

“The thing doesn't have a name”

Learning from emergent real-world interventions in smart home security

Bouwmeester, Brennen; Turcios Rodriguez, E.R.; Gañán, Carlos; van Eeten, Michel; Parkin, Simon

Publication date

2021

Document Version

Final published version

Published in

Proceedings of the 17th Symposium on Usable Privacy and Security, SOUPS 2021

Citation (APA)

Bouwmeester, B., Turcios Rodriguez, E. R., Gañán, C., van Eeten, M., & Parkin, S. (2021). “The thing doesn't have a name”: Learning from emergent real-world interventions in smart home security. In *Proceedings of the 17th Symposium on Usable Privacy and Security, SOUPS 2021* (pp. 493-512). (Proceedings of the 17th Symposium on Usable Privacy and Security, SOUPS 2021). USENIX Association.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



“The Thing Doesn’t Have a Name”: Learning from Emergent Real-World Interventions in Smart Home Security

Brennen Bouwmeester, Elsa Rodríguez, Carlos Gañán, Michel van Eeten, and Simon Parkin, *Delft University of Technology*

<https://www.usenix.org/conference/soups2021/presentation/bouwmeester>

This paper is included in the Proceedings of the Seventeenth Symposium on Usable Privacy and Security.

August 9–10, 2021

978-1-939133-25-0

Open access to the Proceedings of the Seventeenth Symposium on Usable Privacy and Security is sponsored by



“The Thing Doesn’t Have a Name”: Learning from Emergent Real-World Interventions in Smart Home Security

Brennen Bouwmeester, Elsa Rodríguez, Carlos Gañán, Michel van Eeten, Simon Parkin
Department of Multi-Actor Systems (MAS), Delft University of Technology
b.j.bouwmeester@student.tudelft.nl
{e.r.turciosrodriguez, c.hernandezganan, m.j.g.vaneeten, s.e.parkin}@tudelft.nl

Abstract

Many consumer Internet-of-Things (IoT) devices are, and will remain, subject to compromise, often without the owner’s knowledge. Internet Service Providers (ISPs) are among the actors best-placed to coordinate the remediation of these problems. They receive infection data and can notify customers of recommended remediation actions. There is insufficient understanding of what happens in peoples’ homes and businesses during attempts to remediate infected IoT devices. We coordinate with an ISP and conduct remote think-aloud observations with 17 customers who have an infected device, capturing their initial efforts to follow best-practice remediation steps. We identify real, personal consequences from wide-scale interventions which lack situated guidance for applying advice. Combining observations and thematic analysis, we synthesize the personal stories of the successes and struggles of these customers. Most participants think they were able to pinpoint the infected device; however, there were common issues such as not knowing how to comply with the recommended actions, remediations regarded as requiring excessive effort, a lack of feedback on success, and a perceived lack of support from device manufacturers. Only 4 of 17 participants were able to successfully complete all remediation steps. We provide recommendations relevant to various stakeholders, to focus where emergent interventions can be improved.

1 Introduction

The use of “smart” Internet-of-Things (IoT) home devices amongst consumers is growing, where this can include

internet-connected home appliances, entertainment systems, and home fittings such as smart doorbells or locks. The connectivity of these devices has historically lacked sufficient security [1, 23]. Many commonly-used IoT devices have not only technical vulnerabilities, but also ineffective configuration options for password and access permissions [3, 17]. This means that a range of consumer IoT devices continue to be susceptible to malware infections, facilitating various forms of abuse, from recruiting them into botnets to personal stalking and harassment [51].

There is a direction of travel to ensure that consumers purchase secure devices, e.g., increased awareness [48], labels indicating security properties [22, 47], and improved standards of device design [11]. However, for the foreseeable future, insufficiently secure devices continue to enter the consumer market. The brunt of the efforts to clean up infected IoT falls on both the end-users who own the devices and Internet Service Providers (ISPs), where more than 80% of the devices are located [14].

RFC6561 states that ISPs should notify users and ask them to remediate the threat [44]. Helping users protect their computer systems and remove infections has proven to be difficult for PC-based malware, even where users are more likely to have workable, effective tools available to them (for instance, automatic OS update mechanisms [74]). In the consumer IoT space, the conditions for user advice and remediation can be much more constrained when it is an ISP contacting a customer with advice; it is usually unclear what exact device, or even general device type, has been infected, forcing the advice to be highly generic. The lack of accessible user interfaces makes it difficult for users to perform the required security actions on the device they suspect is infected.

Prior work has found that notifying a user about an IoT infection can lead to cleanup [14]. Much less is known about the processes which take place in end-users’ homes after receiving a message with remediation advice. When technical experts are approached to clean a ‘smart’ personal device of suspected malware or unwanted code, they may not be able to confirm it is infected or prove removal of malware [33].

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2021.
August 8–10, 2021, Virtual Conference.

We conduct our study by partnering with an ISP which has sent notifications with remediation advice to customers infected with Mirai malware. We specifically report on the experiences of 17 ISP customers in their efforts to apply the advice. Mirai is a malware family that came to prominence in late 2016 [6], and has been referred to as the “king of IoT malware” [49]. It continues to be the leading malware family [39]. Following the notifications, we approached customers to conduct remote think-aloud observations of their attempts to follow the advice in their home, surrounded by a variety of potentially affected devices.

We focus on the following question: *How do end-users act on remediation advice about their infected Internet of Things device(s)?* To answer this question, we documented the end-to-end story of botnet remediation which included network measurements to identify affected users, and device owner engagement. Infection data received by the ISP allows us to identify users with an infection, but also to gauge the remediation success after the intervention. We combine this with qualitative data collected during the think-aloud observations. We make the following contributions:

- We report on the real-world, in situ experiences of 17 customers acting on advice for IoT devices suspected to be infected with malware. We step out of controlled lab conditions where advice that has a known outcome is directly provided to participants. This allows us to collect data with higher ecological validity.
- We show that users are motivated, yet the advice is constrained by what can be known about the location of the infection on a home network. Many recommended actions are in practice outcomes which users must find a way to reach based on behaviours familiar to them. This adds detail to the shortfalls in the last part of advice communication for smart home users – the implications of the best-placed stakeholders (the ISP) intervening to communicate advice which is the best-available practice or which has been consolidated from manufacturers, to context-expert end-users.
- We capture the importance of advice signal design for effective behaviour change relating to smart home security hygiene. For this we relate our results to the Fogg Behaviour Model [27]. We find that where the Activation Threshold for supporting an individual to reach a target behaviour is often treated as if it were a line to cross, with home IoT it is more akin to an ‘Action Diffraction’. The user is not *able to do enough* in a direct path to the goal, due to limitations inherent in the environment, such that advocated best-practice behaviours are non-deterministic. Participants applied a range of behaviours in an approach that appeared to have a good chance of working but which were not definitely going to be successful, or be confirmed as having been successful.

The context of malware infections of consumer IoT devices is discussed in the Background (Section 2), including how users are typically engaged to remedy consumer IoT infections. We describe our Methodology in Section 3, and Results from our in situ sessions with participants in Section 4. The implications of our participants’ experiences are discussed in Section 5 and contrasted with Related Work in Section 6. Concluding remarks and directions for future work close the paper in Section 7.

2 Background

Many devices enter the market that lack even basic security precautions [3]. The existence of a botnet such as Mirai starts with the manufacturing of IoT devices, which are then shipped, bought by retailers and later by consumers. Once a device has been infected, it is also unclear which of these stakeholders carries the responsibility for cleaning the device, but manufacturers generally lack incentives to prevent and remediate this problem [65].

2.1 Attacks on consumer IoT devices

Different malware families use different vectors to infect vulnerable devices (such as routers, cameras and digital video recorders) [6, 14]. In the case of Mirai, there are four stages [6, 19, 38, 45, 68]. The first stage is to perform a brute-force attempt to access the device using a sequence of entries from a list of standard known username/password combinations. If this brute-force succeeds, the newly infected device sends its IP and username/password combination to the attacker. In the third stage, the report server informs the loader, which loads the malware binaries onto the device. After the binaries have been executed successfully, they are deleted, and the device is now part of the botnet.

Many IoT devices do not support standard user interfaces, which makes it difficult for customers to change the standard passwords (assuming a device has such a feature to begin with, which may not be the case [26]). Even where a device has an adequate interface, many users prefer having a working device as soon as possible over going through security-related installation steps (such as replacing the standard password) thoroughly [40] (where the inter-connected nature of smart homes means this may include securing the entire home network). End-users who do care about security may lack knowledge to perform the right actions, due to the heterogeneity of IoT devices [5, 78].

2.2 Improving consumer IoT security

Information about the security qualities of IoT devices can potentially be difficult to find. One avenue of research focuses on supporting consumers to make informed choices about the smart home devices they buy in the first instance (e.g., security

labels [22, 47] and consumer guides [48]) Another area of focus has been to ensure that device design matches user needs; this has been noted regarding specific requirements for access control [34] and privacy in a shared environment [77], for instance.

Most vendors of IoT devices do not deliver a comprehensive manual or support page with their product. Where information is provided, details relating to security are often absent or not adequate [8, 30]. This means that even for those consumers who do care about security [9, 50, 64], the ‘transaction costs’ of ensuring purchase of the most secure device are simply too high [2, 8].

As the Internet increasingly connects end-users and their devices globally, it becomes complex for governments across the world to organise clear responsibilities and liabilities for security. As the IoT is still relatively new and evolving, it could take some time before governments are able to clean the market of insufficiently secure devices and exert pressure on responsible parties. Simple improvements such as labelling the level of security of devices could improve the purchasing environment [37], but even for such small improvements, incentives are lacking. As present, the most viable mitigation techniques mostly come from Internet Service Providers (ISPs) intervening when customers’ devices are compromised, or information campaigns to realise prevention through consumer awareness. However, levels of remediation are far from perfect. The content of a notification should be understandable and clear for target users, but there is a balance to be struck. Research has found that detailed steps can strengthen the effect of the notification [21, 43, 72]. On the other hand, messages should be plain and simple [29].

Even where users are aware of a security problem and activated to act, there can be uncertainty about which device is infected, or how to take the required action [60]. Users may instead rely on familiar techniques to solve problems on ‘unfamiliar’ devices, which often is not the correct approach for new types of devices and infections [76].

