

This work is on a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license, <https://creativecommons.org/licenses/by-nc-nd/4.0/>. Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Peter Story*, Daniel Smullen, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh*, and Florian Schaub

Awareness, Adoption, and Misconceptions of Web Privacy Tools

Abstract: Privacy and security tools can help users protect themselves online. Unfortunately, people are often unaware of such tools, and have potentially harmful misconceptions about the protections provided by the tools they know about. Effectively encouraging the adoption of privacy tools requires insights into people’s tool awareness and understanding. Towards that end, we conducted a demographically-stratified survey of 500 US participants to measure their use of and perceptions about five web browsing-related tools: private browsing, VPNs, Tor Browser, ad blockers, and antivirus software. We asked about participants’ perceptions of the protections provided by these tools across twelve realistic scenarios. Our thematic analysis of participants’ responses revealed diverse forms of misconceptions. Some types of misconceptions were common across tools and scenarios, while others were associated with particular combinations of tools and scenarios. For example, some participants suggested that the privacy protections offered by private browsing, VPNs, and Tor Browser would also protect them from security threats – a misconception that might expose them to preventable risks. We anticipate that our findings will help researchers, tool designers, and privacy advocates educate the public about privacy- and security-enhancing technologies.

Keywords: Privacy, technology adoption, mental models, private browsing, VPN, Tor, ad blocking, antivirus.

DOI 10.2478/popets-2021-0000

Received 2020-11-30; revised 2021-03-15; accepted 2021-03-16.

***Corresponding Author: Peter Story:** School of Computer Science, Carnegie Mellon University, E-mail: pstory@andrew.cmu.edu

Daniel Smullen, Lorrie Faith Cranor: School of Computer Science, Carnegie Mellon University

Yaxing Yao: Department of Information Systems, University of Maryland, Baltimore County. Previously a postdoctoral associate at Carnegie Mellon University’s School of Computer Science.

Alessandro Acquisti: Heinz College of Information Systems and Public Policy, Carnegie Mellon University

***Corresponding Author: Norman Sadeh:** School of Computer Science, Carnegie Mellon University, E-mail: sadeh@cs.cmu.edu

1 Introduction

A majority of Americans are concerned about their privacy [5]. Yet, despite expressed interest in privacy-enhancing solutions, adoption of tools that offer privacy protection remains low [100]. A variety of tools are available, but it is unclear how much end users know about them. Nudging interventions [2] offer the possibility of increasing adoption, since they have successfully helped people adopt tools in other contexts [3, 85]. Such interventions should target barriers to adoption and correct potentially dangerous misconceptions [2]. Misconceptions should also be addressed in the design and marketing of privacy tools themselves.

To inform the design of nudging interventions and privacy tools, we conducted a demographically-stratified survey of US participants and measured their use of and perceptions about five web browsing-related tools: private browsing, VPNs, Tor Browser, ad blockers, and antivirus software. We focused on answering four research questions:

1. To what extent are people aware of these tools, and how frequently do they use them? (§ 4.2)
2. How interested are people in preventing specific privacy and security threats? (§ 4.3)
3. How accurately can people determine whether these tools afford protection from specific privacy and security threats? (§ 4.4)
4. What misconceptions, if any, do people have about these tools? (§ 4.7)

Our data show a substantial number of misconceptions across all scenarios and tools. This is true even for tools that have widespread adoption, such as antivirus software. In fact, we show that greater experience with VPNs and Tor Browser is associated with confusion about these tools’ protections (§ 4.5). In addition, our thematic analysis of participants’ responses reveals po-

Florian Schaub: School of Information, University of Michigan

tentially harmful misconceptions that exist for all the tools we studied. For example, a number of participants conflated the privacy protections offered by private browsing, VPNs, and Tor Browser with security protections. Therefore, it is vital that when attempting to increase adoption of privacy-enhancing technologies, special care should be taken to help people form accurate mental models. Based on the misconceptions we identified, we offer recommendations for the design of nudging interventions (§ 6.1) and for the design of privacy tools themselves (§ 6.2).

2 Related Work

Privacy and security advice has the potential to help people protect themselves from digital threats (§ 2.1). Such advice is most effective when it is informed by research into people’s mental models (§ 2.2) and their reasons for not already adopting recommended practices (§ 2.3).

2.1 Privacy and Security Advice

The privacy and security community is often eager to give the public advice about how to protect themselves from digital threats [14, 16, 41, 47, 76–78]. However, the amount of time people can dedicate to protecting themselves is limited [9, 37]. This has led to a growing consensus that advice must be prioritized and well-designed to be effective [6, 76, 77]. For example, advice should be relevant to recipients [36] and should address incorrect beliefs and other barriers to adoption [84]. Researchers have shown that well-designed interventions can elicit real world behavior change [3, 4, 49, 85]. Interventions that employ nudging techniques are particularly promising [2, 89]. Nudging interventions are designed to help people act in accordance with their true preferences. Since preferences for security and privacy are highly subjective, nudges encourage but do not enforce the adoption of certain behaviors. Nudges can take many forms, but our focus is on aiding the design of information-based nudges [2]. In particular, we seek to guide future research in this area by providing contextualized insights into participants’ beliefs about privacy tools and threats.

2.2 Mental Models

Research shows that experts and non-experts have different beliefs and behaviors about privacy and security [14, 29, 73]. Thus, experts should draw on research into non-experts’ mental models when they design interventions and tools for non-experts. Wash investigated home computer users’ mental models of viruses, hackers, and security protections [92]. He identified ways in which these folk models leave people vulnerable to botnets. Ion et al. compared experts’ and non-experts’ self-described most important security practices [41]. Two of the biggest differences were experts describing the importance of updating their systems, and non-experts reporting the importance of using antivirus software. Wash et al. show that different segments of the population have different beliefs about security, and suggest that targeted interventions may be most effective [93].

A number of studies have addressed mental models associated with privacy tools. Schaub et al. performed a usability evaluation of three tracker-blocking browser plugins [81]. They found that the plugins increased participants’ awareness of tracking, but did not help participants understand the implications of tracking. Habib et al. conducted a study of private browsing usage [35]. Some participants incorrectly believed that private browsing disabled cookies, allowed anonymous browsing of the web, and that it protected from malware. Dutkowska-Zuk et al. studied university students’ use of VPNs [23]. They found that 40% of participants used VPNs for security and privacy, and that about one-third of participants thought VPNs guaranteed privacy, anonymity, and safety from tracking. Gallagher et al. compared expert and non-expert beliefs about Tor [29]. They found that many non-experts believed that Tor provided protections it did not, such as hiding oneself even while logged in to a service. Similarly, we seek to identify tool-related misconceptions. However, previous work focused on individual tools, whereas we measure misconceptions across different combinations of tools and usage scenarios. This data allows us to identify patterns across tools and scenarios, and to quantify the prevalence of different kinds of misconceptions. In addition, participants’ explanations for their responses lend additional support for themes identified in prior work and reveal several previously unreported themes.

2.3 Adoption of Tools

One of our contributions is an estimate of people’s adoption of a series of privacy-enhancing tools. Zou et al. surveyed crowdworkers about their adoption and abandonment of 30 security and privacy practices [100]. The authors found that the most common reasons for abandoning a practice were perceptions that the practice was no longer needed, that associated risks had decreased, or that the practice was inconvenient. While Zou et al. focused on reasons for adoption and abandonment, our goal was to understand how adoption affects participants’ mental models of tools. Other studies have included measures of adoption of individual tools, including private browsing [22, 35], VPNs [23, 60], and Tor Browser [29]. Unlike those studies, our inclusion of multiple tools facilitates a relative comparison between levels of adoption of tools.

Kang et al. interviewed people about privacy and security risks, and identified reasons people don’t take privacy-protective actions [42]. Reasons included lack of concern, protective actions being too costly or difficult, and lack of knowledge. These factors are similar to components of protection motivation theory (i.e., threat appraisal and coping appraisal) [55, 79, 80], which have been used in effective computer security interventions [3, 85]. We recommend addressing these factors to encourage effective adoption of browsing privacy tools (§ 6.1).

3 Method

We gathered data using an online survey instrument with a demographically-stratified sample of US participants. We used Prolific’s “representative sample” option, which yields representative samples stratified across age, sex, and ethnicity, as compared to US Census data [70]. See Table 4 in the appendix for our participants’ demographics. Past studies have found that crowdworker participants are more internet-savvy than the general US population [75]. Thus, our findings about the usage of different tools might be considered an upper-bound for the general population.

Our survey included four parts. First, we asked participants questions about their general perceptions of online privacy. For example, we asked participants to estimate the likelihood of others observing their web browsing activity, and how concerned they would be if others observed their web browsing activity. Second, we

asked participants questions about the tools we studied, such as whether they had heard of or used each tool. Here we included a fake tool, PrivacyDog, to check for participants’ honesty. Third, we asked participants how effective they thought each tool would be at preventing different scenarios from happening. We asked each participant about six scenarios, which were randomly assigned from twelve total scenarios. See § 3.1 for more details about our selection of tools and scenarios, and how we asked these questions. Finally, we asked participants demographic questions, such as about their education and device usage patterns. Our survey instrument is included in § A.2 in the appendix. Our study was approved by Carnegie Mellon University’s IRB.

We conducted a pilot to test our survey instrument ($n = 20$). We determined the number of participants to recruit for our study by using a bootstrapped power analysis on our pilot data. We had several quantitative research questions, so we conducted multiple power analyses. We conducted power analyses for both Kruskal-Wallis tests and the associated post-hoc Dunn tests. Based on our power analysis of the post-hoc tests for whether self-rated knowledge about privacy tools is associated with answering assessment questions correctly, we decided to recruit 500 participants. This number gave us at least 95% power at $\alpha = 0.05$ for the research questions supported by our exploratory data analysis. In addition, since we randomized which tool-scenario combinations we asked participants to explain with free-text, it was important to get a sufficient number of free-text responses for each combination. A simulation showed that with 500 participants, we would be very likely (> 99% chance) to get at least 20 free-text responses for each combination.

Our goal was to compensate participants \$12 per hour. Based on our pilot, we estimated the survey to take 18 minutes, so we compensated participants \$3.60. We collected data in August 2020. In adherence to Prolific’s rules, we only rejected six participants who wrote low effort free-text responses [71]. This was our only criteria for excluding participants’ responses from our analyses. Since Prolific replaces rejected participants, our final sample contained 500 participants.

3.1 Tools and Assessment Scenarios

An important aspect of our study design was our selection of tools and assessment scenarios. We selected four privacy-centric browsing tools that have been discussed in the literature [41, 54, 77, 100] and which offer

a diverse set of protections: private browsing, VPNs, Tor Browser, and ad blockers.¹ These tools all broadly help people protect their privacy while browsing, but with varying effectiveness depending on the use case. Although we associate antivirus software with security more than privacy, we also included it because we were interested in whether participants would ascribe privacy protections to it. Each participant was asked about all tools, in a random order.

We designed our scenarios based on entities people might want to protect themselves from and information people might want to keep hidden [5, 74], focusing on realistic scenarios in which at least some of our tools would be effective. However, we intentionally included one scenario in which no tools were effective, to see how participants would respond. Each participant was shown six scenarios randomly selected from a total of twelve. See Figure 3 for a list of these scenarios.

Each scenario was introduced as a question in the form of: “When you browse the web, how effective are the tools below at preventing *advertisers from seeing the websites you visit?*” This was followed by a response matrix containing each of the tools and four answer options: “Unsure,” “Not at all effective,” “Somewhat effective,” and “Very effective.” After submitting their responses in the matrix, participants were asked to explain their answer for one randomly selected tool with a free-text response. We chose to ask about only one tool for each scenario in order to reduce participant fatigue.

