Design and Evaluation of Security and Privacy Nudges: From Protection Motivation Theory to Implementation Intentions

Peter Story

CMU-ISR-21-107
August 2021

Institute for Software Research
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Norman Sadeh (Chair)
Lorrie Faith Cranor
Alessandro Acquisti
Florian Schaub (University of Michigan)
Yaxing Yao (University of Maryland, Baltimore County)

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Societal Computing*

This research was supported in part by a grant from the National Science Foundation's Secure and Trustworthy Computing program (CNS-1330596, SES-1513957, CNS-1801316) and in part through a fellowship from the CyLab Security and Privacy Institute at Carnegie Mellon University. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of any sponsoring institution, the U.S. Government, or any other entity.

Keywords: Nudges, privacy, security, technology adoption, protection motivation theory, implementation intentions, action plans, coping plans

Dedicated to my wife, Madeleine, and to my parents, Craig and Kelly.

Abstract

Americans often express concern about their digital security and privacy, yet adoption of security and privacy tools and best practices remains inconsistent. The fields of psychology and behavioral economics offer explanations for this apparent discrepancy, and suggest nudging interventions as a potential solution. Nudges can take many forms, but what nudges have in common is that they should help people make decisions that align with their stated preferences.

My research centers on designing nudges to encourage the adoption of security and privacy tools. My major contribution is the introduction of implementation intention nudges to the field of computer security and privacy. Implementation intentions are plans which help people initiate behaviors (action plans) and overcome obstacles (coping plans). The effectiveness of implementation intentions has been demonstrated in many other contexts, but my work is the first to test them in the context of computer security and privacy. By studying implementation intentions in this context, I offer security and privacy advocates a greater understanding of how this type of nudge can help the public protect themselves from digital threats.

In my first chapter of completed work, I describe my study of nudges designed to encourage adoption of secure mobile payment systems. I tested nudges based on both action planning implementation intentions and protection motivation theory (PMT). I found that participants in both my treatment conditions used Apple Pay more than those in my control condition. Encouraged by these findings, I sought to identify other technologies which might benefit from similar nudging interventions. Thus, I conducted a survey of people's use of and beliefs about web browsing-related privacy tools, which I describe in my next chapter. I found that the most commonly adopted tools did little to address participants' greatest privacy concerns. Based on these findings, I conducted a study of implementation intention nudges designed to help people adopt Tor Browser, which is the subject of my final chapter of completed work. In this study, I tested nudges based on PMT, action planning implementation intentions, and coping planning implementation intentions. These nudges incorporated the recommendations from my second chapter study. I found that my coping planning nudge increased use of Tor Browser in the short-term, while my PMT-based nudge increased use of Tor Browser in both the short- and long-term. In my final chapter, I summarize my research, describe ethical considerations when deploying nudges, and enumerate open research questions relevant to large-scale deployment of nudges.

Acknowledgments

Many others contributed to this thesis. I especially want to thank my advisor, Dr. Norman Sadeh, for the opportunity to work in his lab, his mentorship, and his feedback on my work. My other committee members, Dr. Lorrie Faith Cranor, Dr. Alessandro Acquisti, Dr. Florian Schaub, and Dr. Yaxing Yao, also contributed invaluable feedback.

I am also thankful to my other collaborators, including Daniel Smullen, Rex Chen, Yixin Zou, Abhilasha Ravichander, Hana Habib, Shikun Aerin Zhang, Siddhant Arora, Jongkyu (Talan) Baek, Sushain Cherivirala, Aditya Kamatar, Vinayshekhar Kumar, Pingshan Li, Shaoyan (Sam) Li, Sandy Sun, Yigeng Wang, Dr. Yuanyuan Feng, Dr. Joel Reidenberg, Dr. Shomir Wilson, and Dr. Sebastian Zimmeck.

Finally, this work would not have been possible without the support of ISR staff, particularly Linda Moreci and Connie Herold.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Most Americans express a desire for digital security and privacy [18, 99, 121, 137]. Americans feel a lack of control over their data, and express interest in consumer tools to protect their personal information [18]. However, the limited adoption of security and privacy tools appears inconsistent with these preferences [121, 191]. The fields of psychology and behavioral economics offer explanations for this apparent discrepancy, with concepts such as information asymmetry (e.g., unawareness of the presence of threats to security and privacy) [8], bounded rationality (e.g., failure to fully consider the consequences of threats, unawareness of effective protections from threats, inconsistency in adopting the protections, etc.) [148], and various cognitive and behavioral biases (e.g., optimistically underestimating the chance of being affected by a threat [92]). Thankfully, they also suggest a potential solution to these challenges, in the form of nudging interventions [166]. Nudges can take many forms [13, 54, 93], but what nudges have in common is that they should help people make decisions that align with their stated preferences [6].

The literature on nudging is rich and varied, so it should come as no surprise that some types of nudges have never before been tested in the field of computer security and privacy: my major contribution is the introduction of implementation intention nudges to this field. Implementation intentions are contextually activated plans which help people initiate behaviors and overcome obstacles [22, 56]. The effectiveness of implementation intentions has been demonstrated in many other contexts [91, 107, 108, 109, 112, 122, 123], but my work is the first to test them in the context of computer security and privacy. By studying implementation intentions in this context, I offer security and privacy advocates a greater understanding of how this type of nudge can help the public protect themselves from digital threats.

In Chapter 3, I describe my study of nudges designed to encourage adoption of secure mobile payment systems. As part of the study, I conducted a longitudinal, between-subjects field experiment to test nudges based on protection motivation theory (PMT) and action planning implementation intentions (AP). The experiment included three treatment conditions: a PMT-only condition, a PMT+AP condition, and a control condition. My results showed that participants in the PMT-only and PMT+AP treatment groups were 2.3x ($p = 0.020$) and 3.9x ($p < 0.001$) more likely to use Apple Pay than were participants in the control group, respectively. My findings further suggest that adding an action planning implementation intention to the PMT-only treatment increased its efficacy (1.7x more likely to use, $p = 0.085$). This study showed the promise of nudges based on PMT and action plans to increase real-world adoption of secure technologies.

Encouraged by these findings, I sought to identify other technologies which might benefit from similar nudging interventions. Thus, I conducted a survey of people's use of and beliefs about browsing-related privacy tools, which I describe in Chapter 4. Different than prior work, I asked participants what protections they thought several different tools provided against varied privacy threats. This design allowed me to make comparisons across tools and privacy threats. I found that a substantial number of participants were already using some tools to protect themselves from digital threats. For example, ad blockers, private browsing, and VPNs were widely adopted. Unfortunately, these tools are somewhat or completely ineffective against the privacy threats participants expressed the greatest concern about, such as online observation by advertisers. Adoption of Tor Browser, which can protect against these privacy threats, was significantly lower. Thus, I decided to test using nudges to encourage adoption of Tor Browser in my Chapter 5 study. My results also showed that people have misconceptions about the protections offered by tools, even when they are already using those tools. My analysis of participants' free-text responses characterizes the diverse forms of these misconceptions. Based on the misconceptions I identified, I proposed design recommendations for nudging interventions. For example, I found that participants cited true aspects of tool functionality when explaining incorrect answers about tools' protections. Consequently, I recommend that promoters of privacy tools carefully enumerate the specific threats the tools can protect against, and warn people not to assume the tools can protect against other threats. I incorporated these design recommendations into the nudges in my Chapter 5 study.

Based on these findings, I conducted a study of implementation intention nudges designed to help people adopt Tor Browser, a browsing anonymity system. I describe this study in Chapter 5. I conducted a longitudinal, between-subjects field experiment to test nudges based on protection motivation theory (PMT), action planning implementation intentions (AP), and coping planning implementation intentions (CP). The experiment included four treatment conditions: a PMT-only condition, a PMT+AP condition, a PMT+AP+CP condition, and a control condition. My results show that the PMT-only treatment made participants 1.8x more likely to use Tor Browser in the week immediately following the intervention ($p = 0.026$), and 2.1x more likely to use Tor Browser when I followed up five weeks later ($p = 0.011$), as compared to the control group. I found that those in the PMT+AP+CP group who reported encountering challenges using Tor Browser, and hence were given opportunities to form coping plans, were 2.6x more likely to use Tor Browser in the following week ($p = 0.027$), as compared to the PMT+AP group. However, I did not find statistically significant evidence of the coping planning nudge increasing use of Tor Browser when I followed up four weeks later ($p = 0.678$). In addition, I did not find evidence of the action planning nudge increasing use of Tor Browser at either the one week follow-up ($p = 0.125$), or at the five week follow-up ($p = 0.211$), when comparing the PMT-only and PMT+AP groups. My results suggest that there are opportunities to increase adoption of Tor Browser using nudging techniques, particularly those based on protection motivation theory and coping planning implementation intentions. Together, the results of my Chapter 3 and Chapter 5 experiments show how different types of nudges might be used to encourage adoption of privacy- and security-enhancing tools.

In summary, my thesis is that:

Nudging interventions can motivate people to adopt security and privacy tools, and can help people start using those tools in the real world. By quantifying and comparing the effect of nudges based on implementation intentions and protection motivation theory, we inform their use in the field of computer security and privacy.

Chapter 2

Related Work

As ever more aspects of our lives are mediated by technology, the importance of digital security and privacy is increasingly obvious. The recurrence of security and privacy breaches shows that much progress is still needed. Security and privacy can be protected through regulation and technical measures, but changes to organizational culture and individuals' behavior may be equally important. Researchers in the field of usable security and privacy study ways to help people protect themselves from digital threats. For example, researchers have studied how to improve privacy notices [81], how to guide people towards choices that fit their preferences [93], and how a lack of usability can inhibit adoption of protective technologies [89, 182].

My research is centered around helping people adopt security and privacy tools. For a variety of reasons, public adoption of such tools often lags behind the technical state of the art (Section 2.1). In recent years, security and privacy researchers have begun testing various nudging techniques to guide people towards more secure and privacy protective behavior (Section 2.2). However, to the best of our knowledge, we are the first to study nudges based on implementation intentions in the field of security and privacy. Specifically, we test using implementation intentions to help people adopt security and privacy tools. In the following sections, we describe existing literature on implementation intentions (Section 2.3) and the related protection motivation theory (Section 2.4).

2.1 Obstacles to Adopting Security and Privacy Tools

Various software tools have been designed to protect people's security and privacy. However, adoption of such tools remains inconsistent [191]. In some cases, it might be rational for people not to adopt tools, based on the amount of protection provided and people's limited time and cognitive resources [21, 67, 68]. In other cases, it may be beneficial for users to adopt certain tools, but usability issues may impede adoption. In Whitten and Tygar's seminal paper, they describe end-users' struggles using PGP 5.0 for secure email [182]. Although the software had an aesthetically pleasing user-interface, it did not sufficiently convey the underlying concepts of public-key encryption. As a result, most users were unable to effectively use the software, and some made dangerous mistakes. The authors conclude that for security tools to be usable, they should convey a succinct mental model to their users. Researchers have evaluated the usability of many different security and privacy tools, including private browsing mode [3, 66, 190], password managers [11, 125], Signal [176, 189], and Tor Browser [51, 52, 89, 113]. The importance of conveying accurate mental models about the tools' protections is a recurring theme in this line of research. In accordance with these findings, we strive to help users form accurate mental models as we encourage them to adopt security and privacy tools.

In addition to tool-specific usability issues, adoption may be deterred by more general ways people think about risks and protections [5]. As we describe in Section 2.2, behavioral nudges may help counter these factors, and thereby increase adoption of protective technologies.

2.2 Nudges for Security and Privacy

Many types of behavioral biases may lead people to make decisions about security and privacy they may subsequently regret. For example, people may be unaware of the presence of threats to security and privacy, like the posibility of one's IP address being used to infer one's physical location, which may discourage them from taking protective actions. This is an example of the broader concept of information asymmetry, which describes cases where people don't have access to certain information which might inform their decision-making [8]. As a related example, consider someone who knows about the possibility of IP address-based inferences. They still might fail to fully consider the consequences of this practice, and so not take any protective actions. Alternatively, they may want to protect themselves, but might not know how to use a VPN or Tor Browser to hide their IP address, or they might not remember to consistently use such protective technologies. These are examples of the concept of bounded rationality, which describes how limited time or cognitive resources may impact people's decision-making [148]. The literature describes a variety of other cognitive and behavioral biases which may also cause people to act against their best interests. For example, optimism bias may lead people to underestimate their chance of being affected by a threat, and so discourage them from taking protective actions [92]. The potential negative impact of these biases is further compounded by the fact that security and privacy are not usually people's primary tasks, so people are likely to prioritize achieving their primary tasks (e.g., browsing the web) over protecting their security and privacy.

Work in psychology and behavioral economics has identified different ways these biases can be counteracted, while still allowing individual freedom of choice [166]. In their highly influential book, Thaler and Sunstein popularized the term *nudge* to refer to such interventions. They define nudges as:

> "...any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates."

Using nudges, influencers can guide people toward decisions that align with people's best interests, without prohibiting people from acting otherwise (e.g., not requiring that people use a VPN). Thaler and Sunstein refer to their approach as *libertarian paternalism.*

Researchers are increasingly studying how nudges can improve design for security and privacy [6]. For example, Almuhimedi et al. used nudges to mitigate the information asymmetry between users and the behaviors of apps on their devices [12, 13]. Their nudges were successful at encouraging users to reassess and restrict permissions settings. Frik et at. tested using nudges to overcome present bias [49]. They found that users given the option to be reminded later were less likely to completely dismiss prompts for security updates and 2FA configuration. Albayram et al. and Al Qahtani et al. used educational videos to motivate participants to enable lock screens on their smartphones [9, 10]. In both studies, the videos successfully motivated many participants to enable secure lock screens. However, nudges are not always effective [6], which shows the value of empirical research like ours. In our research, we introduce implementation intention nudges to the field of computer security and privacy, and compare them to the existing state of the art.

2.3 Implementation Intentions

An implementation intention is a concrete plan to achieve some goal, which is triggered by situational factors like time or location [56]. For example, if one has the goal of exercising to reduce one's risk of heart disease, one might form an implementation intention to facilitate exercise: "If it is Wednesday at 5:30pm, then I will jog home from work" [145]. Particularly in the medical domain, there is strong evidence that encouraging people to form implementation intentions has a powerful effect on people achieving their goals. Implementation intentions have been effective in encouraging people to perform breast cancer self exams (an increase from 53% to 100%) [122], in encouraging people to exercise for 20 minutes a week (from 39% to 91%) [109], and in getting people to take actions in many other contexts [107, 108, 112, 123].

In their review of health behavior-related implementation intentions, Sheeran et al. identify potential mediators and moderators of implementation intentions, and make suggestions for operationalizing implementation intentions [145]. In particular, the literature suggests that implementation intentions work by helping people recognize opportunities for action [1, 179] and by helping people perform the action automatically when the opportunity does arise [24, 90]. Also, implementation intentions are most likely to be effective when a person has a strong commitment to both their plan [145] and to the goal which motivates the plan [146]. Implementation intentions offer the greatest benefit when the goal to be achieved is challenging [57].

There are many different ways to design implementation intentions to help people achieve their goals. In recent years, evidence has mounted for the importance of helping people plan how to overcome anticipated obstacles, a form of planning referred to as *coping planning* [22, 152]. For example, a person whose goal is exercising more might anticipate inclement weather interfering with their ordinary exercise routine, so they could make an indoor exercise plan for rainy days. Traditional implementation intentions, concerned with helping people initiate actions without special consideration to obstacles, are referred to as *action plans*. Coping plans and action plans can be used individually [4, 109, 173] or together [95, 151].

2.4 Protection Motivation Theory

Implementation intentions are designed to help people achieve their goals, and the efficacy of implementation intentions depends on people's motivation to achieve those goals. Based on examples from the literature [109], we drew on Protection Motivation Theory (PMT) [100, 140, 141] to motivate participants in our studies. PMT has been widely applied in the medical field [110, 188] and in computer security [9, 10, 27, 147, 149, 153]. PMT proposes that people are more likely to take action to protect themselves from a threat when they perceive that the threat is severe (i.e., greater perception of *threat severity*), that they are susceptible to the threat (i.e., greater perception of *threat susceptibility*), they they are afraid of the threat (i.e., greater *fear arousal*), that the action they could take is not too difficult to perform (i.e., greater perception of *self-efficacy*), that the action they could take will be effective in protecting against the threat (i.e., greater perception of *response efficacy*), and that the costs of taking the action are low (i.e., lower perception of *response costs*) [110, 188]. An intervention drawing on PMT can help people form accurate perceptions of these factors, and may thereby help motivate people to act.

Chapter 3

Nudges to Increase Adoption of Secure Mobile Payments

3.1 Overview

Our primary goal was to determine whether nudges based on action planning implementation intentions (AP) and protection motivation theory (PMT) can increase real-world adoption of secure mobile payment systems. A secondary goal was to measure the effect of our nudging interventions on participants' attitudes about mobile payment systems.

Action plans are relevant to this domain because people must remember to use secure mobile payments at the time of checkout. PMT is relevant because it is often paired with implementation intentions in order to motivate people [109], since implementation intentions are most effective when people are highly motivated [146].

This study was published at SOUPS 2020 [160].

3.2 Methodology

To measure the effect of action plans on behavior, we would ultimately conduct a randomized controlled experiment to test our nudging interventions. However, we began by running an qualitative interview study to refine the design of our nudges. Our findings from these interviews helped increase the validity of our subsequent experiment. In our qualitative interviews, we interviewed users of Apple Pay, Google Pay, and Samsung Pay. In our controlled experiment, we focused on users of Apple Pay.

3.2.1 Qualitative Interviews

This portion of our study included three surveys and two interviews (illustrated in Figure A.1 in the appendix). We recruited participants from Craigslist and Carnegie Mellon University's participant pool. Survey #1 gathered information about users' devices, prior use of payment methods, perceptions of the likelihood and severity of card information theft and fraud, prior experience with card information theft and fraud, and demographic information. We reasoned that our nudges would have the largest impact on people who were not already using Pay,[1] but whose phones were compatible with Pay and were likely to have opportunities to use Pay. Thus, we screened out participants who reported having used Pay in a physical location in the past month, we required that participants had made at least one payment with a credit or debit card in a physical location in the past month, and we required that their smartphone be compatible with either Apple Pay, Google Pay, or Samsung Pay.