For structuring interventions, identifying critical points in life cycle of devices is useful [41]. *Opportune moments* for intervention then emerge [27], which are important for focusing resources toward enacting a behaviour at a specific point where it is more viable. Where purchase of new devices is one such point [22, 54], the notification to a customer of a suspected malware infection is another opportune moment. However, There are challenges inherent to deploying behaviour interventions where the ‘influencer’ does not manage the environment. In managed environments (including the artificial/controlled environment of a lab study), the influencer can know who the target is and how to reach them. Here, we study an environment where that knowledge is not immediately available. We then leverage technical tools to approximate where the intervention is needed, by triangulating across datasets to identify devices which are vulnerable. Simply put, we have to find a way to go to the participant,

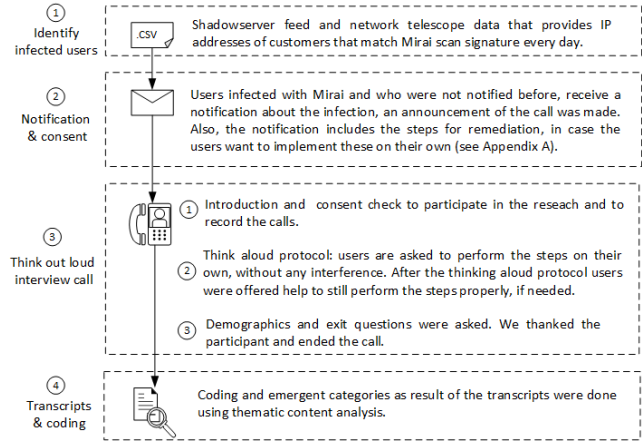


Figure 1: Approach and data collection.

whereas normally in a study the participant comes to us.

3 Methodology

In this section we describe our approach to answering the main research question. This involved partnering with an Internet Service Provider (ISP) and studying customer responses to remediation instructions.

3.1 Overall approach

Our study starts with identifying ISP customers who suffer from an active Mirai malware infection. For this, we used two data sources. One was the Shadowserver drone report [67]. The ISP receives from Shadowserver a daily list of IP addresses of customers that match the Mirai fingerprint. Mirai scans have a particular signature, where an artefact of the malware’s stateless scanning approach is that each probe includes a TCP sequence number equal to the destination IP address that the malware is targeting to attack [6]. This is conventionally used to detect the malware.

A network telescope was then employed. This is a set of unused IP addresses [46], where the traffic targeting this IP set is usually unsolicited. The network telescope of 300K IP addresses logs the IP addresses of hosts that were scanning with the Mirai fingerprint, as described in [6].

This is Phase 1 in the overall approach (as in Figure 1). The ISP is in a unique position to know which customer is associated with an IP address, so that we could identify which customers were suffering from a Mirai infection.

If the identified owner had not yet been notified, the ISP would notify the user about the infection via email (Phase 2, Figure 1). Included in this email would be an explanation of the research, and an invitation to participate in a call to understand better the process that users follow to execute the steps, as part of the standard service. It is also mentioned

that users are free to execute the steps themselves (see [Appendix A](#) for more details on the notification) without opting in to the study. During the call, each customer was asked explicit consent to participate in the research and record the call (see [Appendix B](#)). Minimal data of customers who did not consent to be part of this research was received in advance to be able to contact the customer, but it was not included in the results of this research.

To further ensure that the email notification could be understood by those end-users who received it, several communication experts from the communication department of the ISP transcribed the text to B1 level of the Common European Framework of Reference of Languages (CEFR) [25]. This is an international standard to describe language proficiency, in which B1 indicates basic level. The email notification was written in both English and Dutch (as the main language where the study was carried out).

A day after the email notification, users would be called (Phase 3, [Figure 1](#)). Three users did not answer during three attempts to call them and were left out of the study. Our protocol has a check at the beginning to ensure we talk to the device owner. We then asked users whether they wanted to opt into participating in the study, asked for explicit consent to record the interview, and explained that the participant could end the call at any moment ([Appendix B](#), part 1).

After concluding a call, a transcript was created. We used thematic analysis (Phase 4, [Figure 1](#)) to code transcribed copies of the interactions (from audio recordings). For performing the thematic analysis, the step-wise approach listed by [4] is used. Two of the researchers coded the transcripts to identify themes. Dedicated code review discussions took place between coders (to address emergent themes and conflicts), which happened in stages before arriving at the final set of themes. A balance in themes was found through iterative merging and splitting existing themes until convergence was reached into the most important themes (where the subsection in our Results represent theme families, Section 4). Saturation of themes was reached after 17 calls.

3.2 Think-aloud protocol

Originally we had planned to visit customers' homes/premises, to interact with them in an a natural and comfortable environment, and be physically present when users execute the recommended remediation advice. There was a need to instead develop a novel phone-based protocol for interacting with the customers of the partner ISP, foremost due to social distancing measures (Section 3.6). A positive aspect of this was that all participants were at the appropriate location when they were contacted.

To prepare, experience was gained in managing cases where remediation was not possible. One of the researchers accompanied a senior mechanic from the ISP for a day, and gained insights from the ISP customer support staff regarding

how to build trust with customers. In cases where the engineer is not successful in helping users, the most important step was seen as informing the consumer of the situation and to let them know about the possible ways forward. In such cases, also a supervisor should be informed about the issue. It can reach a point where informing the customer of an issue is the best one can do. This reflects the reality that the ISP is not technically responsible for the device, even though it has the opportunity to intervene.

The think-aloud protocol (Phase 3, [Figure 1](#)) consisted of three stages:

- **Stage 1: Consent and notification:** First, we obtained consent to conduct the study, asking then for approval to record the interview. Next, we checked whether participants received the notification and, if not, we sent it again and provided the participant time to read it.
- **Stage 2: Acting on the advice:** We allowed the participants opportunity to perform the actions and verbalize their thoughts, without direct input from the researcher. This think-aloud activity was transcribed and analysed.
- **Stage 3: Demographics and support:** We collected demographics and, if the researcher saw an action during Stage 2 as incomplete or incorrect, suggestions were offered for performing actions correctly, to the extent that this was possible (see 3.7). Last, we thanked the customer for their participation as well as provide e-mail details for future contact with the researcher in case they had any questions.

See [Appendix A](#) for complete details on the think-aloud protocol. The technical advice provided to customers (in the email and in the second step of the protocol) are steps used by the partner ISP, so it is what the ISP considered best advice. For comparison/reference, these steps are comparable to what is advised in online sources, as found on the Krebs on Security blog¹ and Symantec/Norton website².

During a call with a participant, they would try to implement the 5 recommended actions from the email: (1) determine which devices are connected to the internet that could potentially be infected with Mirai; (2) change the password of these devices; (3) restart the devices by turning them off and on; (4) reset the modem/router to the factory settings, and; (5) change the password of the modem/router ([Appendix A](#) contains the message in full).

3.3 Pilot

The study protocol was tested with 7 customers. These pilot sessions were especially important for refining the protocol,

¹<https://krebsonsecurity.com/2018/01/some-basic-rules-for-securing-your-iot-stuff/>

²<https://us.norton.com/internetsecurity-iot-smart-home-security-core.html>

Table 1: Summary of participants demographics, devices, actions, and outcomes. No. of users refers to the number of people in the household of the participant. Some connections were part of a small business rather than a home. Steps 1-5 refer respectively to actions relating to Device Identification, Device Password, Device Reset, Router Reset, and Router Password. Boxes highlighted in gray refer to an outcome classed as a failure to complete the associated Step, otherwise the action was a variation on a successful outcome. The letter-specific codes for each step are detailed in Figure 2.

Index	Age	Gender	No. Users	Suspected device	Step 1	Step 2	Step 3	Step 4	Step 5	Remediated?	Reinfection?
1	53	M	6	Router	1B	n.a.	n.a.	4A	5A	Yes	No
2	55	F	1	IP camera	1A	2D	3A	4C	5C	Yes	No
3	43	M	2	IP camera	1A	2D	3A	4C	5A	Yes	No
4	49	M	3	IP camera	1A	2D	3A	4C	5A	No	Yes
5	65	M	2	IP camera	1A	2C	3A	4D	5D	Yes	No
6	21	M	Business	IP camera	1A	2B	3C	4C	5A	Yes	No
7	45	M	4	Router	1B	n.a.	n.a.	4C	5C	Yes	No
8	65	M	2	NAS	1A	2C	3A	4C	5A	No	Yes
9	61	M	2	Smart printer	1A	2C	3A	4C	5A	Yes	Yes
10	34	M	Business	IP camera	1A	2A	3B	4A	5A	Yes	No
11	55	M	Business	NAS	1A	2A	3A	4A	5A	Yes	No
12	80	M	2	Doorbell	1A	2A	3A	4C	5A	Yes	Yes
13	49	M	1	IP camera	1A	2D	3A	4A	5A	Yes	No
14	43	M	2	-	1C	2E	3D	4A	5A	Yes	Yes
15	53	M	5	Router	1B	n.a.	n.a.	4B	5B	Yes	No
16	41	M	3	IP camera	1A	2B	3C	4C	5A	Yes	No
17	42	M	4	Smart TV	1A	2C	3A	4A	5A	No	No

as the main study would also involve interacting with real customers of the ISP and an intervention that has not been studied directly in a real-world setting. We could also evaluate the think-aloud protocol, accounting for not being present in the room with the users.

Similar to the insights from the ISP customer support staff, trust was found to be important: 5 of 7 customers were cautious about the call, 4 wanted a more detailed explanation of the research, and one called back to the service desk to confirm the authenticity of the research and email.

The pilot resulted in a check being added at the beginning of the protocol to talk to the person who takes care of security issues (as pilots included cases where the person who set up the devices did not live in the household); issues of delegation to informal technical support are discussed in [56]. The most significant change in the protocol was the inclusion of more upfront information about the purpose of both the call and research, to bolster trust.