Based on research literature and other resources, our team decided on realistic threat models for each scenario. We used these threat models to estimate the true effectiveness of each tool. In evaluating participants’ responses, we allowed them to slightly underestimate the effectiveness of a tool, but we counted any overestimate of a tool’s effectiveness as incorrect. We allowed slight underestimates of the effectiveness of tools because all tools have edge-cases in which they do not provide their optimal level of protection (e.g., if the tool is misused). For example, in our government observation scenario, we consider Tor Browser “Very effective” and VPNs “Somewhat effective.” If a participant indicated that Tor Browser was “Very effective” or “Somewhat effective,”

we counted that as correct, but we counted “Not at all effective” as incorrect. If a participant indicated VPNs as “Somewhat effective” or “Not at all effective,” we counted that as correct, but “Very effective” as incorrect. We counted “Unsure” answers as neither correct nor incorrect. We describe the threat models for each of our twelve assessment scenarios in the paragraphs below. We focus on explaining why certain tools offer some level of protection — tools which are not mentioned should be considered “Not at all effective.”

Preventing hackers from gaining access to your device. Consistent with experts’ advice [44], we suggest that the most realistic threats are from software downloaded and executed by users and from browser exploits [31, 32]. Software offers little protection against certain attacks (e.g., those using novel malware [30, 86] or legitimate software [11, 28]), but antivirus software and ad blockers can help in some cases [46]. For example, antivirus can block some malware from executing [86], and ad blockers can block fake download buttons [83] and potentially malvertising [17]. Thus, we consider these tools “Somewhat effective.”

Preventing online stores from misusing your credit card information. This is the only scenario in which none of the tools we listed provide any protection. The only way to prevent a store from misusing a person’s card information is to not give it to them in the first place, by either avoiding the merchant altogether or using a tokenized payment method like PayPal.

Preventing advertisers from seeing the websites you visit. Advertisers like Google, Facebook, and ComScore have visibility into many websites that people visit because of tracking scripts and other resources that websites choose to embed in their pages [59]. Advertisers can connect different web requests to the same user through cookies and browser fingerprinting [15, 57]. We categorize Tor Browser as the only “Very effective” tool, because it is designed to comprehensively resist fingerprinting [51]. In some cases, private browsing and ad blockers can reduce the amount of tracking taking place by erasing cookies and blocking scripts, respectively, but neither provide comprehensive protection [59]. Thus, we consider them “Somewhat effective.” Although a VPN can hide one’s IP address, which can be used for browser fingerprinting, it provides no protection against other methods of tracking, so we categorize it as “Not at all effective.”

Preventing advertisers from showing you targeted ads based on the websites you visit. The threat model for this scenario is the same as the other advertiser-related scenario, except that the goal is not to

¹ We asked about another tool in our survey: DuckDuckGo. Unfortunately, we did not clarify that we meant the search engine. This led to ambiguity in participants’ responses due to DuckDuckGo’s multiple products: search engine, browser, and browser plugin. As a result, we decided to exclude DuckDuckGo from our analysis.

avoid observation, but simply to avoid seeing targeted ads. Thus, we categorize ad blockers as “Very effective,” since they are capable of blocking many ads [59].

Preventing the websites you visit from seeing what physical location you are browsing from. Websites can see the general geographic location of visitors based on their IP addresses [95, 97]. Both VPNs and Tor Browser provide the ability to hide one’s IP address by passing traffic through another internet connection, so we categorize them as “Very effective” in this scenario.

Preventing your search engine from personalizing the search results you see based on the websites you visit. In this scenario, we assume that search result personalization is tied to a browser cookie, as described by Google in their description of search personalization [39, 96]. Private browsing and Tor Browser disassociate users from their cookies, so we consider those tools “Very effective” at preventing search personalization. We consider ad blockers to be “Somewhat effective,” because they can hide personalized ads from search results, but do not prevent personalization of non-ad results.

Preventing your internet service provider from seeing the websites you visit. An internet service provider (ISP) can observe all traffic that passes through their network. Although SSL/TLS can prevent the ISP from observing the exact pages visited, websites’ IP addresses are not hidden by SSL/TLS. In order to hide the websites visited, one must establish a secure connection to an intermediary, such as a VPN provider or the Tor network. Therefore, we only consider VPNs and Tor Browser “Very effective” in this scenario.

Preventing the government from seeing the websites you visit. In this scenario, we consider two threat models. In the first, the government issues subpoenas for data from internet companies, similar to the PRISM surveillance program [34, 98]. Thus, protection requires preventing one’s web requests from being associated with one’s identity. VPNs hide users’ IP addresses, but other sources of information can still identify users. Also, a VPN provider itself could be the subject of a subpoena, and despite some VPN providers claiming not to log user activity, many VPN providers are known to be untrustworthy [25, 33, 40, 48, 94]. In contrast, Tor Browser is designed for anonymity, though of course it is possible to compromise that anonymity (e.g., by logging into websites or through browser exploits [20]). Thus, we consider Tor Browser “Very effective” and VPNs “Somewhat effective” under this threat model. In the second threat model, the government can

both issue subpoenas to companies and can conduct a forensic analysis of one’s physical device, perhaps obtained by warrant. In this threat model, VPNs are “Not at all effective,” since physical access would allow the government to read one’s browser history. Since Tor Browser automatically erases browser history, we still consider it “Very effective.” In light of both threat models, we consider Tor Browser “Very effective” and VPNs “Somewhat effective” or “Not at all effective.”

Preventing friends or family with physical access to your device from seeing the websites you visit in your browser history. We explicitly described this scenario’s threat model by mentioning browser history in the text we showed participants. We did this because many disparate threat models are associated with physical access, from shoulder surfing to keyloggers. In this scenario, only private browsing and Tor Browser are “Very effective,” because they are the only tools which erase browser history.

Preventing your employer from seeing the websites you visit on your personal device while connected to your work’s WiFi. We adopt the same threat model for this scenario as for our ISP scenario, in which both VPNs and Tor Browser are “Very effective” at preventing observation.

Preventing law enforcement from seeing the websites you visit. We adopt the same threat models for this scenario as for our government observation scenario. As in that scenario, overall Tor Browser is “Very effective” at preventing observation, and VPNs are either “Somewhat effective” or “Not at all effective,” depending on whether law enforcement has physical access to one’s device.

Preventing companies who own movies from seeing if you illegally stream a movie. In practice, the operators of streaming websites are the ones targeted by lawsuits. However, illegally streaming movies can be classified as a misdemeanor [12, 88], so rights-holders could prosecute those who use illegal streaming websites. Similar to the government and law enforcement scenarios, in this scenario we assume that movie rights-holders have the ability to subpoena information from websites and companies. Tor Browser would be “Very effective” at hiding one’s identity, but the efficacy of VPNs would depend on their logging practices, which are impossible to verify, so we consider them only “Somewhat effective.”

3.2 Thematic Analysis

To better understand participants' misconceptions about the tools, we asked participants to explain their answer for one randomly selected tool in each scenario. We used thematic coding to analyze these free-text responses. Since each participant was shown six scenarios, we collected 2,500 free-text responses in total.²

We used a two-pass coding process. In the first pass, the annotators reviewed the free-text responses associated with “Correct” and “Unsure” answers to the multiple-choice assessment scenario questions. The annotators marked whether these free-text responses contained any form of misconception. Our intuition was that responses associated with “Correct” and “Unsure” answers would contain few misconceptions; since we wanted to analyze misconceptions in more detail, this approach allowed us to identify relevant instances for thematic analysis in our second pass. To ensure high quality, two annotators performed this task and reconciled their codes after completing them for each tool. Our intuitions were confirmed, as we found that only 17% of the “Correct” and 13% of the “Unsure” responses contained misconceptions. This greatly reduced the size of our second pass coding task.

After reaching consensus on the first pass coding for a given tool, the lead annotator began the process of second pass coding. For each tool, the lead annotator first coded the “Incorrect” responses as containing misconceptions or not. Next, the lead annotator reviewed all the responses containing misconceptions and created thematic codes. The codebook was finalized after this process was completed for all tools. Our completed codebook contains 23 thematic codes. Using the completed codebook, two annotators independently coded the responses for each tool and then reconciled codes. See Tables 5 and 6 in the appendix for detailed descriptions of our first and second pass codes, respectively.

The annotators eliminated 26 low-quality answers (e.g., unintelligible, clearly about the wrong scenario, etc.) as they encountered them.

Due to the complex interactions between tools and scenarios, our process of reaching consensus was necessarily a collaborative one. For example, several cases arose in which one annotator was unaware of a specific tool behavior. We identified such cases while reconciling codes and researched literature and documentation

to determine whether a response contained a misconception or not. For this reason, we consider measures of annotator reliability to be inappropriate [58]. Because two expert annotators read each response and reached consensus on any differences, we have high confidence in the quality of our data.

4 Results

First, we supply descriptive statistics about participants' general privacy perceptions (§ 4.1), their adoption of browsing-related tools (§ 4.2), and their responses to our assessment scenario questions (Sections 4.3 and 4.4). Next, we convey the results of two exploratory statistical analyses: an analysis of factors associated with correctly identifying the protections offered by the tools (§ 4.5), and an analysis of demographic factors associated with tool use (§ 4.6) Finally, we describe the results of our thematic analysis of participants' misconceptions (§ 4.7).

4.1 General Privacy Perceptions

To measure participants' general perceptions of online privacy, we began by asking broad questions. In response, most participants indicated that their web browsing activity was likely to be observed by others (62%), and most participants were at least slightly concerned about this (83%). Also, most participants thought they knew how to use privacy tools (72%), yet nearly all participants still expressed at least slight interest in learning how to use tools to protect their privacy (96%). These responses suggest that some participants would be receptive to learning how to use privacy-enhancing browsing tools. Our findings are in line with the Pew survey “Americans and Privacy” [5].

4.2 Tool Adoption

Our first tool-specific questions asked whether participants had heard of each tool, and if so, whether they had used the tool before. As shown in Figure 1, nearly all participants had used antivirus software, but less than half had even heard of Tor Browser. Note that only 3% of participants said they had heard of or used PrivacyDog, a fake tool, giving us high confidence in the numbers for the other tools.

² We initially collected 3,000 responses, before eliminating the responses about DuckDuckGo.




Fig. 1. Percentage of participants who reported having used or heard of each tool. Our questions included “Yes,” “No,” and “Unsure” options. For this graph, we grouped “Unsure” and “No” answers together (e.g., if the participant indicated they were unsure whether they’d heard of a tool, we counted them as having not heard of it).




Fig. 2. For each tool, we asked participants who said they had used it before when they had most recently used it. “[Never]” responses belong to participants who were not shown the question because they reported having never used the tool.

We also asked tool users when they had last used each tool. As shown in Figure 2, some of these tools are already widely used, especially antivirus software and ad blockers. Furthermore, 59% of participants had used at least one of the privacy-focused tools in the past day (i.e., a tool other than antivirus software), and 74% had used at least one of the privacy-focused tools in the past week. We see this as further evidence that there is widespread interest in privacy-enhancing tools.

4.3 Interest in Assessment Scenarios

We asked each participant questions about six randomly selected scenarios (from a total of twelve). Figure 3 shows participants’ expressed interest in each scenario. For all scenarios, over half of participants expressed some interest in preventing it. However, participants’ level of interest varied considerably between scenarios. First, it is interesting that the two more security-focused scenarios about hackers and card fraud were of greatest interest to participants. However, participants also showed strong interest in preventing the two advertiser-related scenarios. The difference in wanting to prevent the government and law enforcement observation sce-

narios is notable; 82% of participants had some interest in preventing the government from seeing the websites they visit, while only 67% of participants were interested in preventing law enforcement from doing the same.

4.4 Assessment Scenario Correctness

We asked participants to rate how effective they thought each tool would be at preventing each of the six scenarios they were shown. We evaluated participants’ responses based on the threat models we described in § 3.1. In Sections 4.4.1 and 4.4.2, we explain participants’ responses for two scenarios in detail. For details about the remaining ten scenarios, see Figure 7 in the appendix. In § 4.4.3, we summarize participants’ responses across scenarios and tools.