We invited a diverse subset of qualifying participants to participate in a semi-structured interview (Interview #1). In accordance with purposive sampling, our attempt to balance several factors of interest (e.g., phone type, age, occupation, fraud-related perceptions, etc.) influenced our choice of who to invite to the interview. 20 participants attended Interview #1. 75% of our participants were female, their median age was 26.5, 55% had iPhones, 25% had Samsung phones, and 20% had other Android phones. The interview started with a discussion of prior experiences with card fraud, card information theft, and prior experiences with Pay. Next, the interviewer described recent cases of card information theft from major retailers, and the potential consequences of such theft for the participant. This intervention was included in order to help participants develop an accurate perception of their *susceptibility* to card fraud and the potential *severity* of card fraud, two elements of *threat appraisal* that protection motivation theory (PMT) suggests are associated with protective behavior [110]. Then, the interviewer described how Pay may protect against card information theft, presented the participant with instructions for setting up and using Pay, and gave the participant the opportunity to set up Pay if they wanted to. This intervention was included in order to help participants understand how Pay may help protect them from card fraud and to give them confidence that they can use Pay, influencing perceptions of *response efficacy* and *self-efficacy*, two additional elements of PMT. Next, participants were given an opportunity to form an action planning implementation intention by filling out a paper template. The template encouraged participants to plan where they might use Pay in the coming week and to mentally rehearse using Pay in these locations. These activities were designed to help mentally activate participants' plans to use Pay when they were in these locations [56]. Finally, participants were given the opportunity to express a strong commitment to their plan, which prior work suggests increases the efficacy of implementation intentions [56, 155]. The template was similar in content to the template in our controlled experiment (see Figure 3.4).

One week after completing Interview #1, participants were sent Survey #2, which asked whether participants had set up Pay after the interview, whether they had tried to use Pay, and whether they had successfully used Pay. Participants who completed Interview #1 and Survey #2 were compensated with a $15 Amazon e-gift card.

Participants who completed Survey #2 and who had set up Pay on their phones were invited to Interview #2, which was designed to understand people's experiences using Pay or their reasons for not using it. We also asked questions about whether participants followed their action plans and whether they found the plans to be helpful. Participants in Interview #2 were compensated with an additional $15 Amazon e-gift card.

Four weeks after completing Survey #2, participants who had set up Pay on their phones were invited to take Survey #3, which asked whether participants had used Pay in the past week. We also asked whether participants thought they were likely to use Pay in the future. Our surveys

[1] In the rest of this manuscript, we use *Pay* to refer to Apple Pay, Google Pay, and Samsung Pay generically.

and interview scripts are included in the appendix (Section A.1.1-A.1.5).

Limitations

To protect external validity, it was important that participants understood that they were not required to set up Pay, use Pay, form an action plan, or follow their action plan. We iterated on the design of our interview protocol until we arrived at language which we thought communicated this clearly to participants. However, although setting up Pay was not required to receive compensation for Interview #1 and Survey #2, participants who never set up Pay were not invited to Interview #2 or Survey #3. Although we tried to disguise the qualification criteria for Interview #2 and Survey #3 from participants, participants may have inferred that some action on their part would be required to qualify, and some asked us directly in the interview. To ensure this was not a threat to validity in our controlled experiment, we emphasized that participants' compensation would not be affected by their use or non-use of Pay.

The generalizability of our findings might be impaired by our relatively small sample size ($n = 20$) and recruitment from the geographic area around our institution. To mitigate this, we used purposive sampling to recruit a diverse set of participants. Further, we recruited a much larger set of participants in our controlled experiment.

3.2.2 Controlled Experiment

In the second part of our study, we measured the effect of our nudging interventions using a randomized controlled experiment with a sufficient number of participants ($n = 411$) to determine statistical significance. For ease of recruitment and to reduce the complexity of our protocol, we choose to focus on Apple Pay.

Our design included three experimental conditions. In our control group, we did not try to motivate participants to use Apple Pay. In our PMT group, we presented participants with information about the threat of card fraud (Figure 3.1) and the mitigation of using Apple Pay (Figure 3.2 and Figure 3.3) in order to motivate them to use Apple Pay. This motivational intervention was based on protection motivation theory [110], as described in Section 3.2.1. In our PMT+AP group, we presented participants with the motivational intervention of the PMT group in addition to an opportunity to form an action planning implementation intention. This opportunity took the form of a template we designed to help participants plan where they could use Apple Pay, as shown in Figure 3.4. We did not test an action planning nudge without a PMT nudge because the literature suggests that implementation intentions are only effective when participants are motivated [56].

Our study consisted of three surveys hosted on Qualtrics using recruitment from Prolific (see Figure A.2 and Sections A.2.1–A.2.3 in the appendix). Survey #1 was designed to determine eligibility for Survey #2 and Survey #3. The only requirements for taking Survey #1 were that participants live in the United States, speak English, be at least 18 years old, and have an iPhone. We thought our nudges would have the largest impact on people who were not actively using Apple Pay, but whose phones were compatible with Apple Pay and who were likely to have opportunities to use Apple Pay in the week ahead. Thus, to be eligible for participation in Survey #2 and #3, participants must have purchased their iPhone in the United States, owned an iPhone model compatible with Apple Pay (iPhone 6 or newer), must have had a version of iOS compatible with Apple Pay (iOS 12.2 or higher), in the last week must have made an in-person payment in a physical location using their credit or debit card, in the last week must not have made an in-person payment in a physical location using Apple Pay, and they must have passed a simple attention check.

Shortly after completing Survey #1, participants were invited to Survey #2, which contained our randomly assigned experimental conditions. The control group saw only a short description of Apple Pay. The PMT group was provided with a description of the threat of credit and debit card information theft and fraud, and information about the mitigation of using Apple Pay. This information included written instructions about how to set up and use Apple Pay, a short video showing how to use Apple Pay, and an FAQ addressing questions participants asked in our qualitative interviews. We encouraged participants to set up Apple Pay if they wanted to, but we reassured participants that their compensation would not be affected if they did not set it up. The PMT+AP group received the same information as the PMT group, but was also given a chance to form a plan to use Apple Pay. Near the end of the survey, participants in the treatment groups were given links to the information about Apple Pay and their plan for using Apple Pay, with the option to request that these links be sent to them via Prolific. Participants in all treatment groups were asked demographic questions and questions related to their perceptions of Apple Pay and card fraud.

Survey #3 was sent to participants one week after they completed Survey #2 in order to measure whether they had used Apple Pay. We asked participants whether they had registered a card in Apple Pay, whether they had made an in-person payment using Apple Pay, and about other details related to their use of Apple Pay and other payment methods.

Our goal was to pay participants $12/hour, so compensation was determined based on estimated duration of our surveys. Survey #1 was estimated to take five minutes, so compensation was $1. Survey #2 was estimated to take up to 30 minutes (accounting for time potentially spent outside the survey setting up Apple Pay), so compensation was $6. Survey #3 was estimated to take five minutes, so compensation was $1. Participants only received compensation for Survey #2 and Survey #3 if they completed both surveys within three days of being invited.

We conducted an a priori power analysis using G*Power to determine our target number of participants [47]. We planned three chi-square tests of independence to compare the use of Apple Pay between the three treatment groups. In order to detect a small to medium effect size ($w = 0.2357$, informed by the effect size seen in our interviews), with a Bonferroni corrected $\alpha = 0.05 \div 3 = 0.01667$, power=0.9, and df=1, we determined that we needed 122 participants in each treatment.

We preregistered our protocol on The Open Science Framework prior to collecting any data [159]. After preregistering but before collecting any data, we made two small edits to the survey text. Also, before collecting any Survey #3 data, we added a "using another payment method" option to Q18, Q19, and Q20 in Survey #3. In our preregistration, we described using a Bonferroni correction, but switched to the Holm-Bonferroni method as it controls the experiment's Type I error rate at the same level as a Bonferroni correction, while having a lower Type II error rate [2]. Otherwise, we conducted our study as preregistered.

> There have been many big hacks where credit and debit card information was stolen from retailers. For example, Target [177] was hacked in 2013, Home Depot [135] was hacked in 2014, and Saks Fifth Avenue [55] was hacked in 2018. Information about millions of cards was stolen in these hacks. If criminals get your credit or debit card information, they might use that information to make fraudulent purchases. If you notice fraudulent purchases on your credit card, you can probably get refunded. But if the purchases are made on your debit card, you might not be able to get your money back [32]. In any case, you would need to get a replacement card with a new number, which would be inconvenient.

Figure 3.1: In our experiment, participants in the PMT and PMT+AP groups were shown this text to inform them about the threat of card fraud. This text was included in order to help participants develop accurate perceptions of threat susceptibility and threat severity, two elements of PMT [110].

> Thankfully, there are steps you can take to prevent your card information from being stolen and to protect yourself from card fraud. One of the best things you can do is to start using Apple Pay. Instead of paying by swiping or inserting your card, you can make payments through your phone, which adds an extra layer of security. Payments made with Apple Pay will still be charged to your credit or debit card, but because the payments go through Apple Pay, your card number is not shown to or recorded by retailers [15]. This means that your card number cannot be stolen from transactions made with Apple Pay. If your phone is stolen, the thief will not be able to make payments because Apple Pay is protected by your fingerprint and lock screen PIN. Although no system is perfectly secure, security experts generally agree that Apple Pay is more secure than paying with credit or debit cards [96]. Apple Pay takes just a few minutes to set up, and is widely accepted. As of this year, Apple Pay is accepted in 65% of retail locations [14] in the United States. For example, ALDI grocery, CVS pharmacy, and Starbucks all accept Apple Pay.

Figure 3.2: In our experiment, participants in the PMT and PMT+AP groups were shown this text to inform them about how using Apple Pay can protect them from card fraud. This text was included in order to help participants understand how Apple Pay may help protect them from card fraud and to give them confidence that they would have opportunities to use Apple Pay, influencing perceptions of response efficacy and self-efficacy, two additional elements of PMT [110].

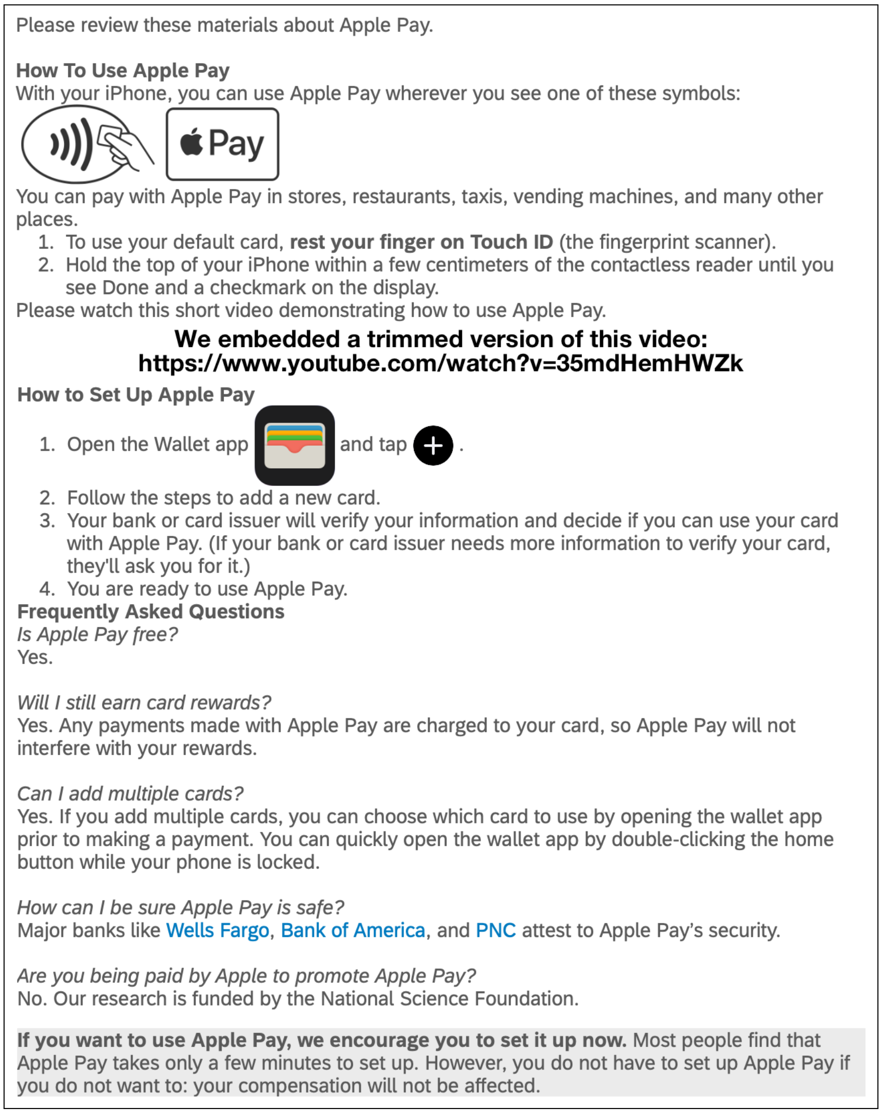

Please review these materials about Apple Pay.

How To Use Apple Pay
With your iPhone, you can use Apple Pay wherever you see one of these symbols:

You can pay with Apple Pay in stores, restaurants, taxis, vending machines, and many other places.

1. To use your default card, **rest your finger on Touch ID** (the fingerprint scanner).
2. Hold the top of your iPhone within a few centimeters of the contactless reader until you see Done and a checkmark on the display.

Please watch this short video demonstrating how to use Apple Pay.

We embedded a trimmed version of this video:
https://www.youtube.com/watch?v=35mdHemHWZk

How to Set Up Apple Pay

1. Open the Wallet app and tap +.
2. Follow the steps to add a new card.
3. Your bank or card issuer will verify your information and decide if you can use your card with Apple Pay. (If your bank or card issuer needs more information to verify your card, they'll ask you for it.)
4. You are ready to use Apple Pay.

Frequently Asked Questions
Is Apple Pay free?
Yes.

Will I still earn card rewards?
Yes. Any payments made with Apple Pay are charged to your card, so Apple Pay will not interfere with your rewards.

Can I add multiple cards?
Yes. If you add multiple cards, you can choose which card to use by opening the wallet app prior to making a payment. You can quickly open the wallet app by double-clicking the home button while your phone is locked.

How can I be sure Apple Pay is safe?
Major banks like Wells Fargo, Bank of America, and PNC attest to Apple Pay's security.

Are you being paid by Apple to promote Apple Pay?
No. Our research is funded by the National Science Foundation.

If you want to use Apple Pay, we encourage you to set it up now. Most people find that Apple Pay takes only a few minutes to set up. However, you do not have to set up Apple Pay if you do not want to: your compensation will not be affected.

Figure 3.3: In our experiment, participants in the PMT and PMT+AP groups were shown these details about Apple Pay. The instructions contained information about either Touch ID or Face ID, based on which technology the participant's phone was compatible with. These instructions were designed to positively influence perceptions of self-efficacy.

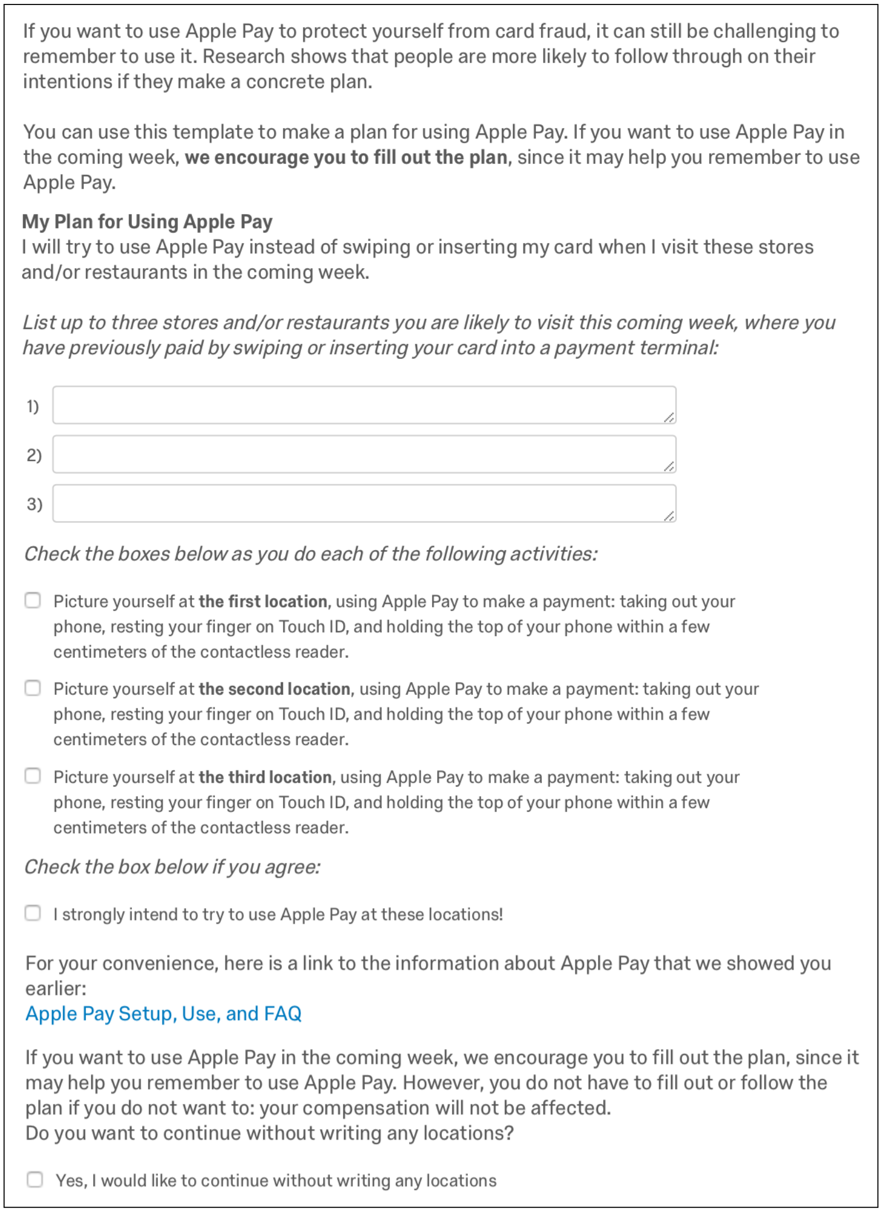

If you want to use Apple Pay to protect yourself from card fraud, it can still be challenging to remember to use it. Research shows that people are more likely to follow through on their intentions if they make a concrete plan.

You can use this template to make a plan for using Apple Pay. If you want to use Apple Pay in the coming week, **we encourage you to fill out the plan**, since it may help you remember to use Apple Pay.

My Plan for Using Apple Pay
I will try to use Apple Pay instead of swiping or inserting my card when I visit these stores and/or restaurants in the coming week.

List up to three stores and/or restaurants you are likely to visit this coming week, where you have previously paid by swiping or inserting your card into a payment terminal:

1)

2)

3)

Check the boxes below as you do each of the following activities:

□ Picture yourself at **the first location**, using Apple Pay to make a payment: taking out your phone, resting your finger on Touch ID, and holding the top of your phone within a few centimeters of the contactless reader.

□ Picture yourself at **the second location**, using Apple Pay to make a payment: taking out your phone, resting your finger on Touch ID, and holding the top of your phone within a few centimeters of the contactless reader.

□ Picture yourself at **the third location**, using Apple Pay to make a payment: taking out your phone, resting your finger on Touch ID, and holding the top of your phone within a few centimeters of the contactless reader.

Check the box below if you agree:

□ I strongly intend to try to use Apple Pay at these locations!