3.4 Participants

All customers with a diagnosed Mirai infection in the period between May and July 2020 were notified by email about the infection and the study. If they did not opt out of the ISP’s support process, they were called the next day. During the experiment period, 37 unique IP addresses corresponded to 37 customers with Mirai infections. 12 were observed during the weekend, where the helpdesk at the ISP does not notify these users as they cannot provide support over the weekend. Of the 25 remaining IP addresses, 3 could not be notified due to

technical issues within the ISP, 2 did not respond to attempts to contact them after being notified, and 3 were not willing to take part in the experiment (did not opt-in to the study). There were think-aloud observations with 17 customers. The age of the participants was between 21 and 80 years old with a median age of 49. We interviewed 16 males and 1 female, and from the 17 participants, 3 used their internet connection to run their own businesses. Table 1 shows the participants’ demographics. As was also the case during the study pilot, sessions each took approximately 30 minutes in total (15 minutes of which was the think-aloud protocol).

No incentive was provided to users to participate, beyond the possibility of providing the technical support detailed in the participant-facing study materials (see Appendix).

3.5 Measuring cleanup

From the two data sources described in subsection 3.1, we received daily lists of IP addresses where infected Mirai hosts were located. This led to the initial identification of the customers and the recruitment of participants. We kept monitoring this data for an additional two weeks after the call.

Mirai reinfection can occur within a few minutes, or for some devices within 48 hours [14]. We chose a conservative 4-day window to determine remediation. Since Mirai attacks involve aggressively scanning the IP space for devices, we presumed a two-week window to measure reinfections as related to the state of participants’ home network. We illustrate this way of measuring outcomes in Table 1. We should note that this observation method is not perfect. While false positives

are highly unlikely, because of the specific Mirai fingerprint, false negatives might occur (an infected host might not show up in the data, even though it is still infected).

3.6 Ethics

The study protocol was approved by our institution’s human research ethics committee (TPM project 1083). The study design followed the principles for ICT human research as detailed in the Menlo Report [20] (as indicated also in the design of the think-aloud protocol). To make sure the end-users feel that they are in a safe environment, the think-aloud protocol is built around ensuring that the participant feels they are in a safe space and have not done anything wrong, and can state their feelings and actions without any judgement.

The first part of the call is about informed consent. This consent involves both taking part in this research anonymously, as well as the call taking place and the recording of it. Users were reminded that they could stop the study at any time. If they did not wish to participate, they were informed that they would be processed as usual by the partner ISP.

3.7 Limitations

In adherence with national social distancing measures related to the Covid-19 pandemic, in-person data collection was avoided. In-person home visits may have allowed for opportunistic observation of relevant details outside of our protocol, or differences between stated and actual behaviour. We compensated for this with a think-aloud protocol. We cannot rule out, however, that users may not have accurately described what they did via the call. Even though the researcher is trying to stay at the side-line, their presence influences the participants [36, 42, 71], who will typically pay more attention and effort to the tasks within the study. This does not detract from the context of the interaction, which would naturally require the individual to focus on the instructions regardless.

The research may have engaged with device owners who were unable to knowingly secure their devices. In such cases, at the end of the protocol they were helped to execute the steps they missed properly (after the think-aloud protocol). Also, an e-mail for future questions or contact was provided. The researcher helped the participants with any unsuccessful steps in accordance with the study protocol. Although infections could have plausibly been remediated, participants were carrying out actions themselves within the online ‘interview call’ format, and outcomes were based on customers’ reported actions. For instance, users may have changed passwords though we may not have been able to corroborate the outcome, or whether the advice absolutely caused the outcome.

Our work is based on users’ data from a single ISP. Hence, more research will be necessary to validate these results across multiple ISPs and different countries. Similarly, we focus our design and analysis on a single malware family, Mirai.

The recommended steps might differ from those for other malware families. We see trends of advice only becoming more complicated, see Section 5.2).

A final point is that our measurements of remediation and reinfection is not perfect. The infection data suffers from a small rate of false negatives. We compensate for this by working with longer time windows. Only when participant’s IP addresses are not seen in the infection data for four consecutive days, do we conclude they successfully remediated.

4 Results

Participant sessions were transcribed and analyzed to understand the ‘journey’ of remediation, following the steps of advice. We present our findings by following this journey. No participants reported having attempted to apply the steps before the session. We describe how participants attempted to: first, identify the infected device (Step 1, subsection 4.1); implement the recommended actions on that device and on their router (Step 2-5, subsection 4.2); infer the success of their actions (subsection 4.3), including their motivation to work through what transpired to be an arduous process for almost all participants (subsection 4.4). Finally, we connect the customer experiences with our measurement data on whether the infection was remediated (subsection 4.5).

Figure 2 provides an overview of reported participant actions. Each labelled box represents a particular action. To illustrate: 13 users took action 1A and identified a specific device as infected. White boxes indicate a successful action in terms of enacting advice, grey indicates no success.

4.1 Identifying suspect devices in the home

The first remediation action is to identify which devices are connected to the internet and could be infected with Mirai. The notification email informed participants that Mirai would not be present on a regular PC, laptop, tablet or phone. The subsequent actions (changing the password and turning the device off and on) are meant to be applied to all the devices that could potentially contain Mirai. A cautious approach is then to remediate and secure all potential victim devices.

Thinking aloud, four participants immediately focused on the device that they thought was the most likely culprit. All other participants started enumerating their devices, e.g., P12: *“I have 22 devices connected to the internet. Cameras, a garden sprinkler, a doorbell, the list goes on.”*

Whether multiple devices were enumerated or not, all participants focused on identifying one suspect – no participant ended up identifying multiple suspect devices. We observed participants using three heuristics to reason about the likely culprit. The first heuristic, used by the majority of participants, was a process of elimination, as with P04: *“I have a laptop, two mobile phones, no three mobile phones. I have a camera, a security cam, and the solar energy is also connected to the*

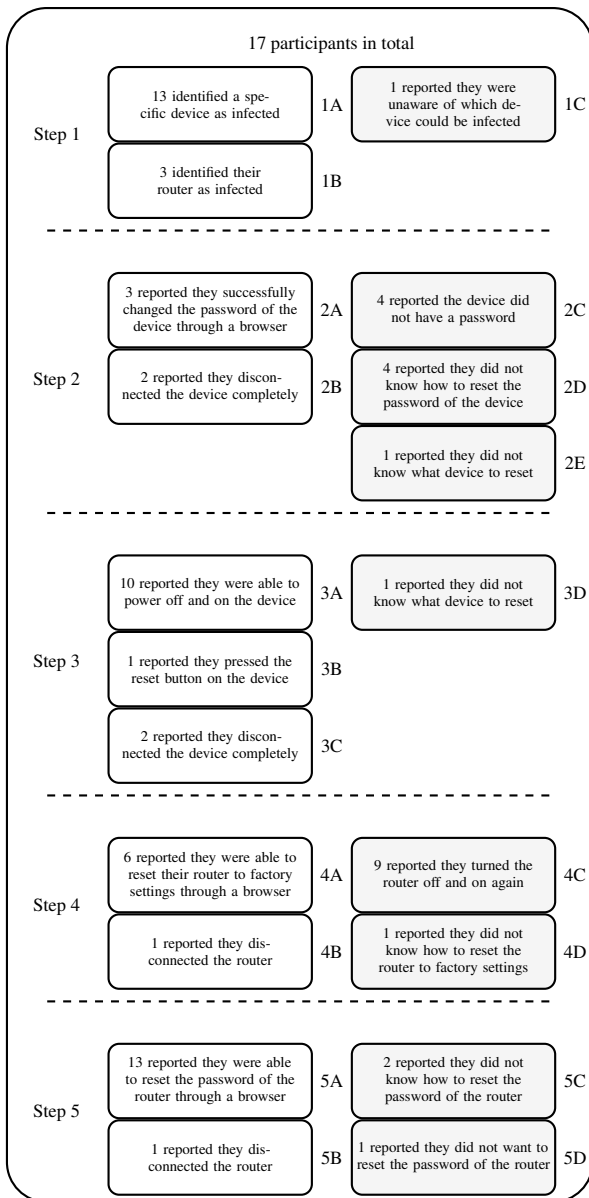


Figure 2: Overview of outcomes of actions by participants, while attempting to execute the remediation advice. Steps correspond to those found in [Appendix B](#).

internet. I run anti-virus on everything. I just bought that for five devices, also for my wife's iPad. According to that email, it would have to be the security camera."

This first heuristic might not lead to a confident identification, as seen with P01: "OK, in the email you write that it can't be phones, laptops, or really anything with Android on it. That leaves us with printers and cameras and the like. But I don't have those. Yeah, I have a printer, one of those all-in-one types, but that isn't even switched on at the moment [...] So that doesn't make sense."

The second heuristic, used by eight participants, was honing

in on a device that the person recently experienced problems with. This occurred for instance with P02: "I think it is the camera. [...] It says there is a system error and it needs a restart. But only the company can do this remotely.", and P06: "There are 4 phones connected to the wifi and a computer. And the security camera, but that doesn't work properly anymore. It actually seems likely that this camera is misbehaving."

A third heuristic was only employed by one person: conducting an Internet search. P15: "I have one all-in-one printer, that is never turned on, a beamer connected to the internet, an Xbox, Nintendo Switch, a smart TV, 2 laptops with Windows 10, a laptop with Windows 8 and a [routerModel] [...] Now, I saw in the email that it can't really be one of these devices, so I searched on Google for all my devices [...] then I found that [routerModel] has been having problems in [another country], so that was really the only clue I could find."

In one case, the participant enumerated the devices they owned, but felt uncertainty around finding the offending device made the whole process meaningless. It is interesting to note that all participants experienced this kind of uncertainty, but only P14, who indicated they had technical expertise, felt it invalidated the remediation path: "Can you see something useful, like an IP or MAC address or something? [...] I have no idea [what device could be the problem], so half of these steps I can't execute. That makes this process kind of useless."

4.2 Taking action with a suspect device

Only three participants reported that they were able to change the password of the suspect device (Fig 2). In these cases, the device either had an associated app or an interface on the device itself that allowed the user to initiate the password change. For, P11, who owned a Network Accessible Storage device (NAS): "Yeah, resetting the password, you can do that via a small screen [...]. It worked, now with a slightly more difficult password."

Four other participants indicated that they thought the device did not have a password, e.g., P09: "This [printer] has no password, does it? I can search on the internet, but I think the printer just appears on screen when I want to print. Other than that, there isn't much to it. I don't get any hits when I search for something related to passwords."