4.4.1 Preventing Hackers from Gaining Access to Your Device

Of the twelve scenarios we asked about, participants indicated that they were most interested in preventing hackers from gaining access to their device. As shown in Figure 4, many participants incorrectly evaluated the tools’ security protections in this scenario. Notably, more than half of participants thought that VPNs would prevent hackers from gaining access to their device.

4.4.2 Preventing the Websites You Visit from Seeing What Physical Location You Are Browsing From

Many participants were also interested in preventing websites from seeing their physical location. As shown in Figure 4, 76% of participants successfully identified that VPNs can provide this protection, but only 36% recognized that Tor Browser provides this protection as well. This discrepancy may be partly explained by participants’ greater familiarity with VPNs. To determine whether this was the case, we tested for an association between participants’ experience with tools and the correctness of their answers (§ 4.5).

4.4.3 Summary of Response Correctness

Figures 5 and 6 show significant numbers of unsure and incorrect responses across tools and scenarios. Figure 5

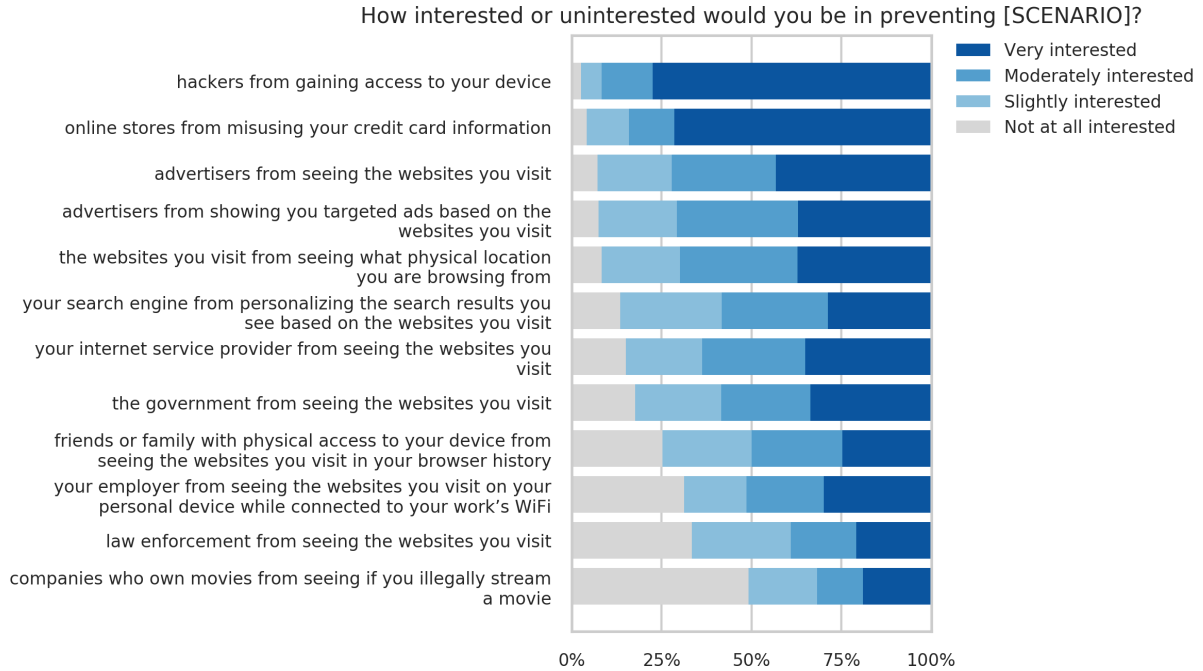


Fig. 3. Participants' interest in preventing each scenario, sorted by the percent of "Not at all interested" responses. Note that each participant was shown six randomly selected scenarios, so percentages are calculated for the participants who did see a given scenario.




Fig. 4. Responses consistent with our threat model are indicated with a star. Tools are sorted by the percent of correct responses.




Fig. 5. The correctness of participants' responses to the scenario-based assessment questions about each tool.

shows that participants answered more questions correctly for tools that are more widely adopted. To explore this apparent relationship, in § 4.5 we test for an association between participants' experience and their answers' correctness. Also, as shown in Figure 6, for all but one scenario participants answered fewer than half of the assessment questions correctly.

4.5 Experience with a Tool Is Not Necessarily Associated with an Accurate Understanding of It

We were interested in whether participants' level of experience with each tool was associated with their ability to answer questions about each tool correctly. For ex-

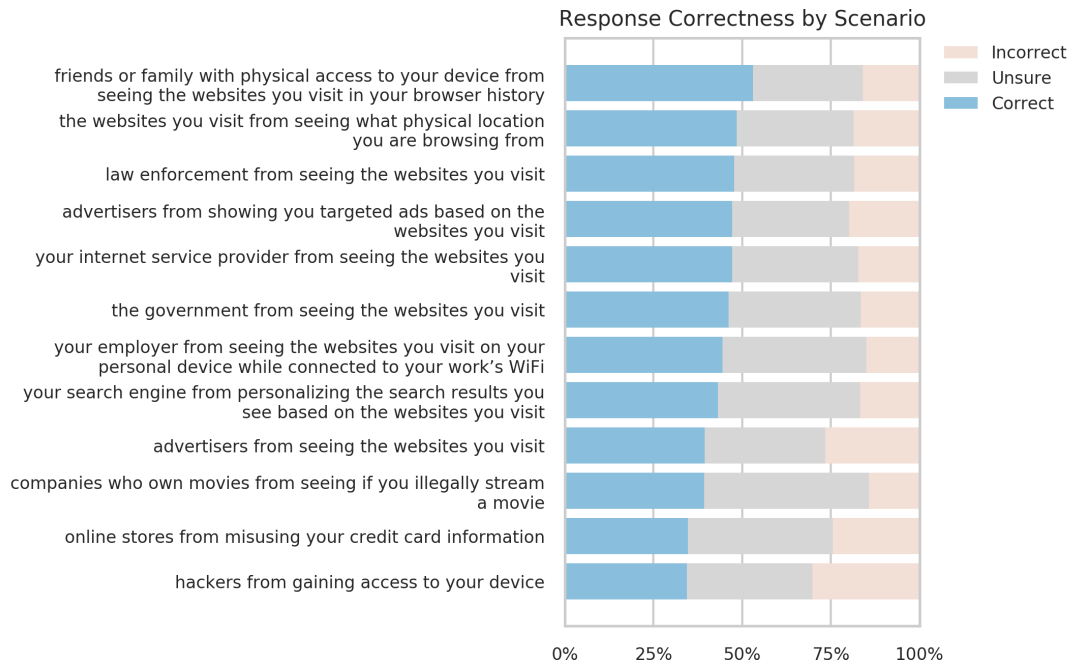


Fig. 6. The correctness of participants' responses to assessment questions about each scenario. Note that each participant was shown six randomly selected scenarios, so percentages are calculated for the participants who did see a given scenario.

ample, are those who have used private browsing more likely to answer questions about private browsing correctly? Ideally, users of a tool would have an accurate understanding of the protections it provides, which would help them use the tool appropriately. We tested for these associations using Kruskal-Wallis tests between level of experience (i.e., have used, haven't used, or haven't heard of the tool) and number of correct answers about the tool. As shown in Table 1, we found statistically significant evidence of an association for all tools at $\alpha = 0.05$. Holm corrected Dunn post-hoc tests show that greater levels of experience are typically associated with answering more questions correctly.

This is an intuitive finding, but when we dug deeper we found something surprising. We conducted similar tests for associations between level of experience and number of incorrect responses, number of unsure responses, and scores (i.e., correct minus incorrect). We found that for VPNs and Tor Browser, greater levels of experience were generally associated with greater numbers of *incorrect* responses. This may be partly due to the tendency of those with greater levels of experience to mark fewer responses as “Unsure.” Subtracting the number of incorrect responses from the number of correct responses to calculate “scores,” we see that those with greater levels of experience only have statistically significantly higher scores for private browsing, Tor Browser, and ad blockers.

Tool	Experience	Mean (Kruskal-Wallis p-value)			
		Correct	Incorrect	Unsure	Score
Private browsing	Have used	2.95	1.70	1.35	1.25
	Haven't used	1.81	1.65	2.54	0.16
	Haven't heard of	0.92	1.58	3.50	-0.66
		(<0.001)	(0.657)	(<0.001)	(<0.001)
VPNs	Have used	2.58	2.25	1.17	0.32
	Haven't used	1.67	1.36	2.98	0.31
	Haven't heard of	0.90	0.34	4.76	0.56
		(<0.001)	(<0.001)	(<0.001)	(0.765)
Tor Browser	Have used	4.26	0.79	0.95	3.47
	Haven't used	2.40	0.67	2.93	1.73
	Haven't heard of	0.53	0.49	4.98	0.04
		(<0.001)	(<0.001)	(<0.001)	(<0.001)
Ad blockers	Have used	3.80	0.81	1.39	2.98
	Haven't used	2.72	0.58	2.71	2.14
	Haven't heard of	2.00	0.82	3.18	1.18
		(<0.001)	(0.182)	(<0.001)	(<0.001)
Antivirus software	Have used	3.37	1.09	1.54	2.28
	Haven't used	2.04	0.71	3.25	1.33
	Haven't heard of	2.50	3.00	0.50	-0.50
		(0.018)	(0.122)	(0.001)	(0.197)

Table 1. The mean number of correct, incorrect, etc. responses by experience with each tool. Bolded cells indicate Kruskal-Wallis tests significant at $\alpha = 0.05$, with p-values shown in parentheses. For example, tool experience was shown to be associated with the number of correct responses to questions about private browsing, so that cell is bolded; we did not find an association between experience and the number of incorrect responses to questions about private browsing, so that cell isn't bolded. Due to space constraints, post-hoc test significance is not shown.

Variable	Model									
	Private browsing		VPNs		Tor Browser		Ad blockers		Antivirus software	
	p-value	e^β	p-value	e^β	p-value	e^β	p-value	e^β	p-value	e^β
Age	<0.001	0.951	0.109	0.986	0.156	0.984	0.322	0.990	0.780	1.006
Non-female	0.008	1.898	<0.001	2.510	<0.001	3.810	0.014	1.837	0.709	1.183
Income: \$10,000 - \$19,999	0.176	0.349	0.824	1.143	0.433	1.750	0.751	0.784	0.998	<0.001
Income: \$20,000 - \$39,999	0.515	0.618	0.474	0.679	0.861	1.120	0.352	0.522	0.998	<0.001
Income: \$40,000 - \$59,999	0.318	0.465	0.304	0.556	0.232	0.413	0.726	1.301	0.998	<0.001
Income: \$60,000 - \$79,999	0.974	1.027	0.900	0.929	0.793	0.830	0.678	1.380	0.998	<0.001
Income: \$80,000 - \$99,999	0.174	0.327	0.949	0.960	0.531	1.600	0.429	0.537	0.998	<0.001
Income: \$100,000 or more	0.395	0.507	0.919	1.062	0.720	1.292	0.924	1.077	0.998	<0.001
Employment: Self-employed	0.424	1.320	0.166	0.666	0.927	0.967	0.539	1.246	0.648	0.740
Employment: Student	0.483	0.647	0.965	1.021	0.584	0.725	0.901	1.077	0.061	0.204
Employment: Not employed	0.684	1.177	0.165	0.636	0.559	1.274	0.343	1.464	0.590	1.595
Employment: Retired	0.685	0.857	0.065	0.503	0.977	0.985	0.985	1.008	0.301	0.443
Education: College or associate degree	0.400	0.769	0.122	1.496	0.660	0.865	0.752	1.102	0.299	1.736
Education: Graduate degree	0.935	0.969	0.050	1.898	0.801	0.902	0.695	0.862	0.676	1.322
Computer-related background	0.015	2.000	0.010	1.814	0.023	1.832	0.068	1.707	0.121	2.737
Living with: Domestic partner	0.036	1.868	0.419	1.237	0.860	0.941	0.277	0.717	0.724	0.810
Living with: Children	0.253	0.733	0.647	0.898	0.512	1.219	0.698	1.113	0.978	0.985
Living with: Parents	0.184	1.838	0.646	1.167	0.415	1.362	0.881	0.940	0.631	1.481
Living with: Other family	0.211	0.599	0.616	1.188	0.695	0.849	0.511	0.765	0.361	0.540
Living with: Roommates	0.621	1.405	0.200	1.891	0.210	2.024	0.692	0.782	0.967	1.053
Intercept	<0.001	34.420	0.961	0.968	0.031	0.158	0.048	5.377	0.998	1.402E+9

Table 2. The variables in our regression models for predicting use of each tool. e^β indicates the change in odds of using the tool for a one unit change in the variable (or when the variable is true). p-values significant at $\alpha = 0.05$ are bolded.

