manage the group's membership. For example, both WhatsApp and Facebook Messenger allow groups to be configured so that only administrators of the group can add or remove members.

Still, there is room for improvement. While group chat tools often inform respondents when a new member is added to a group, this is done as a message to the group and not as a persistent notification. As such, if the user is not actively monitoring the group, then the new member notification can be missed in the flood of other messages. This makes it difficult for users to track who is in a group, impeding their ability to accurately self-filter what they share.

We identify two possible solutions to this problem. First, the group management notifications could be made persistent, requiring them to be acknowledged separately from notifications about chat messages. Second, annotations could be added next to usernames in the group membership roster or in the group message log. These annotations could indicate how recently the member had been added to the group and whether the user had acknowledged this change.

Helping users trust new group members. When a new member is added to the group, it can be challenging to know whether they should trust that user. Most commonly, they need to rely on their knowledge of the new group member. Unfortunately, this approach does not scale well for larger groups.

Instead, tools could help users understand how much they should trust new group members. This could be accomplished using annotations next to usernames, similar to what we recommended for indicating newly added group members. These annotations could be used to quickly identify group members who: (a) the user has explicitly identified as a contact, (b) the user has interacted within another group, or (c) the group member is a contact of one of the user's contacts. This last annotation is reminiscent of PGP's web-of-trust. While the web-of-trust largely failed in e-mail [36], it may be that it provides a reasonable way to establish trust in group chat tools due to the more closed nature of these tools and their support for instantaneous communication. This approach could also help address the key verification problem identified above.

Identifying sensitive information for users. To help users avoid accidental disclosure of sensitive information, tools could try to help users identify when they are about to share sensitive information and warn users, similar to Thunderbird's attachment reminder that asks for forgotten attachments based on certain key words [15]. To aid this process, users could mark which groups are intended to contain sensitive information, allowing tools to only warn users for groups that are not supposed to contain sensitive information. To address potential privacy concerns arising from automated text analysis, the tool could analyze the data locally without storing either the analyzed text or the results of the analysis. More research is needed to identify the best method for preserving the user's privacy while simultaneously helping them avoid sharing sensitive information in the wrong contexts.

Addressing alert fatigue. Vance et al. [45] showed that habituation to non-security-related notifications causes people to disregard actual security warnings. Alert fatigue was a common problem identified by respondents in our study, stemming from the fact that in group chat, group members often receive many notifications regarding messages relevant to a specific member (or specific members), but not all members. This differs from message alerts in one-to-one communication, which are always intended for one user. To address this alert fatigue, some of our respondents mentioned disabling alerts entirely, which would clearly impact their ability to receive security notifications.

There are several potential approaches to addressing this problem. First, applications could show security-related notifications in a different way than other notifications. For example, security notifications could require explicit action to be dismissed, not just swiping the notification away. Second, the number of message notifications could be reduced. This reduction could consider whether the user had viewed existing notifications, how often they viewed those notifications, and how they responded to those notifications. By reducing message notifications, we believe it is more likely that users will pay attention to security notifications. Lastly, tools could display security notifications as interstitial dialogues, preventing the application from being used until the user acknowledged any security-related notifications. Research and development will be needed to identify the benefits and drawbacks of these and other similar approaches, along with which approach is most effective.

Educating app designers rather than users. Dechand et al. [9] make the strong statement that educating users about encryption is not going to change their behavior. Based on our results, we agree. Users choose applications based on what their peers are using, not security. However, the actual app designers do have both the technical knowledge and motivation to improve the security of their applications. Building a strong focus on usability and security within the app-building community is a logical way researchers can help keep group

chat users safe. Such an effort could include creating libraries, sharing knowledge at coding conferences, and establishing partnerships with companies.

7.4 Similarities between the US and UK

Users in the US and UK defined privacy and security in the context of group chat similarly. This does not mean that their broader views on privacy and security are necessarily the same, but only that their views are similar in the context of group chat. These similarities suggest that it would be possible to create tools that broadly meet users' security and privacy needs in the US and UK without customizing tools for different localities.

However, more research is needed to establish this fact for other nations. For example, prior work [22, 27] has shown that users in Asia use group chat in ways that differ significantly from those in Western nations. Our work could be replicated with Asian populations to see how these differences affect perceptions and requirements related to group chat.