For your convenience, here is a link to the information about Apple Pay that we showed you earlier:
Apple Pay Setup, Use, and FAQ

If you want to use Apple Pay in the coming week, we encourage you to fill out the plan, since it may help you remember to use Apple Pay. However, you do not have to fill out or follow the plan if you do not want to: your compensation will not be affected.
Do you want to continue without writing any locations?

□ Yes, I would like to continue without writing any locations

Figure 3.4: In our experiment, participants in the PMT+AP group were shown this action planning template. The template encourages participants to plan where they might use Apple Pay in the coming week and to mentally rehearse using Apple Pay in these locations. These activities should help mentally activate participants' plans to use Apple Pay when they are in these locations [56]. Finally, participants are given the opportunity to strongly commit to their plan [56, 155].

Limitations

One limitation of our study is our reliance on self-reported data. In particular, it is possible that participants did not accurately report whether they used Apple Pay between taking Survey #2 and #3. To encourage honesty, at the beginning of Survey #2 and Survey #3 we included text which encouraged participants to answer honestly and reassured them that there were no right or wrong answers. We also included attention checks in all our surveys. Fifteen participants (2%) failed our Survey #1 attention check and so were not invited to the subsequent surveys, but no participants failed our Survey #2 or Survey #3 attention checks. Another threat to validity is the possibility that some participants may have thought that setting up or using Apple Pay was not optional. To avoid that misconception, we included text assuring participants that setting up or using Apple Pay was not required and would have no effect on their compensation. One threat to the generalizability of our findings is the fact that crowd workers have been shown to differ from the general population. Our use of Prolific was informed by recent findings that Prolific workers are more diverse and honest than Mechanical Turk workers [126]. See Table B.1 in the appendix for a summary of demographic information about our participants.

3.3 Analysis

3.3.1 Qualitative Interviews

We used thematic analysis to analyze transcripts of our interviews and our survey's open-text responses [25]. Two of the authors reviewed these materials together and collaboratively developed a codebook. To ensure that the codes we developed later were consistently applied to the materials we analyzed earlier, one author then re-reviewed all the materials. Since our goal for this portion of our study was to gather rich, qualitative data, we did not attempt to calculate measures of annotator reliability [105]. Table A.1 in the appendix contains our final codebook and the frequencies of our codes.

3.3.2 Controlled Experiment

We collected 670 valid responses to Survey #1, and invited 430 qualifying participants to participate in Survey #2. Of the 430 participants invited to Survey #2, 418 completed Survey #2 and 411 went on to complete Survey #3, for an overall dropout rate of 4%.

After completing data collection, we conducted our preregistered hypothesis tests to compare use of Apple Pay between our three treatment conditions. We also conducted a series of exploratory analyses.

3.4 Results

3.4.1 Qualitative Interviews

Below we summarize key takeaways from our survey and interview data. Although in some cases we report the frequency of codes, due to our use of purposive sampling in selecting participants, it would be inappropriate to assume that these frequencies correspond to the frequencies that might be observed in the general population.

Use of Pay

We received 288 complete responses to Survey #1. Among these respondents, only 34.7% reported using Pay sometime in the past, and a mere 23.6% reported using Pay in the past month. We recruited only respondents who had not used Pay in the last month for Interview #1. In Interview #1, nearly all participants (19/20) said they had heard of Pay before our study, but only one participant reported using it to pay in a physical location before. Multiple participants mentioned seeing Pay in advertisements, seeing it on their phone, seeing it as a payment option, using it for digital purchases, or knowing that friends or family use it. This widespread awareness of Pay makes sense, considering that many smartphones come with Pay preinstalled [61, 143] and that iPhones include persistent reminders to set up Apple Pay [80].

Prior to Interview #2, 11/20 participants had Pay set up on their phone, and so could have used it before Interview #2. One participant had set up Pay prior to Interview #1, four set it up in Interview #1, and six set it up after Interview #1. Participants gave a variety of different reasons for not setting up Pay including being too busy, not thinking they needed it, wanting to do more research, and wanting to consult their partner. Between Interview #1 and Interview #2, seven participants used Pay. Three of these participants used Pay successfully at at least one of the locations in their plan.

To understand whether participants were likely to continue using Pay after our study, four weeks after completing Survey #2 we sent Survey #3 to all participants who had set up Pay. In response to Survey #3, three participants indicated that they had successfully used Pay to make a payment in a physical location in the prior week.

Our results suggest that despite widespread awareness of Pay, most people are not using it regularly. However, after being exposed to our nudges in Interview #1, a substantial percentage of participants (35%) used Pay at least once during the remainder of our study. Furthermore, responses to Survey #3 indicate that our nudges may increase use of Pay long after the initial intervention. These results encouraged us to move forward with the controlled experiment.

Perceptions of Threat Susceptibility and Severity

All but one participant recounted their own or others' experiences with card fraud. Fewer participants (10/20) recounted experiences with card information theft. When asked to describe experiences with card information theft, seven participants instead described cases of card fraud or theft. This makes sense, given that card information theft can be difficult for individuals to detect directly.

After we described recent hacks in which credit and debit card information was stolen and the possible consequences of having one's information stolen, we asked questions to gauge participants' levels of concern about and perceptions of susceptibility to card fraud and information theft.

Participants expressed varied opinions about their susceptibility to these threats. Eleven participants expressed that information theft happened frequently (P14: "It just seems like it does happen so frequently..."), but three participants said such occurrences were infrequent (P6: "Cause even like ... the hacking things you mentioned, I mean they're once in a blue moon."). Three participants said their behavior made theft or fraud more likely, but nine others thought their behavior lowered their likelihood of suffering from card information theft or fraud.

Participants described a number of negative outcomes associated with card theft and fraud, including the hassle and stress of dealing with it, feelings of anger and helplessness, loss of money due to theft or overdraft fees, and the fear that additional bad things might happen to them. Participants also mentioned that their level of concern would depend on the size of the fraudulent purchase and whether the purchase was on their credit or debit card. Ten participants expressed confidence that their card issuer would help them resolve fraudulent purchases, and two even thought they would be refunded under all circumstances. It is potentially a misconception to believe that fraudulent charges will be refunded in all cases, since U.S. law does not require this of card issuers [32].

Our takeaway is that while most participants have a high level of awareness of the possibility of card fraud, some people remain under-informed and might benefit from additional information.

Perceptions of Self-efficacy

Some participants thought Pay setup was easy, but others encountered difficulties. In particular, two participants were confused by Apple Pay's ability to automatically add card details using the phone's camera and three mentioned interacting with their bank to approve registering their card as a challenge. Additionally, two participants found that certain cards simply could not be added to Pay. Seven participants said that setup or use would be a challenge, and would require practice, learning, or attention to detail.

Eleven participants said they did not (or might not) have opportunities to use Pay because they did not go shopping, did not have enough money, or due to other reasons.

Participants described different challenges they might (or did) encounter in stores using Pay. First, stores might or might not accept Pay. Second, participants might not remember to use Pay, suggesting an opportunity for action plans to help in this area. Third, participants might experience difficulty using Pay. Despite our written instructions, some participants still had questions about how to use Pay. Thus, we included a short video alongside written instructions in our controlled experiment. Participants also described positive aspects of Pay. Some participants expressed that Pay was easy to use, that it would allow them to not carry or take out their cards or wallet, that it would be a good backup option if they didn't have a card, and that it would be fun to try something new.

Two usability challenges in particular may be of interest, due to their potential generalizability: the case of accidental activation and the case of failure to activate. Four of our participants who set up Apple Pay described accidentally activating it and not knowing why this was happening. Not understanding this accidental activation alarmed at least one of our participants (P19: "The credit thing keeps popping up whenever I angle my phone a certain direction. I wonder where it's sending my credit info each time."). It is possible to open Apple Pay by either double-clicking the home button when the phone is locked or by bringing the phone in proximity to an NFC reader (even if the NFC reader is not a payment terminal). To address some of this confusion, we added the double-clicking functionality to our instructions in subsequent interviews and in our experiment. One of these same participants (P11) also experienced the problem of Apple Pay not activating. At one location, this participant reported having to scan their phone twice before it worked. At another location, the participant was ultimately unsuccessful using Apple Pay, concluding that it must not have been supported and expressing frustration with this failure mode: "What happens when it doesn't work is nothing happens. It just sits there. And it doesn't even apologize. You know it doesn't say anything on it. 'Oops, sorry. Try again.' Nothing like that." Unfortunately, due to the lack of an NFC signal in the case when a terminal does not support NFC payments, it is hard to imagine a technical solution to this kind of silent failure mode. Thus, while some of these usability challenges may be addressable through education, some may be inherent to the technology.

Perceptions of Response Efficacy

Most participants (14/20) expressed some confidence in the security properties of Pay that we described. However, nine participants also expressed concern about Apple, Google, or their phone being hacked. P16 cited their previous experience having their iTunes account hacked as a reason for not believing that Apple Pay would protect their card information: "[T]he only time I've been hacked was with an Apple product. That's the only reason. ... [T]he only time I had a fraudulent charge was when I was with an Apple product." Interestingly, this participant also recognized that the hack was likely due to their choice of a weak password, saying: "I guess my password wasn't as secure as I thought it was." P11 said that "I feel as if the phone is more vulnerable than the computer." P8 expressed a more concrete concern about NFC signal skimming, expressing concern that "...in a physical store ... the person behind you can actually take your information if they know what they're doing on the phone."

Despite participants' concerns about hacking, Apple Pay is designed to be resistant to hacking: card information is not stored with Apple after the initial enrollment process, mitigating the risk of data breach, device-specific Device Account Numbers are stored on each phone's Secure Element, protecting against phones being compromised, and user interaction is required before making payments [15]. Google Pay and Samsung Pay employ similar protections [62, 142]. Of course, attacks that can thwart these protections are possible (e.g., a persistent threat on Apple's servers), but such attacks would require substantially more resources than simply adding card skimmers to point of sale terminals. Communicating useful mental models to non-technical users remains an open research area [178]. Our participants' responses point to the challenge of communicating complex threat models to a general audience.

Awareness of Protection Actions

Participants demonstrated awareness of many different ways they could protect themselves from credit and debit card information theft and fraud. The most prevalent actions involved working with one's card issuer, such as reporting fraudulent purchases or receiving a new card. Actions involving physical awareness (e.g., looking for card skimmers), monitoring card statements for unauthorized transactions, protecting access to one's account (e.g., with a strong password), using cash, or using a credit card (e.g., due to liability protections) were also common. Interestingly, two participants brought up the possibility of using Apple Pay to protect themselves before we had described it as being a secure payment method (but after we had asked them whether they had used it). P18 even gave an accurate explanation of why Apple Pay might be more secure: "Maybe I could use Apple Pay or something. Then if I don't give my card information directly to these companies or grocery stores, if I go via a secure party like Apple Pay, it should be a good option."

Our overall takeaway is that most participants are aware of some ways they can protect themselves from card information theft and fraud. Unfortunately, prior work and the continued profitability of card fraud suggest that people's ability to protect themselves is limited (e.g., password re-use is prevalent [124]). In addition, most participants seemed unaware that Pay could protect them before we explained that it could, suggesting our information about Pay may be helpful.

Effectiveness of Action Plans

All participants were given the opportunity to form an action planning implementation intention to help them remember to use Pay. 16/20 participants wrote or described at least one location where they might use Pay. About half of participants checked or otherwise indicated that they performed at least one mental rehearsal activity. As we conducted interviews, we refined the way we introduced the plan to communicate that filling out and following the plan were not mandatory, but that filling out the plan was encouraged if the participant wanted to remember to use Pay. Participants described several obstacles to forming an action plan, including not being able to think of places they would visit, not having decided whether they wanted to use Pay, and simply thinking the plan wouldn't be helpful for them. In addition, four participants had at least some difficulty remembering their plan in Interview #2. The act of forming a plan seemed to help four participants understand where Pay could be used. For example, P11 realized that it would be difficult to use Pay at a restaurant where waiters collect cards for payment processing. Participants also described other things that could remind them to setup or use Pay, including receiving notifications from Google Pay about availability, adding Pay to their shopping list, and putting the Pay app on their home screen.

Of the seven participants who used Pay between Interview #1 and Interview #2, three used Pay successfully at at least one of the locations in their plan. As the majority of participants were able to form plans, and some of the participants who formed plans went on to use Pay at their planned locations, we thought that our action plan template was worth testing in our controlled experiment. At the same time, our plan may be unhelpful to participants who have difficulty thinking of locations they are likely to make payments in the coming week, and is almost certain to be unhelpful for participants who simply decide not to use Pay. Since Pay is not available in all locations, it is unsurprising that many participants had questions about where they could use Pay. As part of our description of Pay, we described just four popular locations where Pay can be used in our city. With a more comprehensive list of locations, it might be possible to develop an interactive plan template which could contribute to greater awareness of where Pay is available.

Misconceptions and Other Concerns

Our interviews helped us identify a number of misconceptions related to Pay. For example, four participants thought Pay might interfere with their credit card rewards (P5, P9, P12, P16), four participants thought our study was affiliated with Apple (P11, P16, P19, P20), three participants wondered if Pay cost something (P5, P8, P15), and one participant thought Pay might prevent them from getting receipts (P12). In addition, P7 thought Samsung Pay was a credit card and two participants confused Apple Pay with iPad-based point of sale terminals (e.g., Square). We addressed several of these misconceptions in a "Frequently Asked Questions" section in our controlled experiment (Figure 3.3).

Pay on watches offers the same level of security as on phones, but with potentially greater convenience. Thus, we were surprised that all three of the participants we spoke with about their smartwatches expressed skepticism about using Pay with their watches. P1 thought they would start using Apple Pay on their iPhone, because they thought they would need to practice the motion of making payments with their Apple Watch. P12 thought Apple Pay would be less secure on their Apple Watch than on their iPhone because their Apple Watch did not have a fingerprint reader. Neither P1 nor P12 set up Apple Pay during our study. In Interview #1, P15 was worried that setting up Samsung Pay might allow transactions to be made through their Samsung Watch without their knowledge, due to the fact that their watch did not have a PIN. In Interview #2, P15 said they had figured out how to add a PIN to their watch, and after doing so they proceeded with the setup of Samsung Pay.

3.4.2 Controlled Experiment

Use of Apple Pay

We conducted three chi-square tests of independence to compare the use of Apple Pay between our three treatment groups, as shown in Table 3.1. We used the Holm-Bonferroni method to control Type I error.[2] Participants in the PMT+AP group, who saw our PMT with an action planning nudge, were 3.91x more likely to use Apple Pay than our control group ($p < 0.001$). Participants in the PMT group, who saw only the PMT nudge, were 2.34x more likely to use Apple Pay than our control group ($p = 0.020$). Both of these differences were statistically significant at $\alpha = 0.05$. However, we did not find a statistically significant difference in use of Apple Pay between the PMT and PMT+AP groups ($p = 0.085$).

[2] In our preregistration, we described using a Bonferroni correction. We switched to the Holm-Bonferroni method because it controls the experiment's Type I error rate at the same level as a Bonferroni correction, while having a lower Type II error rate [2]. Using a simple Bonferroni correction, only our Control vs PMT+AP comparison would have been found significant. See [17] for further discussion of the Bonferroni correction.

Comparison	**Use of Apple Pay**	**Odds Ratio**	**p-value**
Control vs PMT+AP	8.7% vs 27.2%	3.91	<0.001
Control vs PMT	8.7% vs 18.3%	2.34	0.020
PMT vs PMT+AP	18.3% vs 27.2%	1.67	0.085

Table 3.1: Comparisons between the percent of participants who reported using Apple Pay in each of our treatment conditions. Per convention, the reported odds ratios correspond to large, medium, and small effect sizes, respectively [163].

Therefore, we have evidence that our interventions in both the PMT+AP and PMT groups had large and medium effects on participants' use of Apple Pay, respectively. Since the treatment conditions only differed in their inclusion of our educational materials (Figures 3.1, 3.2, and 3.3) and our action planning template (Figure 3.4), we can conclude that these differences are what made participants more likely to report using Apple Pay. Although we did not find statistically significant differences between the PMT and PMT+AP groups, our findings suggest ($p = 0.085$) that the inclusion of the action plan had a small effect on increasing the PMT+AP participants' use of Apple Pay.

Effects of Interventions on Attitudes

After testing for the effect of our interventions on participants' use of Apple Pay (Section 3.4.2), we decided to test for other potential effects, as shown in Table 3.2. We used Kruskal-Wallis tests for all variables except whether participants registered a card, where we used a chi-square test of independence. Details of the statistically significant results are shown in Figures 3.5, 3.6, and 3.7. Effect sizes are given as epsilon-squared (ϵ^2) estimates [102, 174]. Insignificant results are included in Figures A.3–A.7 in the appendix. Post-hoc Dunn tests significant at $\alpha = 0.05$ after Holm-Bonferroni correction are bolded.

As shown in Figure 3.5, our treatments had a dramatic effect on participants' agreement that Apple Pay would protect them from card fraud ($\epsilon^2 = 0.241$, $p < 0.001$). In the control group, only 37% of participants agreed that Apple Pay would protect them, whereas in both treatment groups over 84% agreed. Thus, we have strong evidence that our information was effective at correcting people's misconceptions about Apple Pay's security [73]. As illustrated in Figure 3.6, our treatments increased participants' expressed intentions to use Apple Pay, and action plans were even more effective at increasing intention than PMT alone ($\epsilon^2 = 0.172$, $p < 0.001$). Finally, Figure 3.7 shows that our treatments had a small effect on participants' belief that Apple Pay would be useful for making payments ($\epsilon^2 = 0.015$, $p = 0.047$).

Intention vs Behavior

Comparing participants' Survey #2 responses to their Survey #3 responses gave us insight into how participants' stated intentions to act did or did not translate to actual behavior.

First, we measured how stated intention to register a credit or debit card in Apple Pay translated to actually setting up Apple Pay. As shown in Figure 3.8, while half of those who expressed a strong intention to register a card did so, those who expressed weaker intentions were corre-

			Post-hoc Test p-values		
Variable	**p-value**	ϵ^2	**Control vs PMT**	**Control vs PMT+AP**	**PMT vs PMT+AP**
Perception of threat severity	0.932	<0.001			
Perception of threat susceptibility	0.881	0.001			
Perception of self-efficacy	0.523	0.003			
Perception of response-efficacy	**<0.001**	0.241	**<0.001**	**<0.001**	0.880
Expressed intention to use Pay	**<0.001**	0.172	**<0.001**	**<0.001**	**0.001**
Perception of Pay's usefulness	**0.047**	0.015	0.092	0.078	0.856
Self-consciousness using Pay	0.628	0.002			
Variable	**p-value**	V			
Registered card by end of study	0.237	0.084			

Table 3.2: The results of hypothesis tests measuring whether these variables differed between our treatment groups. The p-values in the second column correspond to overall tests of significance; results significant at $\alpha = 0.05$ are bolded, representing rejection of the null hypothesis. The third column contains estimates of effect sizes: ϵ^2 values for Kruskal-Wallis tests, and Cramer's V for the chi-square test of independence [102, 174]. For results which were significant overall, we conducted post-hoc tests to determine which treatment groups were significantly different from each other. The remaining columns contain the results of Holm-Bonferroni corrected post-hoc Dunn tests; results significant at $\alpha = 0.05$ are bolded.

spondingly less likely to register a card. In particular, note that less than half of the participants who responded with "Agree" actually set up Apple Pay by the time of Survey #3.