Model	Cox & Snell R^2
Private browsing	0.171
VPNs	0.145
Tor Browser	0.109
Ad blockers	0.059
Antivirus software	0.043

Table 3. The R^2 values for each of the models in Table 2. R^2 represents the proportion of variance in tool use explained by each of our models.

We performed additional statistical tests to identify associations between self-rated tool knowledge (“I think I know how to use private browsing.”) and participants’ answers, and between having a computer-related background and participants’ answers. In nearly all cases, the significance and direction of our findings were consistent with our analysis of tool experience. For example, we found positive relationships between self-rated tool knowledge and number of correct responses ($p < 0.001$), and between having a computer-related background and number of correct responses ($p = 0.028$). Our only difference in findings was for the association between computer-related background and score, for which we found statistically significant positive relationships only for ad blockers and antivirus software; for tool experience and self-rated knowledge, we found statistically

significant positive relationships for private browsing, Tor Browser, and ad blockers.

Our results suggest that participants who have more experience with tools, who think themselves more knowledgeable about tools, or who have computer-related backgrounds, are more willing to definitively answer questions about the tools. However, these factors are necessarily associated with a more accurate understanding of the tools’ protections.

4.6 Age, Gender, Computer-related Background, and Living Situation Are Associated with Use of Tools

We were interested in how demographic factors like age and education were associated with the use of each tool, so we trained logistic regression models to predict the use of each tool. Our models contain the following seven variables: age, gender (“Female” as baseline), income (“Less than \$10,000” as baseline), employment (“Working (paid employee)” as baseline), education (high school or less as baseline), computer-related background, and living situation (living alone as baseline). We excluded 27 participants who declined to answer questions about income, employment, education, or living situation, leaving us with 473 participants to

train our models. We checked for multicollinearity, and all VIFs were less than 10. We also performed Hosmer-Lemeshow goodness of fit tests for each model, and did not find evidence of poor model fit at $\alpha = 0.05$. Table 2 shows the significance of each model’s coefficients, and Table 3 shows the explanatory power of each model. The lower number of significant variables in our ad blocker and antivirus software models may be due to the broader adoption of these tools (Figure 1). This may also explain the comparatively poor explanatory power of these two models.

Two factors are significant in multiple models. First, our models show that non-female participants are significantly more likely to use private browsing, VPNs, Tor Browser, and ad blockers. For example, our model predicts that non-female participants are 3.8 times more likely to use Tor Browser than female participants, all else being equal. Second, we see that participants with computer-related backgrounds are significantly more likely to use private browsing, VPNs, and Tor Browser. Finally, we see two factors which are only significant for private browsing: the model shows that older participants are less likely to use private browsing, and that those who live with a domestic partner are more likely to use private browsing.

The associations for age, gender, and computer-related background are consistent with the findings of prior work [22, 35, 100], but we are unaware of prior work showing a positive association between living situation and use of privacy-enhancing technologies [82]. The existence of this association makes sense, since participants may want to hide their browsing activity from their partner.

4.7 Thematic Analysis of Misconceptions

We asked participants to explain their responses to our assessment scenarios, and we performed a thematic analysis of these explanations to identify misconception-related themes (§ 3.2). Some themes were associated with particular scenarios (§ 4.7.2) or tools (§ 4.7.3), but others were common across scenarios and tools (§ 4.7.1). Note that we collected 500 free-text responses per tool, and an average of 208 responses per scenario. Based on the misconceptions we discovered, we offer recommendations for the design of nudging interventions (§ 6.1) and privacy tools (§ 6.2).

4.7.1 General Themes

Partial Knowledge, but Incorrect Responses

We collected 501 explanations of participants’ incorrect responses. Participants cited true aspects of tool functionality in 184 of these explanations. For example, P330 indicated that VPNs would be “Very effective” at preventing advertisers from seeing the websites they visit because “VPNs mask one’s IP address...” and P215 indicated that private browsing would be “Very effective” at preventing their employer from seeing browsing done on their employer’s WiFi because “Private browsing does not keep your history...”. We observed similar responses across all tools and scenarios. These responses show that participants know something about the tools, but their knowledge does not prevent them from reaching incorrect conclusions about the protections offered by the tools. This may be due to incomplete mental models about the tools and scenarios.

Resignation

Another theme prevalent across tools and scenarios was that of resignation. Participants frequently wrote that nothing could be done to protect against an entity, or that the entity’s resources were overwhelmingly powerful. We identified this theme 154 times overall, and 92 times in the government and law enforcement observation scenarios. For example, P383 wrote that Tor Browser would be “Not at all effective” at preventing observation by the government because “If the government wants to see what you are doing, they will see it no matter what.” Similarly, P499 wrote that VPNs would be “Not at all effective” at preventing observation by their ISP because “I believe my internet provider can already see everything I do no matter what.” Privacy resignation has been observed in diverse contexts [18, 53, 56, 99], but it is especially striking to see it when effective tools are available, as they are in all but one of our scenarios.

Overconfidence

A final theme prevalent across tools and scenarios was that of overconfidence in tools’ protections. We identified this theme when participants wrote that tools provided total protection or anonymity even though they do not. We observed this theme 69 times overall, across all tools and all scenarios except for observation by friends or family. For example, P312 wrote that antivirus software would be “Very effective” at preventing hackers from gaining access to their device because

“It help prevent any form of virus which might come and affect my data.” Also, P128 wrote that ad blockers would be “Somewhat effective” at preventing the websites they visit from seeing their physical location because “Ad blockers will shield you from getting your information harvested.” Prior work has shown that offering people control over information disclosure can increase people’s willingness to share sensitive information [10]. We worry that overconfidence in tools’ protections will likewise lead users to expose themselves to privacy harms.

4.7.2 Scenario-Related Themes

Conflating Privacy and Security Protections

Among our two security-focused scenarios, we observed 23 instances of participants conflating the privacy protections offered by private browsing, VPNs, and Tor Browser with security protections. In their answers, participants described trying to stay safe from hackers or card fraud by avoiding being noticed or by keeping information hidden. For example, P34 wrote that private browsing would be “Somewhat effective” at preventing hackers from gaining access to their device because “It should make your device hard to find by hackers,” and P158 wrote that VPNs would be “Very effective” at the same because “It is a virtual network that keeps others from your device. Done well, hackers can’t find you.” With respect to preventing online stores from misusing one’s credit card information, P127 wrote that private browsing would be “Very effective” because “[it] allows the user to be undercover and out of reach of basic credit card hackers at online stores,” and P168 wrote that Tor Browser would be “Very effective” because “It would reroute your viewing traffic so they could not see. Might be able to mask it with a different number.” People may conflate privacy and security because they are related concepts, but it is important for them to understand that privacy protections do not necessarily imply security protections. Otherwise, people might expose themselves to undue risk [1, 29].

Citing “Layers” to Justify Incorrect Responses

We observed 13 cases in which participants used language about layers of protection to justify their incorrect responses. Nine of these instances were associated with our hacker-related scenario, and all were associated with either VPNs, Tor Browser, or antivirus software. For example, P268 wrote that Tor Browser would

be “Very effective” at preventing hackers from gaining access to their device “because the onion router is so deep and layered with basic protection it can’t be used to maliciously hack” and P408 indicated the same for VPNs because “... VPN’s give you an extra layer of security that they’d have to hack through.” The security concept of “defense in depth” refers to using multiple protections in case one fails [8], and might be the origin of these references to layers of protection. However, achieving greater protection through layering multiple technologies requires a careful analysis of threat models; it is possible to actually decrease one’s level of protection when using certain technologies together [90, 91]. Thus, the concept of defense in depth might be ultimately misleading for non-expert users.

Referencing Location Permissions

In our scenario about preventing websites from seeing the physical location one is browsing from, we observed five references to location API permissions. Participants explained that “... usually sites ask for your location to be accessed” (P96), “... I do not have location turned on on any devices except certain apps ...” (P325), and that “Location is often a setting on the site, browser, or app that needs to be turned off. I thing the software notifies you if it was accessed but does not prevent it” (P33). These participants seem to assume that websites can only determine their location if websites access it through the location API, possibly revealing unawareness of IP-based location inference.

4.7.3 Tool-Related Themes

Finally, we discuss themes that were associated with particular tools. We collected one free-text response about each tool from each participant, giving us 500 responses for each tool.

Citing Tools’ Names to Justify Incorrect Responses

When answering questions about private browsing and VPNs, a number of participants cited the tools’ names to justify their incorrect responses.

Of the 148 participants who explained their incorrect responses about private browsing, 22 referenced the name of the tool in their explanations. For example, P491 answered that private browsing would be “Somewhat effective” at preventing the government from seeing the websites they visited, explaining that “the name

‘private browsing’ would suggest so.” P389 thought private browsing would prevent websites from seeing their location, writing that “I thought in private browsing you’re incognito which means no one knows what your doing or where you are.” This supports others’ findings that the name “private browsing” can lead users to overestimate its protections [1].

We collected 138 explanations for incorrect responses about VPNs; similar to private browsing, in 16 cases participants referenced the name of the tool in their explanations. For example, P97 indicated that VPNs would be “Very effective” at preventing friends or family from seeing the websites in their browser history because “You have your own private network that others cannot get into.” P383 answered that VPNs would be “Somewhat effective” at preventing advertisers from seeing the websites they visit because “It is a private network, so what you browse is private in the outside.” Also, two participants misunderstood the abbreviation VPN, writing that VPN stands for “V=Very P=Private N=Network” (P257) and “Virtual Processing Networks” (P490). Answers like these suggest that the name “VPN” may be uninformative or misleading.

Tor Browser Is for the Dark Web and File Sharing

Of the 500 free-text responses about Tor Browser, we coded 137 as containing misconceptions. Among these responses, we identified 15 references to the dark web. Some participants seem to believe that Tor Browser is exclusively for use with the dark web: “I thought Tor was just for browsing the darkweb” (P103), “... it is a browser used for illegal activities ...” (P254), “... it is a browser connected with the Dark Web that is hard to use unless you know exactly how to do it or have some type of password that allows you to use it” (P147). Perhaps these beliefs are due to media coverage associating Tor with illegal activity [19, 38, 52].

We also identified four participants who made a connection between Tor Browser and file sharing. For example, P382 wrote that “I know nothing about TOR other than it is Torrent” and P378 wrote that “... Tor Browser was designed from the ground up for very high point-to-point browsing. (The more I think about it, I’m pretty sure I’ve used this a decade or more ago to download large music files.)”. Although these participants didn’t explicitly point to Tor’s name as their reason for making this connection with BitTorrent, the similarity of the words “Tor” and “torrent” suggest that Tor’s name may explain this connection.

If people think Tor Browser is only for illegal activities or torrent downloads, they might think it is less relevant to them, and this might be a potential barrier to adoption [68].

