

User compliance and remediation success after IoT malware notifications

Rodríguez, Elsa; Verstegen, Susanne; Noroozian, Arman; Inoue, Daisuke; Kasama, Takahiro; Van Eeten, Michel; Gañán, Carlos H.

DOI

[10.1093/cybsec/tyab015](https://doi.org/10.1093/cybsec/tyab015)

Publication date

2021

Document Version

Final published version

Published in

Journal of Cybersecurity

Citation (APA)

Rodríguez, E., Verstegen, S., Noroozian, A., Inoue, D., Kasama, T., Van Eeten, M., & Gañán, C. H. (2021). User compliance and remediation success after IoT malware notifications. *Journal of Cybersecurity*, 7(1), Article tyab015. <https://doi.org/10.1093/cybsec/tyab015>

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.

Research paper

User compliance and remediation success after IoT malware notifications

Elsa Rodríguez^{1,*}, Susanne Verstegen¹, Arman Noroozian¹,
Daisuke Inoue², Takahiro Kasama², Michel van Eeten¹
and Carlos H. Gañán¹

¹Organisation and Governance, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands and

²National Institute of Information and Communications Technology, 4-2-1, Nukui-Kitamachi, Koganei, Tokyo 184-8795, Japan

*Correspondence address. Organisation and Governance, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands. Tel: +31-64-54-37915; E-mail: e.r.turciosrodriguez@tudelft.nl

Received 17 September 2020; revised 17 May 2021; accepted 16 June 2021

Abstract

Internet Service Providers (ISPs) are getting involved in remediating Internet of Things (IoT) infections of end users. This endeavor runs into serious usability problems. Given that it is usually unknown what kind of device is infected, they can only provide users with very generic cleanup advice, trying to cover all device types and remediation paths. Does this advice work? To what extent do users comply with the instructions? And does more compliance lead to higher cleanup rates? This study is the first to shed light on these questions. In partnership with an ISP, we designed a randomized control experiment followed up by a user survey. We randomly assigned 177 consumers affected by malware from the Mirai family to three different groups: (i) notified via a walled garden (quarantine network), (ii) notified via email, and (iii) no immediate notification, i.e. a control group. The notification asks the user to take five steps to remediate the infection. We conducted a phone survey with 95 of these customers based on communication–human information processing theory. We model the impact of the treatment, comprehension, and motivation on the compliance rate of each customer, while controlling for differences in demographics and infected device types. We also estimate the extent to which compliance leads to successful cleanup of the infected IoT devices. While only 24% of notified users perform all five remediation steps, 92% of notified users perform at least one action. Compliance increases the probability of successful cleanup by 32%, while the presence of competing malware reduces it by 54%. We provide an empirical basis to shape ISP best practices in the fight against IoT malware.

Key words: IoT security, cleanup IoT malware, user compliance on IoT notifications

Introduction

The number of connected Internet of Things (IoT) devices will soon exceed the world's population [1]. On different continents, more than half of households already have at least one IoT device [2]. Although IoT is bringing convenience to people's lives, the devices also introduce serious security concerns. For a few years now, they have been compromised at scale and recruited into botnets: networks of malware-infected devices under the control of an attacker.

Many of the compromised IoT devices were put on the market without even the most basic security controls in place [3]. This puts the onus of protecting them on their users. Like with regular botnets, most compromised IoT device users are located in Internet Service Provider (ISP) networks [4]. RFC6561 states that ISPs should notify users and ask them to remediate the threat [5]. Researchers [6] also argued that notifying users is an important intervention to diminish the growing number of infected devices.

A core challenge for cleanup of infected IoT is designing usable mitigation advice. Remediating infections has already been proven to be difficult for PC-based malware, where users are more likely to have workable mental models as well as effective tools, most notably antivirus software and automatic update mechanisms. In the IoT space, the conditions for user action are much worse.

First of all, ISPs can typically not ascertain what exact device, or even what general device type, has been infected. Academic research also struggles with this problem. Antonakakis *et al.* [6] could only identify 31.5% of the 1.2 million infected devices and they acknowledge that their method has an unknown error rate. Other approaches rely on intrusive traffic inspection [7] or internal network scanning [2], which are technically or legally infeasible for most ISPs. The lack of visibility into the exact device type will persist for the foreseeable future. Thus, cleanup advice has to fit, by necessity, all potential device types and remediation paths. This restricts ISPs and others to recommending a generic set of steps to the users. Each individual step may or may not be applicable and may or may not be effective in remediating the actual infection at hand.

Second, the absence of accessible user interfaces makes it difficult to perform the recommended actions or apply updates—assuming such updates are even available in the first place, which is often not the case. Combined with the lack of visibility on what device type is affected, this means that the cleanup advice cannot even tell users how to access the device to implement the required steps.

Notwithstanding these challenges, we know from recent work that providing IoT malware notifications with generic cleanup steps does in fact lead to improved remediation rates [4]. It is unknown, however, what users actually did in response to the generic and hard-to-implement instructions. No prior study has measured compliance with the recommended steps.

We present the first empirical study to measure compliance directly and improve our understanding of what users do in response to IoT malware mitigation advice. Thus, our study is able to address three key research gaps: (i) We do not know to what extent users comply with IoT cleanup instructions; (ii) We do not know if notifications cause higher compliance (compared with a control); and (iii) We do not know if compliance causes higher cleanup rates. The latter issue is critical in light of the grave usability problems associated with IoT cleanup advice. We cannot simply assume that trying to follow the advice actually leads to better remediation. To establish evidence-based practices in the field of IoT security, a field with growing societal impact, we need to measure two relationships. First, to what extent does user notification lead to user compliance? And second, to what extent does user compliance lead to user remediation? Prior work could not empirically estimate these relationships, because compliance has never been measured, let alone within a randomized control trial together with notification and remediation.

This paper presents a field study on self-reported user actions following an IoT malware notification. It combines a randomized control trial involving 177 customers of a broadband ISP with a follow-up survey with 95 customers (54% response rate). We studied users' compliance with the suggested actions in the notification and how the amount of compliance affected cleanup. In sum, the contributions of this paper are as follows:

- We present the first empirical analysis of user compliance with a notification asking them to conduct generic remediation steps for infections on any type of IoT device. We find that 92% of all notified users complied with at least one of the recommended five remediation steps. Only 24% of all notified users complied with

all steps. Most users pick and choose their own path from the recommended steps. Many users also reported taking additional actions not mentioned in the notification. Even in the email-only group users comply, while they lack the incentive that quarantined users have.

- We model the impact of notifications and other predictors on user compliance and find that certain user motivations reduce compliance, while the notification comprehension did not seem to have an effect.
- We also model the impact of the amount of compliance on cleanup success. Implementing all five recommended steps increases the probability of cleanup by 32%. The notification itself has a stronger impact on cleanup than the amount of compliance. This suggests that many users chart their own course, rather than following all recommended steps. We also find evidence that the presence of competing malware in the home network reduces the probability of cleanup by 54%.
- We present insights from our survey data on how consumers would like to be approached with notifications regarding IoT infections.

Context

Our study partners with an ISP and its subsidiary in the Netherlands. One of the authors was embedded as an intern in the abuse department in order to conduct the study. The ISP has been mitigating IoT infections of the Mirai family based on the abuse data it receives. We briefly discuss Mirai and then describe the notification mechanisms of the ISP, as well as the remediation steps that the users are asked to perform.

- **Mirai malware.** Mirai emerged in 2016 and became the malware family that demonstrated the threat posed by insecure IoT [6]. Although new families have arisen [8–10], Mirai still has a dominant presence. According to Symantec [11], Mirai was the third most common IoT threat in 2018, accounting for 16% of IoT attacks. Kaspersky mentions that Mirai families were responsible for 21% of the infected devices in that year [12]. A more recent report by IBM X-Force mentions that in the first quarter of 2019, Mirai activity doubled compared with 2018 [13]. In short, Mirai is still a relevant threat and it provides a representative case study for understanding if and how end users can perform remediation.
- **Notification mechanisms.** Our partnering ISP and its subsidiary brand have slightly different user populations and their own abuse handling procedures. Consumers in the subsidiary brand are notified manually on a best-effort basis, while the ISP has an automatic procedure using abuse feeds they receive from third parties to notify consumers. Users can be notified in two ways: walled garden or email-only. (We use the term notification interchangeably with treatment. In other words, notification refers to the whole treatment that users receive.)

Walled Garden. This mechanism moves consumers into a quarantined network, also called a “walled garden,” which controls the internet access of the users. Our partners use a so-called “strict” approach, which limits all internet access except for a set of white-listed domains [5]. Users who want to access the internet get redirected to a landing page. The page tells them about the detected infection and instructs them to take five steps in order to solve the issue and restore internet access. When users are quarantined, the ISP also sends an email containing the same notification content to the user-registered contact email. Apart from notifying consumers, this process also dis-

rupts the communication between the malware command and controls and infected IoT devices.

There are three ways by which consumers can be released from the quarantine environment. First, consumers can release themselves. To achieve this, they can fill a form explaining the steps they have taken to solve the infection and then click a button to leave the walled garden (submitting an empty form also releases them). The self-release option disappears if the customer suffers two subsequent quarantine events in 30 days, to avoid people releasing themselves without taking any action. The second way is to contact the ISP's abuse department and request a release. Finally, if users did not self-release or contact the abuse department, they are automatically released after 30 days.

Email-only. The second mechanism used by our partners is to warn customers only through an email. The email provides consumers with the same notification content and set of five remediation steps. The user's internet connection remains unaffected. The main reason to use this mechanism is that the capacity of the quarantine network and the abuse department support is limited.

Examples of the notifications that consumers received are illustrated in Appendix D. The appendix first shows an example of the landing page users saw when they were notified via walled garden. The same content was also sent as the email notification to consumers in this treatment group. For the email-only group, the appendix shows an example notification sent to consumers in this second group. The email-only notification essentially contains the same content except that it omits statements about consumers having been placed in a quarantine environment.

- **Remediation steps.** Both notification mechanisms, walled garden and email-only, ask the user to comply with the instruction to undertake five generic remediation steps that aim to cover as many remediation paths as possible (Fig. 1; also see Appendix D). Step 1 is to identify the smart device(s) connected to the home network. The explanation mentions that likely candidates for the infected device(s) are IP cameras, digital video recorders (DVRs), or similar devices, not personal computers, laptops, or tablets. Step 2 is to change the password of the smart device(s). Step 3 is to restart the device(s). Since Mirai malware is not persistent, step 2 and step 3 will wipe the malware from the infected device(s) and prevent immediate reinfection via the abuse of factory-default credentials. Step 4 is to reset the modem or router to factory settings. This removes port forwarding that exposes IoT devices to the public internet, as well as a possible infection of the router itself. Finally, step 5 is to change the password of the modem or router. These five steps are generic enough to deal with the fact that the ISP cannot reliably identify what generic device type is infected, let alone the exact model number. Therefore, the ISP cannot explain to the user how to exactly change the password or install an update—or even whether such an update is actually available for their device. Also, these steps are seen as the steps that are most likely to help, but the ISP cannot know whether they are in fact effective for each notified user.