8 CONCLUSIONS

We examined the security and privacy perceptions and requirements of 996 survey respondents regarding group chat. Our results demonstrated that users do share sensitive information in group settings, that they do not choose group tools based on their security properties, and that instead, they rely on non-technical strategies for protecting their privacy, such as self-filtering and monitoring group membership lists. We also find that group chat inundates users with alerts, making it easy to miss important security notifications. We conduct cognitive walkthroughs for five popular group chat tools to identify how these results relate to existing tools, showing that these tools largely fail to satisfy users' security and privacy needs and that these tools need to be improved to better align with user perceptions and requirements. Based on these results, we formulated several suggestions for improving these existing group chat tools, such as improving group membership management, helping users establish trust in new group members, and reducing alert fatigue.

APPENDICES

A CONSENT FORM

Introduction

You are invited to participate in a research study. The purpose of this study is to understand how people are using group instant messaging tools. This study is being conducted collaboratively by the University College of London, the University of Tennessee, and Brigham Young University. This study is open to all respondents 18 years and older.

Respondents' Involvement in the Study

You will complete a survey that asks questions regarding your experience using instant messaging tools for group communication. This can include sending text messages, images, video messages, or voice notes to others using instant messaging tools, such as Facebook Messenger, iMessage, or WhatsApp. These questions include multiple-choice and free-response questions on a range of topics related to using IM tools for group communication.

In this study, we are trying to understand how individuals use instant messaging tools for group chat. As such, there are no right or wrong answers. Please provide as honest of answers as you feel comfortable giving. We will not collect any information that can be used to connect you to the answers you give, such as your name or address.

Completion of the one-time survey, available through Qualtrics, should take approximately 1015 minutes. You will be compensated $2 for your efforts, with payment distributed through Prolific Academic.

Risks
There are no foreseeable risks relative to any procedures in this study other than those encountered in everyday life.

Benefits
This study will identify areas where existing instant messaging tools do not meet user needs. We anticipate that this information will be used by tool developers to improve the utility and usability of instant messaging tools.

Confidentiality
Responses to questions in this survey will be made available to the research community. Data will be sanitized to ensure that all personally identifiable information is removed before the data is shared. No reference will be made in oral or written reports which could link respondents to the study.

Participation
Your participation in this study is voluntary; you may decline to participate without penalty. If you decide to participate, you may withdraw from the study at any time without penalty and without loss of benefits to which you are otherwise entitled. If you withdraw from the study before data collection is completed, your data will be deleted by the researcher from the data collection file. Completion of the survey is all that is required to receive the full payment.

Consent
I have read the above information. I have had the opportunity to print a copy of this form. Clicking on the button to continue and completing the survey constitute my consent to participate.

B QUALTRICS SURVEY

Q1. Please enter your Prolific Academic ID.

Q2. How frequently do you use instant messaging tools for group chat?
○ *Daily* ○ *4-6 times a week* ○ *2-3 times a week* ○ *Once a week* ○ *Rarely* ○ *Never*

B.1 Tools

Q3. Please mark which of the following tools, if any, you have used. *(select all that apply)*
○ *Blackberry Messenger* ○ *Discord* ○ *Facebook Messenger* ○ *iMessage* ○ *IMO* ○ *Instagram Direct* ○ *Kakaotalk* ○ *Kik* ○ *Line* ○ *Marco Polo* ○ *Signal* ○ *Skype* ○ *Slack* ○ *Snapchat* ○ *Telegram* ○ *Viber* ○ *WeChat* ○ *WhatsApp* ○ *N/A*
Q4. For what purposes do you use instant messaging tools for group communication?

Q5. What, if anything, do you like about using instant messaging tools for group communication?

Q6. What, if anything, do you dislike about using instant messaging tools for group communication?

Q7. How do you choose which instant messaging tools to use for group communication?

B.2 Group Dynamics: Participation

Q8. When using instant messaging tools for group communication, who do you talk to? *(select all that apply)*
○ *Immediate/nuclear family members* ○ *Extended family members* ○ *Friends* ○ *Work colleagues* ○ *Other* ○ *I prefer not to answer*
Q9. On average, how large are your instant messaging groups?
○ *3–5 people* ○ *6–10 people* ○ *11–20 people* ○ *20+ people* ○ *My groups vary largely in size* ○ *Unsure / I prefer not to answer*

`Q10 shown if answer to Q9 is ‘‘My groups vary largely in size’’.`

Q10. Why do your instant messaging groups vary largely in size?

Q11. Have you ever been removed from an instant messaging group without your permission?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q12–13 shown if answer to Q11 is ‘‘Yes’’.`

Q12. If you are willing, please share why you were removed from an instant messaging group without your permission.