Tor Browser Should Be Used with a VPN

Three participants suggested that Tor Browser is most effective when used with a VPN. P109, who had used Tor Browser before, wrote “... it’s only completely ‘safe’ if you also use a VPN or have it configured to use a proxy, since your data still goes through your ISP ...” and P300, who hadn’t used it before, wrote “... stories of using Tor usually [recommend] that you have a VPN or something to mask where you are coming from.” Such claims are frequently present in content advertising VPNs [27, 45, 72]. However, experts caution that combining Tor with a VPN can either increase or decrease one’s privacy protections, depending on one’s threat model [90, 91]. Those who think that Tor Browser requires a VPN to be fully effective may perceive adopting Tor Browser to be more challenging than it is in reality. Thus, correcting this misconception may lower a barrier to the adoption of Tor Browser.

Ad Blockers Hide Browser History

We asked 40 participants to explain whether ad blockers would prevent those with physical access to their device from seeing the websites in their browser history. In response, six participants indicated that because ad blockers can block personalized ads, they would be “Somewhat effective” or “Very effective.” For example, P306 wrote that ad blockers “... will stop your family members from seeing ads that were personalized for you.” Note that we intentionally phrased this scenario to draw participants’ attention to the “browser history” function. Although an ad blocker may hide some signs of one’s browsing history, it will do nothing to prevent other users of the computer from viewing the browser history itself, or from seeing other signs of browsing history, like search autocomplete.

Citing Experience to Justify Incorrect Beliefs

Several participants cited their experience with ad blockers and antivirus software when explaining incorrect beliefs they held about those tools.

Interestingly, three participants wrote that ad blockers would be “Not at all effective” at blocking targeted ads because they still saw ads despite using an ad

blocker. Although the efficacy of ad blockers varies [59], we doubt that the ad blockers these participants used were completely ineffective. Instead, perhaps a lack of visual feedback when ads were blocked led these participants to doubt their ad blockers were working.

Three participants incorrectly indicated that antivirus software would be “Very effective” at preventing three different scenarios because they had not yet suffered adverse consequences. For example, P72 claimed that antivirus software would prevent law enforcement from seeing the websites they visited “Because I have never had any indication that law enforcement has been on my computer in 20 years of computer use with the antivirus system I have used.” Similarly, P481 explained that antivirus software would prevent websites from misusing their card information because “This software had been set up for awhile. Looks like nothing goes wrong.” These participants seem to attribute their lack of negative experiences to their use of antivirus software, when external factors are a more likely explanation (e.g., law enforcement not viewing one’s browsing activity because one is not under investigation).

Antivirus Software Blocks Malicious Ads

We asked 54 participants to explain whether antivirus software would prevent advertisers from showing them targeted ads. In response, four participants wrote that antivirus software would specifically block malicious ads. For example, P472 wrote that “...some ads do carry viruses and so I guess this software would block them.” However, we are unaware of any antivirus software that claims to distinguish between regular ads and malvertising; it is concerning that participants thought that antivirus software offered this functionality, since that may lead them to take unnecessary risks.

5 Limitations

Our study is subject to various limitations.

First, our use of the Prolific platform for recruitment means that our participants are not completely representative of the general public. Prolific participants differ from the general public in some obvious ways (e.g., all have access to the internet) and in more nuanced ways [67, 75]. We attempted to mitigate this limitation by collecting a demographically-stratified sample of US participants using Prolific’s “representative sample” feature, similar to the recommendation of Redmiles

et al. [75]. However, we still expect our participants to be more internet-savvy than the general public, so our findings about the usage of different tools might be considered an upper-bound for the general population.

Second, since we relied on self-reported behavior, participants’ responses may be biased [50]. We checked for participants’ honesty by asking whether they had heard of or used a fake tool, PrivacyDog. Only 3% of participants said they had heard of or used PrivacyDog, which suggests that our participants’ were generally honest.

Finally, our choice of threat models for our assessment scenarios represents a possible threat to the validity of our study. We based our threat models on published research, news stories, and our own knowledge as security experts. Notably, we based our threat models on the technologies underpinning the tools, but a potential confounding factor is that some security and privacy products bundle multiple technologies. For example, while NordVPN functions primarily as a VPN, it also includes an optional feature, CyberSec [62]. NordVPN advertises that this feature performs the functions of ad blockers and antivirus software, though we are unaware of independent evaluation of its efficacy. Similarly, while Norton offers traditional antivirus software, they also offer Norton Secure VPN, which in addition to functioning as a VPN is also advertised as “block[ing] unwanted tracking technologies” [66]. Norton also offers Norton Privacy Manager, which among other features blocks ads and trackers, includes a VPN, and includes a privacy-friendly search engine [63, 65]. Thus, participants might answer based on their familiarity with these bundled products, rather the component technologies. When counting the number of correct answers, we choose not to count these bundled functionalities as correct (e.g., not to assume that antivirus functions as a VPN). As we describe in the appendix (§ A.1), our data suggest that most incorrect answers were based on inappropriate mental models, rather than on an awareness of these bundled products.

6 Discussion

In our survey, we asked participants about the protections offered by five different tools in twelve realistic scenarios. The substantial number of incorrect and unsure responses across tools and scenarios (§ 4.4.3) shows that misconceptions are widespread. In addition, our qualitative analysis of participants’ free-text responses

characterizes the diverse ways in which misconceptions are expressed (§ 4.7). Our participants’ misconceptions are cause for concern. For example, if a person mistakenly believes that a tool offers protections that it does not provide in actuality, that person may take unnecessary risks under the belief that the tool is protecting them, and may thereby expose themselves to privacy harms (e.g., unwanted observation). Conversely, if a person doesn’t believe in the protections that tools can actually offer, that person may engage in unnecessary self-censorship to avoid privacy harms. However, our data suggest that people are receptive to learning more about how to protect their privacy (§ 4.1), showing a need for effective interventions to help them protect themselves. Informed by our results, we offer design recommendations for nudging interventions and for the design of privacy tools.

6.1 Recommendations for Designing Nudging Interventions

When designing nudging interventions to encourage the adoption of privacy tools, we suggest that designers adhere to the following recommendations.

First, we recommend that interventions **focus on helping people protect themselves from well-defined threats**. One of the most common themes we observed was participants answering incorrectly despite demonstrating partial knowledge of a tool. Perhaps these participants’ partial knowledge made them confident enough to choose an answer, rather than selecting “Unsure.” We worry that partial knowledge could also lead to inadvertent risk-taking, when a person thinks a tool provides a protection it does not. It seems unrealistic to expect people to make accurate judgments about the protections offered by tools, as doing so would require in-depth technical knowledge. Therefore, we think interventions should warn people not to assume that tools provide protections from threats other than those described in the intervention. Also, some participants seemed to conflate privacy and security concerns, assuming they would be protected from security threats if they browsed anonymously. Therefore, it seems especially prudent to remind people that privacy-focused tools like Tor Browser provide no additional security guarantees (e.g., against malware). To inform the choice of which threats to focus on, we recommend that researchers consult our data on participants’ relative interest in protecting against different threats (Figure 3).

Second, we recommend that interventions address the components of protection motivation theory (PMT), which has informed the design of other effective interventions in the computer security domain [3, 85]. Three relevant components of PMT are perceived *threat susceptibility*, *response efficacy*, and *self-efficacy*. A person’s perception of their threat susceptibility is how likely they think they are to be affected by a given threat (e.g., to be tracked by advertisers). A person’s perception of response efficacy is their belief that the suggested response will protect against the threat (e.g., that using a privacy-enhancing technology will prevent them from being tracked by advertisers). Finally, a person’s perception of self-efficacy is their belief that they will be able to perform the suggested response successfully (e.g., that it will be easy for them to adopt the recommended technology). PMT suggests that people’s motivation to act is influenced by these components. Themes from our participants’ qualitative responses suggest opportunities for helping people form realistic perceptions of threat susceptibility, response efficacy, and self-efficacy.

One common theme was participants expressing that nothing could be done to protect themselves from a given threat (i.e., resignation [18, 53, 56, 99]), which may be associated with a low perception of response efficacy. For example, many of our participants suggested that information could not be hidden from the government or law enforcement. People may not be aware of or believe in the privacy protections that tools can provide against these and other entities. Thus, it might be helpful to **reassure participants of the efficacy of the tool or action being promoted, in order to bolster their perception of response efficacy**. For example, describing the complexity of law enforcement operations against Tor users might reassure people of the protections provided by using Tor Browser [24, 69]; if gaining access to data about Tor users were as simple as issuing a subpoena, law enforcement would have had an easier time shutting down sites like Silk Road [7, 26, 87].

Our participants often misattributed protections to tools. This may correspond to a low perception of threat susceptibility, especially when participants are already using those tools. For example, 42% of participants thought private browsing would prevent websites from seeing their physical location. As another example, some responses suggested that location could only be accessed through the browser location API, rather than inferred from IP address. Misconceptions like these may cause participants to think they are already protected from threats by their existing behavior. Therefore, we think interventions will be more effective if they **emphasize**

the lack of effectiveness of other tools and practices, in order to increase people’s perception of threat susceptibility.

For Tor Browser in particular, we identified several impediments to an accurate perception of self-efficacy. First, some responses suggested that Tor Browser was primarily for accessing the dark web, and one participant thought that users might even need “some type of password that allows you to use it” (P147). These participants might be surprised to learn that Tor Browser can be used like a regular browser to visit ordinary websites, and that it does not require any special credentials or advanced skills. Second, several participants incorrectly thought that Tor Browser needed to be used with a VPN in order to be fully effective. However, it is not necessary to use a VPN to achieve anonymity with Tor Browser. These types of misconceptions portray Tor Browser as difficult to use, which may lead people to think that it would be too difficult for them to use Tor Browser successfully. People should be made aware of the real challenges associated with using Tor Browser (i.e., increased latency), but these misconceptions should not discourage them from trying to use it. Thus, we recommend that interventions **debunk misconceptions which may contribute to a decreased sense of self-efficacy.**

6.2 Recommendations for Designing Privacy Tools

Our design recommendations for nudging interventions also apply to the marketing of privacy tools. Although it might be possible to exaggerate the effectiveness of a tool, responsible marketing should attempt to convey accurate perceptions by following the recommendations we outlined above. In addition, we have several recommendations specifically for tool designers.

First, we recommend that designers **choose a name for their tool which doesn’t mislead users.** We observed name-related misconceptions for both private browsing and VPNs. Pre-testing product names with prospective users seems promising, since it might be difficult to predict misconceptions a priori.

Second, we recommend **testing the tool with non-experts**, since misconceptions can arise while using a tool. For example, some users of ad blockers thought the tool was not working because they still saw some ads. In this case, displaying the number of ads blocked might counter this misconception. Norcie et al.’s

work with the Tor Browser Bundle shows that user testing can yield substantial improvements to usability [61].

7 Conclusions

Privacy-enhancing tools can help address some of the public’s concerns about privacy, and public awareness campaigns employing nudging have the potential to encourage adoption. However, misconceptions about privacy tools are common, and addressing these misconceptions is crucial if the tools are to be adopted effectively. Misconceptions can be addressed as part of nudging interventions, and in the marketing and design of tools themselves. To inform the design of nudges and tools, we conducted a demographically-stratified survey to study people’s use of and perceptions about five tools. First, we collected descriptive data on prevalence and recency of tool use (§ 4.2). Next, we asked participants to indicate which protections they thought the tools provided in twelve realistic scenarios. These questions allowed us to quantify the prevalence of misconceptions about the tools’ protections (§ 4.4) and to understand nuances of these mistaken beliefs (§ 4.7). Especially common were participants answering questions incorrectly despite demonstrating partial knowledge, and participants expressing either resignation or overconfidence. We show that those who have used a tool answer more questions about it correctly, but that those who have used VPNs and Tor Browser also answer more questions incorrectly, suggesting that partial knowledge may lead some participants to make mistaken assumptions about these tools’ protections (§ 4.5). We also identify demographics associated with use of the tools, which may help target nudging interventions to those who would most benefit (§ 4.6). Finally, we offer recommendations for designing both nudges and tools themselves (§ 6). In particular, we suggest that interventions should target well-defined threats and address obstacles to realistic perceptions of threat susceptibility, response efficacy, and self-efficacy. We suggest that tool designers follow these same recommendations and that they test the name of their tool to ensure it is not misleading. They should also test their tools with non-experts to identify emergent misconceptions. We hope our findings will lead to more widespread and effective use of privacy- and security-enhancing technologies.