Next, we compared stated intention to use Apple Pay to actual use of Apple Pay. We performed a chi-square test of independence and found that those who indicated they intended to use Apple Pay in the week ahead were more likely to use Apple Pay than those who did not ($p < 0.001$). However, as shown in Figure 3.9, many participants who expressed an intention to use Apple Pay did not do so. This reinforces our belief that it is important to ask participants about their actual behavior, rather than only measuring their intentions.

Finally, we took a closer look at the behavior of participants in the PMT+AP group, who were given the opportunity to make a plan for using Apple Pay. 96.3% of participants in the PMT+AP group wrote plans in Survey #2. Of those who wrote plans, 88.5% visited a location in their plan, 25.2% used Apple Pay at a location in their plan, and 87% used other payment methods at a location in their plan. Of those who wrote plans, 83.2% checked a box indicating "I strongly intend to try to use Apple Pay at these locations!" Of these participants, 89.9% visited a location in their plan, 30.3% used Apple Pay at a location in their plan, and 87.2% used other payment methods at a location in their plan.

In conclusion, although intention to set up and use Apple Pay was associated with actually doing so, many participants who expressed intentions did not follow through. This suggests nudges like action plans may help participants follow through on their intentions. This also demonstrates the importance of measuring actual behavior in addition to intention when evaluating the effectiveness of nudging techniques.

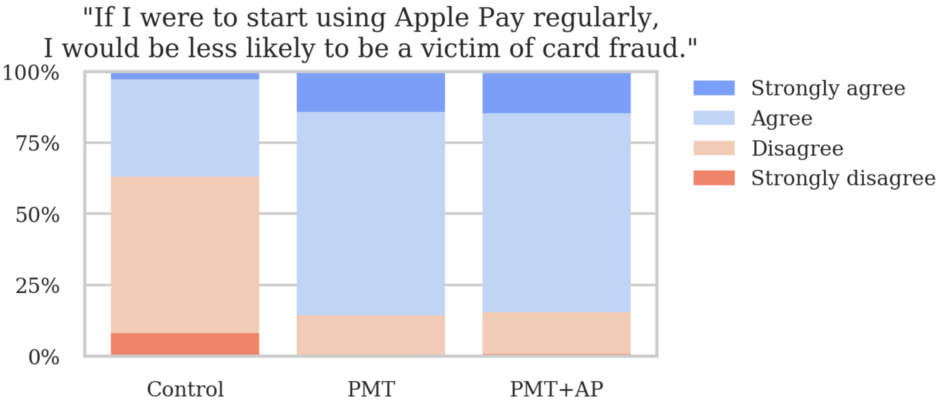

Figure 3.5: Participants in our treatment groups expressed greater agreement that Apple Pay would protect them from card fraud (i.e., response efficacy). Post-hoc tests: **Control vs PMT**, **p < 0.001**; **Control vs PMT+AP**, **p < 0.001**; PMT vs PMT+AP, $p = 0.880$.

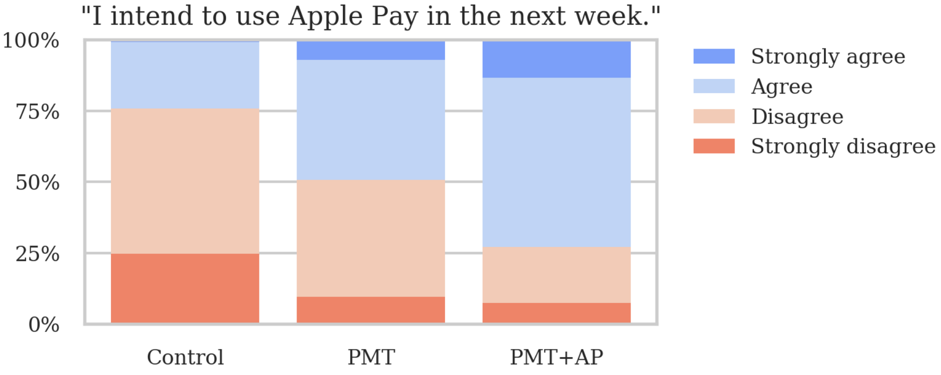

Figure 3.6: Participants in our treatment groups expressed stronger intentions to use Apple Pay. Further, participants who received the action plan treatment expressed even stronger intentions to use Apple Pay than did participants who only received the PMT treatment. Post-hoc tests: **Control vs PMT**, **p < 0.001**; **Control vs PMT+AP**, **p < 0.001**; **PMT vs PMT+AP**, **p $= 0.001$**.

When Did Participants Set Up Apple Pay?

As shown in Figure 3.10, 35% of participants had set up Apple Pay before Survey #2. In Survey #2, we encouraged the participants in our treatment groups to set up Apple Pay, but only 2.9% reported setting it up during Survey #2. However, an additional 10.5% reported setting up Apple Pay when we asked again in Survey #3. Overall, about half of participants had Apple Pay set up by the end of our study.

Note that most of the participants who set up Apple Pay during our study did so after completing Survey #2. The same pattern held in our qualitative interviews (Section 3.4.1). This suggests the importance of an experimental design like ours, in which information is given to participants, but participants are allowed time to think about that information and potentially conduct additional research before taking action.

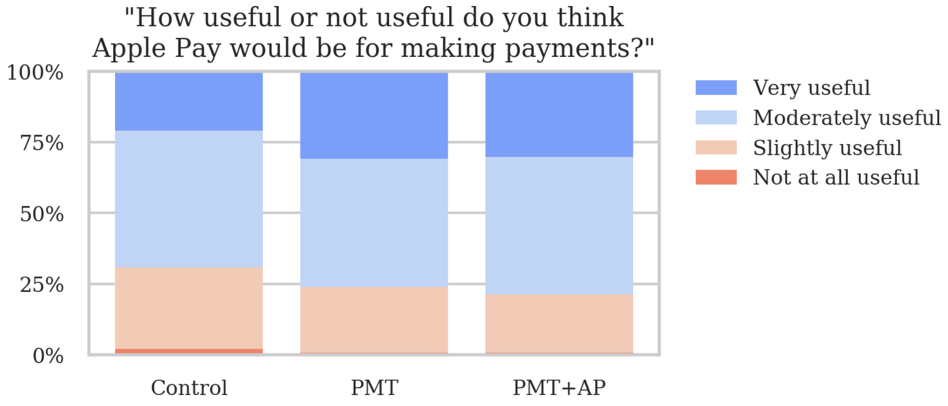

Figure 3.7: Our treatments had an effect on participants' belief that Apple Pay would be useful for making payments. Post-hoc tests: Control vs PMT, $p = 0.092$; Control vs PMT+AP, $p = 0.078$; PMT vs PMT+AP, $p = 0.856$.

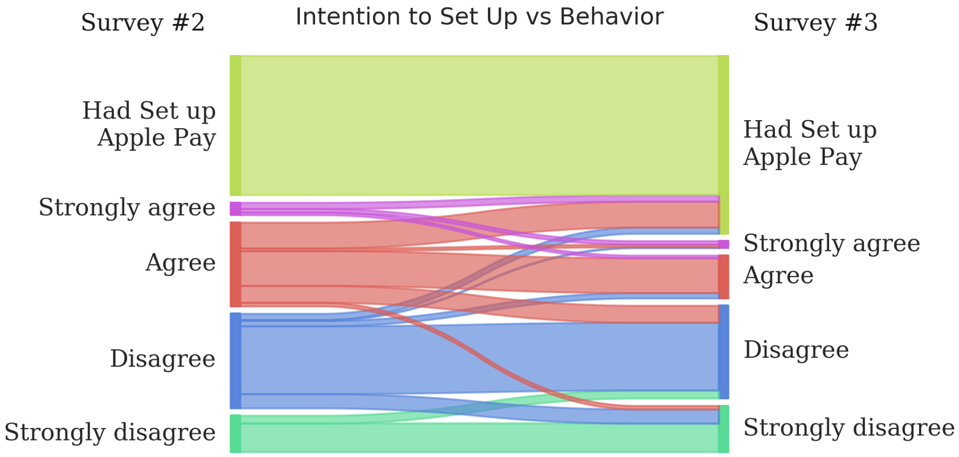

Figure 3.8: In both Survey #2 and Survey #3, we asked participants whether they had registered a card in Apple Pay. Those who had not were asked to rate their level of disagreement or agreement with the statement: "I intend to register a credit or debit card in Apple Pay in the next week."

Factors Associated with Use of Apple Pay

Having found that our treatments were associated with participants using Apple Pay, we trained three logistic regression models to identify additional factors associated with using Apple Pay.

First, we trained a model on all participants who completed all three of our surveys ($n = 411$). Our model contains the following 17 variables: treatment condition, security attitudes (SA-6) [46], age, Computer Science (CS) background, prior experience with card fraud, phone biometric (Face ID or Touch ID), gender, expressed intention to use Apple Pay, whether the participant knew anyone who used Apple Pay, whether the participant owned an Apple Watch, whether the participant had used Apple Pay before the study, and the participants' perceptions of response efficacy, self-efficacy, threat severity, threat susceptibility, self-consciousness, and Apple Pay's usefulness. Our model is shown in Table A.3 in the appendix. The model suggests that those with a computer science background and those who have experienced card fraud before are less

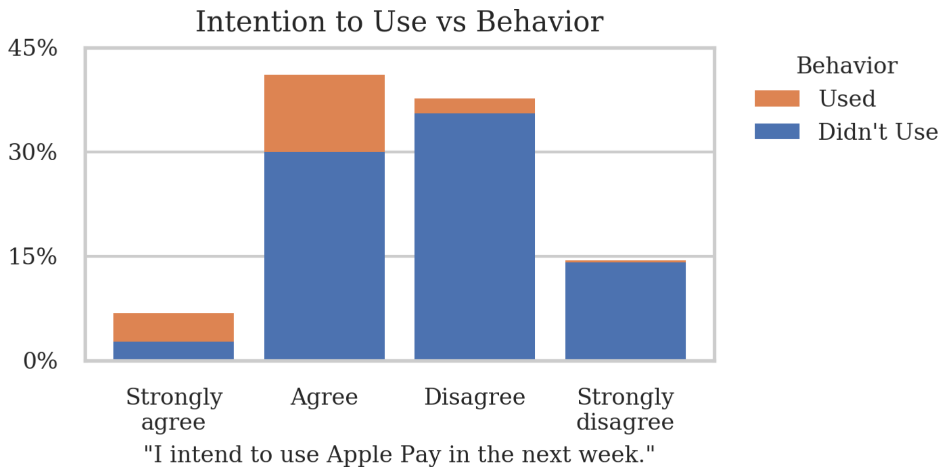

Figure 3.9: In Survey #2 we asked participants to rate their intention to use Apple Pay in the week ahead. We compared those responses to whether participants reported using Apple Pay in Survey #3.

likely to use Apple Pay (0.24x and 0.45x as likely, respectively). Perhaps those with a computer science background generally know more about Apple Pay, making those eligible for our experiment more likely to have consciously decided not to use it in advance of our interventions. This possibility is supported by Survey #1 from the qualitative interviews showing a positive association between having a CS background and having previously used Pay. The model also suggests that those whose phones are compatible with Face ID (2.1x), those who are non-female (2.4x), those who have used Apple Pay before (3.7x), and those who express an intention to use Apple Pay (6.1x) are more likely to use it.

Next, we trained a model on only the participants in the PMT+AP group ($n = 136$). Our model contains the same variables as our first model, with the removal of treatment and the addition of these variables: whether the participant checked the box indicating that they strongly intended to follow their plan, whether the participant requested they be sent information about Apple Pay, whether the participant requested they be sent their plan, and whether the participant visited at least one location in their plan. Our model is shown in Table A.4 in the appendix. Like our first model, this model suggests that those who experienced card fraud before are less likely to use Apple Pay (0.22x), and that those who used Apple Pay before are more likely to use it again (4.0x). Perhaps counterintuitively, the model also suggests that those who express self-consciousness about using Apple Pay in public are *more likely* to use it (5.1x). It is possible that participants' increased self-consciousness was due to their greater engagement with the plan, which could have made them more likely to use Apple Pay. There is also some evidence that whether the participant visited a location in their plan was associated with using Apple Pay (30x, $p = 0.058$).

Finally, we trained a model on the data we collected in Survey #1 to identify factors associated with people having used Apple Pay in the week before our study. We eliminated participants whose phones were incompatible with Apple Pay and who failed our attention check, leaving us with 590 participants. Due to the limited number of questions we asked participants in Survey #1, our model only contains age, phone compatibility with either Face ID or Touch ID, and

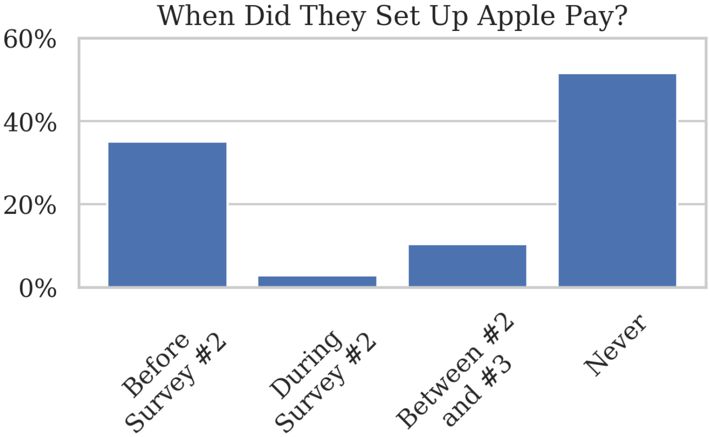

Figure 3.10: Participants could choose to set up Apple Pay at any point in our study, or not at all. More participants set up Apple Pay after Survey #2 than during Survey #2, suggesting the importance of giving participants time to think about the information we gave them.

Variable	β	e^{β}	**p-value**
age	-0.002	0.998	0.813
Face_ID	0.365	1.441	0.080
own_watch	1.035	2.815	<**0.001**
Intercept	-1.539	0.215	<**0.001**

Table 3.3: The variables in our regression model for predicting use of Apple Pay in the week prior to Survey #1 ($n = 590$). e^{β} indicates the change in odds of using Apple Pay for a one unit change in the variable (or when the variable is true). p-values significant at $\alpha = 0.05$ are bolded. Cox & Snell $R^2 = 0.051$.

Apple Watch ownership. The variables in our model are shown in Table 3.3. Overall, 23.7% of participants reported using Apple Pay in the past week. Our model shows a strong association between owning an Apple Watch and using Apple Pay, with Apple Watch owners being more than 2.8x more likely to use Apple Pay than non-owners. It is difficult to know the reason for this association, but one possible explanation might be that it's easier to use Apple Pay with an Apple Watch.

3.5 Discussion and Future Work

Our results have implications for both practitioners and researchers. First, banks, card issuers, and mobile payment operators could use our nudges to increase use of mobile payments instead of less secure, physical card payments. More widespread adoption of secure mobile payments has the potential to reduce card fraud, saving companies and customers both time and money. Second, our findings advance the field of nudging research. Our PMT and action planning implementation intention nudges corrected participants' misconceptions and increased intention to and actual use of mobile payments. In particular, we believe our PMT-inspired description of Apple Pay's security (Figure 3.2) was instrumental in correcting participants' misconception that mobile payments are less secure than physical card payments. Our action planning nudge was designed to help participants remember to use mobile payments when they visited certain locations. Although we did not find sufficient evidence to conclude that our action planning nudge increased *use of Apple Pay* compared to the PMT nudge (Table 3.1), we did find strong evidence that it increased *intention to use Apple Pay* (Figure 3.6). This discrepancy is due to the fact that many participants who expressed an intention to use Apple Pay did not actually use it (Figure 3.9). This shows the importance of an experimental design which measures both intention to use and actual behavior, as we did in our study. Our results also show the need for additional research into techniques that may help people translate their intentions to act in a secure manner into actual behavior.

Our study suggests several possible areas for future work. First, it would be useful to compare our PMT and action planning nudges in an experiment with a larger sample size. This would allow us to conclusively determine whether action planning yields improvements over PMT alone. Second, testing variations of PMT and action planning nudges could yield insight into what exactly makes these nudges effective. Knowing the most essential elements of these nudges could help translate them into a form suitable for large-scale messaging campaigns. Relatedly, people's receptivity to such messaging campaigns may depend on the entities conducting the campaigns, making a study of such messenger effects worthwhile. Third, PMT and action plans should be tested for their potential to increase adoption of other secure technologies and for encouraging adherence to security best practices. Finally, other forms of implementation intentions, such as coping plans, should also be tested.

3.6 Conclusions

Despite the security benefits they offer, adoption of mobile payments in the United States remains low, at least in part due to the belief that mobile payments are less secure than payments with physical cards [73, 128]. Our nudges addressed this misconception and increased adoption of mobile payments: participants in our PMT and PMT with action planning treatment groups were 2.3x and 3.9x more likely, respectively, to use Apple Pay than those in our control group. Our qualitative interviews suggested additional factors which may limit adoption of mobile payments, including lack of availability and usability challenges.

Our findings show that it is possible to increase real-world adoption of security-enhancing technologies using carefully crafted informational interventions. At the same time, many people who express an intention to adopt such technologies may fail to do so. This suggests the need for further research into interventions which can help people translate intention into action. Implementation intentions are designed to do this. In our study, we found only weak evidence of a small improvement (1.67x) from adding action planning to our PMT intervention. However, action plans might become more helpful as mobile payments become more available and other barriers to adoption are removed. Clearly, there is no single solution for increasing adoption of security-enhancing technologies, but PMT and action planning nudges are two tools that may help.

Chapter 4

Measuring Adoption of and Beliefs About Web Browsing Privacy Tools

4.1 Overview

Our Chapter 3 study suggests that nudges based on PMT and implementation intentions can increase real-world adoption of security-enhancing technologies like Apple Pay. Next, we sought to identify privacy-enhancing technologies which might benefit from nudging interventions. In this chapter, our goal was to measure people's awareness, adoption, and understanding of different web browsing-related privacy tools. Our study led us to identify Tor Browser as a tool which would benefit from nudging interventions, and informed the design of those nudges. We tested using nudging interventions to increase adoption of Tor Browser in Chapter 5.

This study was published at PoPETS 2021 [161].

4.2 Methodology

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 [132]. See Table B.1 in the appendix for our participants' demographics. Past studies have found that crowdworker participants are more internet-savvy than the general US population [138]. 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

Section 4.2.1 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 Section B.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 [133]. This was our only criteria for excluding participants' responses from our analyses. Since Prolific replaces rejected participants, our final sample contained 500 participants.

4.2.1 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 [75, 97, 139, 191] 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 [18, 136], 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 4.3 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*?"

[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.