Related Work

As pointed out by Cetin *et al.* [4], infected IoT devices are located in broadband networks managed by ISPs. Our research is motivated by the role that ISPs can take in notifying and warning users infected with IoT malware. Nevertheless, the warnings and notifications only work if users are able to comply and this compliance actually results in remediating the infection. We first look at the literature related

to botnet mitigation by ISPs. Next, we look at work on abuse and notification and security warnings. Finally, we discuss relevant work on security behavior of users.

- **Botnet mitigation by ISPs.** Security literature highlights ISPs as a critical control point against botnets [14], and it highlights that ISPs can make a difference. Although ISPs are a critical actor to fight botnets, Asghari, Ciere and Van Eeten [15] also looked at the impact of anti-botnet initiatives on the cleanup success of botnets and concluded they have no impact. Nevertheless, they conclude that anti-botnet initiatives need to engage ISPs in taking action. Also, Pijpker and Vranken [16] developed a model with measures that ISPs can implement to fight botnets. They found that ISPs mainly focus on prevention and notification. In addition, RFC6561 recommends best practices that can be implemented by ISPs to notify users, so they can remediate botnets [5]. Moreover, the recent RFC8520 [17] proposes to whitelist IoT traffic through a manufacturer usage description (MUD). There are discussions on how MUDs can help ISPs pinpoint abuse when devices show a different behavior than that specified in the MUD [18].

The literature has highlighted that ISPs are a relevant actor to fight botnets. Nevertheless, little attention has been paid to so far to understand whether notifying or warning users about infected IoT devices leads to user compliance with the remediation advice and whether this, in turn, leads to successful mitigation of IoT infected devices.

- **Abuse and vulnerability notifications.** There is a large body of work on the effectiveness of abuse and vulnerability notifications by measuring the presence or absence of the security issue, without actually observing the user's behavior. Vasek and Moore [19] studied the effect of notification content and found that verbose notifications caused more remediation of compromised websites than brief notifications. Li *et al.* [20] studied notification content and mechanisms in terms of webmasters cleaning up compromised servers. They found that direct communication with the webmaster substantially increased the likelihood of cleanup. Cetin *et al.* [21] studied the effect of the reputation of the notification sender and found that a better reputation did not improve cleanup rates. Durumeric *et al.* [22] sent notifications for servers vulnerable for Heartbleed and found a positive impact in patching. In addition, Li *et al.* [23] notified thousands of different network operators about security issues in their networks finding that notification has a positive impact on remediation. In contrast, Stock *et al.* [24] and Cetin *et al.* [25] sent notifications to thousands of domains with vulnerabilities and found very low remediation rates. The experiment highlighted the shortcomings of email notifications and the gap between awareness of the problem and actually taking action.

Our work directly relies on two notifications mechanisms (email and walled garden) to inform infected users about the Mirai infection. Although the related work has shown that such notifications sometimes work, in certain cases low remediation rates are observed [24, 25]. A few studies looked specifically at walled garden mechanisms. Cetin *et al.* [26] found high remediation rates for Windows-based malware cleanup. The only research on end-user remediation of IoT compromise, closest to our current study, found that a walled garden was also effective in cleaning up IoT malware [4]. However, these studies treat the user behavior that translates notification into remediation as a black box. No prior work has tried to observe what users actually do with the cleanup instructions, nor whether better




Figure 1: Recommended steps in the notification mechanisms.

compliance actually results in better cleanup. Our work wants to contribute to this literature by performing a real-world experiment with users who have been notified.

- **Security warnings.** Notifications are also related to the work on security warnings. Previously, Egelman *et al.* [27] studied how tolerant individuals were with delays in their activities when they are informed that they were due to security purposes; it was found that users were likely to not wait when they were not properly informed that the delay was due to security purposes. Also, Egelman and Schechter [28] studied web browser warnings on phishing websites manipulating the background and color of the warning to observe if users obeyed the warning. They found that text and color did not have an effect on users actually following the warning. Felt *et al.* [29] undertook the task of designing a new SSL warning, so that they were not disregarded by users. Moreover, Akhawe and Felt [30] assessed if security warnings were effective for malware and phishing websites, and they demonstrated its effectiveness in practice. Moreover, Krol, Moroz and Sasse [31] looked at how users reacted to PDF download warnings and showed that these might get ignored by the user because of the exposure to false positives, incorrect mental models, or not understanding that PDFs can also contain viruses. Bravo-Lillo *et al.* [32] used mental models to understand how advanced and novice computer users responded to computer warnings. They reported that the groups differ in terms of how they perceive the risk they might be facing.

These studies show that user responses to notifications and warnings are highly variable and we do not yet know all the factors at play. In our study design, we incorporate factors from communication–human information processing (C–HIP) [33] theory to test if they help explaining user compliance and cleanup success (see more details in the “Methodology” section).

- **User security behavior.** Notification mechanisms rely on consumer behavior to be effective. Fagan and Khan studied users’ motivation to follow security advice; they found that individual concern of following advice is rated higher than how this can affect other [34]. Redmiles [35] looked up the immediate response of Facebook users who receive warnings about suspicious login incidents defining the common process of users to respond to the incident as consisting of incident awareness, mental model generation, and behavioral response. Vaniea, Rader and Wash [36] studied how users’ negative experiences of software updates impacted their willingness to update software.

As the literature expresses user’s security practices might be based on wrong mental models, yet we need to rely on the fact that users need to learn on how to react on notifications, especially in the area of IoT devices, which present very different challenges, e.g. because of lack of a web interface on devices.

The related work has shown that notifications might work, but that their effectiveness is highly variable. The work on warnings underlines that users might ignore them. Behavioral research, moreover, has highlighted the gap between awareness and actual behav-

ior. To the best of our knowledge no prior study, including [4], has measured compliance with IoT cleanup instructions send in a notification. As we describe in the introduction, IoT cleanup advice has huge usability problems. So, we cannot assume that following the advice actually leads to better remediation. Also, our study differs from this prior work on abuse notifications by providing the first study that opens up the black box of user behavior, most notably compliance, after receiving a notification in the area of IoT malware.

Methodology

Our data collection was carried out between May and June 2019. To answer our research questions, we combined a randomized control experiment with a survey among participants. We first randomly assigned 177 Mirai-infected customers of our partners to one of the treatments (walled garden or email-only) or to the control group. We then conducted a short phone survey based on C–HIP theory [33]. Of the 177 participants, 95 were reachable via phone within three attempts and accepted to respond the survey. Finally, we tracked the infections of these customers during the experiment and for two additional months, to see if the infected devices were successfully cleaned.

- **Sampling and random assignment.** Our partner ISP receives a daily feed from Shadowserver containing IP addresses of Mirai-infected users in its network and that of its subsidiary brand. In collaboration with both, we used additional infection data by identifying scans that matched the Mirai fingerprint (as described by Antonakakis *et al.* [6]) in a /15 network telescope. All identified infected users were randomly assigned to a treatment or the control group. The latter received a notification delayed by 2 weeks, so as to have a baseline against which to measure the impact of either notification mechanism.

Consumers detected as having infected devices during the weekends were not included in the random assignment to the treatments. This decision was made because the abuse department of the ISP does not work during weekends. So, if users needed immediate support after receiving a notification in the weekend, it would not have been possible to respond to their inquiries.

We also excluded users who had been notified about an IoT infection prior to our study. Their behavior might be different from users who were notified for the first time due to previous exposure to the remediation process. Only nine users were excluded here.

In total, the sample consisted of 177 customers. Of these, 128 have a contract with the ISP and 49 with the subsidiary. Our design was to randomly assign customers to three equal groups: walled garden, email-only and control. However, during the experiment we discovered that there was a malfunction with the mail server at the ISP.

Consequently, users assigned to the email-only group did not get the intended email notification at the ISP. This meant that 43 users in the originally intended email-only group had to instead be assigned

Table 1: Overview of group assignments and survey respondents

| Group | | Control | Email | Walled garden | Total |
|---------------------------|--------------------|----------|----------|---------------|----------|
| Internet Service Provider | Participants | 85 | 0 | 43 | 128 |
| | Survey respondents | 35 (41%) | 0 | 28 (65%) | 63 (49%) |
| Subsidiary | Participants | 17 | 16 | 16 | 49 |
| | Survey respondents | 10 (59%) | 11 (68%) | 11 (65%) | 32 (65%) |
| Total | Participants | 102 | 16 | 59 | 177 |
| | Survey respondents | 45 (44%) | 11 (68%) | 39 (66%) | 95 (54%) |

to the control group. At the subsidiary however, email notifications functioned properly.

This company is smaller however, as is the number of infected users, so the email-only group consisted of 16 users. All in all, this meant that our study had a larger control group than originally intended and an email-only group that was too small to allow for strong statistical inference about its differences with the other groups. We retain the group in our analysis however for qualitative comparisons.

Table 1 provides an overview of the overall group assignments. It also reports the portion of each group that responded to the survey. We had high response rates in all groups.

- **Survey framework.** We used C–HIP theory as a basis to develop our survey. To maximize the response rate, we limited the survey to require only around 10 min to complete. This was tested during 17 pilot interviews, which are not included in the final sample on which this study is based.

C–HIP was originally proposed as a stage model for information processing, allowing for feedback loops among stages, in which an entity tries to communicate a message to change the behavior of the receiver [33]. In our case, the ISP and its subsidiary brand are the sources of the notification which are trying to get their customers to comply with the recommended cleanup steps.¹

We chose the C–HIP theory because it includes the source of the notification. Different sources can have different consequences on how users react. In our case, the email and walled garden seemed likely to be received quite differently. Since we had to make a trade-off between maximizing responses and the length of the survey, we study only the comprehension and motivation of the users to understand their behavior, compliance. The model includes attention, comprehension, beliefs and attitudes, and motivation. Due to the real-life settings, we could not measure the attraction that the notification caused to the users when they received it. We only notified users who were not previously notified, so this reduced the familiarity that users had regarding doing the steps and they did not have an accumulation of knowledge about the tasks. Hence, we did not measure users' attitudes and beliefs either. We cannot assume that all users comprehend the notifications, since the notifications reach users of different backgrounds with different abilities and experiences. So we have to check first whether users understand what they were asked to do. If users do not understand the notifications, they cannot correctly act upon it. In addition, motivation is key because it can activate people to comply with any directive [33]. The cost of compliance should be lower than the benefits that the users perceive by taking the recommended steps.

¹ According to later versions of the model [37], the message needs to create an attention switch and attention maintenance in its receiver. This stage was not included in our adapted theoretical framework, since due to the real-life setting of the experiment, we were not able to measure it. Nevertheless, the notification method can trigger the users' attention.

In our theoretical framework, we also included the type of devices and demographics to control for other variables that could influence behavior that might not be related to comprehension and motivation. For instance, if the device the users' own has a web interface this could influence how easily the user can change the password of the device versus when the device does not have a web interface. Demographics can also play a role in compliance since research has shown that characteristics such as gender and age can influence technology acceptance [38], and thus how users could handle IoT devices. Hence, we want to control for these variables. This model covers two important aspects of the related work: (i) the role of the ISPs as intermediaries and how the different types of notifications can influence compliance and cleanup in the IoT domain; and (ii) drivers of user behavior, in this case comprehension and motivation, to understand the degree of compliance.