Acknowledgments

We thank Yuanyuan Feng and Linda Moreci for their assistance with our research. This research was supported in part by a grant from the National Science Foundation's Secure and Trustworthy Computing program (CNS-1801316) and in part through a fellowship from the CyLab Security and Privacy Institute at Carnegie Mellon University.

References

- [1] Ruba Abu-Salma and Benjamin Livshits. Evaluating the End-User Experience of Private Browsing Mode. *CHI '20: CHI Conference on Human Factors in Computing Systems*, April 2020.
- [2] Alessandro Acquisti, Manya Sleeper, Yang Wang, Shomir Wilson, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, and Florian Schaub. Nudges for Privacy and Security. *ACM Computing Surveys*, 50(3):1–41, August 2017.
- [3] Yusuf Albayram, Mohammad Maifi Hasan Khan, Theodore Jensen, and Nhan Nguyen. "...better to use a lock screen than to worry about saving a few seconds of time" - Effect of Fear Appeal in the Context of Smartphone Locking Behavior. *Symposium on Usable Privacy and Security*, 2017.
- [4] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015.
- [5] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information: Pew Research Center*. November 2019.
- [6] Maria Bada, Angela M. Sasse, and Jason R. C. Nurse. Cyber Security Awareness Campaigns: Why do they fail to change behaviour? *arXiv preprint*, January 2019.
- [7] Matthew Ball, Roderic Broadhurst, Alexander Niven, and Harshit Trivedi. Data Capture and Analysis of Darknet Markets. *SSRN Electronic Journal*, 2019. <https://www.ssrn.com/abstract=3344936>.
- [8] Sean Barnum, Michael Gegick, and C.C. Michael. Defense in Depth | CISA. <https://us-cert.cisa.gov/bsi/articles/knowledge/principles/defense-in-depth>, September 2005.
- [9] Adam Beaument, M. Angela Sasse, and Mike Wonham. The compliance budget: Managing security behaviour in organisations. In *Proceedings of the 2008 Workshop on New Security Paradigms - NSPW '08*, pages 47–58, Lake Tahoe, California, USA, 2008. ACM Press. <http://portal.acm.org/citation.cfm?doid=1595676.1595684>.
- [10] Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science*, 4(3):340–347, May 2013. <https://doi.org/10.1177/1948550612455931>.
- [11] Jon Brodtkin. Fake tech support scam is trouble for legitimate remote help company. <https://arstechnica.com/information-technology/2013/11/fake-tech-support-scam-is-trouble-for-legitimate-remote-help-company/>, November 2013.
- [12] Dalvin Brown. Is streaming video from sketchy websites illegal? <https://www.usatoday.com/story/tech/2019/12/16/can-get-arrested-streaming-illicit-movies-its-complicated/2662072001/>, December 2019.
- [13] United States Census Bureau. Current Population Survey (CPS), 2018. <https://www.census.gov/cps/data/cpstablecreator.html>.
- [14] Karoline Busse, Julia Schäfer, and Matthew Smith. Replication: No One Can Hack My Mind Revisiting a Study on Expert and Non-Expert Security Practices and Advice. *Symposium on Usable Privacy and Security*, pages 116–136, August 2019.
- [15] Yinzhi Cao, Song Li, and Erik Wijmans. (Cross-)Browser Fingerprinting via OS and Hardware Level Features. *Network and Distributed System Security Symposium*, March 2017.
- [16] CISA. Tips | CISA. <https://us-cert.cisa.gov/ncas/tips>.
- [17] Cisco. 2013 Cisco Annual Security Report. January 2013.
- [18] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Honolulu HI USA, April 2020. ACM. <https://dl.acm.org/doi/10.1145/3313831.3376389>.
- [19] Cyrus Farivar and Andrew Blankstein. Feds take down the world's 'largest dark web child porn marketplace'. <https://www.nbcnews.com/news/crime-courts/feds-take-down-world-s-largest-dark-web-child-porn-n1066511>, October 2019.
- [20] Roger Dingledine. [tor-announce] Tor security advisory: Old Tor Browser Bundles vulnerable. <https://lists.torproject.org/pipermail/tor-announce/2013-August/000089.html>, August 2013.
- [21] Disconnect. Best privacy VPN app for iOS and Mac. Powerful protection with one tap., Sep 2020. <https://disconnect.me>.
- [22] DuckDuckGo. A Study on Private Browsing: Consumer Usage, Knowledge, and Thoughts. Whitepaper, January 2017.
- [23] Agnieszka Dutkowska-Zuk, Austin Hounsel, Andre Xiong, Marshini Chetty, Nick Feamster, Molly Roberts, and Brandon Stewart. Practicing Safe Browsing: Understanding How and Why University Students Use Virtual Private Networks. *arXiv preprint*, cs.HC, February 2020.
- [24] EFF. The Playpen Cases: Frequently Asked Questions. <https://www.eff.org/pages/playpen-cases-frequently-asked-questions>, August 2016.

- [25] EFF. Choosing the VPN That's Right for You. <https://ssd.eff.org/en/module/choosing-vpn-thats-right-you>, March 2019.
- [26] E. Erdin, C. Zachor, and M. H. Gunes. How to Find Hidden Users: A Survey of Attacks on Anonymity Networks. *IEEE Communications Surveys Tutorials*, 17(4):2296–2316, Fourthquarter 2015.
- [27] ExpressVPN. How to Combine a VPN and Tor Browser for Online Anonymity. <https://www.expressvpn.com/>, October 2020.
- [28] FTC. Tech Support Scams. <https://www.ftc.gov/tips-advice/business-center/small-businesses/cybersecurity/tech-support-scams>, October 2018.
- [29] Kevin Gallagher, Sameer Patil, and Nasir D. Memon. New Me - Understanding Expert and Non-Expert Perceptions and Usage of the Tor Anonymity Network. *Symposium on Usable Privacy and Security*, 2017.
- [30] Samuel Gibbs. Antivirus software is dead, says security expert at Symantec. <http://www.theguardian.com/technology/2014/may/06/antivirus-software-fails-catch-attacks-security-expert-symantec>, May 2014.
- [31] Dan Goodin. Millions exposed to malvertising that hid attack code in banner pixels. <https://arstechnica.com/information-technology/2016/12/millions-exposed-to-malvertising-that-hid-attack-code-in-banner-pixels/>, December 2016.
- [32] Dan Goodin. Malvertisers target Mac users with steganographic code stashed in images. <https://arstechnica.com/information-technology/2019/01/malvertisers-target-mac-users-with-stenographic-code-stashed-in-images/>, January 2019.
- [33] Yael Grauer. The impossible task of creating a "Best VPNs" list today | Ars Technica. <https://arstechnica.com/information-technology/2016/06/aiming-for-anonymity-ars-assesses-the-state-of-vpns-in-2016/>, June 2016.
- [34] Glenn Greenwald and Ewen MacAskill. NSA Prism program taps in to user data of Apple, Google and others. *The Guardian*, June 2013. <https://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data>.
- [35] Hana Habib, Jessica Colnago, Vidya Gopalakrishnan, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, and Lorrie Faith Cranor. Away From Prying Eyes - Analyzing Usage and Understanding of Private Browsing. *SOUPS @ USENIX Security Symposium*, 2018.
- [36] Cormac Herley. So long, and no thanks for the externalities - the rational rejection of security advice by users. *NSPW*, pages 133–144, 2009.
- [37] Cormac Herley. More Is Not the Answer. *IEEE Security & Privacy*, 12(1):14–19, 2014.
- [38] Aaron Holmes. The dark web turns 20 this month — here's how it changed the world - Business Insider. <https://www.businessinsider.com/dark-web-changed-the-world-black-markets-arab-spring-2020-3>, March 2020.
- [39] Bryan Horling and Matthew Kulick. Personalized Search for everyone. <https://googleblog.blogspot.com/2009/12/personalized-search-for-everyone.html>, December 2009.
- [40] Muhammad Ikram, Narseo Vallina-Rodriguez, Suranga Seneviratne, Mohamed Ali Kaafar, and Vern Paxson. An Analysis of the Privacy and Security Risks of Android VPN Permission-enabled Apps. *IMC 2016*, 2016.
- [41] Iulia Ion, Rob Reeder, and Sunny Consolvo. "...No one Can Hack My Mind" - Comparing Expert and Non-Expert Security Practices. *Symposium on Usable Privacy and Security*, 2015.
- [42] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara B. Kiesler. "My Data Just Goes Everywhere - " User Mental Models of the Internet and Implications for Privacy and Security. *Symposium on Usable Privacy and Security*, 2015.
- [43] Kaspersky. Disk and File Encryption, Sep 2020. <https://www.kaspersky.com/enterprise-security/wiki-section/products/encryption>.
- [44] Kaspersky Lab. The main sources of malware infection. <https://web.archive.org/web/20201127210358/https://support.kaspersky.com/789>, November 2018.
- [45] Katie Kasunic. How To Use Tor Browser: Everything You MUST Know (2020). <https://www.vpnmentor.com/blog/tor-browser-work-relate-using-vpn/>, August 2020.
- [46] Brian Krebs. Tools for a Safer PC — Krebs on Security. <https://krebsonsecurity.com/tools-for-a-safer-pc/>.
- [47] Brian Krebs. Krebs's 3 Basic Rules for Online Safety — Krebs on Security. <https://krebsonsecurity.com/2011/05/krebs-3-basic-rules-for-online-safety/>, May 2011.
- [48] Brian Krebs. Post-FCC Privacy Rules, Should You VPN? — Krebs on Security, March 2017.
- [49] Ponnurangam Kumaraguru, Justin Cranshaw, Alessandro Acquisti, Lorrie Cranor, Jason Hong, Mary Ann Blair, and Theodore Pham. School of phish: A real-word evaluation of anti-phishing training. In *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09*, Mountain View, California, 2009. ACM Press. <http://portal.acm.org/citation.cfm?doid=1572532.1572536>.
- [50] Ozan Kuru and Josh Pasek. Improving social media measurement in surveys: Avoiding acquiescence bias in Facebook research. *Computers in Human Behavior*, 57:82–92, April 2016. <https://linkinghub.elsevier.com/retrieve/pii/S0747563215302788>.
- [51] Pierre Laperdrix. Browser Fingerprinting: An Introduction and the Challenges Ahead | Tor Blog. <https://blog.torproject.org/browser-fingerprinting-introduction-and-challenges-ahead>, September 2019.
- [52] Selena Larson. Infant Social Security numbers are for sale on the dark web. <https://money.cnn.com/2018/01/22/technology/infant-data-dark-web-identity-theft/index.html>, January 2018.
- [53] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–31, November 2018. <https://dl.acm.org/doi/10.1145/3274371>.
- [54] Monica G. Maceli. Encouraging patron adoption of privacy-protection technologies: Challenges for public libraries. *IFLA Journal*, 44(3):195–202, August 2018.
- [55] James E Maddux and Ronald W Rogers. Protection motivation and self-efficacy: A revised theory of fear appeals and attitude change. *Journal of experimental social psychology*, 19(5):469–479, 1983.
- [56] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. "What Can't Data Be Used For?": Privacy Ex-