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 [78], we suggest that the most realistic threats are from software downloaded and executed by users and from browser exploits [58, 60]. Software offers little protection against certain attacks (e.g., those using novel malware [53, 162] or legitimate software [28, 50]), but antivirus software and ad blockers can help in some cases [83]. For example, antivirus can block some malware from executing [162], and ad blockers can block fake download buttons [154] and potentially malvertising [34]. 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 [106]. Advertisers can connect different web requests to the same user through cookies and browser fingerprinting [31, 104]. We categorize Tor Browser as the only "Very effective" tool, because it is designed to comprehensively resist fingerprinting [86]. 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 [106]. 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 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 [106].

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 [181, 184]. 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 [72, 183]. 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 [65, 185]. 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 [43, 64, 74, 84, 180]. 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 [38]). 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 [29, 165], 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."

4.2.2 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 [126, 138]. 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. [138]. 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 [85]. 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 [114]. 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" [118]. Norton also offers Norton Privacy Manager, which among other features blocks ads and trackers, includes a VPN, and includes a privacy-friendly search engine [115, 117]. 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 (Section B.1), our data suggest that most incorrect answers were based on inappropriate mental models, rather than on an awareness of these bundled products.

4.3 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 analysis to analyze these free-text responses [25]. 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 B.2 and B.3 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 [105]. Because two expert annotators read each response and reached consensus on any differences, we have high confidence in the quality of our data.

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

4.4 Results

First, we supply descriptive statistics about participants' general privacy perceptions (Section 4.4.1), their adoption of browsing-related tools (Section 4.4.2), and their responses to our assessment scenario questions (Sections 4.4.3 and 4.4.4). Next, we convey the results of two exploratory statistical analyses: an analysis of factors associated with correctly identifying the protections offered by the tools (Section 4.4.5), and an analysis of demographic factors associated with tool use (Section 4.4.6) Finally, we describe the results of our thematic analysis of participants' misconceptions (Section 4.4.7).

4.4.1 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" [18].

4.4.2 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 4.1, 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.

We also asked tool users when they had last used each tool. As shown in Figure 4.2, 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.

4.4.3 Interest in Assessment Scenarios

We asked each participant questions about six randomly selected scenarios (from a total of twelve). Figure 4.3 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 scenarios is

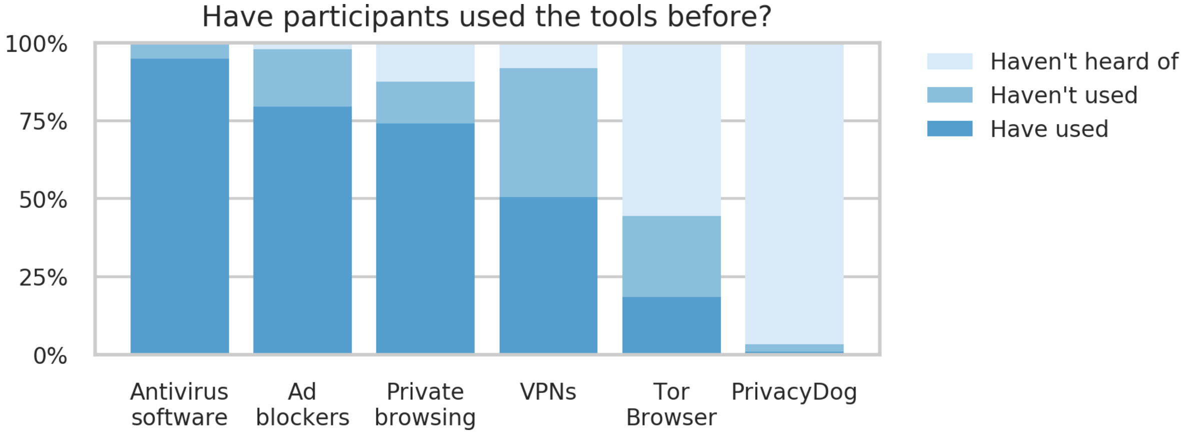

Figure 4.1: Percentage of participants who reported having used or heard of each tool. Our questions included "Yes," "No," and "Unsure" options. For this graph, we grouped "Unsure" and "No" answers together (e.g., if the participant indicated they were unsure whether they'd heard of a tool, we counted them as having not heard of it).

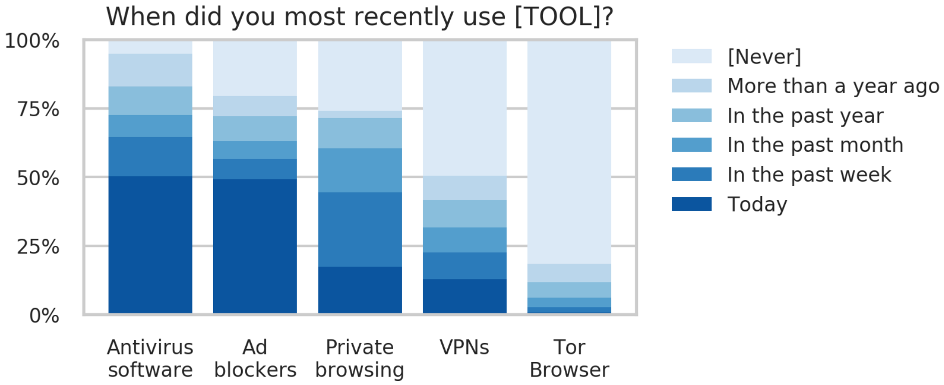

Figure 4.2: For each tool, we asked participants who said they had used it before when they had most recently used it. "[Never]" responses belong to participants who were not shown the question because they reported having never used the tool.

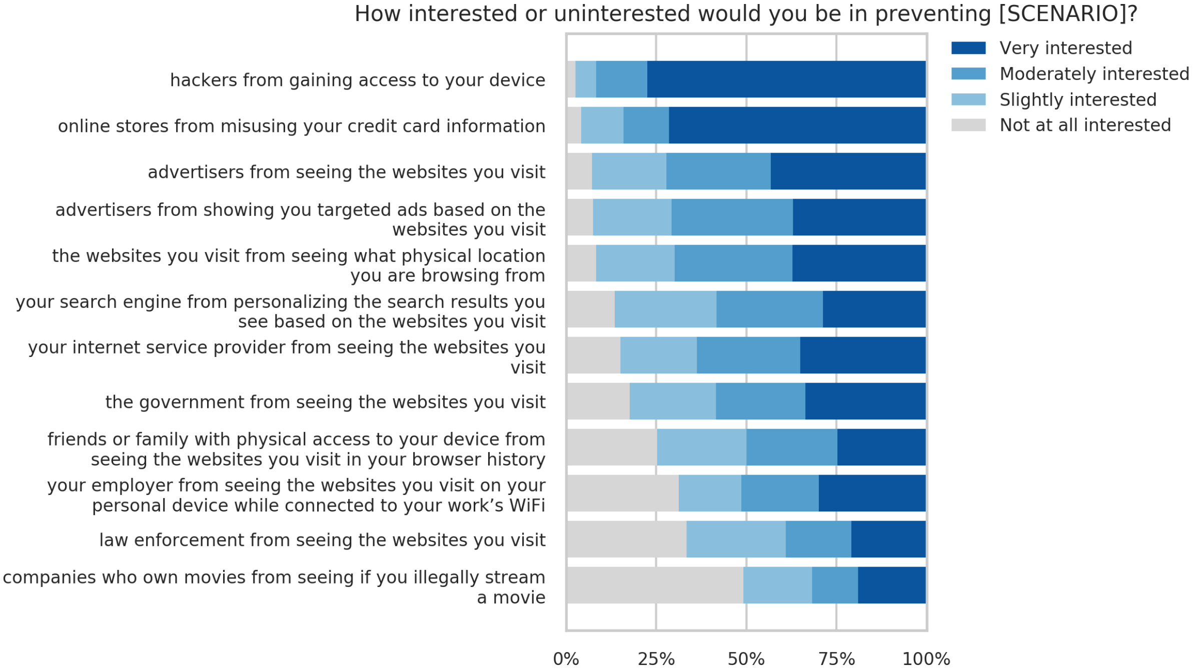

Figure 4.3: Participants' interest in preventing each scenario, sorted by the percent of "Not at all interested" responses. Note that each participant was shown six randomly selected scenarios, so percentages are calculated for the participants who did see a given scenario.

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.

4.4.4 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 Section 4.2.1. In Sections 4.4.4 and 4.4.4, we explain participants' responses for two scenarios in detail. For details about the remaining ten scenarios, see Figure B.1 in the appendix. In Section 4.4.4, we summarize participants' responses across scenarios and tools.

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 4.4, 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.

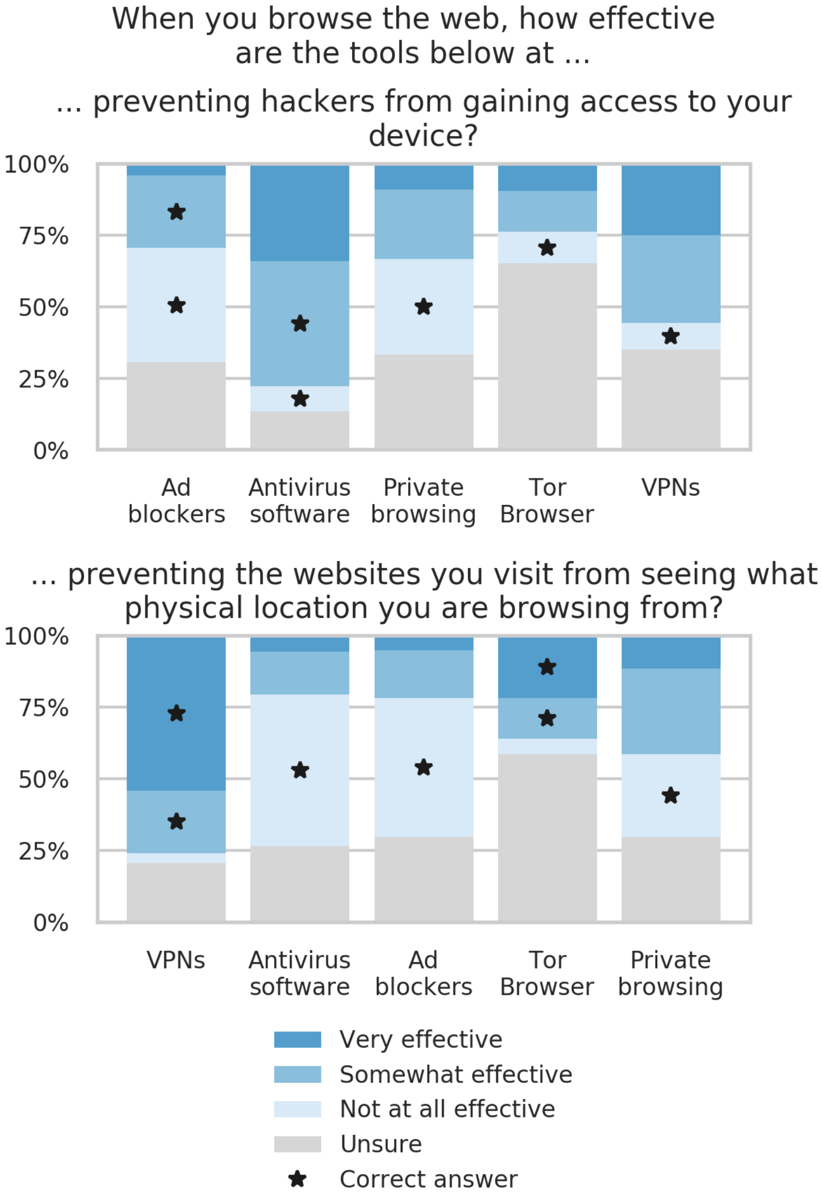

Figure 4.4: Responses consistent with our threat model are indicated with a star. Tools are sorted by the percent of correct responses.

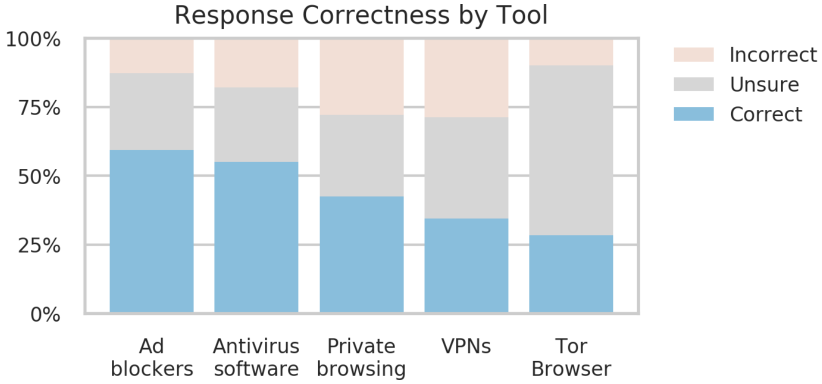

Figure 4.5: The correctness of participants' responses to the scenario-based assessment questions about each tool.

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 4.4, 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 (Section 4.4.5).

Summary of Response Correctness

Figures 4.5 and 4.6 show significant numbers of unsure and incorrect responses across tools and scenarios. Figure 4.5 shows that participants answered more questions correctly for tools that are more widely adopted. To explore this apparent relationship, in Section 4.4.5 we test for an association between participants' experience and their answers' correctness. Also, as shown in Figure 4.6, for all but one scenario participants answered fewer than half of the assessment questions correctly.

4.4.5 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 example, 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 4.1, we found statistically significant evidence of an association for all tools at $\alpha = 0.05$. Holm corrected Dunn post-

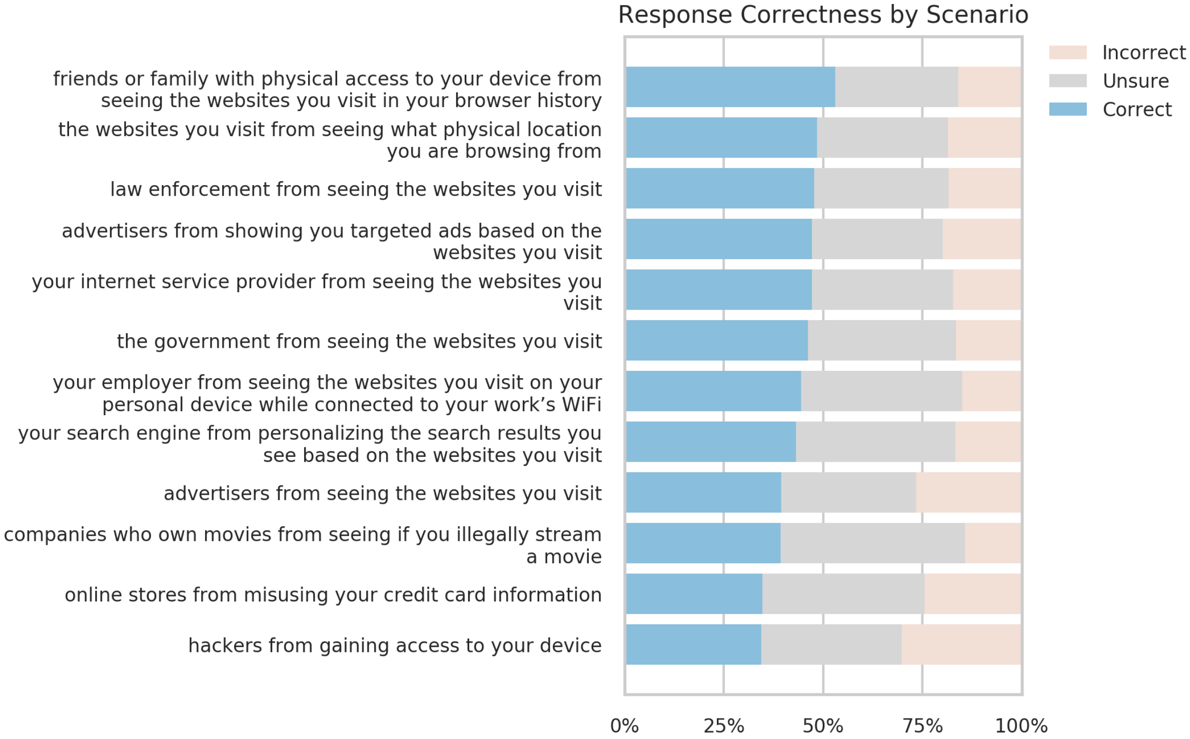

Figure 4.6: The correctness of participants' responses to assessment questions about each scenario. Note that each participant was shown six randomly selected scenarios, so percentages are calculated for the participants who did see a given scenario.

		Mean (Kruskal-Wallis p-value)			
Tool	**Experience**	**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 4.1: 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.

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.

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 not necessarily associated with a more accurate understanding of the tools' protections.

4.4.6 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 4.2 shows the significance of each model's coefficients, and Table 4.3 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 4.1). 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 [41, 66, 191], but we are unaware of prior work showing a positive association between living situation and use of privacy-enhancing technologies [150]. The existence of this association makes sense, since participants may want to hide their browsing activity from

	Model									
	Private browsing		**VPNs**		**Tor Browser**		**Ad blockers**		**Antivirus software**	
Variable	**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 4.2: 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 4.3: The R^2 values for each of the models in Table 4.2. R^2 represents the proportion of variance in tool use explained by each of our models.

their partner.

4.4.7 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 (Section 4.3). Some themes were associated with particular scenarios (Section 4.4.7) or tools (Section 4.4.7), but others were common across scenarios and tools (Section 4.4.7). 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 (Section 4.5.1) and privacy tools (Section 4.5.2).

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 [36, 88, 101, 187], 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 [23]. We worry that overconfidence in tools’ protections will likewise lead users to expose themselves to privacy harms.

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 [3, 52].

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 [20], 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 [167, 175]. 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.

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 [3].

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 [37, 71, 87].

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 [127].

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 [45, 79, 134]. However, experts caution that combining Tor with a VPN can either increase or decrease one's privacy protections, depending on one's threat model [167, 175]. 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 [106], 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.

4.5 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 (Section 4.4.4) shows that misconceptions are widespread. In addition, our qualitative analysis of participants' free-text responses characterizes the diverse ways in which misconceptions are expressed (Section 4.4.7). 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 (Section 4.4.1), 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.

4.5.1 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 4.3).

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 [10, 160]. 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 [36, 88, 101, 187]), 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 [42, 129]; 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 [19, 44, 164].

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**.

4.5.2 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. Pretesting 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 [113].

4.6 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 (Section 4.4.2). 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 (Section 4.4.4) and to understand nuances of these mistaken beliefs (Section 4.4.7). 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 (Section 4.4.5). We also identify demographics associated with use of the tools, which may help target nudging interventions to those who would most benefit (Section 4.4.6). Finally, we offer recommendations for designing both nudges and tools themselves (Section 5.6). 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.