Due to the structure of sequential stages, C–HIP can be an easy tool to pinpoint where an end user drops out of the process of compliance. Each stage can be a potential bottleneck to comply. An interesting notion within the C–HIP model is that notification effectiveness can also be measured based on other stages, in this case comprehension and motivation, than the binary distinction between compliance and noncompliance.

Our survey addressed the following:

Comprehension of the notification. Customers were asked if they recalled receiving the notification and if they read them. Also, they were asked if they understood the notification. The answers related to reading and understanding the notification were coded as dummy variables to measure comprehension of the notifications.

Motivation of users. The notification must motivate customers to perform the desired behavior, in this case, to comply with the five recommended steps. Customers were asked about their motivations to comply, or not, with the notification, and their replies coded as categorical variables. After the pilot survey, these categories were refined based on the most common responses.

Compliance and additional steps. Customers were asked which steps, if any, they followed to resolve the IoT infection. The question was open-ended, so that they could tell us what they remembered doing, rather than prompting their answers by mentioning the recommended steps. We precoded the five steps from the notification. Steps were a binary variable that took a value of 1 if the consumer performed any of the five recommended steps or a value of 0 if the consumer did not perform a step. We define compliance as the number of steps that consumers took to clean their IoT devices, out of the five recommended steps. As such compliance is expressed as a ratio from 0/5 steps to 5/5 steps.

Sometimes customers would mention taking additional steps, that were not mentioned in the notification. This was registered as unstructured text describing an additional step. After the survey was concluded we coded these answers in several recurring additional actions.




Figure 2: Experimental setup drawing on adaptation of C-HIP model described in [33].

Figure 2 depicts the adapted theoretical framework used as a guide to study consumers' compliance and cleanup. We also wanted to control for demographic differences and for different device types when measuring compliance and cleanup. In the survey, we also asked the type of IoT device the user thought was infected, so as to control for how device type might impact compliance and cleanup success. Furthermore, we included demographic characteristics of the customers, recorded in the ISP customer data, to control for differences among users. Finally, we measured cleanup success independently, to see if the self-reported compliance is predictive of remediation.

Device type. Survey respondents were asked if they could identify the infected IoT device(s). Answers might be influenced by speculation or incorrect mental models. While we have no ground truth to compare these answers against, we did lookups for the customer IP addresses in Shodan [39]. Shodan is an IoT search engine, and it indexes IoT devices which means these devices are exposed to the open internet, and compared the results with the answers.

Age and gender. We used the data of the respondent as recorded by the ISP. When we reached someone at the listed phone number, we asked them if the subscription was in their name. In some cases, the respondent reported that they were small businesses. We coded them as such. Six users reported that someone else did the steps for them and so we coded based on their description.

Attitudes and beliefs. We did not address these because of time constraints in the phone call, although we did try to minimize the difference in customers' beliefs and attitudes by not including users who had been notified before.

- **Survey process.** The survey was developed and tested in 17 pilot interviews. The pilot survey was carried out also with real consumers to check if they understood the questions, how long the survey could take, and to refine some potential answers for the open questions. Feedback to improve the protocol was obtained

and incorporated in the final design of the survey. The data of the pilot survey were not included in our results. Finally, the questions were adapted slightly, depending on if the consumer was in the control group or the treatment group. The questionnaires are included in Appendix C.

We conducted the survey 2 weeks after the notification was sent. For the control group, since they did not receive a notification, the 2 weeks were counted from the first day of their detection as infected. The survey call was the first notification the control group received, and for users in the control group that we could not reach by phone, we sent an email. These users, of course, were not included in the survey study, and they were not included in the measurement of the remediation rate of the control group. We set the time to contact all participants to 2 weeks because we want to obtain as much reliable information as possible regarding what actions a consumer took, while also giving the user time to conduct the remediation steps without being prompted to do so by the survey request. To ensure that the protocol would be consistently carried out, one person did the survey.

Survey respondents were explicitly reminded of the right to opt out from the survey. One respondent chose to opt out. The survey respondents did not receive any incentive to participate in the survey. Out of 177 calls we placed, 95 respondents accepted to respond to the survey, one person opted out and 81 customers could not be reached.

Because of privacy concerns, the ISP did not allow us to record the phone survey. A script was developed to log the answers of the survey respondents. For the closed questions, the possible answers were already precoded. For the open questions, we added potential answers that had been given during the pilot survey and had the investigator enter manually any additional information given by the respondent.

Email logs from the abuse department were used to check if consumers contacted them for additional information. Moreover, the quarantine forms that users filled out in order to leave the walled garden were used to check if they were reporting the same device types as mentioned during the response to the survey. We used this information to validate our results.

- **Cleanup and competing malware.** We collected data during the experiment and for 2 additional months (July and August 2019) to see whether the infection was successfully removed after the experiment. We monitored the Shadowserver abuse feeds received by the ISP [40], the Global Cyber Alliance IoT honeypot data [41], IoTPot data [42], and also a network telescope of 300 K IP addresses.

We coded the infection as cleaned when the user's IP address was absent from the abuse reports, honeypot logs and not scanning the network telescope, either with the Mirai fingerprint [6] or without it. We included the latter to measure cleanup conservatively. It suggests there is still an infection on the device(s) in the home network, since we would not expect a normal subscriber to aggressively scan large network blocks. We coded these cases as "no cleanup." This analysis revealed a surprise where sometimes we found both scanning patterns for the same customer. This pattern might reveal the presence of competing malware in the home network. It has been well documented that various IoT malware families actively compete with each other for control over devices [43]. To take this factor into account, we created a dummy variable called "competing malware" to capture when we saw other scanning patterns than Mirai for the same customer. To reiterate: all scanning patterns were coded as "no cleanup."

Table 2: Summary of findings

| | No steps | One or more steps | Odds |
|--|-----------------------|-----------------------------|-------------|
| Control | 33 | 12 | 0.36 |
| Notifications (email-only and walled garden) | 4 | 46 | 11.5 |
| | Still infected | Successfully cleaned | Odds |
| Control | 22 | 23 | 1.04 |
| Notifications (email-only and walled garden) | 7 | 43 | 6.14 |

- One or multiple devices infected.** We were aware that the Mirai infection could be present on just one device, but also on multiple devices in the home network. Also, the “competing malware” that we observed could have been present on the same device as the Mirai infection (but at a different time) or on another device. Since neither we nor the ISP could know if one or more devices are infected, the notification was designed to handle both scenarios. It told users that one or multiple devices could be infected with Mirai. In terms of observing cleanup success, we cannot differentiate partial cleanup from no cleanup, i.e. one device was actually remediated, but another device is still infected. As long as we observed any malware scanning behavior coming from the customer IP address, we coded that case as “not clean” in order to have a conservative measurement of the remediation rate. In sum, while we lack visibility into the number of infected devices in the customer home network, we designed both the notifications as well as the measurement of cleanup to handle both scenarios.

Ethical Considerations

Our study follows the ethical principles set forth within the Menlo Report [44], namely that of respect for persons, respect for the law, justice, and beneficence. We additionally followed legal guidelines and policies set forth by our partner ISP regarding the study and the collection of empirical data to understand consumer behavior with respect to IoT malware cleanup.

In light of the first two ethical principles (respect for persons and law), we operated within the privacy policies of our partner ISP. One of the researchers was embedded as an intern and processed the customer data on the ISP premise.

The survey was also conducted by the intern from within the ISP. Consumer contact details were looked up every time prior to each phone interview and are not part of our collected study data. All respondents were first asked for their consent to respond the survey and for the survey data to be anonymously used for the purpose of this study. The possibility to opt out of the survey was explicitly mentioned. Only one person declined to participate in the survey. (The rest of the nonresponse was caused by not being able to reach the respondent.)

In terms of the latter two ethical principles (justice and beneficence), we believe that our study does not create harm and it treats individuals fairly. Our study follows a randomized control trial design (see more details in the “Methodology” section). All ISP subscribers affected by Mirai-like malware were notified of the infection. The notification for the subscribers in the control group was delayed by 14 days. Since Mirai attacks first and foremost target third parties, not the owners of the infected devices, this delay is unlikely to expose the subscriber to substantial harm. We evaluate the downsides of this delay to be outweighed by the fact that our study aims to improve the

mechanisms for users, and society at large, to combat IoT malware and prevent attacks to third parties in the future.

Findings

To reiterate, our question is: to what extent do users comply with the instructions? And does more compliance lead to higher cleanup rates? We will model both relationships in light of the factors discussed in our adapted theoretical framework (see the “Methodology” section).

Table 2 summarizes the findings; the notifications seem to be extraordinarily effective. We calculate the odds of customers who received the notification and customers in the control group. Then we look for the odds ratio of doing one or more steps and remediation. We can observe that notified customers had 31.9 times the odds ratio of doing more than one step than customers who were not notified. Also, we can observe that customers notified had 5.9 times the odds ratio of successfully cleaning their infected device.

Before turning to the explanatory models, we will discuss these factors more descriptively.

Age and gender

To check for potential bias in the sample of participants who were reached for a survey, we compare the age and gender of survey respondents against the other participants. Table 3 shows the distributions. The groups are very similar across treatment conditions and demographics. Except for a bit lower proportion of female customers among the survey respondents in the control group, we see no evidence for potential bias.

Overall, the age of customers with an infected device ranges from 25 to 87 years old, with a median age of 47.5. As explained in the “Methodology” section, when participants were reached for a survey, we asked them if the subscription was in their name. In seven cases, the survey respondents indicated it was actually owned by a small business. We coded these users separately.


We also compared the age of the Mirai-infected customers versus the total subscriber population of the ISP and the subsidiary brand. We find a right-skewed distribution for the infected customers compared with the distribution of all subscribers (Fig. 3). The mean age of Mirai-infected consumers is 6 years younger ($\mu = 48$) than the mean age in the total subscriber population of the ISP and the subsidiary ($\mu = 54$). Welch’s unequal variance *t*-test estimates this difference to be significant ($P < 0.0001$). In short, Mirai-infected consumers are relatively young. This fits with the speculation that younger consumers are more likely to buy IoT devices.