- pectations about Smart TVs in the U.S. In *Proceedings 3rd European Workshop on Usable Security*, London, England, 2018. Internet Society. https://www.ndss-symposium.org/wp-content/uploads/2018/06/eurosec2018_16_Malkin_paper.pdf.
- [57] Jonathan R. Mayer and John C. Mitchell. Third-Party Web Tracking: Policy and Technology. *2012 IEEE Symposium on Security and Privacy*, 2012.
- [58] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM Human-Computer Interaction*, August 2019.
- [59] Georg Merzdovnik, Markus Huber, Damjan Buhov, Nick Nikiforakis, Sebastian Neuner, Martin Schmiedecker, and Edgar Weippl. Block Me If You Can: A Large-Scale Study of Tracker-Blocking Tools. *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, March 2017.
- [60] Moses Namara, Daricia Wilkinson, Kelly Caine, and Bart P. Knijnenburg. Emotional and Practical Considerations Towards the Adoption and Abandonment of VPNs as a Privacy-Enhancing Technology. *Proceedings on Privacy Enhancing Technologies*, 2020(1):83–102, December 2019.
- [61] Greg Norcie, Jim Blythe, Kelly Caine, and L. Jean Camp. Why Johnny Can't Blow the Whistle: Identifying and Reducing Usability Issues in Anonymity Systems. *Workshop on Usable Security*, February 2014.
- [62] NordVPN. Block ads and malicious websites with CyberSec, Sep 2020. <https://nordvpn.com/features/cybersec/>.
- [63] Norton. Norton Privacy Manager | How it works, Mar 2019. <https://www.youtube.com/watch?v=iKsHl-uzVrU>.
- [64] Norton. Browse the Internet securely with Norton Safe Web, Oct 2020. <https://support.norton.com/sp/en/us/home/current/solutions/v19116982>.
- [65] Norton. Norton Privacy Manager, Sep 2020. <https://us.norton.com/norton-privacy-manager>.
- [66] Norton. Norton Secure VPN, Sep 2020. <https://us.norton.com/products/norton-secure-vpn>.
- [67] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70(C):153–163, May 2017.
- [68] Nicole Perloth. Tor Project, a Digital Privacy Group, Reboots With New Board (Published 2016). *The New York Times*, July 2016. <https://www.nytimes.com/2016/07/14/technology/tor-project-a-digital-privacy-group-reboots-with-new-board.html>.
- [69] Nathaniel Popper. The Tax Sleuth Who Took Down a Drug Lord. *The New York Times*, December 2015. <https://www.nytimes.com/2015/12/27/business/dealbook/the-unsung-tax-agent-who-put-a-face-on-the-silk-road.html>.
- [70] Prolific Team. Representative Samples on Prolific. <https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-Samples-on-Prolific>, March 2019.
- [71] Prolific Team. Reviewing submissions: How do I decide who to accept/reject?, May 2020. <https://researcher-help.prolific.co/hc/en-gb/articles/360009092394-Reviewing-submissions-How-do-I-decide-who-to-accept-reject->.
- [72] ProtonVPN. Why use Tor over VPN. <https://protonvpn.com/blog/tor-vpn/>, July 2018.
- [73] E. Racine, P. Skeba, E. P. S. Baumer, and A. Forte. What are PETs for Privacy Experts and Non-experts? *Symposium on Usable Privacy and Security*, August 2020.
- [74] L. Rainie, S. Kiesler, R. Kang, and Mary Madden. *Anonymity, Privacy, and Security Online*. Pew Research Internet Project. 2013.
- [75] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. *IEEE SP*, 2019.
- [76] Elissa M Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L Mazurek. A Comprehensive Quality Evaluation of Security and Privacy Advice on the Web. *USENIX Security Symposium*, pages 88–108, August 2020.
- [77] Robert W. Reeder, Iulia Ion, and Sunny Consolvo. 152 Simple Steps to Stay Safe Online - Security Advice for Non-Tech-Savvy Users. *IEEE Secur. Priv.*, 15(5):55–64, 2017.
- [78] Consumer Reports. 66 Ways to Protect Your Privacy Right Now, February 2017.
- [79] Ronald W Rogers. A Protection Motivation Theory of Fear Appeals and Attitude Change. *The Journal of Psychology*, 91(1):93–114, 1975.
- [80] Ronald W Rogers and Steven Prentice-Dunn. Protection motivation theory. 1997.
- [81] Florian Schaub, Aditya Marella, Pranshu Kalvani, Blase Ur, Chao Pan, Emily Forney, and Lorrie Faith Cranor. Watching Them Watching Me: Browser Extensions Impact on User Privacy Awareness and Concern. *Workshop on Usable Security*, February 2016.
- [82] Daniel Smullen, Yuanyuan Feng, and Norman Sadeh. The Best of Both Worlds: Mitigating Trade-offs Between Accuracy and User Burden in Capturing Mobile App Privacy Preferences. *Proceedings on Privacy Enhancing Technologies*, 1:195–215, 2020.
- [83] Ben Stegner. How to Avoid Fake Ads Disguised as Fake Download Links. <https://www.makeuseof.com/tag/spot-avoid-ads-disguised-download-buttons/>, July 2019.
- [84] Geordie Stewart and David Lacey. Death by a Thousand Facts - Criticising the Technocratic Approach to Information Security Awareness. *HAI SA*, 2011.
- [85] Peter Story, Daniel Smullen, Alessandro Acquisti, Lorrie Faith Cranor, Norman M Sadeh, and Florian Schaub. From Intent to Action - Nudging Users Towards Secure Mobile Payments. *SOU PS @ USENIX Security Symposium*, 2020.
- [86] O. Sukwong, H. S. Kim, and J. C. Hoe. Commercial Antivirus Software Effectiveness - An Empirical Study. *Computer*, 44(3):63–70, 2011.
- [87] Xiao Hui Tai, Kyle Soska, and Nicolas Christin. Adversarial Matching of Dark Net Market Vendor Accounts. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1871–1880, Anchorage AK USA, July 2019. ACM. <https://dl.acm.org/doi/10.1145/3292500.3330763>.
- [88] Karyn A. Temple. U.S. Copyright Office Responses To Specific Questions. <https://troypoint.com/wp-content/>

- uploads/2019/07/letter-to-senators-tillis-and-coons-on-felony-streaming-from-copyright-office.pdf, July 2019.
- [89] Richard H Thaler and Cass R Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. J. Wiley and Sons, 2008.
- [90] The Tor Project. Doc/TorPlusVPN – Tor Bug Tracker & Wiki. <https://trac.torproject.org/projects/tor/wiki/doc/TorPlusVPN>, October 2019.
- [91] Matt Traudt. VPN + Tor: Not Necessarily a Net Gain - Matt Traudt. <https://matt.traudt.xyz/posts/vpn-tor-not-mRikAa4h.html>, November 2016.
- [92] Rick Wash. Folk models of home computer security. *Symposium on Usable Privacy and Security*, pages 1–16, 2010.
- [93] Rick Wash and Emilee Rader. Too Much Knowledge? Security Beliefs and Protective Behaviors Among United States Internet Users. *Eleventh Symposium On Usable Privacy and Security*, pages 309–325, 2015.
- [94] Zachary Weinberg, Shinyoung Cho, Nicolas Christin, Vyas Sekar, and Phillipa Gill. How to Catch when Proxies Lie. *IMC '18: Internet Measurement Conference*, October 2018.
- [95] WhatIsMyIPAddress.com. How does geolocation work? <https://whatismyipaddress.com/geolocation>, November 2020.
- [96] Wikipedia. Google Personalized Search. *Wikipedia*, April 2020. https://en.wikipedia.org/w/index.php?title=Google_Personalized_Search&oldid=952991366.
- [97] Wikipedia. Internet geolocation. *Wikipedia*, September 2020. https://en.wikipedia.org/w/index.php?title=Internet_geolocation&oldid=980936345.
- [98] Wikipedia. PRISM (surveillance program). *Wikipedia*, September 2020. [https://en.wikipedia.org/w/index.php?title=PRISM_\(surveillance_program\)&oldid=978695563](https://en.wikipedia.org/w/index.php?title=PRISM_(surveillance_program)&oldid=978695563).
- [99] Jakob Wirth, Christian Maier, Sven Laumer, and Friedrich-Alexander Universität Erlangen-Nürnberg. The Influence Of Resignation On The Privacy Calculus In The Context Of Social Networking Sites: An Empirical Analysis. *Twenty-Sixth European Conference on Information Systems*, 2018.
- [100] Yixin Zou, Kevin Roundy, Acar Tamersoy, Saurabh Shintre, Johann Roturier, and Florian Schaub. Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices. *CHI*, pages 1–15, April 2020.

A Appendix

A.1 Bundled Products and Incorrect Answers

As we describe in our limitations section (§ 5), the fact that VPNs and antivirus software are sometimes bundled with additional functionality could be a source of participants’ “incorrect” responses. To estimate the extent of this confounding factor, we examined incorrect answers for VPNs and antivirus software, searching for references to these bundled functionalities.

First, we consider VPNs, which can be bundled with ad blocking and antivirus functionality [62]. We inspected free-text responses from those who had answered “incorrectly” about whether VPNs would prevent advertisers from showing them targeted ads, and whether VPNs would prevent hackers from gaining access to their device. We reasoned that these free-text responses would be the most likely to contain explicit references to ad blocking and antivirus capabilities, respectively, since these were the two scenarios in which these capabilities would be most effective. In all, we inspected 42 such responses. In these responses, the dominant theme seemed to be about IP and location hiding, rather than the VPN blocking ads or malware. For example, P204 wrote that “it still allows ads, just targeted for the area your vpn is located” and P158 wrote that “... hackers can’t find you.” Only one of these incorrect responses clearly described the possibility of a VPN acting as an ad blocker, with P295 writing that “A VPN with ad blocking protects your privacy by preventing third-party ad domains from installing trackers on your device when they display their ads. By blocking the trackers, the VPN prevents the ad domains from collecting data about you.” No participants clearly described the mechanism by which VPNs can protect from malware (i.e., blocking known malware-distributing domains), but four participants described protections from hackers more generally, writing that “VPNs would not allow other programs into your computer system...” (P184), “VPN’s give you an extra layer of security” (P408), and that “... The connection ... blocks unwanted intrusions” (P198). Thus, it seems likely that inappropriate mental models were in fact responsible for most of these “incorrect” answers, rather than participants correctly considering the ways in which optional VPN features can block ads or malware.

Next, we consider antivirus software, which can be bundled with VPNs [66]. We inspected free-text responses from those who had answered incorrectly about the three scenarios in which VPNs are very effective: preventing your employer from seeing the websites you visit, preventing your ISP from seeing the websites you visit, and preventing websites from seeing your physical location. We reasoned that these free-text responses would be the most likely to contain explicit references to VPN capabilities. In all, we inspected 21 such responses. In these responses, the dominant theme seemed to be about virus prevention, rather than antivirus acting as a VPN. A representative answer from P99 reads: “Malicious software would give away my location directly to a hacker or website. Antivirus software eliminates track-

ing malware.” Only one of these 21 incorrect responses clearly alluded to the possibility of antivirus acting as a VPN, with P201 writing “... my free AVG does not block my location but offers to do that for additional cost per year.” P59 gave a more opaque response that hints at an awareness of additional features, but does not go into detail, writing that “Good Antivirus software has many features built-in and I think it is quite effective.” Thus, we think it is likely that inappropriate mental models were in fact responsible for most of these “incorrect” answers, rather than participants considering the possibility of antivirus acting as a VPN. We do wonder whether the availability of these optional features might lead consumers to assume that basic antivirus itself can provide these protections.

A.2 Survey Instrument

All participants are asked to answer the screening questions below.

Based on your answers to the screening questions, we will determine your eligibility for our survey. If you are eligible, the survey will take about 15 minutes to complete.

In what country do you currently reside?
(The United States, Other country)

Do you speak English?
(Yes, No)

What is your age in years? _____

Based on your answers to our screening questions, we have determined that you are eligible for our survey.

Please review the details below:

[Consent Form]

Have you read and understood the information above?
(Yes, No)

Do you want to participate in this research and continue with the survey?
(Yes, No)

Researchers at OMITTED are conducting a study to understand people’s use of web browsing-related tools. Please answer honestly and **take the time to read the information in this survey carefully.**

What do you think is **the likelihood** of others observing your web browsing activity?