Four participants said they did not know how to change the password, as with P03: "Well, I really have no idea how to do this. I do not have a booklet or anything. And the thing has no name, I think. So you tell me how to do this. A friend of mine helped me with installing this thing, but he got killed in a car accident, so I can't ask him." One participant consulted the manual, P17: "There is really nothing useful in the booklet that comes with it. I only see things that prevent us from suing them." Two participants reported visiting the manufacturer website, to no avail, as for P13: "Yeah, I searched for this and I found a website that belongs to the device. But the site is totally unhelpful. I already know it is a camera, can't they put

something more useful on the site?”

Two participants ‘solved’ the problem by completely disconnecting the device, e.g., P06: *“You know what, I will just disconnect it. I have no idea how to change the password, but it is broken anyway, so I will take it offline and then we will buy a better one [...] I don’t want a virus in my network.”*. P16 followed a similar behaviour: *“Well, I thought that [the camera] would hang there as a deterrent. But then I got your email. I threw out the device right away, because I definitely do not want a virus.”* Chalhoub & Flechais [15] considered disconnecting a smart device as a *compensatory behaviour* that owners apply to address security and privacy concerns, regardless of whether it directly addressed the concern.

When it comes to restarting the suspected device, two participants looked for a dedicated reset button. P10: *“I am pressing the reset button for a long time [...] OK, it is turning off and on again.”* The second person looked for such a button but ended up, like nine other participants, disconnecting the power cable: P02: *“I don’t really see a button or anything on the camera. Perhaps just pulling the plug then?”*

The last two steps concerned the modem/router. At least six participants had the standard router issued by the ISP. The email from the ISP contained a link to a help page that described two actions: how to restart the device by disconnecting the power, and how to factory reset the device via a web interface. While the email asked users to factory-reset the router, the presence of both actions on the help page led some participants to take the first listed action: only disconnecting the power. Strengthened by the presence of this action on the help page, participants were convinced their efforts were the requested ones, P02: *“It says here to pull the plug and wait for 10 seconds, I can do that, great”*. Moreover, participants tend to copy the actions they took for earlier steps and implement those for their router, P08: *“Reset? So I will do the same as with the camera. I have disconnected it for 5 seconds and it is back in. I see a green light so I guess that worked”*. Overall, 6 participants reported having enacted a factory reset, while 9 participants removed the power cable to reset the device.

P05, who was running a small business, said they did not want to execute a factory reset: *“The problem is that I would have to set up all port forwarding again and I don’t really want to do that [...] Then I have to let IT come again. [...] Were the previous steps not enough to make the virus disappear?”*

For the final step, 13 participants reported that they successfully set a new password via the ISP web interface of the device, while two said they did not know how to do this. For this step, six participants made use of the URL in the notification (see [Appendix A](#)).

4.3 Inferring the success of remediation

When users manage to complete an action on the suspect device, they receive almost no feedback on the success of their efforts. The exception was when setting a new password

was supported via an interface that the participants are familiar with. The users who managed to reach a web interface for their router, for example, would get a clear confirmation when they successfully completed a password change. Still, all participants experienced actions that lacked feedback on whether they were successfully completed. More importantly, all participants lacked feedback on the success of their actions in terms of the main outcome: removal of the malware. These observations are of interest when compared to Forget et al. [29], and the examination of whether ‘engaged’ or ‘disengaged’ users arrive at secure outcomes to their (in)action to secure a computer – here the problem is that the outcome, secure or not, is not visible.

During the calls, we witnessed a clear desire by many participants to receive confirmation of whether they were doing the right things, as with P02: *“Shall I wait a few seconds? [...] OK, I think 10 seconds is enough, I am putting the plug back in [...] I am waiting for the lights to turn on again. It is supposed to be orange, right? Or green?”*

Some remediation actions were surrounded by uncertainty, while others were more clearly unsuccessful to the participants. In either case, participants regularly requested confirmation that they were successfully removing the virus. For instance, P04: *“Could it be enough if we do not change the password. That we do all other steps?”*, and P08: *“The device is already disconnected. Does that count as a reset if I now reconnect it again? I am really curious whether the virus is really gone. Can I reconnect it now?”*.

4.4 Motivation under uncertainty

All participants were willing, in some cases eager, to undertake the recommended actions, e.g., P09: *“I am now putting the plug of the router back in. What is the next step of this adventure?”* Participant motivation was illustrated by the degree to which they tolerated their uncertainty about what was asked of them, and whether they conducted the actions correctly. Motivation was also visible through the effort that was made. For example, the device or router might be in another part of the house or access to it might be blocked. This was the case for P03: *“You ask quite a bit from me, because then I have to make quite a mess. [...] Let me put the phone down, I need to move a few boxes... OK. What do I do now?”*, and P07: *“Then I will walk to the utility closet [...] I see the cable already, I will pull it out completely.”*

In addition, the factory reset of the router means that users lose their configuration, which might not be trivial to set up again. P10 debated this, *“Ah, so then I have to set up all port-forwarding and port assignments again. Well, I think that is the right thing to do, otherwise the virus will hang around.”*, as did P04: *“Oh, that is complicated. I did the same thing a while ago, but then I need to reconfigure all port forwarding again. But OK, if that helps, then we will do it again.”*

Only a few participants expressed doubts about the effort,

in all cases because they were not clear what problem Mirai posed, as with P01: “*Eh, let’s take a step back. I have no idea whatsoever about how that Mirai virus actually works. I mean, I do not experience any issues, right? So what is the problem?*” After an explanation about how Mirai-infected devices are used for criminal activity against other users and organizations on the internet, P01 concluded: “*Ah, right. That is understandable, I am happy to cooperate.*”. Renaud & Goucher [63] note that the ‘gulf of evaluation’ differentiates between the sense of being able to enact a security behaviour, and the ‘response efficacy’ of whether the behaviour is appropriate.

No participants dropped out before completing the steps. The only case where a participant did not want to conduct a specific step was P14, who felt none of their devices were plausible suspects, and as such did not want to implement a reset and password change on any of those devices. They did, however, proceed with subsequent steps involving the router.

Regarding the evidence for users’ seemingly high motivation, one potential source of bias here (as discussed in Section 3.7) could be an observer effect (a.k.a. the ‘Hawthorne effect’), where the fact that the participants know their actions are being ‘observed’ makes them more motivated than they might have been without the presence of the researcher.

4.5 The end: remediation, and reinfections

Table 1 presents an overview of participant-level actions and outcomes. Again, the coding used in the columns for the remediation steps relate to the boxes in Figure 2. After the intervention, 14 of the 17 participants were observed to be remediated, as measured by the absence of their IP address in the daily data feed of Mirai infections received by the ISP in the four days after the call. This may count as good news. The cumbersome non-deterministic remediation process seems at least probabilistically related to the desired outcome. Three participants remained infected. It is true that they did not fully execute the recommended steps, but the same holds for other participants who were regarded as having managed to remediate. Only four participants could be said to have fully executed the recommended actions (P01, P10, P11, P15). We include P15, because this person took the suspect device permanently offline, so in that sense ‘secured’ it from further harm. We monitored the presence of the IP address in the daily data feed for two more weeks after the remediation period. In 5 cases, we observed a re-infection with Mirai; there was a gap of three consecutive days where the user’s IP address was not reported in the daily data feed, and then it reappeared. Two of these reinfections were non-remediated users, three were users who did manage to remediate at first.

For the two non-remediated cases, the infection disappeared by an unknown cause five or more days after the call. This is consistent with the relatively high ‘natural’ cleanup rate seen elsewhere [14]. One explanation is that the Mirai malware is not persistent on the device, at least not at the time

of the study. This means that a power cycle may have removed the infection, although the device is still in a vulnerable state. It might be discovered and reinfected soon thereafter, because of the aggressive scanning conducted by Mirai bots.

The three cases where we observed an initial remediation, and a later reinfection, can have various explanations, and as such are indicative of avenues for future work. One explanation is that the detection of infections via the daily data feed is not perfect, potentially including false negatives. Another explanation is that these users did manage to get rid of the infections by power cycling the devices, but did not remediate the underlying vulnerability (i.e., set a secure password). This is consistent with our observations, because all three users did not fully execute the recommended actions. As noted from the observations, users may have otherwise had multiple infected devices and only focused on one, or focused their attention on the wrong device.

In the end, the gap we observed between advice and user actions cannot be blamed wholly on either the user or the advice-giver the ISP. It points to the responsibilities of a third actor: the manufacturer. Even when users went online and tried to find manufacturer information about solving security problems, there were complications. This was certainly the case for P16, who was not able to even identify the manufacturer: “*Well, there is no brand name on the device, haha, only IP-camera is printed on the side of it.*”

5 Discussion

Returning to our overarching research question, we provide real-world evidence of the gap between advice and outcomes in IoT [7], but also the impact this gap can have on smart home users. There are two sides to this story – the quality of advice, and the characteristics of the response to that advice.

Successful behaviour for our participants was often unconfirmed and unconfirmable, and neither the users nor the advice-giver can resolve this at present, given the constraints inherent in the situation (in home infection, limited device visibility, etc). This unbridgable gap points to the responsibilities of other actors, notably the manufacturer [32]. We could argue users lack capability, but it is not a lack of user capability, but a design flaw, pointing to the relationship between behaviour support and interface design to provide situational feedback (as highlighted elsewhere for user access control guidance [76]). The lack of ‘normal’ computing interfaces on IoT devices creates an environment fraught with confusion and uncertainty for applying standard security advice.

What we have for network-connected smart home devices is also a multi-party intervention. Participants had to wait for their efforts to be confirmed as worth it (that remediation will be confirmed at some point afterwards via network scans, and a *lack* of capacity for the ISP to follow up). Participants demonstrated despair over not knowing what to do and whether their effort was successful. Remediation is then




Figure 3: Action Diffraction for resetting a smart home device. Users may vary in Motivation, and rely on their Motivation to enumerate over possible solutions (standing in for a lack of knowing the precise Ability they need to apply). The target behaviour may be deterministic (the small circle, top right), but plausible variations surround it, informed in part by Instructions. It can be unclear if the applied Ability has achieved the intentions of the Prompt, even if it has been successful.

non-deterministic (very likely to work, but not definitely going to work). The lack of feedback stands in contrast to, say, removing Windows malware, where a removal tool—such as an anti-virus client—will typically report on what it found and whether it was effective in removing it. This limits the potential of checklists, for instance, if instructions cannot be made specific enough to a particular user’s set of network-connected devices (and are as such, ‘sub-optimally targeted’).