Device type

We asked survey respondents which of their devices they thought was infected. Table 4 shows the type of devices consumers reported as the offender. It is possible to notice that surveillance cameras make up a

Table 3: Study participant demographics

| Group | | Control | | Email | | Walled garden | |
|-------------------|--------|-----------|------------|-----------|----------|---------------|----------|
| | | Yes | No | Yes | No | Yes | No |
| Survey respondent | | | | | | | |
| Age | Range | 29–76 | 25–77 | 30–69 | 26–67 | 26–83 | 29–87 |
| | Median | 47 | 46 | 47 | 45 | 46 | 46 |
| Gender | F | 7 (15.5%) | 14 (24.5%) | 0 (0%) | 0 (0%) | 2 (5%) | 2 (10%) |
| | M | 31 (69%) | 38 (67%) | 11 (100%) | 5 (100%) | 32 (82%) | 14 (70%) |
| | N/A | 2 (4.5%) | 2 (3.5%) | 0 (0%) | 0 (0%) | 3 (8%) | 3 (15%) |
| Business | | 5 (11%) | 3 (5%) | 0 (0%) | 0 (0%) | 2 (5%) | 1 (5%) |

**Figure 3:** Age distribution—infected consumers vs all subscribers.**Table 4:** Infected IoT devices

| Device type | No. of consumers |
|---------------------------|------------------|
| Surveillance camera | 36 (37.8%) |
| Raspberry Pi | 33 (35%) |
| NAS | 9 (9.5%) |
| Unknown device | 8 (8.4%) |
| DVR | 2 (2.10%) |
| Router | 2 (2.10%) |
| Printer | 2 (2.10%) |
| Linux embedded system | 1 (1%) |
| Smart meter | 1 (1%) |
| Power consumption monitor | 1 (1%) |

large portion of devices (36; 37.8%). This is consistent with prior studies [4, 6]. Next, 33 users mentioned a Raspberry Pi (35%). This is different from previous research. The high percentage can be understood by the fact that during our experiment, a new Mirai-based attack vector emerged targeting a known vulnerability in Domoticz software [45]. Domoticz is an open-source software that can manage home automation systems. It is often run on a Raspberry Pi. The Mirai variant exploited an “unauthenticated remote command execution” vulnerability, which allowed the malware to bypass the authentication mechanism of the devices. This was detected in April 2019. Although a new version of the software was released on 9 May 2019 [45], users reported a peak of infected IoT devices with this variant of Mirai during the study.

Nine users (9.5%) reported a network attached storage device (NAS) as the culprit, which is again consistent with other studies. Next, we find a list of devices such as DVRs, routers, printers, Linux embedded systems, smart meters, and power consumption monitors.

Surprisingly, only a small portion of the survey respondents (8; 8.4%) felt unable to identify the offending device. This could mean that most users have a pretty good understanding of their computing environment or it could mean that users are overconfident in their expertise. For example, one survey participant mentioned the “smart meter” as the compromised device. The Dutch smart meters are locked-down devices that have been rolled out and maintained by the distribution grid operators. So far, there is no known attack against these devices. Some of the answers from the survey respondents might be triggered by socially desirable behavior, as they might want to convince the investigator that they are technically savvy.

We have no ground truth against which to test the accuracy of the answers. We did conduct two crosschecks, however. First, we compared the survey answers against the submitted user forms from the walled garden. We found no inconsistencies. Second, we looked up the IP addresses of the infected IoT devices in Shodan [39]. For 36 of the 95 survey respondents (38%), we found a device listed in Shodan. Interestingly enough, 35 of these 36 (97%) survey respondents had reported the same device during the survey as was observed by Shodan. While this is hardly conclusive evidence, it does give credence to the idea that users have honestly answered our question and that they have at least a plausible speculation about the offending device. The fact that Shodan can observe it means it is exposed to the open internet, which implies a high level of risk for poorly secured devices.

Comprehension

In the survey, we asked whether participants received, read and understood the notification. In the walled garden group, 37 out of the 39 users (95%) remember receiving and reading the notification either via the landing page or the corresponding email. However, only 25 (67.5%) indicated they understood the notification. Interestingly, all users who acknowledged that they did not understand the message had emailed the ISP’s abuse department. In other words, even

Table 5: Customer motivations to comply with notifications

| Treatment | Motivation | No. of consumers |
|---------------|--|------------------|
| Email-only | Safe internet is important | 7 (78%) |
| | Malfunctioning device | 1 (11%) |
| | No answer | 1 (11%) |
| Walled garden | Internet back | 19 (51%) |
| | Internet back and safe internet is important | 9 (24%) |
| | Safe internet is important | 3 (8%) |
| | No answer | 3 (8%) |
| | Malfunctioning device | 1 (3%) |
| | Need the device | 1 (3%) |
| | Privacy concern and safe internet | 1 (3%) |

though they did not understand the notification, they all took action to find out how to solve the problem. For example, they asked for more technical information or they stated that they did not understand the cause of the infection. Of the 25 people who did claim to understand the message, 22 also emailed the ISP. Their messages were typically stating the actions they took and then asking for confirmation whether that was enough to solve the problem.

While the email-only group was too small to make robust statements (see the “Methodology” section), it is worth noting that 9 of the 11 (82%) acknowledged receiving and reading the notification. Of these, eight declared that they understood the notification. Again, the one person who did not understand emailed the ISP’s abuse department. The consumer was asking for more technical details.

In total, these results indicate that for 46 out of the 50 notified users (92%), the message was successfully delivered and read. Those recipients who did not understand the message, contacted the ISP and asked for further details and advice. Even among people who said they did understand the message, the majority contacted the abuse department to state the actions they took.

Motivation

We asked users an open question regarding what drove them to comply with the recommended steps.

We found some recurrent topics in the answers to this question. Table 5 presents an overview.

In the walled garden group, 19 users (51%) said that they were driven by the fact that they did not have an internet connection. Nine users (24%) mentioned not only the lack of an internet connection, but also that safe internet is important.

In the email-only group, no one loses their internet connection, which shifts the answers more toward more intrinsic motivations to improve security. Seven consumers in the email-only group (78%) expressed that they complied because a safe internet is important. One consumer said that a malfunctioning device was the motivation.

Similar to Fagan and Khan [34], of all notified customers only 11 (22%) expressed some social motivation to comply. Hence, it is clear that most users were thinking about how the infection affects themselves rather than others. The email-only group differs in this respect. While it is too small to draw firm conclusions, it does hint at the possibility that security practices in the IoT domain would benefit from relying on the users’ social considerations regarding how infections could affect others.

Table 6: Participants’ self-reported compliance (1) or not (0) with each step in the notification (listed in Fig. 1)

| Group | Followed steps | | | | | Freq. |
|---------------|----------------|---|---|---|---|-------|
| | 1 | 2 | 3 | 4 | 5 | |
| Walled garden | 0 | 0 | 0 | 0 | 0 | 2 |
| | 1 | 0 | 0 | 0 | 0 | 9 |
| | 1 | 0 | 0 | 1 | 0 | 1 |
| | 1 | 0 | 0 | 1 | 1 | 4 |
| | 1 | 0 | 1 | 0 | 0 | 1 |
| | 1 | 0 | 1 | 1 | 0 | 3 |
| | 1 | 0 | 1 | 1 | 1 | 1 |
| | 1 | 1 | 0 | 0 | 0 | 2 |
| | 1 | 1 | 0 | 1 | 1 | 1 |
| | 1 | 1 | 1 | 0 | 0 | 3 |
| | 1 | 1 | 1 | 0 | 1 | 1 |
| | 1 | 1 | 1 | 1 | 0 | 2 |
| | 1 | 1 | 1 | 1 | 1 | 9 |
| | 1 | 1 | 1 | 1 | 1 | 2 |
| Email | 0 | 0 | 0 | 0 | 0 | 2 |
| | 1 | 0 | 0 | 0 | 0 | 1 |
| | 1 | 0 | 0 | 1 | 0 | 1 |
| | 1 | 0 | 0 | 1 | 1 | 1 |
| | 1 | 1 | 1 | 0 | 0 | 1 |
| Control | 0 | 0 | 0 | 0 | 0 | 33 |
| | 1 | 0 | 0 | 0 | 0 | 10 |
| | 1 | 1 | 0 | 0 | 0 | 2 |
| | 1 | 1 | 0 | 0 | 0 | 2 |

Compliance

We asked the participants an open-ended question about compliance and then coded the answers in terms of which of the recommended steps were mentioned. We also recorded when users mentioned other steps than those recommended in the notification.

Table 6 displays the results for the recommended steps. Each row is one pattern of steps complied with, or not. The end of each row contains the number of users who reported this pattern. Of the 50 users who were notified and accepted to respond to the survey, 12 notified users (24%) fully complied with all five steps (9 in the walled garden group and 3 in the email-only group). At the other extreme, four people in the treatment groups reported taking none of the recommended actions (two in the walled garden group and two in the email group). Taking no action whatsoever was, for obvious reasons, the dominant pattern in the control group, since they had not been notified of the problem. We will discuss this group later.

The overwhelming majority (92%) of participants in the treatment groups reported taking at least one of the recommended steps (95% of the walled garden group and 82% in the email-only group); 10% took two steps in the walled garden group and 9% in the email-only group; 26% took three steps in the walled garden group and 18% in the email-only group; and 13% took four steps in the walled garden group and 18% in the email-only group. All the steps were taken in various combinations.

Even in the control group, we found that some users also reported having taken certain steps in the 2 weeks before, even though they had not been informed about the infection. Some users with Domoticz devices identified that their device had a security update, which they applied. In total, ten users (22%) followed step 1. Of course, as they had not received any notification, this means that even identifying the device was a step they followed without complying with a notification. We did code it as a compliance step, to capture the degree in which users to take security actions for other reasons. Two users (4.4%) followed step 1 and step 2. In the conversation with these consumers, we learned that they were prompted to take action either because of the malfunctioning of their devices or the type of device they owned.

Some combinations of steps occurred more often than others. We used Spearman's rank correlation to measure the strength and direction of the association among the steps. We observed that there is a high correlation ($r_s = 0.74$, $P < 0.001$) between step 2 (change the password of the device) and step 3 (restart the device). Similarly, there is a correlation ($r_s = 0.73$, $P < 0.001$) between step 4 and step 5: reset the modem to factory settings and change the password of the modem. This might indicate that although not all consumers did the five steps, there is some pattern to how they proceeded to mitigate the infection. Steps 2 and 3 are focused on the compromised device, while 4 and 5 are more oriented at preventing new infections. Some users focus on one, rather than the other. In Appendix A, the complete correlation table is presented.

We also looked at what other actions people reported, beyond the five steps. Table 7 summarizes the extra steps that users mentioned. As with compliance, we also include the actions taken in the control group. Interestingly, 25 (64%) of the consumers who were in the walled garden did extra steps versus 4 (44%) of the consumers in the email-only group.

Even among the users who had fully complied with the notification, some reported taking extra steps. One user, for example, described doing a software update. Among users who did not do all the steps, we found that they did report taking other actions to resolve the issue. For example, one customer reported identifying the device and also doing a software update. Other customers reported more drastic actions. After identifying the offending device, they discon-

Table 7: Additional steps consumers performed

| Treatment | Additional steps | No. of consumers |
|---------------|----------------------------------|------------------|
| Email-only | Only followed notification steps | 5 (55.5%) |
| | Disconnected device | 2 (22.5%) |
| | Software update | 1 (11%) |
| | Disable port forwarding | 1 (11%) |
| | Ask for help | 2 (5.5%) |
| Walled garden | Only followed notification steps | 12 (31%) |
| | Disconnect device | 9 (24%) |
| | Stop using the device | 6 (16%) |
| | Software update | 5 (13.5%) |
| | Disable port forwarding | 3 (8%) |
| Control group | Software update | 8 (18%) |
| | Stop use | 2 (4.4%) |
| | Disconnected device | 1 (2%) |

ected it or stopped using it altogether. One person even mentioning that he had brought the device to the recycling center.

Of the 12 customers who took actions in the control group, some also reported extra steps. Eight customers reported doing a software update, two customers said they stopped using the device, and one customer described disconnecting the device.