Q13. How did you feel after having been removed from an instant messaging group without your permission?

Q14. Should people ask the group for permission before inviting others to join an instant messaging group you are a member of?
○ *Yes* ○ *No* ○ *It depends* ○ *Unsure / I prefer not to answer*

`Q15 shown if answer to Q14 is ‘‘Yes’’ or ‘‘It depends’’.`

Q15. Why do you want other people to ask the group for permission before inviting others to join an instant messaging group you are a member of?

Q16. When do you review the member list of an instant messaging group? *(select all that apply)*
○ *When I first join a group* ○ *When a new member joins the group* ○ *When a member leaves the group* ○ *I check every now and then to see if anything has changed* ○ *I never review the member list* ○ *Unsure / I prefer not to answer*

Q17. Have you ever been in an instant messaging group chat where one or more members of the group chat only rarely participate in the group conversation?
○ *Yes* ○ *No* ○ *Unsure*

`Q18 shown if answer to Q17 is ‘‘Yes’’.`

Q18. How did you feel about having an instant messaging group chat where one or more members of the group chat only rarely participate in the group conversation?

B.3 Group Dynamics: Privacy

Q19. How comfortable are you with other members of an instant messaging group saving and/or sharing your conversations with non-members?
○ *Extremely uncomfortable* ○ *Somewhat uncomfortable* ○ *Neither comfortable nor uncomfortable* ○ *Somewhat comfortable* ○ *Extremely comfortable* ○ *Unsure / I prefer not to answer*

Q20. Are there topics that make you uncomfortable to read or discuss in instant messaging groups you are a member of? *(select all that apply)*
○ *Religion* ○ *Politics* ○ *Medical health* ○ *Mental health* ○ *Sexuality* ○ *Drug use* ○ *Other* ○ *Unsure / I prefer not to answer*

Q21. Has anyone ever shared something in an instant messaging group that placed you in an awkward position?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q22–23 shown if answer to Q21 is ‘‘Yes’’.`

Q22. If you are willing, please share how what was shared put you into an awkward position.

Q23. How did you respond to someone sharing something that placed you in an awkward position?

Q24. Have you ever joined an instant messaging group because you were interested in the topic being discussed and not because of who the group members were?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q25–26 shown if answer to Q24 is ‘‘Yes’’.`

Q25. What topics were discussed in these groups?

Q26. Do you recall a time when privacy was a concern for you when joining or participating in these groups?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q27 shown if answer to Q26 is ‘‘Yes’’.`

Q27. If you are willing, please share what your privacy concerns were when joining or participating in these groups.

B.4 Sensitive Information

Q28. Have you ever shared sensitive information in an instant messaging group?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q29 shown if answer to Q28 is ‘‘Yes’’.`

Q29. If willing, please share the types of sensitive information you have shared in an instant messaging group.

Q30. What does it mean to you that an instant messaging tool is secure for group communication?

Q31. What do you personally do to make sure your instant messaging group communications are secure?

Q32. How do you decide if an instant messaging tool is secure for group communication?

Q33. Are there any instant messaging tools you believe are secure for group communication?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q34 shown if answer to Q33 is ‘‘Yes’’.`

Q34. Please specify which tools you believe to be secure for group communication.

Q35. Have you ever been concerned that someone is not who they say they are when using instant messaging for group communication?
○ *Yes* ○ *No* ○ *Unsure / I prefer not to answer*

`Q36–37 shown if answer to Q35 is ‘‘Yes’’.`

Q36. Why were you concerned that someone is not who they say they are when using instant messaging for group communication?

Q37. How do you verify that someone is who they say they are when using instant messaging for group communication?

B.5 Demographics

Q38. What is your age?
○ *Under 21* ○ *21–34* ○ *35–44* ○ *45–54* ○ *55–64* ○ *65+* ○ *I prefer not to answer*

Q39. What is your gender?
○ *Male* ○ *Female* ○ *Other* ○ *I prefer not to answer*

Q40. Please specify your ethnicity.
○ *White or Caucasian* ○ *Black or African American* ○ *Asian* ○ *Pacific Islander* ○ *Mixed race* ○ *Other (specify)* ○ *I prefer not to answer*

Q41. What is the highest level of school you have completed or the highest degree you have received?
○ *Less than high school degree* ○ *High school graduate (high school diploma or equivalent including GED)* ○ *Some college but no degree* ○ *Associate's degree in college (2-year)* ○ *Bachelor's degree in college (4-year)* ○ *Master's degree* ○ *Professional degree (JD, MD)* ○ *Doctoral degree* ○ *I prefer not to answer*

B.6 Post-Survey Questionnaire

Q42. Please rate the overall difficulty of this survey.
○ *Very difficult* ○ *Somewhat difficult* ○ *Neither easy nor difficult* ○ *Somewhat easy* ○ *Very easy*

Q43. Please rate your overall satisfaction with the survey.
○ *Good* ○ *Neutral* ○ *Bad*

Q44. Please provide any additional comments on the survey overall.