Participants applied one of the heuristics identified in our results, to navigate the gap in specificity, and attempt to identify the target of an advocated behaviour. Applying advice then leans on motivation, in that most participants were willing to try quite convoluted steps (going to another room to unplug the router, coming back to the phone, then back to the router, etc.). Where Redmiles et al. [59] isolate ‘bad advice’, we step back from this to identify ‘ecosystem factors which limit the capacity to construct good advice’. We regard this then as also exploring the limitations of *emergent* interventions for smart home security.

What is remarkable and worthy of further exploration is that our participants demonstrated somewhat correct reasoning in identifying suspect devices, consistent with actual properties of these devices. Mostly the heuristic is to eliminate suspect devices. This further highlights the important of local context to instantiating security advice for the smart home [76], but also making advice specific enough to be actionable [62].

5.1 Informing effective interventions

Where participants felt a need to enumerate over familiar behaviours, many would push back if they did not know how to enact the advocated behaviour. This points to self-efficacy, important for prompting action within various behaviour change approaches [24]. To put our findings in the context of enacting (what appeared to our participants as) a new behaviour, acting on notification of a malware infection is an *opportune moment* or *prompt* to enact a new behaviour, so we refer to the Fogg Behaviour Model (or B=MAP / B=MAT model [27]). In this model, Behaviour = Motivation + Ability + Prompt. The model has been used extensively across areas such as persuasive design and personal development, but also to understand social interventions for security [18], and opportunities for security interventions in a retail environment [54].

A Prompt can be a Facilitator, Spark, or Signal – here it is a Signal, that a device in the home is infected and that actions must be taken to resolve the issue, as a call to Motivation and upon an Ability to act. The ISP carries the Signal to the user (highlighting that ISPs are the *best-placed* party to intervene, but that this does not mean they are the *most appropriate*) – this relies on sufficient Motivation and Ability already being present. We found that participants were *over-investing* Motivation to make up for an insufficient definition of the target behaviour or outcome (a lack of capability to identify or confirm the appropriate Ability). Among our participants, there was uncertainty as to what was right to do, to the extent that a user may enact a behaviour which removes malware, but continue with further actions for lack of indication that they had already succeeded. This even includes where some of our participants chose to permanently disconnect or dispose of a suspect device (representing an *unintended harm* of unclear advice [16]). ‘Actionable choices’, with clear outcomes, are regarded as feasible in areas such as smart home privacy [66], and in supporting a user-defined ‘recovery state’ [35].

We show the gulf where these harms manifest as what we refer to here as ‘Action Diffraction’ (Figure 3). Where Renaud & Goucher refer to the ‘Gulf of Execution’ [63] (including knowing what needs to be done, but not how to do it), here we find a gulf created by restrictions in the vehicle of the intervention itself which makes the target behaviour indistinct. This applies to both knowing what the target behaviour is, and knowing whether it has been reached. Where the Activation Threshold is the point of realising a target Behaviour, and a user being activated to try to get over the Threshold, our results show efforts being ‘diffracted’, splitting off in many directions as participants find themselves exploring non-deterministic and potentially inapplicable behaviours (this includes where they have Ability to do something, but are not willing to try everything unless they can be Motivated to do so).

Renaud & Goucher [63] frame a ‘Gulf of Evaluation’ in formulating an intention to adopt a secure behaviour, and Redmiles et al. [61] identify dimensions of advice quality.

We note in reference to the latter that the specificity – and actionability – of advice, including the capacity to evaluate the efficacy of the behaviour [63], are also impacted by the specificity of the target behaviour and its confirmation. Our findings showed also that, as with other forms of security advice [62], multiple sources of instructions can potentially confuse users further.

One contributory factor to this problem is best articulated through the Behaviour Wizard of Fogg & Hreha [28]. The best-practice advice seen by smart home users is an ‘unfamiliar’ task (requiring a link to existing practices), but framed more like a ‘familiar’ task (one that does not need explanation), and so we saw participants replacing an unfamiliar action with familiar behaviour(s). This is a complex world of Things, where enacting the wrong behaviour can result in ‘proxy changes’ [53], regardless of whether the intended outcome is reached. A user may turn on and off many devices, or the wrong one and not the right one, or achieve the goal but lose tailored configuration settings in the effort, all while not knowing in the moment whether they have succeeded.

5.2 Implications for evolving IoT threats

If users only apply some of the advice they are given, or devices have inherent security weaknesses, they may *continue* to be vulnerable and require *regular* intervention. Users can follow advice but still suffer the same consequences again, if IoT infrastructure does not help them to stay recovered, or malware evolves. There are parallels to the Transtheoretical Model [57], where understanding specific stages of behaviour can identify security improvements [55]. Inherent weaknesses in the design of many smart home devices put a user back into an ‘unhealthy’ situation (e.g., a device repeatedly falling back into an insecure state), requiring repeated cycles of *contemplating* and *acting* on advice, to *maintain* secure devices.

New malware variants are moving away from short-lived infections, and becoming persistent and resistant to current interventions [12]. More efforts of the type we have observed for Mirai infections would be required where, for instance, thousands of QNAP network access storage devices have been targeted by persistent malware [75], and the direction of travel shows that advice from manufacturers is requiring users to follow 20 or more steps completely and successfully to resolve these issues [58]. Moreover, some of these variants are also starting to include countermeasures to make detection difficult. For instance, malware leveraging blockchain DNS or TOR makes it even harder for the interveners to assess the efficacy of the user’s actions [10, 69, 73]. This is all within the context of increasing use of smart home devices, which itself already increases the complexity of remediation when there are problems (as we saw evidence of here).

5.3 Recommendations

Here we describe recommendations emerging from our Results and consideration of behaviour change approaches, associating recommendations to specific stakeholders.

- **Confirmation of settings changes.** Visibility of changes to system status is a crucial design principle [52]. Here this applies to both Apps and Interfaces, as created and maintained by the *manufacturer*. This was seen among our participants as already happening for some devices and interfaces, but should be enshrined as a consistent design choice, to reduce the ‘diffraction’ of remediation efforts. This would then serve as a visible ‘security outcome’ [29], to *then* be able to consider whether the visible outcome was the correct step to follow. This may be necessary for future security issues if resetting / unplugging a device actually runs the risk of reinstating default credentials, for instance. This would complement efforts to standardise smart home device functionality (as in e.g., the UK [70] and US [26]) which aim to have *manufacturers* reduce the scope for misconfiguration as a vector for device compromise (as with e.g., easily-guessed ‘default’ settings).
- **Settings logs.** A log of settings changes can help both *users* and *ISPs* (or indeed anyone ‘helping’ users) to see and refer to a clear record of changes. This could also include notifying users of security settings which need to be changed at setup but have not been, or which have been changed but not by a registered user. Ideally, there would be some signalling to users when a security issue is suspected, where there is a general lack of event logging related to security [26].
- **Assisted remediation.** Our study showed that not all participants were able to follow the advice, or needed confirmation that they had followed it. For lack of being able to move incrementally toward a clearly focused outcome (Figure 3), having a helper on-call or on-site would increase chances of a successful outcome, if the previous steps cannot be achieved. This would be a low-bar in terms of ensuring that there is an intervention for all levels of Motivation and Ability – if *users* are as keen to follow advice as our participants, they cannot be blamed if they are trapped in a cycle of trying advice without confirmation of actions or visible evidence of success. This relates to having actionable choices to begin with. It also aligns with the incentives of *ISPs*, which could commercially offer such services, though this brings the risk of users distrusting notifications as a ploy to sell a service – *ISPs* might only offer the service if the user asks for it.

6 Related Work

Chalhoub & Flechais [15] studied real-world users of smart home devices, where limitations in device features and transparency were seen to frustrate privacy-related decisions. The authors characterised *compensatory behaviours* in response to concerns (such as disconnecting a device). We saw participants defaulting to ‘familiar’ behaviours as a strategy to approach the uncertain process of situating generic advice. Geeng & Roesner [31] studied multi-user smart homes, noting that when devices fail to function properly, alternative paths to using a device are needed. We saw a parallel, where participants sought a viable solution to critical security issues, but were at times reluctant to dismantle their smart home device configurations to achieve it.

In terms of supporting behaviour change, Forget et al. [29] studied the security attitudes, behaviours, and understanding of active computer users from device activity and interviews. The authors characterised ineffectively proactive users, who exerted too much effort for security or regularly performed familiar behaviours even if they did not match the security concern. Where the authors saw information-seeking behaviours, our participants felt challenged in determining what to seek information about (lacking both clarity as to what was the target device, and available diagnostic information). Crucially, Forget et al. highlighted the importance of tangible outcomes to user actions, where here there was a lack of clear outcomes; the authors identified ‘problematic knowledge gaps’, where for consumer IoT environments these gaps are constraints in advice and user support.

Reeder et al. [62] identify a range of criteria for good home security advice, including that it must be actionable. We identify a gap that requires the recipient of smart home security advice to be able to complete advice and relate advice received from others to their personal context. The authors also discuss the potential need to enumerate over possible versions of generic advice to reach specific advice, considering “offering the generic advice followed by specific instructions on how to implement it” – similarly, our participants applied strategies to do this themselves.

Redmiles et al. [61] identify ‘perceived efficacy’ of advice as important, where here there is an element of efficacy in being able to localise advice received from others. The advice the authors reviewed was regarded as mostly ‘actionable’, where here we explore the implications of advice which, at least for our participants, was not immediately actionable. Redmiles et al. regard network security as amongst the least actionable and most general security advice (e.g., “Secure your router”), raising questions of whether non-actionable advice should be given to users in the first place, and we provide real-world evidence informing this discussion.

Çetin et al. [13] studied a ‘walled garden’ approach of limiting users’ capacity to access the Internet while a device is infected. Here we learned about the remediation journey

while users were acting on suggested remediation actions locally themselves, rather than checking the effectiveness of the notification method alone.