Modeling compliance

All users in the treatment groups (92%) took some steps, though in many different combinations. Figure 4 shows the distribution of the count of steps taken by the users. When notified users do take action, they report on average 2.9 steps recommended by the notification, while the control group report on average 0.3 steps on their own initiative, without being notified.

Before turning to the models, we also did a chi-square test to validate that customers who responded were not more diligent or motivated than those who did not respond to our survey. The test result was $X^2(4, N = 177) = 0.032$, $P = 0.99$. The test suggests that there is no relationship between interviewed/non interviewed and clean/no clean outcomes. Also, we carried out a chi-square test to compare if the ISP and the subsidiary had differences in the compliance steps. We only checked the walled garden group of the ISP and subsidiary. The test result was $X^2(12, N = 39) = 5.69$, $P = 0.93$. The test suggests that ISP and the subsidiary have no significant differences in terms of performed steps.

We want to understand which causal factors are associated with user compliance. We operationalized compliance as a ratio of the number of steps divided by five, the number of steps recommended

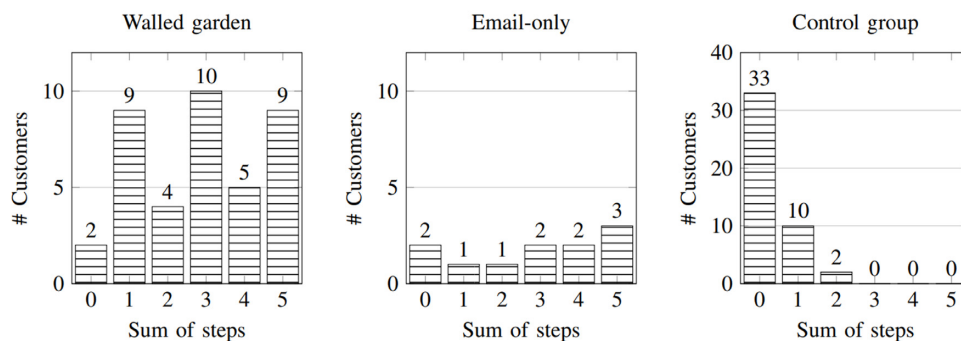


Figure 4: Distribution of the count of steps taken by the users.

Table 8: Summary of variables

| Reference category | Variables | Explanation |
|--------------------|-------------------------|--|
| Control group | Walled garden | True if in walled garden group and not in email and control group |
| | Email-only | True if in email-only group and not in walled garden group and control group |
| Female | :: Age | Discrete variable |
| | Small business | True if it is a small business and not male and female |
| | Male | True if male and not small business and female |
| No Domoticz | Domoticz | True if the device type is Domoticz |
| Did not understand | Understood notification | True if consumer understood the notification |
| Internet back | Safe internet | True if motivation is not to get internet back and other motivations |
| | Other motivations | True if motivations are others motivation and not internet back and safe internet. |

Note for Model (4) and Model (5) in Table 9, the reference category for the walled garden group is the email-only group. For Models 1–3, $N = 95$ and for Models 4–5, $N = 50$. The vertical bars are to visually group the independent variables with their reference category (since “Age” does not have a reference category we used :: as symbol).

in the notification (see the “Methodology” section). Since it is a proportion from 0 to 1, this type of data can be analyzed with a beta regression. However, beta regressions assume that the ratio is between 0 and 1, excluding the extremes. Since we also have scores of 0 and 1 in our data, we have to transform these extreme values as suggested in [46, 47]. The distribution of the dependent variable did not change.

We model the driving factors of performance using the explanatory variables from our adapted theoretical framework. Hence, we have five groups of independent variables. The first category is the treatment consumers received: walled garden notification, email-only notification, no notification (control). Second, we include as control variables the user characteristics age, gender and status as a small business. Since there were three observations in the walled garden group and two observations in the control group with missing values for gender, we used as imputation method the most frequent value, meaning we replaced the missing values with the most common value.

Next, we control for device types. In the subsection “Device type” of the section “Findings,” we discussed the range of devices reported by the user. We cannot use the reported device types as explanatory factors, since many of them are used by only a few people, so the samples would be way too small to register any effect on compliance. The key difference in the population of device types lies between the Raspberry Pis and the other IoT devices like cameras and DVRs. The Raspberry Pis were specifically targeted by attackers via a known vulnerability in Domoticz (CVE-2019-10664, CVE-2019-10678). Hence, we created a categorical variable called Domoticz to distinguish between device types.

Next, we have comprehension of the notification, coded as: Understood or Did not understand. When we asked this question, there were two missing values from the walled garden group, so similar to gender, we used the most frequent value as imputation method. We ran the model with and without using the most frequent value imputation method for gender and comprehension variables, and the results did not change. And finally, we include the different motivations that were reported by users to comply. Similarly to device type, many of the motivations had a small size, so they would be way too small to register any effect on compliance. Therefore, we grouped motivations into three categories. The first category was users who wanted their internet back. The second category was composed by users whose motivations were to have the internet back and safe internet, only safe internet, and privacy concern and safe internet. Finally, other motivations include malfunctioning device, the need of the device, and no answers. Table 8 provides a summary of the variables that will be included in the regression model as well as the corresponding reference categories.

Table 9 presents the estimated coefficient values, significance levels, and additional goodness-of-fit indicators of interest. We decided to take a stepwise approach in adding each group of variables, so we can assess their effects on compliance. Model (1) shows that the treatments—that is, the fact that users were notified—already explain 50% of the variance in compliance ($R^2 = 0.501$). Simply put, notifications do get many people to take action. This holds even for email-only. This is somewhat surprising, as earlier work [4] found that sending an email was indistinguishable from the control group, in terms of cleanup at least. In contrast, we find that emails are not ignored by users, even though they easily could do so.

From Model (4), it is also possible to observe that understanding the notification does not have a significant impact on compliance compared with consumers who did not understand since this variable only explains 6.4% of the variance in compliance ($R^2 = 0.064$). As visible in Model (5), comprehension does not have a significant impact on compliance compared with consumers who did not understand the message, though the positive coefficient is in the expected direction. In terms of the different motivations, “other motivations” have a significant negative impact on compliance compared with users who want their internet back. Users whose motivations are related to the need to use a device or to the malfunctioning of it, or users who did not give an answer to this question, comply less. Note that Models (4) and (5) do not include the control group, as comprehension cannot be measured for this group because they did not receive a notification. For these models, the email-only group is the reference group. Other goodness-of-fit indicators, such as log likelihood, are reported for all models. Higher log-likelihood values are preferred, although they alone cannot be used to determine the fit of the model. See Appendix B for more details on the likelihood ratio test of the models. We will proceed to interpret our final model, Model (3), as the best fit for the data.


To interpret intuitively the coefficients of Model (3), the coefficients were converted to average marginal effects, Fig. 5 presents a summary of the average marginal effects of the predictor variables on the compliance ratio, which is to say, the average expected change in compliance ratio for a change in a predictor.

We will interpret only the significant coefficients of Model (3). Model (3) suggests that being in the walled garden increases the average compliance ratio by 0.39. Since the dependent variable is a proportion of the five steps that users took, we should multiply the coefficient 0.39 times five. Meaning that consumers in the walled garden do 1.95 steps more on average relative to the control group, which compliance ratio is on average 0.3. Similarly, receiving an email increases the average expected compliance ratio by 0.36 respective to the control group. Meaning consumers in the email group do 1.8

Table 9: Estimated coefficients beta regression on compliance ratio

| | Dependent variable: Compliance Ratio (Transformed) | | | | |
|-------------------------|--|---------------------|---------------------|-------------------|----------------------|
| | Beta Regression—link = 'logit' | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Walled garden | 2.037*** (0.287) | 2.036*** (0.292) | 2.035*** (0.292) | 0.129 (0.449) | −0.295 (0.461) |
| Email-only | 1.928*** (0.416) | 1.897*** (0.429) | 1.874*** (0.433) | | |
| Age | | −0.004 (0.009) | −0.004 (0.010) | −0.007 (0.014) | −0.007 (0.013) |
| Small business | | −0.270 (0.754) | −0.233 (0.759) | −0.769 (1.475) | −1.028 (1.443) |
| Male | | 0.008 (0.412) | −0.023 (0.421) | −0.878 (0.948) | −1.082 (0.892) |
| Domoticz | | | 0.095 (0.258) | −0.044 (0.379) | 0.099 (0.380) |
| Understood notification | | | | 0.610 (0.396) | 0.340 (0.378) |
| Safe internet | | | | | −0.303 (0.451) |
| Other motivation | | | | | −1.807*** (0.515) |
| Constant | −1.608*** (0.195) | −1.393** (0.624) | −1.438** (0.635) | 1.035 (1.213) | 2.278* (1.201) |
| Observations | 95 | 95 | 95 | 50 | 50 |
| Pseudo R-squared | 0.501 | 0.503 | 0.505 | 0.064 | 0.277 |
| Log likelihood | 95.155 | 95.267 | 95.332 | 15.630 | 22.185 |

Note: * $P < 0.1$; ** $P < 0.05$; *** $P < 0.01$.

**Figure 5:** Average marginal effect of each predictor variable.

steps more on average respective to the control group. Although Model (5) explain less variance having other motivations rather than wanting the internet back decreases the average compliance ratio by -0.38 . Meaning that consumers in the group with other motivation do 1.9 steps less on average than consumers who want their internet back.

In summary, our model finds clear evidence for the impact of the notification and of user motivation. Comprehension seems to have less effect, which is somewhat puzzling, since the notifications have to rely on rather generic advice, rather than clear-cut and actionable instructions. Perhaps the generic advice is easy to understand so users do not see the subsequent questions and difficulties (“how can I actually change the password on my IP camera of brand X?”) as a part of the message itself, but rather as a challenge separate from understanding the message. In that case, they would answer that they un-

derstood the message, even if they had trouble understanding how to comply with it.

Modeling cleanup

Now we turn to the actual goal of the notifications and compliance: cleanup of the infected devices. Figure 6 shows how many devices were cleaned up after 2 weeks of being notified or assigned to the control group, distributed over the number of steps the user reportedly took. As expected, cleanup rates are higher when the number of compliance steps increases.

An important finding is that cleanup also happens in the control group—mostly concentrated in the column with zero steps. In line with earlier work [4], we also found that 33% of the survey respondent users in the control group, who reported not taking any step, also got clean. It is unclear how this happens. We did find that around 26% of the users in that group also undertook action, even though they were not informed. Certain security behaviors are triggered by other mechanisms, such as update notifications. While our study added a new piece to this puzzle because users reported no action, we still cannot present a satisfying answer.

Compared with the control group, remediation rates in the two treatment groups are significantly higher. In the walled garden group, 90% got cleaned up versus 73% in the email group.

The final part of our research question is to estimate the effect of compliance on cleanup. We do this via a binomial logistic regression model. Binomial logistic regression is used when the dependent variable is binary—in this case, whether a device has been cleaned up or not.