AVAILABILITY

Our sanitized survey data is available for download at https://bitbucket.org/user-lab/oesch2020understanding/.

REFERENCES

[1] Ruba Abu-Salma, Kat Krol, Simon Parkin, Victoria Koh, Kevin Kwan, Jazib Mahboob, Zahra Traboulsi, and M. Angela Sasse. 2017. The security blanket of the chat world: An analytic evaluation and a user study of Telegram. In *Proceedings of the European Workshop and Usable Security*. Internet Society.

[2] Ruba Abu-Salma, Elissa M. Redmiles, Blase Ur, and Miranda Wei. 2018. Exploring user mental models of end-to-end encrypted communication tools. In *Proceedings of the 8th USENIX Workshop on Free and Open Communications on the Internet*.

[3] Ruba Abu-Salma, M. Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. 2017. Obstacles to the adoption of secure communication tools. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy*. 137–153.

[4] Wei Bai, Moses Namara, Yichen Qian, Patrick Gage Kelley, Michelle L. Mazurek, and Doowon Kim. 2016. An inconvenient trust: User attitudes toward security and usability tradeoffs for key-directory encryption systems. In *Proceedings of the 12th USENIX Symposium on Usable Privacy and Security*. 113–130.

[5] David L. Baumer, Julia B. Earp, and J. C. Poindexter. 2004. Internet privacy law: A comparison between the United States and the European Union. *Computers & Security* 23, 5 (2004), 400–412.

[6] John M. Carroll and Mary Beth Rosson. 1987. *Paradox of the Active User*. The MIT Press.

[7] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. 2018. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In *Proceedings of the 2018 ACM Conference on Computer and Communications Security*. 1802–1819.

[8] Alexander De Luca, Sauvik Das, Martin Ortlieb, Iulia Ion, and Ben Laurie. 2016. Expert and non-expert attitudes towards (secure) instant messaging. In *Proceedings of the 12th USENIX Symposium on Usable Privacy and Security*. 147–157.

[9] Sergej Dechand, Alena Naiakshina, Anastasia Danilova, and Matthew Smith. 2019. In encryption we don't trust: The effect of end-to-end encryption to the masses on user perception. In *Proceedings of the 2019 IEEE European Symposium on Security and Privacy*. 401–415.

[10] Heather Desurvire, Jim Kondziela, and Michael E. Atwood. 1992. What is gained and lost when using methods other than empirical testing. In *Proceedings of the Posters and Short Talks of the 1992 SIGCHI Conference on Human Factors in Computing Systems*.

[11] Steve Dodier-Lazaro, Ruba Abu-Salma, Ingolf Becker, and M. Angela Sasse. 2017. From paternalistic to user-centred security: Putting users first with value-sensitive design. In *Proceedings of the ACM Workshop on Values in Computing*.

[12] EverybodyWiki. 2021. K-modes Clustering. Retrieved August 30, 2021 from https://en.everybodywiki.com/K-modes_clustering.

[13] Rachel L. Finn, David Wright, and Michael Friedewald. 2013. Seven types of privacy. In *Proceedings of the European Data Protection: Coming of Age*. S. Gutwirth, R. Leenes, P. de Hert, and Y. Poullet (Eds.), Springer, 3–32.

[14] Gerhand Fischer. 1991. Supporting learning on demand with design environments. In *Proceedings of the International Conference on the Learning Sciences*.

[15] Daniel Folkinshteyn. [n.d.]. Attachment Reminder :: Add-ons for Thunderbird. Retrieved August 30, 2021 from https://addons.thunderbird.net/en-US/thunderbird/addon/attachment-reminder/.

[16] Nina Gerber, Verena Zimmermann, Birgit Henhapl, Sinem Emeröz, and Melanie Volkamer. 2018. Finally Johnny can encrypt. But does this make him feel more secure? In *Proceedings of the 13th International Conference on Availability, Reliability and Security*. 11.

[17] Mark Handel and James D. Herbsleb. 2002. What is chat doing in the workplace? In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*. 1–10.