7 Conclusion

Here we studied user efforts to apply advice provided to them by their ISP. We found that the advice was not specific enough to ensure that it was applicable to participants’ own smart home context. Critically, constraints to the specificity of advice limited how it was produced, communicated, and put into practice in a real-world setting. Only 4 of 17 participants completed all applicable advice steps successfully. Action typically went wrong at the second step (changing the password of the suspected device), or at the fourth step (resetting the router to its factory settings). 16 participants were able to pinpoint a plausible infected device, using one of three strategies we identified (including by process of elimination).

Our work informs the understanding of interventions for real-world IoT settings. The construction, communication, and enactment of technical advice to home users is both complex and collaborative. It involves end-users, their ISPs, device manufacturers, and technical experts to support successful outcomes. Putting our findings into perspective with the continuing need for technical support for home computers and mobile devices, the need to fix security issues of smart home devices can be expected to persist. Given the complexity and role of local context, this can be expected to require analysis of the smart home in situ, including return visits to users of reinfected devices. Future work will explore the capacity of intervention approaches which include multiple relevant stakeholders. For instance, a list of known vulnerable device models could aid both ISPs in informing end-users, and end-users themselves in identifying problematic devices which they use or are considering for purchase.

Acknowledgments

We thank the partner ISP company for permitting access to network data and knowledgeable staff, and facilitating engagement with their customers. This publication is part of the MINIONS project (number 628.001.033) of the “Joint U.S.-Netherlands Cyber Security Research Programme” which is (partly) financed by the Dutch Research Council (NWO). We also wish to thank our paper reviewers for their comments.

References

- [1] Fadele Ayotunde Alaba, Mazliza Othman, Ibrahim Abaker Targio Hashem, and Faiz Alotaibi. Internet of Things security: A survey. *Journal of Network and Computer Applications*, 88:10–28, 2017.

- [2] D Allen. Transaction costs. *Encyclopedia of Law and Economics*, 1999.
- [3] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. SoK: Security evaluation of home-based IoT deployments. In *2019 IEEE symposium on Security and Privacy (S&P)*, pages 1362–1380. IEEE, 2019.
- [4] Rosemarie Anderson. Thematic content analysis (TCA). *Descriptive presentation of qualitative data*, pages 1–4, 2007.
- [5] Eirini Anthi, Shazaib Ahmad, Omer Rana, George Theodorakopoulos, and Pete Burnap. EclipseIoT: A secure and adaptive hub for the Internet of Things. *Computers & Security*, 78:477–490, 2018.
- [6] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. Understanding the Mirai botnet. In *26th USENIX security symposium (USENIX Security '17)*, pages 1093–1110, 2017.
- [7] Christopher Bellman and Paul C van Oorschot. Best practices for IoT security: What does that even mean? *arXiv preprint arXiv:2004.12179*, 2020.
- [8] JM Blythe and SD Johnson. The consumer security index for IoT: A protocol for developing an index to improve consumer decision making and to incentivize greater security provision in IoT devices. *Living in the Internet of Things: Cybersecurity of the IoT*, 2018.
- [9] John M Blythe, Shane D Johnson, and Matthew Manning. What is security worth to consumers? investigating willingness to pay for secure Internet of Things devices. *Crime Science*, 9(1):1–9, 2020.
- [10] Leon Böck, Nikolaos Alexopoulos, Emine Saracoglu, Max Mühlhäuser, and Emmanouil Vasilomanolakis. Assessing the threat of blockchain-based botnets. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–11. IEEE, 2019.
- [11] Irina Brass, Leonie Tanczer, Madeline Carr, Miles Elsdon, and Jason Blackstock. Standardising a moving target: The development and evolution of IoT security standards. *IET*, 2018.
- [12] Calvin Brierley, Jamie Pont, Budi Arief, David J Barnes, and Julio C Hernandez-Castro. Persistence in Linux-based IoT malware. *NordSec 2020*, 2020.
- [13] Orçun Çetin, Lisette Altena, Carlos Gañán, and Michel Van Eeten. Let me out! Evaluating the effectiveness of quarantining compromised users in walled gardens. *Fourteenth Symposium on Usable Privacy and Security (SOUPS)*, 2018.
- [14] Orçun Çetin, Carlos Gañán, Lisette Altena, Takahiro Kasama, Daisuke Inoue, Kazuki Tamiya, Ying Tie, Katsunari Yoshioka, and Michel Van Eeten. Cleaning up the internet of evil things: Real-world evidence on ISP and consumer efforts to remove Mirai. In *Network and Distributed System Security Symposium (NDSS)*, 2019.
- [15] George Chalhoub and Ivan Flechais. “Alexa, are you spying on me?”: Exploring the effect of user experience on the security and privacy of smart speaker users. In *International Conference on Human-Computer Interaction*, pages 305–325. Springer, 2020.
- [16] Yi Ting Chua, Simon Parkin, Matthew Edwards, Daniela Oliveira, Stefan Schiffner, Gareth Tyson, and Alice Hutchings. Identifying unintended harms of cybersecurity countermeasures. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–15. IEEE, 2019.
- [17] Andrei Costin, Jonas Zaddach, Aurélien Francillon, and Davide Balzarotti. A large-scale analysis of the security of embedded firmwares. In *23rd USENIX Security Symposium (USENIX Security '14)*, pages 95–110, 2014.
- [18] Sauvik Das, Laura A Dabbish, and Jason I Hong. A typology of perceived triggers for end-user security and privacy behaviors. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, 2019.
- [19] Michele De Donno, Nicola Dragoni, Alberto Giaretta, and Angelo Spognardi. DDoS-capable IoT malwares: Comparative analysis and Mirai investigation. *Security and Communication Networks*, 2018, 2018.
- [20] David Dittrich, Erin Kenneally, et al. The Menlo report: Ethical principles guiding information and communication technology research. http://www.caida.org/publications/papers/2012/menlo_report_actual_formatted, 2012. Accessed:2021-05-25.
- [21] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, et al. The matter of Heartbleed. In *Proceedings of the 2014 conference on internet measurement conference*, pages 475–488, 2014.
- [22] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. Ask the experts: What should be on an IoT privacy and security label? In *2020 IEEE Symposium on Security and Privacy (S & P)*, pages 447–464. IEEE, 2020.

- [23] European Union & Agency for Network and Information Security. Baseline security recommendations for IoT in the context of critical information infrastructures - Publications Office of the EU. <https://op.europa.eu/en/publication-detail/-/publication/c37f8196-d96f-11e7-a506-01aa75ed71a1/language-en>, 2017. Accessed:2021-05-25.
- [24] European Union Agency for Cybersecurity (ENISA). Cybersecurity culture guidelines: Behavioural aspects of cybersecurity. <https://www.enisa.europa.eu/publications/cybersecurity-culture-guidelines-behavioural-aspects-of-cybersecurity>, 2018. Accessed:2021-05-25.
- [25] European Union and Council of Europe. Document Library | Europass. <https://europa.eu/europass/en/document-library>, 2004. Accessed: 2021-05-25.
- [26] Michael Fagan, Mary Yang, Allen Tan, Lora Randolph, and Karen Scarfone. Security review of consumer home Internet of Things (IoT) products. <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8267-draft.pdf>, 2019.
- [27] Brian J Fogg. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*, pages 1–7, 2009.
- [28] Brian J Fogg and Jason Hreha. Behavior wizard: A method for matching target behaviors with solutions. In *International Conference on Persuasive Technology*, pages 117–131. Springer, 2010.
- [29] Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang. Do or do not, there is no try: user engagement may not improve security outcomes. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 97–111, 2016.
- [30] Steven Furnell. Making security usable: Are things improving? *Computers & Security*, 26(6):434–443, 2007.
- [31] Christine Geeng and Franziska Roesner. Who’s in control? interactions in multi-user smart homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [32] Julie M Haney, Yasemin Acar, and Susanne M Furman. "it’s the company, the government, you and I": User perceptions of responsibility for smart home privacy and security. In *30th USENIX Security Symposium (USENIX Security ‘21)*, Vancouver, B.C., August 2021. USENIX Association.
- [33] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium (USENIX Security ‘19)*, pages 105–122, 2019.
- [34] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlene Fernandes, and Blase Ur. Rethinking access control and authentication for the home internet of things (IoT). In *27th USENIX Security Symposium (USENIX Security ‘18)*, pages 255–272, 2018.
- [35] Weijia He, Jesse Martinez, Roshni Padhi, Lefan Zhang, and Blase Ur. When smart devices are stupid: negative experiences using home smart devices. In *2019 IEEE Security and Privacy Workshops (SPW)*, pages 150–155. IEEE, 2019.
- [36] Bonnie E John and Steven J Marks. Tracking the effectiveness of usability evaluation methods. *Behaviour & Information Technology*, 16(4-5):188–202, 1997.
- [37] Shane D Johnson, John M Blythe, Matthew Manning, and Gabriel TW Wong. The impact of IoT security labelling on consumer product choice and willingness to pay. *PLoS one*, 15(1):e0227800, 2020.
- [38] G. Kambourakis, C. Koliass, and A. Stavrou. The Mirai botnet and the IoT zombie armies. In *MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM)*, pages 267–272, 2017.
- [39] Kaspersky. New Mirai botnet is targeting enterprise IoT | Kaspersky official blog. <https://www.kaspersky.com/blog/mirai-enterprise/26032/>, 2019. Accessed: 2021-05-25.
- [40] Constantinos Koliass, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. DDoS in the IoT: Mirai and other botnets. *Computer*, 2017.
- [41] David Kotz and Travis Peters. Challenges to ensuring human safety throughout the life-cycle of smart environments. In *Proceedings of the 1st ACM Workshop on the Internet of Safe Things*, pages 1–7, 2017.
- [42] Clayton Lewis. *Using the "thinking-aloud" method in cognitive interface design*. IBM TJ Watson Research Center Yorktown Heights, NY, 1982.
- [43] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. Remediating web hijacking: Notification effectiveness and webmaster comprehension. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1009–1019, 2016.