Table 10 presents the estimated coefficient values, significance levels, and additional goodness-of-fit indicators. The primary focus is on the relationship between compliance and cleanup. We also look at the

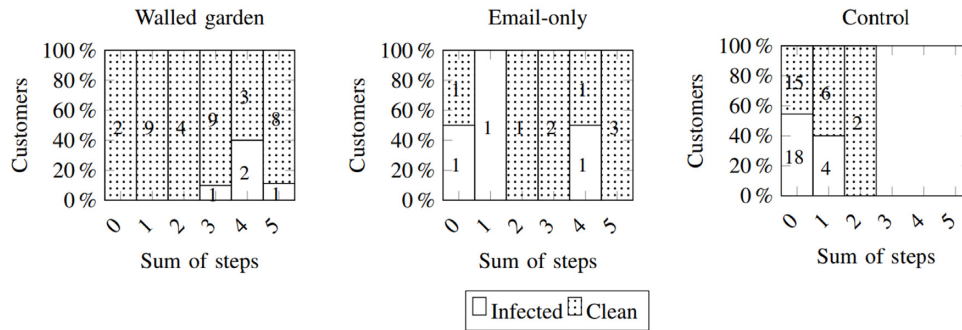


Figure 6: Cleaned versus infected devices after 14 days.

Table 10: Estimated coefficients binomial logistic regression on cleanup

| | Dependent variable: clean | | |
|-------------------|--|-------------------|----------------------|
| | Binomial Logistic Regression*—link = 'logit' | | |
| | (1) | (2) | (3) |
| Compliance ratio | 2.197*** (0.806) | 1.524* (0.849) | 1.627* (0.859) |
| Domoticz | | 0.316 (0.498) | 0.648 (0.541) |
| Did extra step | | 0.869 (0.592) | 0.802 (0.602) |
| Competing malware | | | -1.576*** (0.552) |
| Constant | 0.220 (0.292) | -0.037 (0.348) | 0.776 (0.474) |
| Observations | 95 | 95 | 95 |
| Log likelihood | -53.782 | -52.219 | -47.618 |
| Akaike Inf. Crit. | 111.565 | 112.437 | 105.237 |
| McFadden R2 | 0.07 | 0.10 | 0.18 |

Note: *P < 0.1; **P < 0.05; ***P < 0.01.
*Also known as binary logistic regression.

effect of the extra steps that consumers reported performing, at the device type, and at the issue of whether we observed scanning activity from competing malware variants for the customer. We define three models in which we estimate the effects of each additional variable on remediation.

An intuitive way to represent the results of binary logistic regression models is converting the coefficients into a relative risk (RR). This will capture the change of the probability of remediation after the exposure to each predictor variable. From Model (1), once converting the coefficient (2.197) to RR, we can observe that an increase in the compliance ratio increases the probability of remediation by 37% as compared with the control group. In Model (2), we checked the influence of device type, and it does not have a significant effect. Figure 7 shows the relative risk of the coefficients of our final Model (3). An increase in compliance ratio increases the probability of remediation by 32%. Extra steps and the device type have no significant effect. Competing malware presence in the home network decreases the probability of remediation by 54%.

Customer Experience

The survey ended with two questions about their experience as customers of the ISP and the subsidiary brand. There was an open ques-




Figure 7: Relative Risk Model (3) on remediation.

tion asking about what consumers thought of ISP reaching out to infected customers, and 24 (61%) of the survey respondents in the walled garden group were satisfied with this approach versus 11 (100%) in the email-only group. These results are more encouraging than in [4], in which only 17 respondents out of 76 expressed satisfaction. A possible explanation for this difference is that in our study, we asked specifically about a customer's opinion of the service, rather than analyzing the logs of people contacting the support center. The latter is likely biased toward customers being frustrated and struggling with resolving the infection. In our study, some consumers expressed frustration with losing their internet access, but they were also glad to be contacted.

Customers were also asked for suggestions to improve the notification and remediation mechanism. Five customers in the email-only group and 24 in the walled garden group gave an answer. From the email-only group, two customers suggested that more information on the offender device is needed. One customer expressed that a more personalized email would help to avoid users thinking it is a phishing email. Another customer expressed the necessity of a higher availability of the abuse team, since they do not work during weekends. Finally, a customer suggested giving more publicity to the abuse team, so users would be aware of their role. From the walled garden group, 12 customers suggested that a warning prior to being in a walled garden was necessary. Along the same lines, five customers expressed that a call before putting them in a walled garden was a way to improve. Seven customers expressed that more availability of the abuse team was necessary. Other suggestions from the walled garden respondents were to explain more clearly the quarantine process and how to get out, to provide more information on the malware, to

work on the authenticity of the warning, and to include information on what device type was actually infected.

Limitations and Future Work

We discuss the main limitations of our study. First, there is the issue that there could potentially be multiple infections in the same home network. The ISP notification did tell users that there could be more than one infection. As long as we saw signs of an infection, we coded the user as “not clean,” though the user might have cleaned up one of the infected devices. This means we cannot measure partial cleanup, only full cleanup. Second, our data is on self-reported actions. Users might have forgotten what they did or give socially desirable answers. We cannot rule out these effects, but we did see that the devices that participants mentioned as being the culprits were, in fact, the same ones were found by Shodan at those IP addresses. We also observed that more than half of all users reported taking no action or only one step (excluding the control group, the count is one in four users). At the other extreme, only around one in four people stated doing all the steps. This pattern suggests that the tendency to provide socially desirable answers was limited.

A third limitation is the limited sample size: 177 participants in the whole study and 95 participants in the survey sample. This sample is large enough to find robust results for certain effects and causal factors. That being said, we were still left with a large portion of unobserved effects in the study on the impact of compliance on cleanup.

Finally, the experiment was carried out in one ISP and its subsidiary brand in the Netherlands. It is unclear how well these results will generalize beyond this ISP and country.

Future work might pursue a study with a larger sample size and in other ISPs and countries. Laboratory experiments might be an alternative, but they have their own methodological weaknesses compared with a field study with a sample of real and heterogeneous users. An important direction for future work is also to test various approaches in terms of how to actually provide usable as well as effective cleanup advice or understand why users do not take some of the suggested steps in the notification. This might need future work to collect actual ground-truth on the infected devices on customer premises, in order to have an empirical basis for remote device identification and identifying the best cleanup advice, as well as better understanding of users’ mental models.

Discussion and Conclusion

ISPs are asked to implement best practices to notify consumers about IoT infections. Is cleaning IoT something that consumers can actually do? While earlier work [4] suggests that the answer is Yes, we actually knew little about the underlying mechanism. Without that understanding, we cannot design better interventions. For this reason, we measured, first, whether ISP customers complied with the cleanup advice and, second, whether this compliance improved cleanup rates.

We identified that only 24% of all survey respondents and notified participants succeeded in performing all remediation steps. The overwhelming majority of notified users, however, took at least some action upon receiving the notification. Even in the email-only group, which only received an email and had no further incentive to act, over 80% took some action. This finding suggests, differently than [24, 25], that a less intrusive notification could be effective. However, due to the sample size, more research is needed. In short, we found significant evidence that when consumers are informed about compromised IoT, they are willing to act. Users notified via email do 1.8 steps on average, while users in the walled garden do 1.95 steps on

average, both compared with the control group, where users only do 0.31 steps.

When analyzing the impact on cleanup, an increase in the compliance ratio increase the probability of remediation by 32%. However, if the home network was infected with competing malware, this reduced the probability of remediation by 54%. It suggests that user compliance with the recommended steps might not apply to all types of malware. Some devices remain infected or are being reinfected. IoT malware analysis has confirmed that some families fight for control over vulnerable devices. Another explanation for the effect of competing malware might be that the user owned more than one infected device. Both explanations are consistent with our finding of that competing malware are correlated with worse remediation rates.

If the impact of compliance is limited, it does not mean that the notifications as such are ineffective. Rather, it signals that the recommended remediation steps are not a sure way to get rid of the infection. Users who receive the notification might comprehend their IoT devices well enough to chart their own course out of the problem. This is supported by the fact that the impact of the notification on cleanup is higher than the impact of compliance. Cleanup rates are high in both treatment groups: 90% in the walled garden group and 73% in the email-only group. This suggests that users, once aware of the problem, are often able to resolve it, irrespective of the grave usability problems plaguing the recommended steps and IoT security in general. Putting this into the context of the C–HIP model, the notification acts as an attention switch that triggers users to comply. Comprehension does not play a role in changing user behavior (compliance), while the type of motivation that users expressed can negatively influence compliance compared with users who want their internet back. Its effect is not as big as the notifications. Perhaps we are seeing an effect of early IoT adopters being also more technically competent than average users. In that case, we would expect to see diminishing cleanup rates with the wider adoption of IoT.

Consistent with Fagan and Khan [34], we have observed how users’ motivations are related to how the infection could affect themselves rather than how the infected IoT devices could affect others. Similar to Redmiles [35], we take a step forward on understanding compliance with users fixing a real infection in their home network, giving ecological validity to these findings.

These findings clearly underline the recommended best practice for ISPs to notify infected users. Walled gardens perform the best in terms of cleanup. However, they have achieved only limited adoption among ISPs, because of cost considerations and the fear of customer pushback. Bad luck caused our email-only group to end up too small to make robust inferences. That being said, contrary to [4], users in this group had high compliance rates and high remediation rates. Since email is a cheap and easily available option for ISPs, this could be a good second-best notification mechanism. Future work should test whether our findings for this group hold up with larger samples. In the end, though, the lion’s share of the burden is not borne by the ISP. The good news from our study is that consumers are willing and able to take action, even in the absence of usable security advice and solutions.

Acknowledgments

We would like to thank our anonymous reviewers for their feedback and suggestions to improve the quality of our manuscript.

Funding

This publication is part of the MitigatNG IoT-based DDoS attacks via DNS (MINIONS) project [number 628.001.033] of the “Joint U.S.–Netherlands Cy-

ber Security Research Programme,” which is (partly) financed by the Dutch Research Council (NWO); and of the related MINIONS-TLD project, which is financed by the Stichting Internet Domeinnaamregistratie Nederland (SIDN), the .nl registry.

Conflict of interest statement. None declared.