[18] Michael M. Harris, Greet Van Hoye, and Filip Lievens. 2003. Privacy and attitudes towards Internet-based selection systems: A cross-cultural comparison. *International Journal of Selection and Assessment* 11, 2–3 (2003), 230–236.

[19] Anat Hashavit, Naama Tepper, Inbal Ronen, Lior Leiba, and Amir D. N. Cohen. 2018. Implicit user modeling in group chat. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*. 275–280.

[20] Richard M. Hessler. 1995. Privacy ethics in the age of disclosure: Sweden and America compared. *The American Sociologist* 26, 2 (1995), 35–53.

[21] Claire-Marie Karat, Robert Campbell, and Tarra Fiegel. 1992. Comparison of empirical testing and walkthrough methods in user interface evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

[22] Shipra Kayan, Susan R. Fussell, and Leslie D. Setlock. 2006. Cultural differences in the use of instant messaging in Asia and north America. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work*. 525–528.

[23] Jane E. Kirtley. 1999. Is implementing the EU data protection directive in the United States irreconcilable with the first amendment? *Government Information Quarterly* 16, 2 (1999), 87–91.

[24] Clayton Lewis, Peter G. Polson, Cathleen Wharton, and John Rieman. 1990. Testing a walkthrough methodology for theory-based design of walk-up-and-use interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

[25] Clayton Lewis and John Rieman. 1993. *Task-Centered User Interface Design: A Practical Introduction*.

[26] Rich Ling and Chih-Hui Lai. 2016. Microcoordination 2.0: Social coordination in the age of smartphones and messaging apps. *Journal of Communication* 66, 5 (2016), 834–856.

[27] Paul Benjamin Lowry, Jinwei Cao, and Andrea Everard. 2011. Privacy concerns versus desire for interpersonal awareness in driving the use of self-disclosure technologies: The case of instant messaging in two cultures. *Journal of Management Information Systems* 27, 4 (2011), 163–200.

[28] Juan Ramón Ponce Mauriés, Kat Krol, Simon Parkin, Ruba Abu-Salma, and M. Angela Sasse. 2017. Dead on arrival: Recovering from fatal flaws in email encryption tools. In *Proceedings of the LASER 2017 Learning from Authoritative Security Experiment Results Workshop*. 49–57.

[29] Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner. 2015. Investigating the computer security practices and needs of journalists. In *Proceedings of the 24th USENIX Conference on Security Symposium*. 399–414.

[30] Sandra J. Milberg, H. Jeff Smith, and Sandra J. Burke. 2000. Information privacy: Corporate management and national regulation. *Organization Science* 11, 1 (2000), 35–57.

[31] Deirdre K. Mulligan, Colin Koopman, and Nick Doty. 2016. Privacy is an essentially contested concept: A multi-dimensional analytic for mapping privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 2083 (2016), 20160118.

[32] Jakob Nielsen. 1994. Usability inspection methods. In *Proceedings of the Conference on Human Factors in Computing Systems*.

[33] Peter G. Polson, Clayton Lewis, John Rieman, and Cathleen Wharton. 1992. Cognitive walkthroughs: A method for theory-based evaluation of user interfaces. *Journal of Man-Machine Studies* 36, 5 (1992), 741–773.

[34] Stanley Presser, Mick P. Couper, Judith T. Lessler, Elizabeth Martin, Jean Martin, Jennifer M. Rothgeb, and Eleanor Singer. 2004. Methods for testing and evaluating survey questions. *Public Opinion Quarterly* 68, 1 (2004), 109–130.

[35] Paul Rösler, Christian Mainka, and Jörg Schwenk. 2018. More is less: On the end-to-end security of group chats in Signal, Whatsapp, and Threema. In *Proceedings of the IEEE European Symposium on Security and Privacy*. 415–429.

[36] Scott Ruoti, Jeff Andersen, Luke Dickinson, Scott Heidbrink, Tyler Monson, Mark O'neill, Ken Reese, Brad Spendlove, Elham Vaziripour, Justin Wu, Daniel Zappala, and Kent Seamons. 2019. A usability study of four secure email tools using paired participants. *Transactions on Privacy and Security* 22, 2 (2019), 1–33.

[37] Scott Ruoti, Jeff Andersen, Scott Heidbrink, Mark O'Neill, Elham Vaziripour, Justin Wu, Daniel Zappala, and Kent Seamons. 2016. "We're on the same page": A usability study of secure email using pairs of novice users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 4298–4308.