- [44] Jason Livingood, Nirmal Mody, and Mike O'Reirdan. Recommendations for the Remediation of Bots in ISP Networks. <https://tools.ietf.org/html/rfc6561>, 2012. Accessed:2021-05-25.
- [45] J. Margolis, T. T. Oh, S. Jadhav, Y. H. Kim, and J. N. Kim. An in-depth analysis of the Mirai botnet. In *2017 International Conference on Software Security and Assurance (ICSSA)*, pages 6–12, 2017.
- [46] David Moore. Network telescopes: Observing small or distant security events. In *11th USENIX Security Symposium (USENIX Security '02)*, San Francisco, CA, August 2002. USENIX Association.
- [47] Philipp Morgner, Christoph Mai, Nicole Koschate-Fischer, Felix Freiling, and Zinaida Benenson. Security update labels: Establishing economic incentives for security patching of IoT consumer products. In *2020 IEEE Symposium on Security and Privacy (S&P)*, pages 429–446. IEEE, 2020.
- [48] Mozilla. *Privacy not included. <https://foundation.mozilla.org/en/privacynotincluded/>, 2021. Accessed: 2021-05-25.
- [49] Netscout. Dawn of the terrorbit era. https://www.netscout.com/sites/default/files/2019-02/SECR_001_EN-1901%20-%20NETSCOUT%20Threat%20Intelligence%20Report%20H%202018.pdf, 2018. Accessed: 2021-05-25.
- [50] Kenneth D Nguyen, Heather Rosoff, and Richard S John. Valuing information security from a phishing attack. *Journal of Cybersecurity*, 3(3):159–171, 2017.
- [51] Larissa Nicholls, Yolande Strengers, and Jathan Sadowski. Social impacts and control in the smart home. *Nature Energy*, 5(3):180–182, 2020.
- [52] Jakob Nielsen. Enhancing the explanatory power of usability heuristics. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 152–158, 1994.
- [53] Magda Osman, Scott McLachlan, Norman Fenton, Martin Neil, Ragnar Löfstedt, and Björn Meder. Learning from behavioural changes that fail. *Trends in Cognitive Sciences*, 2020.
- [54] Simon Parkin, Elissa M Redmiles, Lynne Coventry, and M Angela Sasse. Security when it is welcome: Exploring device purchase as an opportune moment for security behavior change. In *Proceedings of the Workshop on Usable Security and Privacy (USEC'19)*. Internet Society, 2019.
- [55] Shari Lawrence Pfleeger, M Angela Sasse, and Adrian Furnham. From weakest link to security hero: Transforming staff security behavior. *Journal of Homeland Security and Emergency Management*, 11(4):489–510, 2014.
- [56] Erika Shehan Poole, Marshini Chetty, Tom Morgan, Rebecca E Grinter, and W Keith Edwards. Computer help at home: methods and motivations for informal technical support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 739–748, 2009.
- [57] James O Prochaska, Colleen A Redding, Kerry E Evers, et al. The transtheoretical model and stages of change. *Health behavior: Theory, research, and practice*, 97, 2015.
- [58] QNAP. Security Advisory for Malware QSnatch - Security Advisory | QNAP. <https://www.qnap.com/en/security-advisory/nas-201911-01>, 2021. Accessed: 2021-05-25.
- [59] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. Where is the digital divide? a survey of security, privacy, and socioeconomic. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 931–936, 2017.
- [60] Elissa M Redmiles, Amelia R Malone, and Michelle L Mazurek. I think they're trying to tell me something: Advice sources and selection for digital security. In *2016 IEEE Symposium on Security and Privacy (S&P)*, pages 272–288. IEEE, 2016.
- [61] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. A comprehensive quality evaluation of security and privacy advice on the web. In *29th USENIX Security Symposium (USENIX Security '20)*, pages 89–108. USENIX Association, August 2020.
- [62] Robert W Reeder, Iulia Ion, and Sunny Consolvo. 152 simple steps to stay safe online: Security advice for non-tech-savvy users. *IEEE Security & Privacy*, 15(5):55–64, 2017.
- [63] Karen Renaud and Wendy Goucher. The curious incidence of security breaches by knowledgeable employees and the pivotal role a of security culture. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, pages 361–372. Springer, 2014.
- [64] Brent Rowe and Dallas Wood. Are home internet users willing to pay ISPs for improvements in cyber security? In *Economics of information security and privacy III*, pages 193–212. Springer, 2013.

- [65] Bruce Schneier. *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company, 2018.
- [66] William Seymour, Martin J Kraemer, Reuben Binns, and Max Van Kleek. Informing the design of privacy-empowering tools for the connected home. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020.
- [67] Shadowserver. Drone/Botnet-Drone Report | Shadowserver. <https://www.shadowserver.org/what-we-do/network-reporting/drone-botnet-drone-report/>, 2021. Accessed: 2021-05-25.
- [68] H. Sinanović and S. Mrdovic. Analysis of Mirai malicious software. In *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–5, 2017.
- [69] Alex Turing, Hui Wang, and Liu Yang. New threat: Matryosh botnet is spreading. <https://blog.netlab.360.com/matryosh-botnet-is-spreading-en/>, 2021. Accessed: 2021-05-25.
- [70] UK Department for Digital, Culture, Media & Sport (DCMS). Code of practice for consumer IoT security. <https://www.gov.uk/government/publications/code-of-practice-for-consumer-iot-security>, 2018.
- [71] Maaïke Van Den Haak, Menno De Jong, and Peter Jan Schellens. Retrospective vs. concurrent think-aloud protocols: testing the usability of an online library catalogue. *Behaviour & information technology*, 22(5):339–351, 2003.
- [72] Marie Vasek and Tyler Moore. Do malware reports expedite cleanup? an experimental study. In *5th Workshop on Cyber Security Experimentation and Test (CSET '12)*, 2012.
- [73] Hui Wang. Fbot, a Satori related botnet using block-chain DNS system. <https://blog.netlab.360.com/threat-alert-a-new-worm-fbot-cleaning-adbminer-is-using-a-blockchain-based-dns-en/>, 2018.
- [74] Rick Wash, Emilee Rader, Kami Vaniea, and Michelle Rizor. Out of the loop: How automated software updates cause unintended security consequences. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 89–104, 2014.
- [75] ZDNet. Thousands of QNAP NAS devices have been infected with the QSnatch malware | ZDNet. <https://www.zdnet.com/article/thousands-of-qnap-nas-devices-have-been-infected-with-the-qsnatch-malware/>, 2019. Accessed: 2021-05-25.
- [76] Eric Zeng, Shrirang Mare, and Franziska Roesner. End user security and privacy concerns with smart homes. In *thirteenth symposium on usable privacy and security (SOUPS 2017)*, pages 65–80, 2017.
- [77] Eric Zeng and Franziska Roesner. Understanding and improving security and privacy in multi-user smart homes: a design exploration and in-home user study. In *28th USENIX Security Symposium (USENIX Security '19)*, pages 159–176, 2019.
- [78] Verena Zimmermann, Paul Gerber, Karola Marky, Leon Böck, and Florian Kirchbuchner. Assessing users' privacy and security concerns of smart home technologies. *i-com*, 18(3):197–216, 2019.

A Appendix A – Notification Message and Instructions

Dear Sir/Madam [name],

We have discovered a security issue on your internet connection. We would like to resolve this issue together with you. The following sections explain how.

What is going on?

We have noticed that one or more internet-connected devices in your home have been infected with the mirai virus. While we cannot exactly detect which one of your connected devices has been infected, it is most likely a device such as a digital video recorder (DVR), security camera or printer connected to the Internet. Devices infected with the Mirai virus are typically **not** computers, laptops, tablets or mobile phones. The infection means that right now criminals have access to your infected device. This is putting you and other internet users at risk.

Tomorrow we will call you to resolve the issue

Our colleague, Mr. _____, will call you within a day to help you remove the virus. We gladly help you with this, as customers find it difficult to resolve the issue on their own. Moreover, the call will be a part of a scientific research that is executed together with _____ about the virus. This means we will ask you several questions to be able to help our customers better in the future.

Do you wish to remove the virus on your own?

Please let us know by a reply to this email or during the phone call. After that, please execute the following steps.

These are the steps needed to remove the virus

Step 1. Determine which devices are connected to your Internet connection. The Mirai virus mainly infects Internet connected devices such as a digital video recorder (DVR), security camera or printer connected to the Internet (not computers, laptops, tablets or mobile phones).

Step 2. Change the password of the Internet connected devices. Choose a password that is hard to guess. If you do not know the current password, please refer to the manual. By following these steps, you have prevented future infections.

Step 3. Restart the Internet connected devices by turning them off and on again. Hereafter, the Mirai virus has been removed from the memory of the devices.

Now that your Internet connected devices are safe, the last steps are to protect your router/modem.

Step 4. Reset your modem/router to the factory settings. On https://www._____.htm it is described how you can do this for an _____.

Step 5. Change the password of your modem/router. On https://www._____ it is described how you can do this for an _____.

NOTE: If remote access to a certain device is absolutely necessary, manually define port forwards in your router for the device. On https://_____/internet-9/port-forwarding-upnp-_____ it is described how you can do this for an _____.

Do you have any questions?

Please ask them in a reply to this email or during the phone call.