References

1. Tung L. IoT devices will outnumber the world's population this year for the first time. 2017. [Online]. <https://www.zdnet.com/article/iot-devices-will-outnumber-the-worlds-population-this-year-for-the-first-time/>.
2. Kumar D, Shen K, Case B. *et al.* All things considered: an analysis of IoT devices on home networks. In: *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, 2019, pp. 1169–85. [Online]. <https://www.usenix.org/conference/usenixsecurity19/presentation/kumar-deepak>.
3. HP. HP news: HP study reveals 70 percent of internet of things devices vulnerable to attack. 2014. [Online]. <https://www8.hp.com/us/en/hp-news/press-release.html?id=1744676>.
4. Cetin O, Gañán C, Altena L. *et al.* Cleaning up the internet of evil things: real-world evidence on ISP and consumer efforts to remove mira. In: *NDSS*, 2019. [Online]. <http://dx.doi.org/10.14722/ndss.2019.23438>.
5. Livingood J, Mody N, O'Reirdan M. Recommendations for the remediation of bots in ISP networks. 2012. [Online]. <https://tools.ietf.org/html/rfc6561>.
6. Antonakakis M, April T, Bailey M. *et al.* Understanding the mirai botnet. In: *Proceedings of the 26th USENIX Security Symposium*, pp. 1093–110. 2017. [Online]. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis>.
7. Song Y, Huang Q, Yang J. *et al.* IoT device fingerprinting for relieving pressure in the access control. In: *Proceedings of the ACM Turing Celebration Conference, China*. New York, NY: Association for Computing Machinery, 2019. [Online]. <https://doi.org/10.1145/3321408.3326671>.
8. Koliás C, Kambourakis G, Stavrou A. *et al.* DDoS in the IoT: Mirai and other botnets. *Computer*, 2017;50:80–4.
9. De Donno M, Dragoni N, Giaretta A. *et al.* Analysis of DDoS-capable IoT malwares. In: *Proceedings of the 2017 Federated Conference on Computer Science and Information Systems, FedCSIS 2017*, 2017.
10. Vlajic N, Zhou D. IoT as a land of opportunity for DDoS hackers. *Computer*, 2018;51:26–34.
11. Symantec. Internet security threat report, volume 24. Tech. Rep., 2019.
12. Mikhail Kuzin VK, Shmelev Y. New trends in the world of IoT threats. 2018. [Online]. <https://securelist.com/new-trends-in-the-world-of-iot-threats/87991/>.
13. DeBeck C, Chung J, McMillen D. I can't believe mirais: tracking the infamous IoT malware. [Online]. <https://securityintelligence.com/posts/i-cant-believe-mirais-tracking-the-infamous-iot-malware-2/>.
14. Ashgari H, van Eeten MJ, Bauer JM. Economics of fighting botnets: lessons from a decade of mitigation. *IEEE Secur Priv* 2015;13:16–23. [Online]. <http://ieeexplore.ieee.org/document/7310846/>.
15. Ashgari H, Ciere M, Van Eeten MJG. Post-mortem of a zombie: conficker cleanup after six years. In: J. Jung (ed.), *Proceedings of the 24th USENIX Security Symposium (pp. 1-16)*. USENIX Association. 2015.
16. Pijpker J, Vranken H. The role of internet service providers in botnet mitigation. In: *2016 European Intelligence and Security Informatics Conference (EISIC)*. IEEE, 2016, pp. 24–31. [Online]. <http://ieeexplore.ieee.org/document/7870186/>.
17. Lear E, Droms R, Romascanu D. Manufacturer usage description specification. Tech. Rep., 2019. [Online]. Available: <https://www.rfc-editor.org/info/rfc8520>.
18. Richardson M, Ranganathan M. Manufacturer usage description for quarantined access to firmware. Internet Engineering Task Force, Tech. Rep., 2019. [Online]. <https://datatracker.ietf.org/doc/draft-richardson-hg-mud-quarantined-access/>.
19. Vasek M, Moore T. Do malware reports expedite cleanup? An experimental study. In: *Presented as part of the 5th Workshop on Cyber Security Experimentation and Test*. Bellevue, WA: USENIX, 2012. [Online]. <https://www.usenix.org/conference/csset12/workshop-program/presentation/Vasek>.
20. Li F, Ho G, Kuan E. *et al.* Remediating web hijacking: notification effectiveness and webmaster comprehension. In: *25th International World Wide Web Conference, WWW 2016*, 2016.
21. Cetin O, Jhaveri MH, Gañán C. *et al.* Understanding the role of sender reputation in abuse reporting and cleanup. *J Cybersecurity* 2016;2:83–98.
22. Durumeric Z, Kasten J, Adrian D. *et al.* The matter of Heartbleed. In: *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC*, 2014.
23. Li F, Bailey M, Durumeric Z. *et al.* You've got vulnerability: exploring effective vulnerability notifications. In: *USENIX Security Symposium*, 2016.
24. Stock B, Pellegrino G, Li F. *et al.* Didn't you hear me? Towards more successful web vulnerability notifications. In: *Proceedings of the 25th Annual Symposium on Network and Distributed System Security (NDSS '18)*, February 2018. [Online]. <https://publications.cispa.saarland/1190/>.
25. Cetin O, Gañán C, Altena L. *et al.* Tell me you fixed it: evaluating vulnerability notifications via quarantine networks. In: *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 326–39.
26. Cetin O, Altena L, Gañán C. *et al.* Let me out! Evaluating the effectiveness of quarantining compromised users in walled gardens. In: *Fourteenth Symposium on Usable Privacy and Security*, 2018.
27. Egelman S, Molnar D, Christin N. *et al.* Please continue to hold: an empirical study on user tolerance of security delays. In: *Workshop on the Economics of Information Security (WEIS)*, 2010.
28. Egelman S, Schechter S. The importance of being earnest [in security warnings]. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2013.
29. Felt AP, Ainslie A, Reeder RW. *et al.* Improving SSL warnings: comprehension and adherence. In: *Conference on Human Factors in Computing Systems, Proceedings*, 2015.
30. Akhawe D, Felt AP. Alice in warningland: a large-scale field study of browser security warning effectiveness. In: *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*. Washington, DC: USENIX, 2013, pp. 257–72. [Online]. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe>.
31. Krol K, Moroz M, Sasse MA. Don't work. Can't work? Why it's time to rethink security warnings. In: *7th International Conference on Risks and Security of Internet and Systems, CRiSIS 2012*, 2012.
32. Bravo-Lillo C, Cranor LF, Downs J. *et al.* Bridging the gap in computer security warnings: a mental model approach. *IEEE Secur Priv*, 2011;9:18–26.
33. Wogalter MS, Laughery KR. Warning! sign and label effectiveness. *Curr Dir Psychol Sci*, 1996;5:33–7.
34. Fagan M, Khan MMH. Why do they do what they do?: a study of what motivates users to (not) follow computer security advice. In: *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association, 2016, pp. 59–75. [Online]. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/fagan>.
35. Redmiles EM. “Should I worry?” A cross-cultural examination of account security incident response. In: *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 920–34.
36. Vaniea K, Rader E, Wash R. Betrayed by updates: how negative experiences affect future security. In: *Conference on Human Factors in Computing Systems, Proceedings*, 2014.
37. Conzola VC, Wogalter MS. A communication–human information processing (C–HIP) approach to warning effectiveness in the workplace. *J Risk Res* 2001;4:309–22. [Online]. <https://www.tandfonline.com/doi/full/10.1080/13669870110062712>.

38. Sun H, Zhang P. The role of moderating factors in user technology acceptance. *Int J Hum Comput Stud* 2006;64:53–78.
39. Shodan. Shodan. 2019. [Online]. <https://www.shodan.io/>.
40. Shadowserver. Drone/botnet-drone report. 2019. [Online]. <https://www.shadowserver.org/what-we-do/network-reporting/drone-botnet-drone-report/>.
41. GCA. GCA—global cyber alliance—working to eradicate cyber risk. 2019. [Online]. Available: <https://www.globalcyberalliance.org/>.
42. Pa YMP, Suzuki S, Yoshioka K. et al. IoT POT: analysing the rise of IoT compromises. In: *9th USENIX Workshop on Offensive Technologies (WOOT '15)*. Washington, DC: USENIX Association, 2015. [Online]. <https://www.usenix.org/conference/woot15/workshop-program/presentation/pa>.
43. Malware Wiki. Mirai: malware wiki. 2019. [Online]. <https://malwiki.org/index.php?title=Mirai>.
44. Dittrich D, Kenneally E. The menlo report: ethical principles guiding information and communication technology research. 2012. [Online]. <http://www.caida.org/publications/papers/2012/menlo-report-actual-formatted>.
45. Domoticz. Domoticz downloads. 2019. [Online]. <https://www.domoticz.com/downloads/>.
46. Smithson M, Verkuilen J. A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychol Methods* 2006;11:54–71.
47. Cribari-Neto F, Zeileis A. Beta regression in R. *J Stat Softw* 2010;34:1–24.

Appendix A: Correlation between the steps performed by customers

| r_s | Step 1 | Step 2 | Step 3 | Step 4 | Step 5 |
|--------|--------|--------|--------|--------|--------|
| Step 1 | 1 | | | | |
| Step 2 | 0.49 | 1 | | | |
| Step 3 | 0.49 | 0.74 | 1 | | |
| Step 4 | 0.49 | 0.42 | 0.58 | 1 | |
| Step 5 | 0.44 | 0.56 | 0.56 | 0.73 | 1 |

Figure A1: Correlation between the steps performed by customers.

Appendix B: Likelihood ratio test compliance models

Table B1:

Likelihood ratio test Models 1–3

Model 1: Compliance ratio ~ Walled Garden + Email-only

Model 2: Compliance ratio ~ Walled Garden + Email-only + Age + Small business + Male

Model 3: Compliance ratio ~ Walled Garden + Email-only + Age + Small business + Male + Domoticz #Df LogLik Df Chisq Pr(>Chisq)

1 4 95.155

2 7 95.267 3 0.2240 0.9736

3 8 95.332 1 0.1319 0.7165

Likelihood ratio test Models 4–5

Model 1: Compliance ratio ~ Walled Garden + Age + Small business + Male + Domoticz + Understood notification

Model 2: Compliance ratio ~ Walled Garden + Age + Small business + Male + Domoticz + Understood notification + Safe

internet + Other motivation #Df LogLik Df Chisq Pr(>Chisq)

4 8 15.630

5 10 22.185 2 13.11 0.001423**

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Note: Models 1–3 are not different with respect to likelihood value, and Model 5 shows improvement with respect to likelihood value of Model 4.

Appendix C: Survey protocol

These are the survey protocols that were used to conduct the survey with the users in the different treatment groups. The survey was conducted in Dutch. We translated the questions as accurately as possible to English.