Chapter 5

Nudges to Increase Adoption of Tor Browser

5.1 Overview

Our Chapter 4 study revealed that Tor Browser was not yet widely adopted, yet was effective at protecting against many of the privacy threats people were most concerned about. Furthermore, many people were unaware of Tor Browser or had misconceptions about it. In this study, our goal was to determine whether nudges based on action planning implementation intentions (AP), coping planning implementation intentions (CP), and protection motivation theory (PMT) can increase real-world adoption of privacy-enhancing technologies like Tor Browser. In particular, we wanted to compare the relative effects of these nudges to each other.

Action plans are relevant to this domain because users must remember to switch to Tor Browser for particular privacy-sensitive activities in order to protect their privacy. Our earlier study of implementation intentions for secure mobile payments only included action plans (Chapter 3), but in this study we also included coping plans. Coping plans are relevant because there are well-documented usability challenges associated with Tor Browser [51], and coping plans may help users overcome those challenges. PMT is relevant because it can help motivate participants' implementation intention plans [109, 146]. In addition, our PMT nudge may help correct the types of misconceptions we identified in our study of browsing privacy tools (Chapter 4).

This research also allowed us to study three additional questions. First, we measured the effect of these nudging interventions on participants' attitudes about Tor Browser. Second, we identified and quantified obstacles to widespread adoption of Tor Browser. Third, we compared the mechanism of this study's interventions to those from our earlier study of secure mobile payments (Chapter 3).

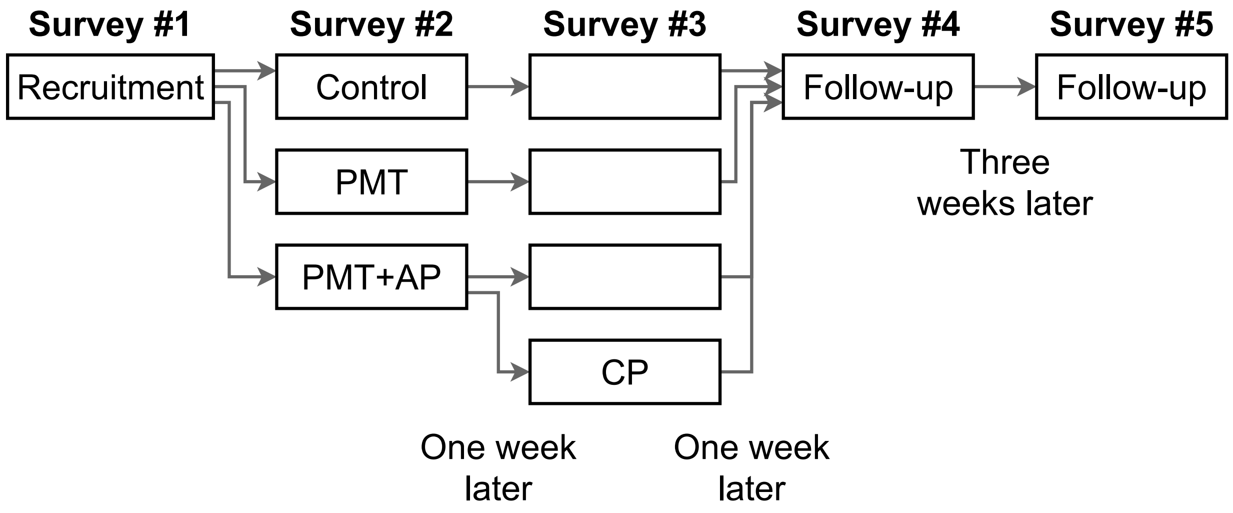

Figure 5.1: An overview of the surveys in our study.

5.2 Methodology

5.2.1 Overview

The goal of our study was to test whether nudges based on protection motivation theory (PMT), action planning (AP) implementation intentions, and coping planning (CP) implementation intentions could increase real-world adoption of Tor Browser. In total, we had four treatment conditions: Control, PMT, PMT+AP, and PMT+AP+CP. Comparing use of Tor Browser between the treatment conditions allowed us to see the effects of our interventions. The PMT nudge was designed to motivate participants to use Tor Browser, the action planning nudge to help participants identify opportunities to use Tor Browser, and the coping planning nudge to help participants overcome challenges associated with using Tor Browser. The literature suggests that implementation intention plans are most effective when people are strongly motivated [109, 146], so we tested our implementation intentions together with our PMT nudge. We administered our coping planning nudge one week after the initial interventions, to give participants time to encounter challenges using Tor Browser. Note that the PMT+AP and PMT+AP+CP conditions did not diverge until Survey #3, so for the purposes of describing them in our protocol and in our data analyses we refer to them as the same condition until that point.

Our study used a longitudinal design because we needed to give participants time to use Tor Browser in their everyday lives. After administering each treatment, we checked back with participants one week later to see whether they had used Tor Browser in the intervening week. A week gave participants time to perform activities they might only do on certain days (e.g., weekends). Each type of nudge was administered only once to each participant. Participants could request a link to their nudges.

Figure 5.1 shows an overview of the surveys in our study, and complete survey materials are included in Appendix C.1. Next, we describe the contents of each survey in detail.

5.2.2 Survey Design

Survey #1 We recruited participants from the Prolific crowdsourcing platform [126]. We sought to recruit participants who we thought would be motivated to adopt Tor Browser and who would have the ability to install it on their devices. To identify these participants, we employed a screening survey, Survey #1. To qualify for Survey #1, participants had to live in the United States, speak English, be at least 18 years old, and have a Windows, macOS, or Ubuntu operating-system[1] running on their computer. In Survey #1, we asked about people's use of privacy enhancing technologies, devices, and web browsers. We also asked whether they felt comfortable installing software on their devices and how interested they would be in preventing four threats to their online privacy. Participants had to meet multiple criteria to qualify for our experiment. First, in the past week they must have used either private browsing mode or a VPN, as long as the VPN usage wasn't primarily for work. Second, in the past week they must not have used Tor Browser. Third, on multiple days in the past week, they must have used a web browser on a laptop or desktop. Additionally, we asked which devices they had used at least once in the past week, and compared their responses to those about web browser usage; we required their responses to be consistent, and this served as our attention check. Fourth, participants must have indicated that they were generally comfortable installing software on their laptop or desktop. Finally, participants must have indicated they were "Very interested" in preventing at least one of the privacy threats we described. These criteria were designed to help us recruit participants who we thought would be motivated and able to install and use Tor Browser. Based on their responses to Survey #1, we invited all qualifying participants to our experiment. Our experiment began in Survey #2, and continued in Surveys #3 and #4, which we invited participants to one week after they completed the previous survey.

Survey #2 In Survey #2, we randomly assigned participants to our treatment conditions. Those in the control group only saw a short description of Tor Browser: "Tor Browser is an alternative web browser." Those in the PMT treatment were shown a description of privacy threats (Figure 5.2), the protection offered by Tor Browser (Figure 5.3), and instructions for installing and using Tor Browser (Figure 5.4). We also addressed common misconceptions (Figure 5.6) and usability issues (Figure 5.7), and offered technical details to those who were interested (Figure 5.5). Participants in the PMT+AP treatment were given the same information as the PMT treatment, but were also given a chance to form an action plan to help them remember to use Tor Browser for privacy-sensitive browsing activities (Figure 5.8). Note that the fourth treatment, the PMT+AP+CP treatment, did not diverge from the PMT+AP treatment until Survey #3. Finally, we asked demographic questions and questions related to perceptions of Tor Browser and privacy threats.

Survey #3 We invited participants to Survey #3 one week later. In this survey, we measured whether people set up and used Tor Browser following Survey #2, and whether they encountered any challenges when trying to use it. We also asked those in the PMT+AP and PMT+AP+CP

[1] Our goal was to measure use of Tor Browser, irrespective of device type. However, we wanted to ensure all participants had devices compatible with Tor Browser. We selected these three operating systems because they were available as prescreening criteria on the Prolific platform.

Many different organizations can gather information about your browsing activity. Here are just a few examples:

- **Advertisers** can see which websites you visit [186]. By tracking your browsing, advertisers can learn about your interests, and they may show you annoying or embarrassing ads
- **Every website you visit** receives information about you which can be used to infer the city or even neighborhood in which you live [120]
- **Your internet service provider** sees every website you visit, and there are few laws preventing them from selling that information [111]
- **The government** can request that companies give them information [63] about your online activities

And unfortunately, **most browsing tools offer only partial protection** against these privacy threats. For example:

- **Private browsing** only partially hides your browsing from advertisers, and does nothing to hide your location from websites or your browsing from your internet service provider or the government.
- Most **VPNs** do nothing to hide your browsing from advertisers, many VPNs keep logs which can be accessed by the government [69], and some VPNs even spy on their users [59]
- **Ad blockers** only partially hide your browsing from advertisers, and do nothing to protect against other privacy threats

Figure 5.2: As part of our PMT-based intervention we informed participants about threats to their browsing privacy. We primarily focused on threat susceptibility, although we also touched on threat severity [110]. In accordance with our recommendations from Chapter 4, we addressed well-defined threats and common misconceptions about other tools' protections.

Thankfully, there is a tool called **Tor Browser** which is effective at protecting against these kinds of privacy threats. Tor Browser is a web browser which makes web browsing anonymous. It does this by making each user's browsing indistinguishable from the browsing of thousands of other users around the world. If you use Tor Browser correctly, you can be confident your browsing is hidden from advertisers, your internet service provider, and even the government. Tor Browser also hides your location from the websites you visit. Tor Browser is available for free [169] and is simple to use.

Figure 5.3: As part of our PMT-based intervention we informed participants about the protections offered by Tor Browser. In this text, we addressed response efficacy and response cost [110].

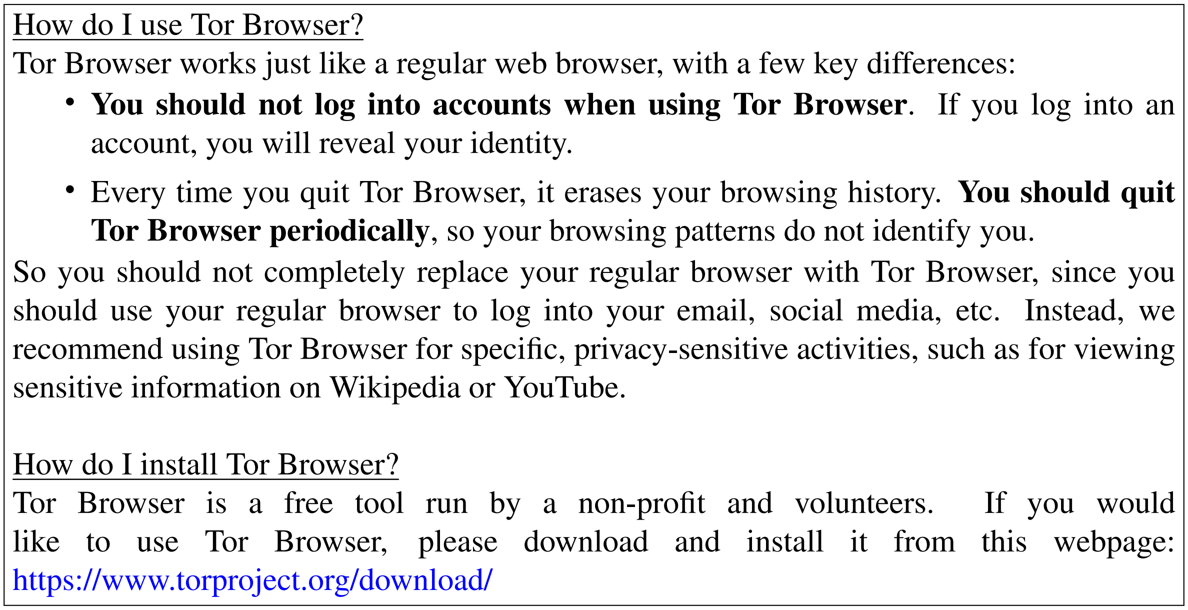

How do I use Tor Browser?

Tor Browser works just like a regular web browser, with a few key differences:

- **You should not log into accounts when using Tor Browser**. If you log into an account, you will reveal your identity.
- Every time you quit Tor Browser, it erases your browsing history. **You should quit Tor Browser periodically**, so your browsing patterns do not identify you.

So you should not completely replace your regular browser with Tor Browser, since you should use your regular browser to log into your email, social media, etc. Instead, we recommend using Tor Browser for specific, privacy-sensitive activities, such as for viewing sensitive information on Wikipedia or YouTube.

How do I install Tor Browser?

Tor Browser is a free tool run by a non-profit and volunteers. If you would like to use Tor Browser, please download and install it from this webpage: https://www.torproject.org/download/

Figure 5.4: As part of our PMT-based intervention we informed participants about how to use and install Tor Browser. We also reminded participants that Tor Browser is free. This text was designed to increase participants' perceptions of self-efficacy and to reduce perceptions of response cost [110].

conditions whether they had followed their action plans for using Tor Browser. Those in the PMT+AP+CP condition who reported encountering challenges using Tor Browser were given the opportunity to form a coping plan to overcome the challenges. We included two pre-defined plan templates corresponding to two challenges identified by Gallagher et al. [51]. Figure 5.9 shows our coping plan for participants who reported encountering extremely slow websites, which recommends that participants use the "New Circuit" button to fix this problem. Figure 5.9 also shows our coping plan for participants who reported encountering websites which didn't work in Tor Browser; in this case, we recommend that participants use alternative websites, and we suggest alternatives for YouTube and Reddit. Both of these challenges were encountered by participants in our pilot study. Finally, Figure 5.10 shows the open-ended template we showed participants who reported encountering other challenges.

Survey #4 One week later, we invited participants to Survey #4. Again, we asked whether participants had set up and used Tor Browser following Survey #3. We also remeasured perceptions of Tor Browser and privacy threats. In addition, we asked those in the PMT+AP and PMT+AP+CP conditions whether they followed their plans, and whether their plans were helpful to them. Finally, we asked whether participants were interested in an optional follow-up survey.

Survey #5 Three weeks after completing Survey #4, we invited participants who expressed interest to Survey #5. Survey #5 was similar to Survey #4, remeasuring usage of Tor Browser and Tor Browser-related perceptions.

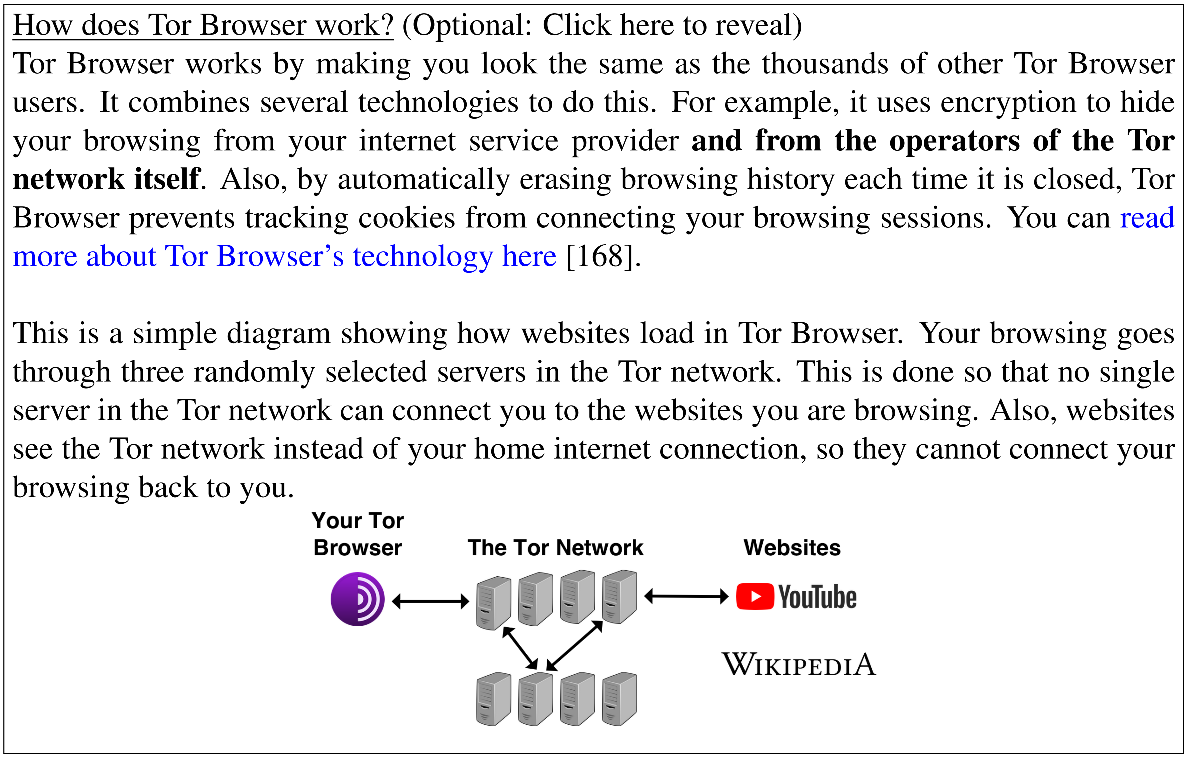

How does Tor Browser work? (Optional: Click here to reveal)

Tor Browser works by making you look the same as the thousands of other Tor Browser users. It combines several technologies to do this. For example, it uses encryption to hide your browsing from your internet service provider **and from the operators of the Tor network itself**. Also, by automatically erasing browsing history each time it is closed, Tor Browser prevents tracking cookies from connecting your browsing sessions. You can read more about Tor Browser's technology here [168].

This is a simple diagram showing how websites load in Tor Browser. Your browsing goes through three randomly selected servers in the Tor network. This is done so that no single server in the Tor network can connect you to the websites you are browsing. Also, websites see the Tor network instead of your home internet connection, so they cannot connect your browsing back to you.

Figure 5.5: For more technically inclined participants, we offered technical details about how Tor Browser works. To avoid overwhelming participants, this information was hidden until the heading was clicked.

Frequently Asked Questions

Who uses Tor Browser?

Citizens avoiding government censorship [48], journalists [144], and many other people [156] use Tor Browser.

Is it legal to use Tor Browser?

Yes: In the United States, free speech laws mean that it is completely legal to use Tor Browser. However, Tor Browser is blocked in countries which employ censorship, like China.

Is Tor Browser useful for torrenting files?

No: Tor Browser is intended for loading websites, and the similarity in name of Tor and BitTorrent is purely coincidental. Torrenting files over Tor is not recommended [16].

Does using Tor Browser ***protect me from malware or hackers****?*

No: Tor Browser provides no additional protections against malware or hackers.

Does using Tor Browser ***guarantee*** *that I will be* ***anonymous?***

No: Tor Browser initially provides anonymity, but if you log into internet accounts (e.g., your email account) or identify yourself through other ways (e.g., Googling your name) in Tor Browser, you will reveal your identity. But when used correctly, Tor Browser provides strong privacy protections: law enforcement has successfully caught some criminals who commit crimes using Tor Browser [129], but such investigations are time-consuming and expensive.

What if I ***accidentally log into an account*** *using Tor Browser?*

To become anonymous again, you should clear Tor Browser of all account-related data by either quitting Tor Browser or by clicking the "New Identity" button.