Kind regards,

Abuse Team
abuse@_____

The _____ Abuse department deals with security incidents for _____. You can find more information about the Abuse department on: https://www._____/abuse

Figure 4: Notification and opt-out invitation

B Appendix B – Think-Aloud Protocol

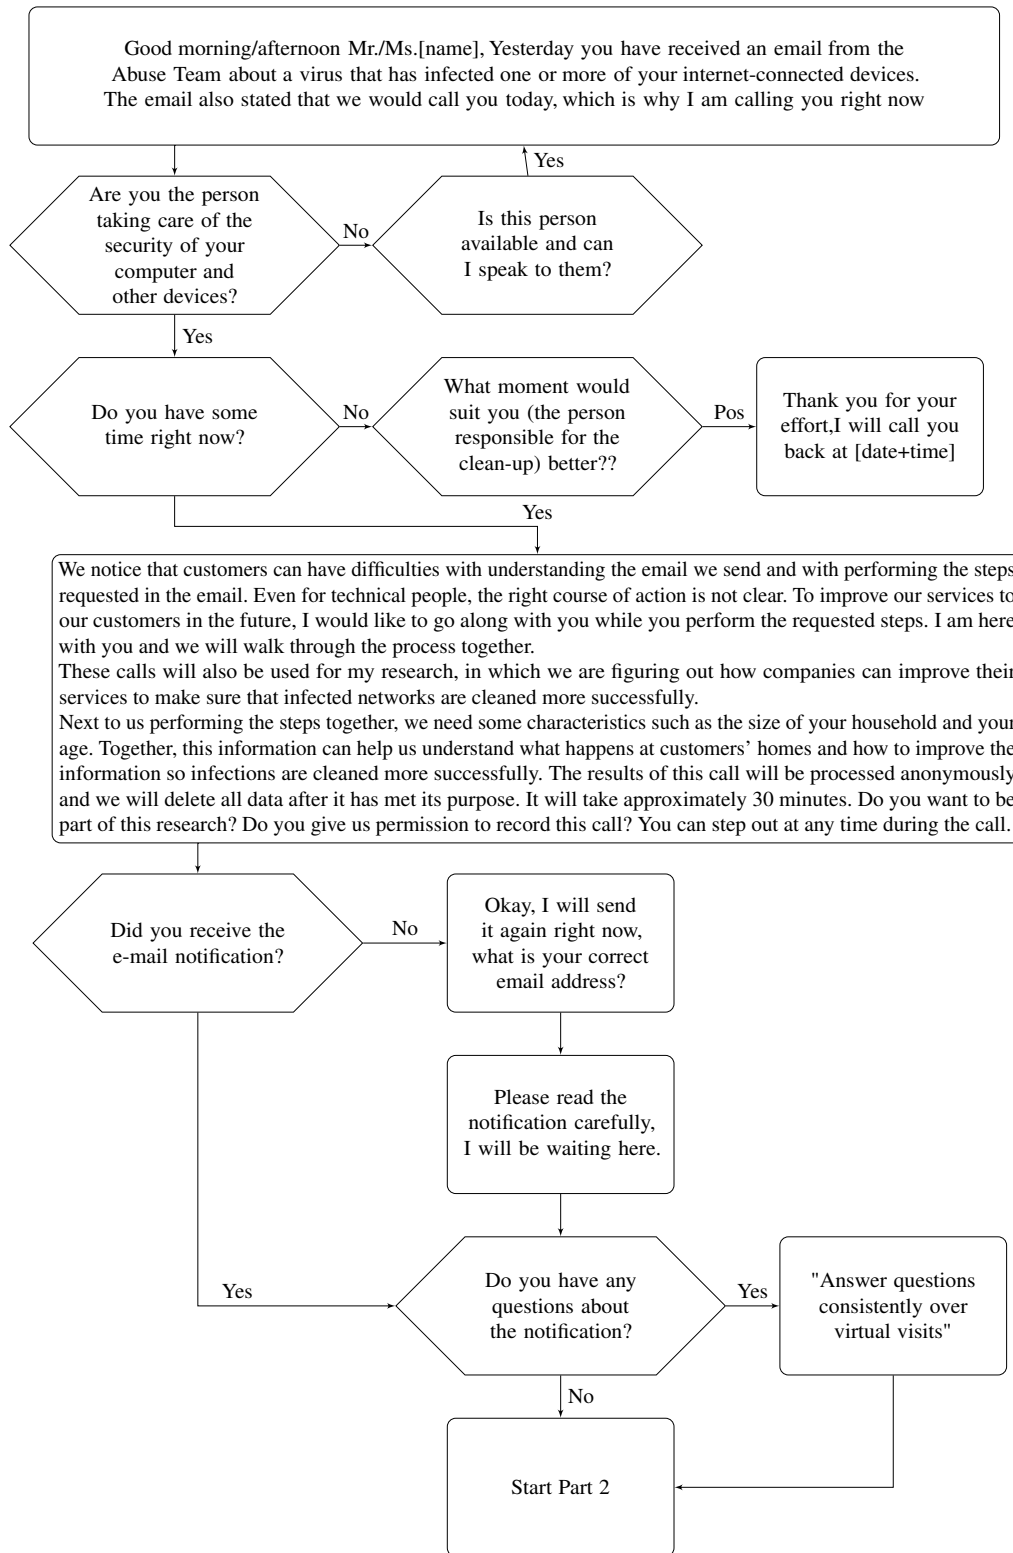


Figure 5: Think-aloud protocol - Part 1

I would like to go along with you through the steps that are described in the notification. Could you look them up? We will do this step by step, and I would like to ask you to share with me clearly what actions you are taking. The idea is that you will perform the steps as if we were not calling, except that you continuously think aloud while you take actions.

Note: At the end of each step users were told "Just tell me every thought, that goes through your mind, there are no wrong thoughts"

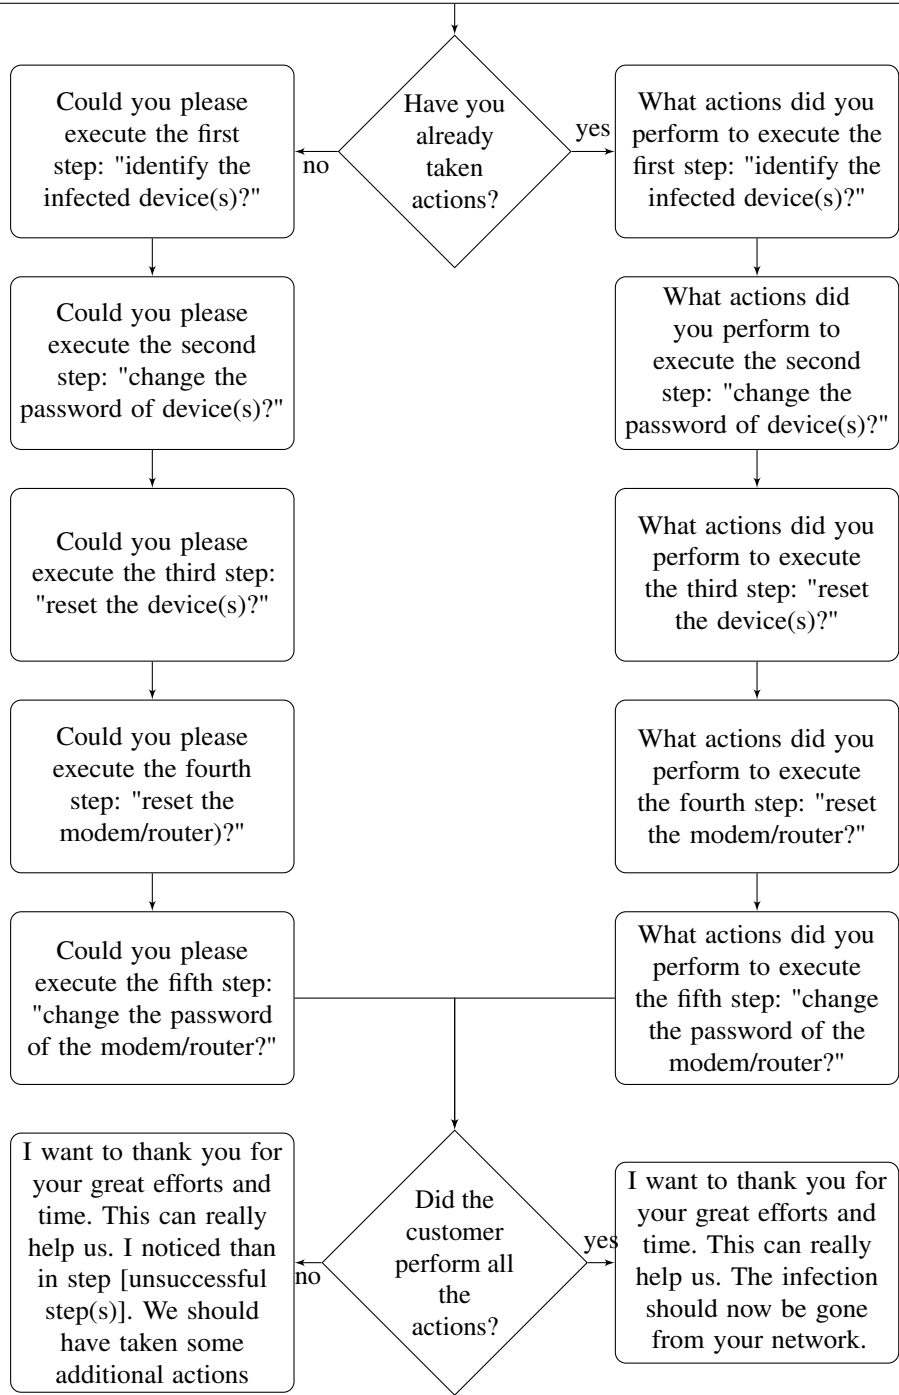


Figure 6: Think-aloud protocol - Part 2

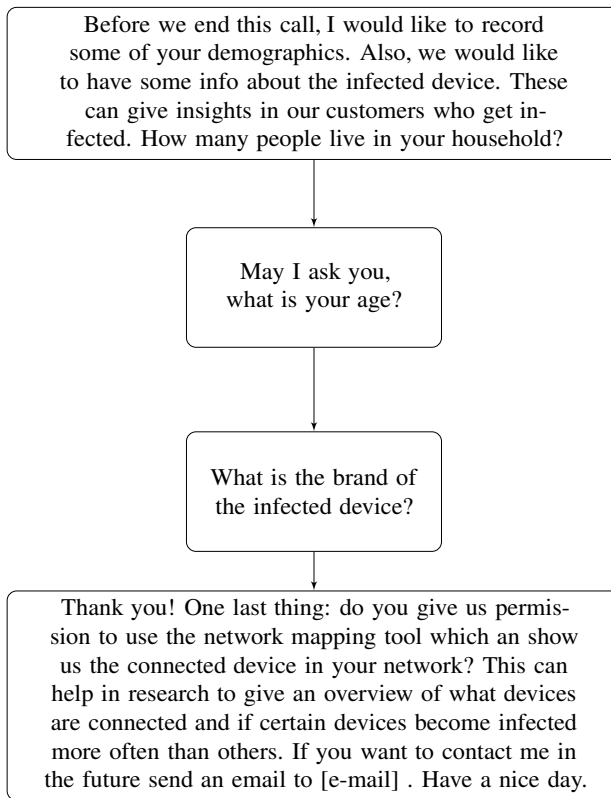


Figure 7: Think-aloud protocol - Part 3