[38] Scott Ruoti, Tyler Monson, Justin Wu, Daniel Zappala, and Kent Seamons. 2017. Weighing context and trade-offs: How suburban adults selected their online security posture. In *Proceedings of the 13th USENIX Symposium on Usable Privacy and Security*. 211–228.

[39] Michael Schliep and Nicholas Hopper. 2019. End-to-end secure mobile group messaging with conversation integrity and deniability. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society*. 55–73.

[40] Brian Shackel. 1985. Human factors and usability-whence and whither? In *Software-Erqonomie'85: Mensch-Computer-lnteraktion*. H.-J. Bullinger (Ed.), B. G. Teubner, 13-31.

[41] H. Jeff Smith. 2001. Information privacy and marketing: What the US should (and shouldn't) learn from Europe. *California Management Review* 43, 2 (2001), 8–33.

[42] Eva Thulin. 2018. Always on my mind: How smartphones are transforming social contact among young Swedes. *Young* 26, 5 (2018), 465–483.

[43] Nik Unger, Sergej Dechand, Joseph Bonneau, Sascha Fahl, Henning Perl, Ian Goldberg, and Matthew Smith. (2015). SoK: Secure messaging. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*.

[44] Willemijn M. Van Dolen, Pratibha A. Dabholkar, and Ko De Ruyter. 2007. Satisfaction with online commercial group chat: The influence of perceived technology attributes, chat group characteristics, and advisor communication style. *Journal of Retailing* 83, 3 (2007), 339–358.

[45] Anthony Vance, David Eargle, Jeffrey L. Jenkins, C. Brock Kirwan, and Bonnie Brinton Anderson. 2019. The fog of warnings: How non-essential notifications blur with security warnings. In *Proceedings of the 15th USENIX Symposium on Usable Privacy and Security*.

[46] Elham Vaziripour, Justin Wu, Reza Farahbakhsh, Kent Seamons, Mark O'Neill, and Daniel Zappala. 2018. A survey of the privacy preferences and practices of Iranian users of telegram. In *Proceedings of the Workshop on Usable Security*.

[47] Elham Vaziripour, Justin Wu, Mark O'Neill, Daniel Metro, Josh Cockrell, Timothy Moffett, Jordan Whitehead, Nick Bonner, Kent Seamons, and Daniel Zappala. 2018. Action needed! helping users find and complete the authentication ceremony in signal. In *Proceedings of the 14th USENIX Symposium on Usable Privacy and Security*. 47–62.

[48] Elham Vaziripour, Justin Wu, Mark O'Neill, Jordan Whitehead, Scott Heidbrink, Kent Seamons, and Daniel Zappala. 2017. Is that you, Alice? A usability study of the authentication ceremony of secure messaging applications. In *Proceedings of the 13th USENIX Symposium on Usable Privacy and Security*. 29–47.

[49] Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the 6th USENIX Symposium on Usable Privacy and Security*. 1–16.

[50] Cathleen Wharton, John Rieman, Clayton Lewis, and Peter Polson. 1994. The cognitive walkthrough method: A practitioner's guide. In *Usability Inspection Methods*. J. Nielsen and R. Mack. (Eds.), John Wiley & Sons, Inc., 105–140.

[51] Wikipedia. Elbow Method (Clustering). [n. d.]. Retrieved August 30, 2021 from https://en.wikipedia.org/wiki/Elbow_method_(clustering).

[52] Wikipedia. K-means Clustering. [n. d.]. Retrieved August 30, 2021 from https://en.wikipedia.org/wiki/K-means_clustering.

[53] Wikipedia. Silhouette (Clustering). [n. d.]. Retrieved August 30, 2021 from https://en.wikipedia.org/wiki/Silhouette_(clustering).

[54] Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a privacy fundamentalist sell their DNA for $1000... if nothing bad happened as a result? The westin categories, behavioral intentions, and consequences. In *Proceedings of the 10th USENIX Symposium on Usable Privacy and Security*. 1–18.

[55] Justin Wu, Cyrus Gattrell, Devon Howard, Jake Tyler, Elham Vaziripour, Daniel Zappala, and Kent Seamons. 2019. "Something isn't secure, but I'm not sure how that translates into a problem": Promoting autonomy by designing for understanding in Signal. In *Proceedings of the 15th USENIX Symposium on Usable Privacy and Security*.

[56] Justin Wu and Daniel Zappala. 2018. When is a tree really a truck? Exploring mental models of encryption. In *Proceedings of the 14th USENIX Symposium on Usable Privacy and Security*. 395–409.

Received May 2021; accepted August 2021