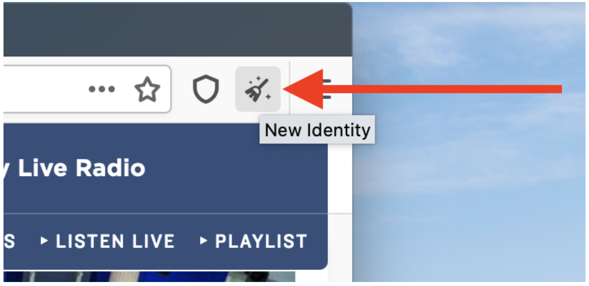

Figure 5.6: We use an FAQ to address the misconceptions about Tor Browser we discovered in our Chapter 4 study.

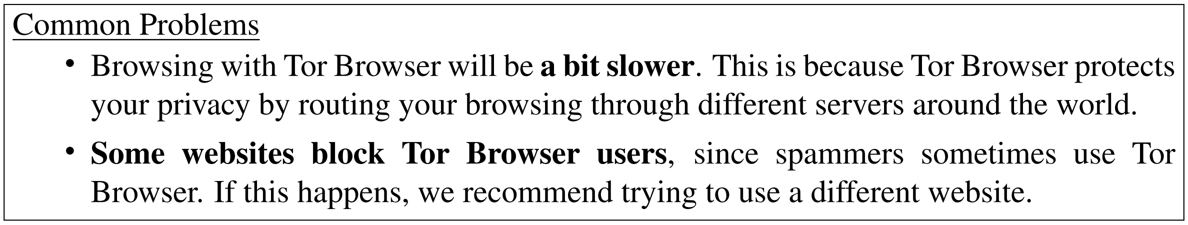

Common Problems

- Browsing with Tor Browser will be **a bit slower**. This is because Tor Browser protects your privacy by routing your browsing through different servers around the world.
- **Some websites block Tor Browser users**, since spammers sometimes use Tor Browser. If this happens, we recommend trying to use a different website.

Figure 5.7: We briefly addressed two common challenges to using Tor Browser [51, 191]. Norcie et al. and Gallagher et al. suggest that making users aware of such usability issues may make users more willing to tolerate them in exchange for greater privacy [51, 113]. Awareness of these issues may also help participants form accurate perceptions of response cost [110].

5.2.3 Compensation

We estimated survey durations based on the longest treatment (PMT+AP+CP). We estimated Surveys #1, #2, #3, #4, and #5 to take four, eight, six, three, and three minutes, respectively. The median times taken by our participants for each survey were 2.4, 6.8, 2.2, 3.0, and 2.9 minutes, respectively. We aimed to compensate participants at least $12/hour. Thus, we paid $0.80 for Survey #1, $3.50 for successful completion of the experiment, and $1.00 for Survey #5. The actual median rates of compensation were $20.00/hour for Survey #1, $16.48/hour for the experiment, and $20.87/hour for Survey #5. Since our survey questions were time sensitive, we required participants to answer Surveys #2, #3, and #4 within two days of being invited. We allowed up to one week for Survey #5.

5.2.4 Hypothesis Tests

We pre-planned four one-tailed tests of two independent proportions: First, comparing usage of Tor Browser reported in Survey #3, between the control and PMT groups. Second, comparing usage of Tor Browser reported in Survey #3, between the PMT and PMT+AP groups. Note that this uses data collected before the PMT+AP+CP group diverged from the PMT+AP group. Third, comparing usage of Tor Browser reported in Survey #4, between the PMT+AP and PMT+AP+CP groups. Finally, comparing usage of Tor Browser reported in Survey #4, between the PMT+AP and PMT+AP+CP groups, including only those who reported encountering challenges, since only they were presented with opportunities to form coping plans. Our hypothesis was that each treatment would progressively increase usage of Tor Browser (e.g., that PMT+AP would increase usage to a greater extent than PMT alone).

We conducted a small pilot study ($n = 116$ completed Survey #1) to test our surveys and to gather data for our power analysis. Based on effect sizes observed in our pilot and budgetary constraints, we determined effect sizes of interest, and used these to determine our sample size. We only describe power analysis for the final test listed above, since this showed the greatest number of required participants. Our pilot showed that of those in the PMT+AP+CP treatment who reported encountering challenges in Survey #3, 71.4% went on to use Tor Browser in the following week, as reported in Survey #4. Our effect size of interest was 30% (i.e., if 71.4% of those in the PMT+AP+CP treatment use Tor Browser, we want to detect if 41.4% or fewer

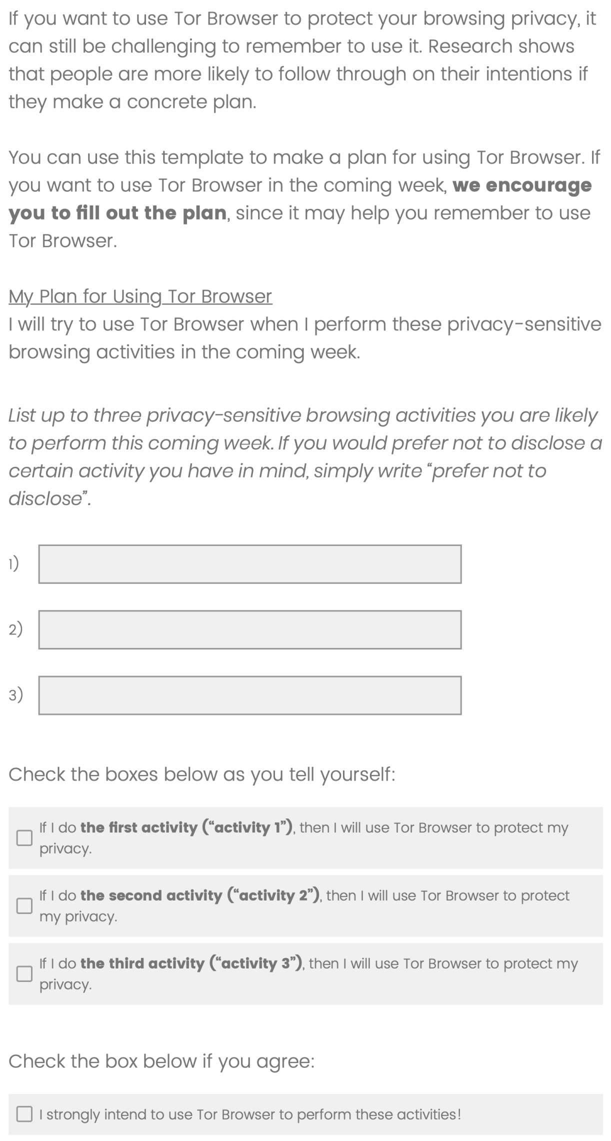

If you want to use Tor Browser to protect your browsing privacy, it can still be challenging to remember to use it. Research shows that people are more likely to follow through on their intentions if they make a concrete plan.

You can use this template to make a plan for using Tor Browser. If you want to use Tor Browser in the coming week, **we encourage you to fill out the plan**, since it may help you remember to use Tor Browser.

My Plan for Using Tor Browser
I will try to use Tor Browser when I perform these privacy-sensitive browsing activities in the coming week.

List up to three privacy-sensitive browsing activities you are likely to perform this coming week. If you would prefer not to disclose a certain activity you have in mind, simply write "prefer not to disclose".

1)

2)

3)

Check the boxes below as you tell yourself:

□ If I do **the first activity ("activity 1")**, then I will use Tor Browser to protect my privacy.

□ If I do **the second activity ("activity 2")**, then I will use Tor Browser to protect my privacy.

□ If I do **the third activity ("activity 3")**, then I will use Tor Browser to protect my privacy.

Check the box below if you agree:

□ I strongly intend to use Tor Browser to perform these activities!

Figure 5.8: We encouraged participants in our PMT+AP condition to form an action plan to help themselves use Tor Browser in the coming week. The template is designed to help participants formulate a plan and then mentally rehearse it in an "if-then" format [56, 119, 145]. The template also includes an opportunity for participants to form a strong commitment to their plan [56].

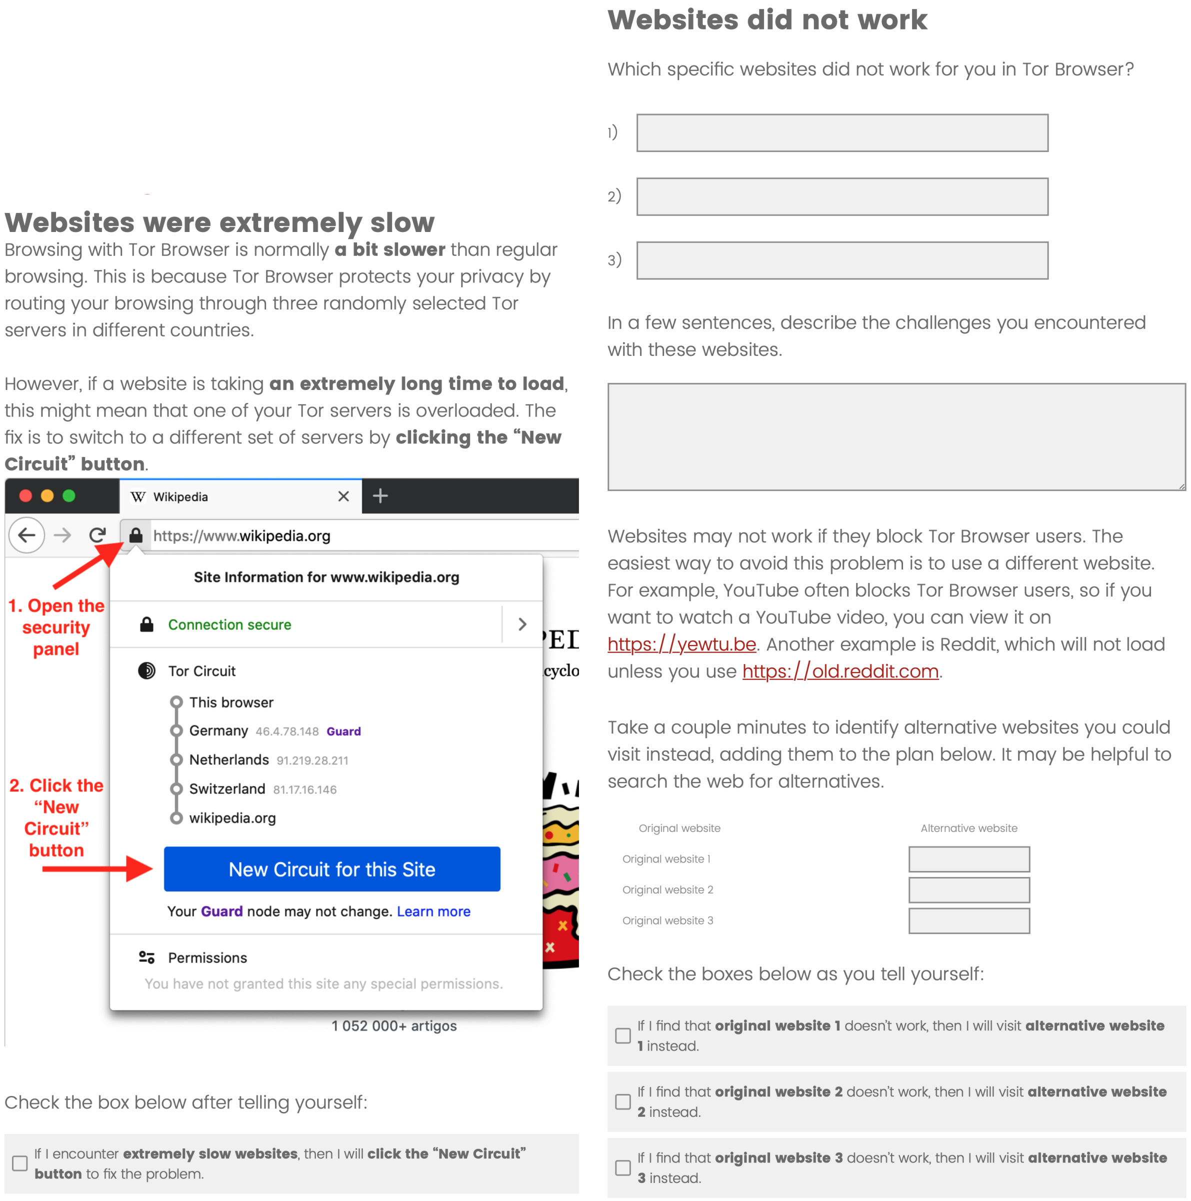

Websites were extremely slow

Browsing with Tor Browser is normally **a bit slower** than regular browsing. This is because Tor Browser protects your privacy by routing your browsing through three randomly selected Tor servers in different countries.

However, if a website is taking **an extremely long time to load**, this might mean that one of your Tor servers is overloaded. The fix is to switch to a different set of servers by **clicking the "New Circuit" button**.

Wikipedia
https://www.wikipedia.org
1. Open the security panel
Site Information for www.wikipedia.org
Connection secure
Tor Circuit
This browser
Germany 46.4.78.148 Guard
Netherlands 91.219.28.211
Switzerland 81.17.16.146
wikipedia.org
2. Click the "New Circuit" button
New Circuit for this Site
Your Guard node may not change. Learn more
Permissions
You have not granted this site any special permissions.
1 052 000+ artigos

Check the box below after telling yourself:

□ If I encounter **extremely slow websites**, then I will **click the "New Circuit" button** to fix the problem.

Websites did not work

Which specific websites did not work for you in Tor Browser?

1)

2)

3)

In a few sentences, describe the challenges you encountered with these websites.

Websites may not work if they block Tor Browser users. The easiest way to avoid this problem is to use a different website. For example, YouTube often blocks Tor Browser users, so if you want to watch a YouTube video, you can view it on https://yewtu.be. Another example is Reddit, which will not load unless you use https://old.reddit.com.

Take a couple minutes to identify alternative websites you could visit instead, adding them to the plan below. It may be helpful to search the web for alternatives.

Original website	Alternative website
Original website 1	
Original website 2	
Original website 3	

Check the boxes below as you tell yourself:

□ If I find that **original website 1** doesn't work, then I will visit **alternative website 1** instead.

□ If I find that **original website 2** doesn't work, then I will visit **alternative website 2** instead.

□ If I find that **original website 3** doesn't work, then I will visit **alternative website 3** instead.

Figure 5.9: We encouraged participants in our PMT+AP+CP condition who encountered challenges using Tor Browser to form coping plans to overcome the challenges [22, 152]. The plan on the left was shown to participants who reported encountering extremely slow websites. This plan suggests that in this case, people use the "New Circuit" button to fix the problem. The plan on the right was shown to participants who reported encountering websites that did not work. This plan suggests that participants identify alternative websites and visit those if they encounter problems again. Both plans give participants the opportunity to mentally rehearse the solution in an "if-then" format [56, 119, 145].

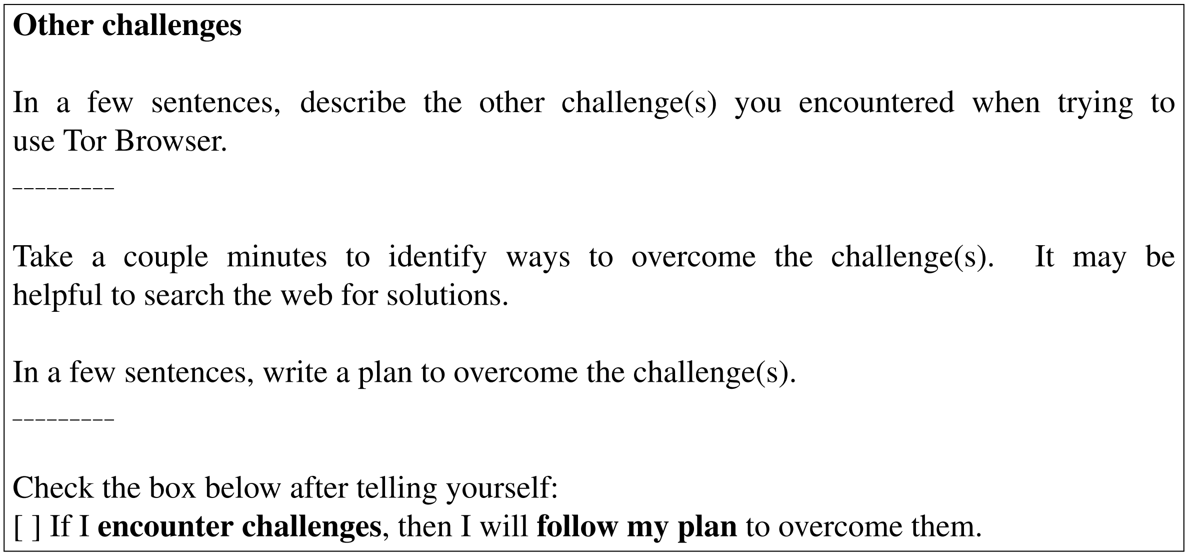

Other challenges

In a few sentences, describe the other challenge(s) you encountered when trying to use Tor Browser.

Take a couple minutes to identify ways to overcome the challenge(s). It may be helpful to search the web for solutions.

In a few sentences, write a plan to overcome the challenge(s).

Check the box below after telling yourself:
[ ] If I **encounter challenges**, then I will **follow my plan** to overcome them.

Figure 5.10: We encouraged participants in our PMT+AP+CP condition who encountered challenges using Tor Browser to form coping plans to overcome the challenges [22, 152]. This plan template was shown to participants who reported encountering a challenge other than those we listed. The template gives participants the opportunity to mentally rehearse their plan in an "if-then" format [56, 119, 145].

of those in the PMT+AP treatment use it). This corresponds to $h = 0.62$, a medium to large effect. For 80% power at $\alpha = 0.05$, G*Power showed we need 33 participants in each group. In our pilot, only 28% reported encountering challenges, which suggests that 118 participants are needed in each group, in order to have an estimated 33 participants in each group when running the tests.

We pre-registered our protocol on Open Science Framework prior to collecting the data used for our analysis [158].

5.2.5 Data Collection

We began collecting data on March 24th, 2021 and completed collecting data for our experiment on May 6th, 2021. We collected our final long-term follow-up response on May 30th, 2021. We spread recruitment across multiple days of the week, since participants' behavior might vary by day (e.g., weekday vs weekend). Our goal was for at least 118 participants to complete the experiment in each treatment group. Our pre-registration described weekly recruitment of the *minimum* number of participants needed to replace dropouts. We followed this procedure for two weeks, then used data about our dropout rates to estimate the size of a final batch of replacement participants, sized so that additional batches would not be needed.