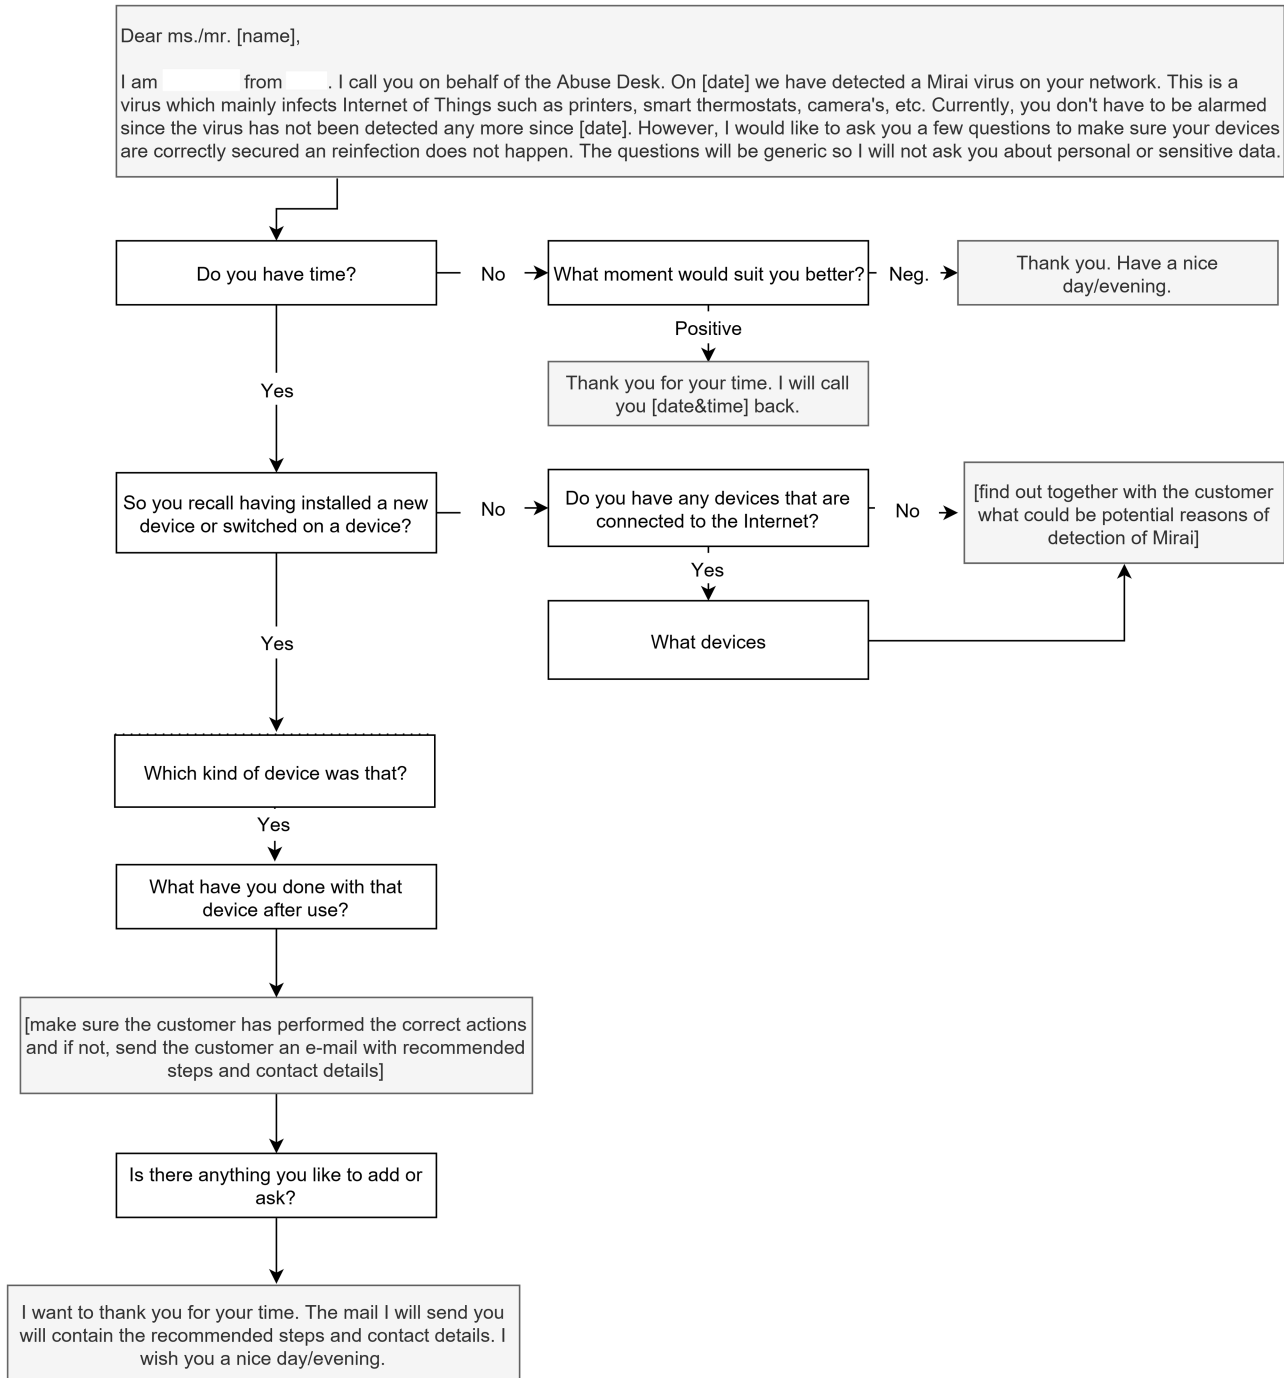


Figure C1: Survey protocol control group.

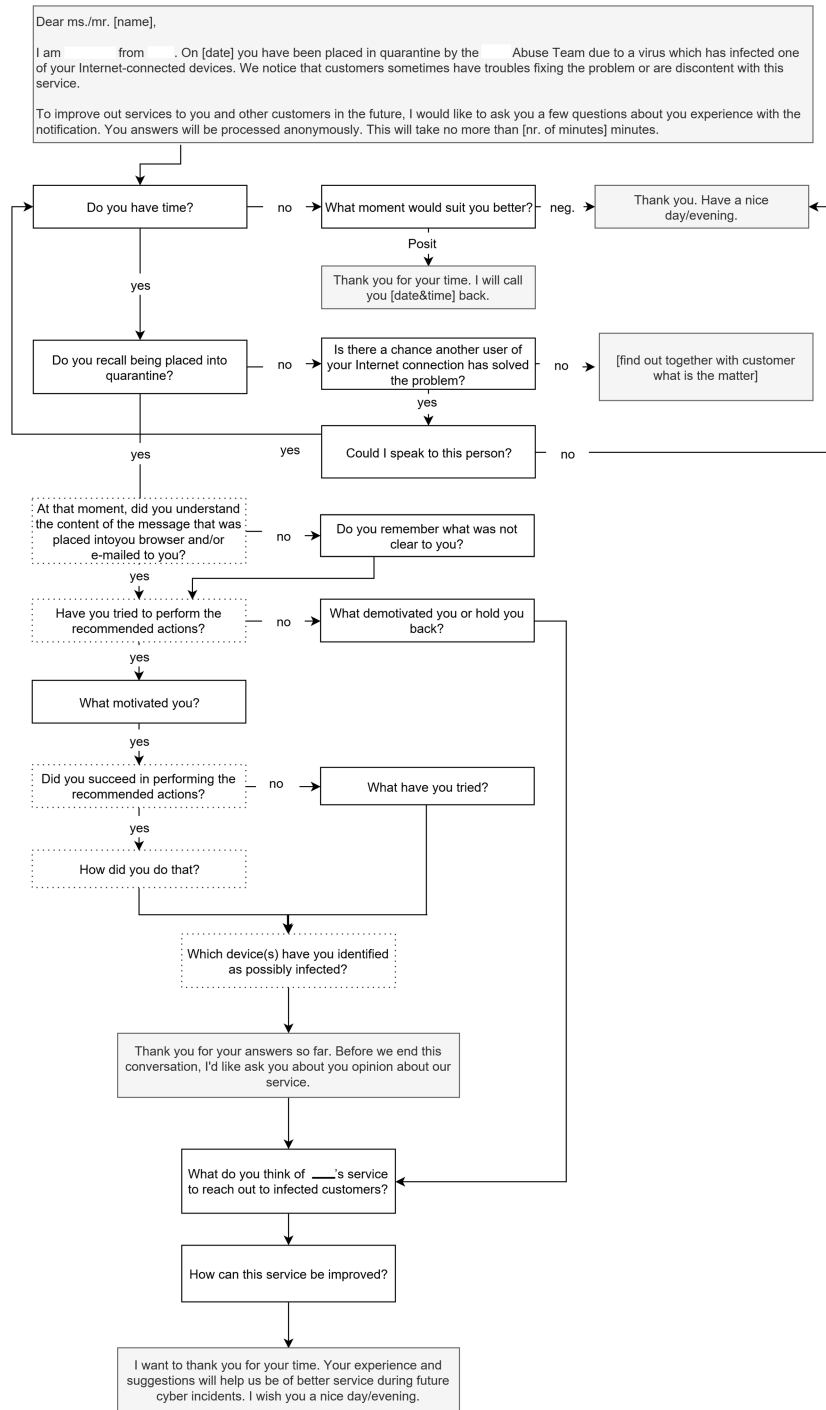


Figure C2: Survey protocol walled garden group.




Figure C3: Survey protocol email-only group.

Appendix D: Notifications

Walled garden

Illustration of walled garden landing page displayed to consumers who were randomly assigned to the walled garden treatment group. The same content was also sent to consumers via email.

Quarantainenet

Secure environment

A safe Internet is in everyone's interest. We, [Quarantainenet](#), strongly care about protecting your (confidential) information.

We have received information from one of our partners that a security issue has been detected on your Internet connection. You probably have not noticed anything yet.

Don't worry. To protect you against the security risks we have placed your Internet connection in our secure environment. In this environment you can safely solve the security issues. We are willing to help you to do so.

What is the problem and how can you solve it?

One or more Internet connected devices in your home have been infected with the Mirai virus. We cannot detect which Internet connected device has been infected. Most likely it is a digital video recorder (DVR), security camera or printer connected to the Internet rather than a computer, laptop, tablet or mobile phone.

What should you do to remove the Mirai virus and prevent future infections?
Please follow the steps below. If you cannot complete a step, please proceed to the next one.

- Determine which devices are connected to your Internet connection.
Reminder: The Mirai virus mainly infects Internet connected devices such as a DVR, security camera or printer connected to the Internet.
- 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.
- Restart the Internet connected devices by turning it 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.

- Reset your modem/router to the factory settings. On [this page](#) it is described how you do this for an Experia Box.
- Set the password of your modem/router. On [this page](#) it is described how you do this for an Experia Box.

Warning! If remote access to a certain device is absolutely necessary, manually define port forwards in your router for this device. On [this page](#) it is described how you do this for an Experia Box.

Necessary steps

- Take the measures stated above.
- Fill in our [form](#) (and restore your Internet Connection).

General security tips

- * Use an up-to-date virus scanner to keep out potential hazards.
- * Keep computer software, like your operating system, up to date.
- * Do not open messages and unknown files that you do not expect or trust.
- * Secure your wireless connection with a unique and strong password

Figure D1: Landing page of walled garden.

Email-only

Example of notification email sent to consumers randomly assigned to the email-only treatment group. The notification content essentially only differs with the previous example in that it omits statements about placing the recipient in a quarantine environment.

Dear Sir/Madam,

A safe internet is in everyone's interest. We, _____, strongly care about protecting your (confidential) information.

We have observed a security issue on your internet connection. You probably have not noticed anything, because it's about processes that run in the background.

One or more Internet connected devices in your home have been infected with the Mirai virus. We cannot detect which Internet connected device has been infected. Most likely it is a digital video recorder (DVR), security camera or printer connected to the Internet rather than a computer, laptop, tablet or mobile phone.

What should you do to remove the Mirai virus and prevent future infections? Please follow the steps below. If you cannot complete a step, please proceed to the next one.

1. Determine which devices are connected to your Internet connection. Reminder: The Mirai virus mainly infects Internet connected devices such as a DVR, security camera or printer connected to the Internet.
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.
3. Restart the Internet connected devices by turning it 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.
4. Reset your modem/router to the factory settings. On _____ it is described how you do this for an Experia Box.
5. Set the password of your modem/router. On _____ it is described how you do this for an Experia Box.

Warning! If remote access to a certain device is absolutely necessary, manually define port forwards in your router for this device. On _____ it is described how you do this for an Experia Box.

We ask you to take above steps within a day and to respond to this message. You can also ask additional questions in a reply to this email.

Kind regards,

____ Abuse Team

Abuse team email

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

Figure D2: Notification email sent to consumers in email-only treatment group.