(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

How **concerned or unconcerned** would you be if others observed your web browsing activity?

(Not at all concerned, Slightly concerned, Moderately concerned, Very concerned)

Rate your level of **disagreement or agreement** with the following statement:

“I think I know how to use privacy tools to prevent others from observing my web browsing activity.”

(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How **interested or uninterested** would you be in learning to use privacy tools to prevent others from observing your web browsing activity?

(Not at all interested, Slightly interested, Moderately interested, Very interested)

How **easy or difficult** do you think it would be for you to use privacy tools to prevent others from observing your web browsing activity?

(Very difficult, Somewhat difficult, Somewhat easy, Very easy)

Rate your level of **disagreement or agreement** with the following statement:

“If I were to start using privacy tools, in general I would prevent others from observing my web browsing activity.”

(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

The following set of questions are about web browsing-related tools:

[The real tools are displayed in a random order, with the fake tool last (i.e., PrivacyDog)]

- Private browsing
- VPNs
- Tor Browser
- DuckDuckGo
- Ad blockers
- Antivirus software
- PrivacyDog

If you’ve never heard of some or all of these tools, that’s okay! Please simply answer the questions to the best

of your ability, without searching for the answers online.

[The following block of questions is displayed once for each real tool. We used an abbreviated block of questions for PrivacyDog. The blocks were shown in a random order. For brevity, we show only the blocks for private browsing and PrivacyDog.]

Private Browsing

Note that "private browsing" is referred to as "Incognito" in Google Chrome and "InPrivate" in Microsoft Edge.

[This kind of explanatory text was only included for private browsing.]

Have you **heard of** private browsing before?

(Yes, No, Unsure)

[If Yes, has heard of]

Have you **used** private browsing before?

(Yes, No, Unsure)

[If Yes, has used]

When did you most recently use private browsing?

(Today, In the past week, In the past month, In the past year, More than a year ago)

Do you **know anyone else** who has used private browsing?

(Yes, No, Unsure)

[If No, has not used]

Have you **tried to use** private browsing?

(Yes, No, Unsure)

Rate your level of **disagreement or agreement** with the following statement:

"I think I know how to use **private browsing**."

(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How **easy or difficult** do you think it would be for you to use **private browsing**?

(Very difficult, Somewhat difficult, Somewhat easy, Very easy)

Rate your level of **disagreement or agreement** with the following statement:

"If I were to start using **private browsing**, in general I would prevent others from observing my web browsing activity."

(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

When, if ever, do you think you will next use **private browsing**?

(Today, Sometime in the next week, Sometime in the next month, Sometime in the next year, More than a year from now, Never, I don't know)

PrivacyDog

Have you **heard of** PrivacyDog before?

(Yes, No, Unsure)

[If Yes, has heard of]

Have you **used** PrivacyDog before?

(Yes, No, Unsure)

Which tool did we ask you about in the most recent set of questions?

[The real tools are displayed in a random order, with the fake tool last (i.e., PrivacyDog)]

- Private browsing
- VPNs
- Tor Browser
- DuckDuckGo
- Ad blockers
- Antivirus software
- PrivacyDog

[The following block of questions is displayed six times, each time populated with a different randomly selected scenario, drawn from a pool of twelve possible scenarios.]

When you browse the web, how effective are the tools below at **preventing hackers from gaining access to your device**?

[Answers options are shown in a response matrix, where each row is labeled with a tool, and the columns are labeled with the answers options: Unsure, Not at all effective, Somewhat effective, Very effective]

[For each block, we ask the following follow-up questions for a single randomly selected tool. The tools are selected without replacement, so the follow-up questions are only asked one time for each tool.]

[If Unsure]

In a few sentences, please explain why you indicated that you were **unsure whether Private browsing would be effective at preventing hackers from**

gaining access to your device. _____

[If not Unsure]

In a few sentences, please explain why you indicated that **Private browsing would be [SELECTED_EFFECTIVENESS] at preventing hackers from gaining access to your device.**

How interested or uninterested would you be in **preventing hackers from gaining access to your device?**

(Not at all interested, Slightly interested, Moderately interested, Very interested)

Please answer the following questions about your use of devices **in the past week.**

In the past week, which of the following types of devices did you **use at least once?**

(Smartphone, Tablet, Laptop computer, Desktop computer)

In the past week, which of the following types of devices, if any, did you **share with other people?**

(Smartphone, Tablet, Laptop computer, Desktop computer)

In the past week, how often did you **use a web browser** on each of the following devices?

[Answer options are shown in a response matrix. Rows are labeled with device types: Smartphone, Tablet, Laptop computer, Desktop computer, Other device(s). Columns are labeled with the answer options: Every day, On multiple days, On one day, Never.]

[If Never is not selected for Other device(s)]

Please briefly describe the other device(s) you used to browse the web, and how often you used them to browse the web.

What gender do you identify with?

(Male, Female, Non-binary, Other: _____, Prefer not to answer)

What best describes your employment status?

(Working, paid employee; Working, self employed; Student; Not employed; Retired; Prefer not to answer)

Have you ever worked in or studied in a computer-

related field? (Computer Science, IT support, etc.)
(Yes, No)

What is the highest level of school you have completed or degree you have earned?

(Less than high school, High school or equivalent, College or associate degree, Master's degree, Doctoral degree, Professional degree, Other: _____, Prefer not to answer)

Please estimate what your total household income will be for this year:

(Less than \$10,000; \$10,000 - \$19,999; \$20,000 - \$39,999; \$40,000 - \$59,999; \$60,000 - \$79,999; \$80,000 - \$99,999; \$100,000 or more; Prefer not to answer)

Please indicate which other people, if any, live in your household.

(Domestic partner, e.g., spouse, boyfriend/girlfriend, etc.; Children; Parents; Other family; Unrelated roommates; I live alone; Other: _____, Prefer not to answer)

When you browse the web, how effective are the tools below at ...




Fig. 7. Responses consistent with our threat model are indicated with a star. Tools are sorted by the percent of correct responses.

Demographic Factor		Survey	Census
Age	18-27	16.8%	17.4%
	28-37	18.6%	17.6%
	38-47	16.2%	16.1%
	48-57	17.2%	16.9%
	58+	31.2%	32.1%
Gender	Female	50.6%	51.6%
	Male	48.4%	48.4%
	Other	1%	
Ethnicity	White	72.4%	78.0%
	Black	12.8%	12.6%
	Asian	7.4%	6.2%
	Mixed	3.8%	1.8%
	Other	3.6%	1.4%
Employment	Working (paid employee)	45.0%	
	Working (self employed)	17.4%	
	Student	6.8%	
	Not employed	15.0%	
	Retired	15.0%	
	Prefer not to answer	0.8%	
Education	High school or less	24.0%	
	College or associate	52.4%	
	Graduate degree	18.0%	
	Professional degree	3.2%	
	Other	2.2%	
	Prefer not to answer	0.2%	
Worked or studied in a computer-related field	Yes	28.4%	
	No	71.6%	
Living situation	Domestic partner	50.6%	
	Children	30.8%	
	Parents	16.4%	
	Other family	12.2%	
	Unrelated roommates	4.8%	
	I live alone	21.6%	
	Other	0.8%	
Prefer not to answer	0.8%		
Household income	Less than \$10,000	4.6%	
	\$10,000 - \$19,999	8.0%	
	\$20,000 - \$39,999	23.2%	
	\$40,000 - \$59,999	16.4%	
	\$60,000 - \$79,999	13.8%	
	\$80,000 - \$99,999	10.4%	
	\$100,000 or more	21.0%	
	Prefer not to answer	2.6%	

Table 4. Our participants' demographics ($n = 500$). For ethnicity, we report data received from the Prolific platform about our participants, since we did not ask about ethnicity in our survey instrument. We collected data about the other demographic factors using our own survey instrument. We requested a demographically representative sample, so Prolific stratified across age, sex, and ethnicity, in an attempt to match proportions from the US Census Bureau [70]. We include data from the US Census Bureau for comparison [13].

First Pass Code	Description	Number of Occurrences
MISCONCEPTION	Describes an incorrect belief about the tool/scenario (e.g., “private browsing hides your location”). We classify thinking that an entity can see things no matter what you/others do as a misconception. We classify wondering if a tool is fake as a misconception. We classify referencing related products (e.g., DDG browser instead of search, VPNs acting as ad blockers, and antivirus acting as a VPN) as misconceptions.	796
NO_MISCONCEPTION	No incorrect belief about the tool/scenario is described	1678
POOR	A low-quality answer. Incomprehensible, clearly about the wrong tool/scenario, etc.	26

Table 5. We used a multi-step coding process to make our analysis more efficient. We applied these first pass codes to all free-text responses (n=2500), before applying the second pass codes shown in Table 6 to only those responses which contained any kind of misconception (n=796).

Second Pass Code	Description	Number of Occurrences
DANGEROUS_ADS	The tool tries to stop dangerous ads in particular	4
DARK_WEB	Mentioning the dark web	15
EXPERIENCE	Citing one’s own experiences as evidence	35
EXTRAS	Writing that the tool offers optional extra features (simply mentioning a feature that isn’t normally in the tool doesn’t count)	13
HIDING	Trying to stay secure by avoiding being noticed, or by keeping information hidden (not as much about privacy as security, so not applicable every time hiding is mentioned)	24
LAYERS	Having more layers of protection keeps you secure/private	13
NAME	Referencing the name of the tool as justification for a belief (e.g., “private”, “incognito”)	49
NOTHING	Nothing can be done to provide protection (e.g., “nothing can stop advertisers from seeing everything you do”), the resources of the entity are too great to overcome, etc.	154
OTHER_AD_BLOCKER	Mentions of this tool when it was not the tool the participant was asked about	4
OTHER_ANTIVIRUS	Mentions of this tool when it was not the tool the participant was asked about	4
OTHER_BRAVE	Mentions of this tool when it was not the tool the participant was asked about	1
OTHER_BROWSER	Mentions of this tool when it was not the tool the participant was asked about	11
OTHER_DISK_ENCRYPTION	Mentions of this tool when it was not the tool the participant was asked about	1
OTHER_DUCKDUCKGO	Mentions of this tool when it was not the tool the participant was asked about	1
OTHER_FIREWALL	Mentions of this tool when it was not the tool the participant was asked about	1
OTHER_PRIVATE_BROWSING	Mentions of this tool when it was not the tool the participant was asked about	7
OTHER_TORRENT	Mentions of this tool when it was not the tool the participant was asked about	4
OTHER_VPN	Mentions of this tool when it was not the tool the participant was asked about	17
PERMISSIONS	Referencing permissions (e.g., the location permission)	6
SEARCH_ADS	Mentioning ads in search results	6
SHOULDER_SURFING	Mentioning or implying a shoulder surfing threat model (e.g., someone watching you use your device, or someone else using your device and seeing information without seeking it out)	7
TOTAL	Writing that the tool provides total protection, hides things from everyone, provides total anonymity, etc.	69
TRUE	Accurately describing a true function of the tool (e.g., not retaining cookies). Excessively vague responses aren’t counted. Some edge cases: For private browsing: We don’t count “blocking” cookies. For ad blockers: We don’t count blocking cookies generally, but we do count blocking advertisers’ cookies and blocking tracking (e.g., Google Analytics, other ad networks, etc.). For VPNs: We don’t count generic “giving privacy” or “masking info”. We do count extra features of VPNs: review these marketing materials [21, 62]. For Tor Browser: We don’t count vague references to the “dark web”. We count writing that Tor provides anonymity and encrypts traffic. For Antivirus software: We don’t count generic “staying safe”. Since antivirus software can be bundled with extra features, review some examples of marketing materials [43, 63, 64, 66]: we count these extra features as true functions.	262

Table 6. Our final set of thematic codes, and their frequencies of occurrence. We only applied these thematic codes to responses we identified as containing any form of misconception (n=796), since we only wanted to analyze misconceptions in greater detail.