Of the 1870 participants who responded to Survey #1, 689 qualified for our experiment. To ensure high quality data, we reviewed participants' free text responses. We rejected one participant who gave a low-effort response. In total, 537 participants completed our experiment. Of these participants, 148 were in the control group, 124 were in the PMT group, 125 were in

Comparison	Use of Tor Browser	Odds Ratio	p-value
Control vs PMT	S3: 14.9% vs 24.2%	1.83	**0.026**
PMT vs PMT+AP	S3: 24.2% vs 29.8%	1.33	0.125
PMT+AP vs PMT+AP+CP	S4: 34.4% vs 40.0%	1.27	0.173
Comparison, for those who encountered challenges			
PMT+AP vs PMT+AP+CP	S4: 42.3% vs 65.9%	2.64	**0.027**

Table 5.1: Our pre-planned tests for the effect of our treatments on participants' self-reported use of Tor Browser. "S3" and "S4" indicate that the Tor Browser usage data came from Surveys #3 and #4, respectively. For odds ratios, 1.5, 2, and 3 are the conventional thresholds for small, medium, and large effect sizes, respectively [163]. Results significant at $\alpha = 0.05$ are bolded.

the PMT+AP group, and 140 were in the PMT+AP+CP group.

5.2.6 Thematic Coding

We analyzed participants' free text responses as part of several exploratory analyses (Section 5.3.3, Section 5.3.4, and Section 5.3.5) [25]. For each analysis, the lead annotator began by developing a draft codebook. Next, the lead annotator and another annotator coded a batch of responses from a set of randomly selected participants. Then, the annotators reconciled their codes, and potentially refined the codebook. If they made any changes to the codebook, they reapplied the codes to any earlier batches. The annotators repeated this process until the coding task was complete. The numbers we report in our paper are based on dual-coding, so we have high confidence that we applied our codes consistently.

5.3 Results

5.3.1 Effect of Nudges on Use of Tor Browser

To determine the effect of our treatments on participants' use of Tor Browser, we conducted four one-tailed tests of two independent proportions. The results are shown in Table 5.1. Note that our PMT and action planning (AP) interventions were administered in Survey #2, and our coping planning (CP) intervention was administered in Survey #3. Treatments were layered, such that those in the PMT+AP+CP condition saw all three interventions. Also, in each survey we asked about use of Tor Browser since the previous survey.

The results show that our PMT-based informational treatment made participants 1.8x more likely to report using Tor Browser in the week between Surveys #2 and #3 than those in our control condition ($p = 0.026$). Our action planning intervention did not significantly increase use of Tor Browser relative to the PMT-only treatment ($p = 0.125$). For participants who reported encountering challenges using Tor Browser, our coping planning intervention made them 2.6x more likely to report using Tor Browser in the following week ($p = 0.027$). When all participants are analyzed, we do not see a significant effect from our coping planning intervention ($p =$

Survey	Variable	p-value	ϵ^2
2	Perception of threat susceptibility	**0.001**	0.027
2	Perception of threat severity	**0.008**	0.018
2	Perception of self-efficacy	<**0.001**	0.041
2	Perception of response efficacy	<**0.001**	0.030
2	Knowledge of how to use Tor Browser	<**0.001**	0.217
2	Expressed intention to use Tor Browser	<**0.001**	0.117
3	Expressed intention to use Tor Browser	**0.001**	0.032
4	Expressed intention to use Tor Browser	0.142	0.010
4	Perception of threat susceptibility	0.661	0.003
4	Perception of threat severity	**0.035**	0.016
4	Perception of self-efficacy	0.490	0.005
4	Perception of response efficacy	0.179	0.009
4	Knowledge of how to use Tor Browser	<**0.001**	0.033
4	Perception of privacy control	0.874	0.001

Table 5.2: The results of hypothesis tests measuring whether these variables differed between our treatment groups. The survey numbers in which the data were collected are shown in the leftmost column. p-values significant at $\alpha = 0.05$ are bolded, representing tests where the null hypothesis was rejected. Effect sizes are estimated as ϵ^2 values [102, 174].

0.173), which is unsurprising since only those who reported encountering challenges were given the opportunity to form coping plans. In summary, we have statistically significant evidence of a small effect from our PMT-based intervention and a medium effect from our coping planning intervention.

5.3.2 Perceptions of Tor Browser

We were also interested in the effect of our interventions on participants' perceptions of Tor Browser. We measured participants' perceptions using Likert scale questions. We analyzed these questions using Kruskal-Wallis tests, testing whether perceptions differed between our treatment groups. The results of our tests are shown in Table 5.2. For the significant results, we performed pairwise comparisons between the treatment groups using Dunn tests, employing the Holm-Bonferroni method to control Type I error.

The results suggest that our interventions affected participants' perceptions. In Survey #2, after administering the PMT and action plan nudges, these factors differed significantly between our treatment groups: threat susceptibility (Figure 5.11), threat severity (Figure 5.12), self-efficacy (Figure 5.13), response efficacy (Figure 5.14), self-reported knowledge of how to use Tor Browser (Figure 5.15), and intentions to use Tor Browser (Figure 5.16). In Survey #3, we administered the coping plan nudge and we only remeasured expressed intention to use, finding it still significant (Figure 5.17). In Survey #4, we remeasured all these variables at the end of our experiment. In Survey #4, only changes to threat severity (Figure 5.18) and knowledge of Tor Browser (Figure 5.19) remained significant. Graphs of non-significant results are shown in Figures C.1-C.5 in the appendix.

Survey, Variable	Control vs PMT	Control vs PMT+AP	Control vs PMT+AP+CP	PMT vs PMT+AP	PMT vs PMT+AP+CP	PMT+AP vs PMT+AP+CP
2, Threat susceptibility	**<0.001**	0.072		**0.023**		
2, Threat severity	**0.017**	**0.014**		0.788		
2, Self-efficacy	**<0.001**	**<0.001**		0.991		
2, Response efficacy	**0.002**	**0.001**		0.887		
2, Knowledge of Tor Browser	**<0.001**	**<0.001**		0.947		
2, Intention to use Tor Browser	**<0.001**	**<0.001**		0.096		
3, Intention to use Tor Browser	0.652	**0.039**	**0.001**	0.200	**0.021**	0.652
4, Threat severity	0.335	0.412	0.821	0.821	0.076	0.121
4, Knowledge of Tor Browser	**0.003**	**0.008**	**0.002**	1.000	1.000	1.000

Table 5.3: For results which were significant overall (Table 5.2), we conducted post-hoc tests to determine which treatment groups were significantly different from each other. We performed pairwise comparisons between the treatment groups using Dunn tests, employing the Holm-Bonferroni method to control Type I error. This table contains Holm-Bonferroni corrected p-values. Corrected p-values significant at $\alpha = 0.05$ are bolded.

As expected, the results show that our PMT intervention increased perceptions of self-efficacy, response efficacy, knowledge of how to use Tor Browser, and intention to use Tor Browser. Surprisingly, our action planning nudge appeared to negate the increase in perceptions of threat susceptibility from our PMT nudge (Figure 5.11). Perhaps our participants' plans to use Tor Browser made them feel more protected against online observation. We further discuss this in our limitations section (Section 5.4). Also, our PMT nudge reduced perceptions of threat severity (Figure 5.12). This might be because our descriptions of privacy threats did not emphasize the most severe possibilities (Figure 5.2), and perhaps participants' fears in the abstract are greater than those pertaining to the threats we described. This is not necessarily a problem, since our PMT-based nudge is designed to help participants form accurate perceptions of threats and protective responses, rather than to motivate participants to the greatest extent possible (e.g., by exaggerating threats). It is notable that by Survey #4, we no longer observe significant differences in intention to use Tor Browser or in perceptions of threat susceptibility, self-efficacy, or response efficacy. This suggests that some of our nudges' effects diminish over time, but as we discuss in Section 5.3.8, our Survey #5 data suggest that use of Tor Browser may persist long-term. We did not ask about perceptions of privacy control in Survey #2; since we do not see differences in Survey #4, it is unclear whether our nudges ever had an effect on these perceptions (further discussed in Section 5.4).

5.3.3 Why Do or Don't People Use Tor Browser?

At multiple points throughout our study, we asked participants about their reasons for either installing or using Tor Browser, or for not doing so. We collected multiple responses from all 537 participants who completed our experiment. We coded these responses to identify common themes, stopping after reaching code saturation. In total, we coded 558 free text responses from 150 randomly selected participants. Table C.1 in the appendix shows our codebook.

Participants most commonly explained that they used or installed Tor Browser because they wanted to test it out. For example, P33 wrote that they installed Tor Browser "To try it out, to

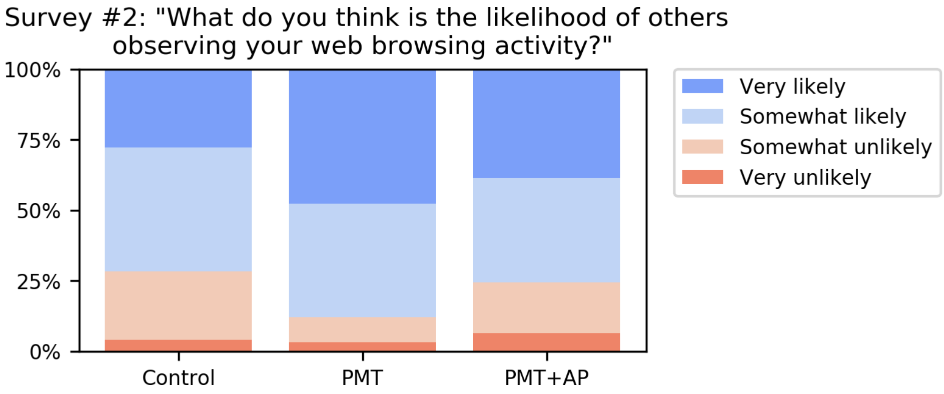

Figure 5.11: This question measured perceptions of threat susceptibility. Our PMT nudge increased perceptions of threat susceptibility, but our action planning nudge appears to negate this increase. Post-hoc tests: **Control vs PMT**, $\mathbf{p < 0.001}$; Control vs PMT+AP, $p = 0.072$; **PMT vs PMT+AP**, $\mathbf{p = 0.023}$.

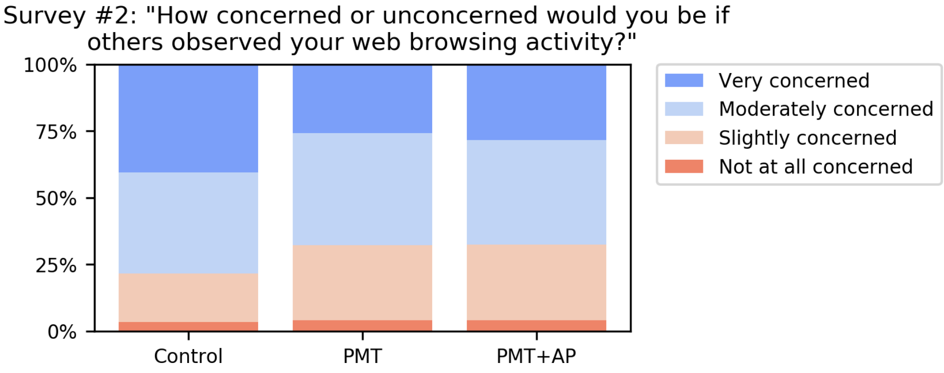

Figure 5.12: The question measured perceptions of threat severity. Our PMT nudge reduced perceptions of threat severity. Post-hoc tests: **Control vs PMT**, $\mathbf{p = 0.017}$; **Control vs PMT+AP**, $\mathbf{p = 0.014}$; PMT vs PMT+AP, $p = 0.788$.

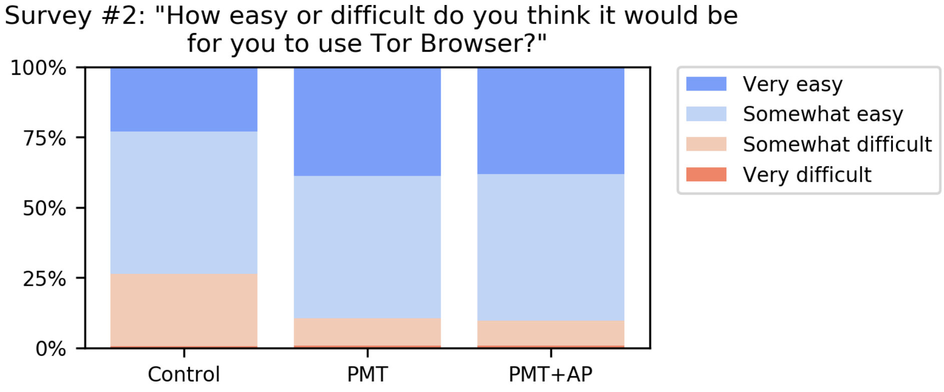

Figure 5.13: This question measured perceptions of self-efficacy. Our PMT nudge increased perceptions of self-efficacy. Post-hoc tests: **Control vs PMT**, $\mathbf{p < 0.001}$; **Control vs PMT+AP**, $\mathbf{p < 0.001}$; PMT vs PMT+AP, $p = 0.991$.

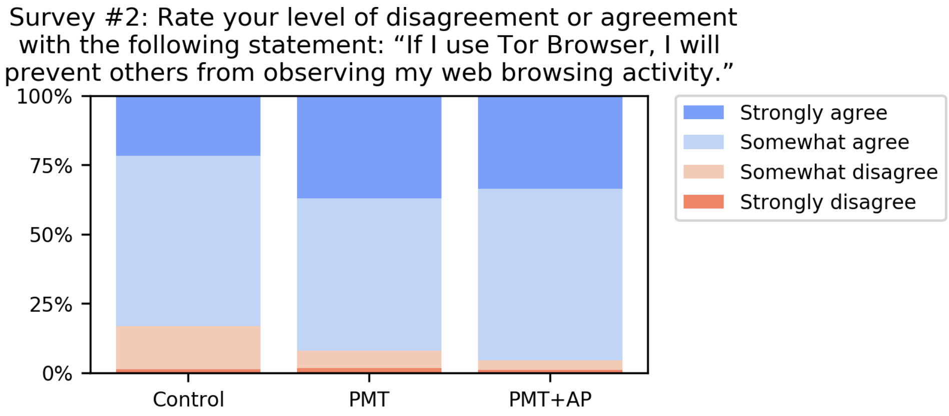

Figure 5.14: This question measured perceptions of response efficacy. Our PMT nudge increased perceptions of response efficacy. Post-hoc tests: **Control vs PMT**, **p $= 0.002$**; **Control vs PMT+AP**, **p $= 0.001$**; PMT vs PMT+AP, $p = 0.887$.

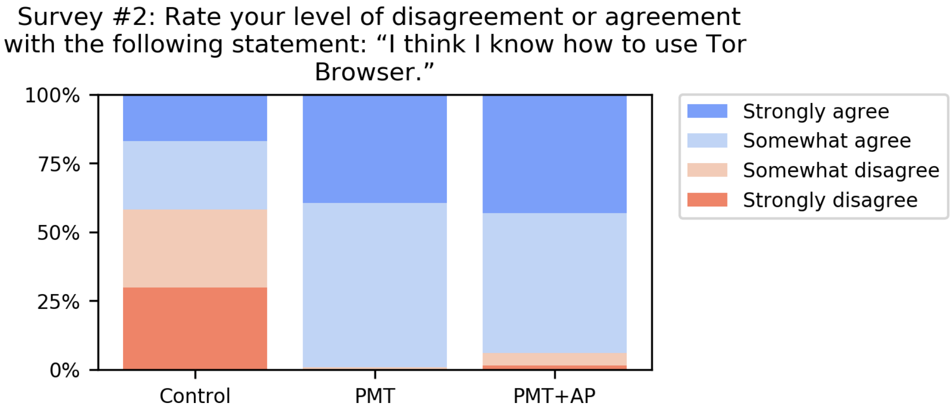

Figure 5.15: This question measured self-reported knowledge of how to use Tor Browser. Our PMT nudge increased knowledge of how to use Tor Browser. Post-hoc tests: **Control vs PMT**, **p < 0.001**; **Control vs PMT+AP**, **p < 0.001**; PMT vs PMT+AP, $p = 0.947$.

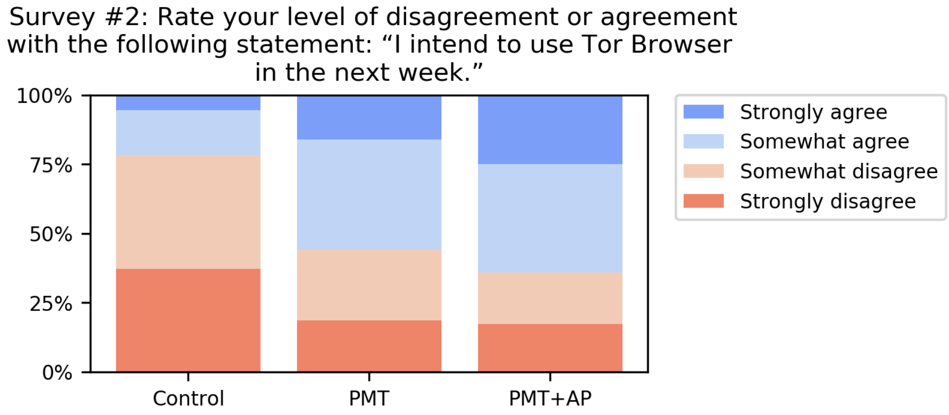

Figure 5.16: This question measured intention to use Tor Browser. Our PMT nudge increased intentions to use Tor Browser. Post-hoc tests: **Control vs PMT**, **p < 0.001**; **Control vs PMT+AP**, **p < 0.001**; PMT vs PMT+AP, $p = 0.096$.

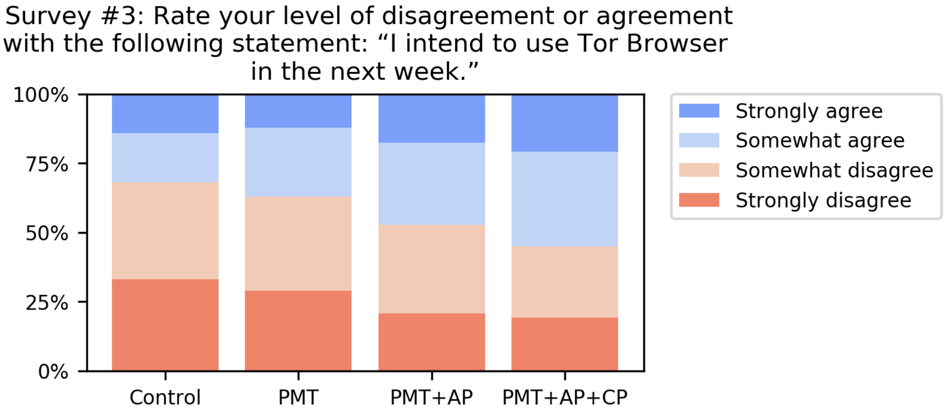

Figure 5.17: In Survey #3, we remeasured intention to use Tor Browser. Post-hoc tests: Control vs PMT, $p = 0.652$; **Control vs PMT+AP**, **p** $= \mathbf{0.039}$; **Control vs PMT+AP+CP**, **p** $= \mathbf{0.001}$; PMT vs PMT+AP, $p = 0.200$; **PMT vs PMT+AP+CP**, **p** $= \mathbf{0.021}$; PMT+AP vs PMT+AP+CP, $p = 0.652$.

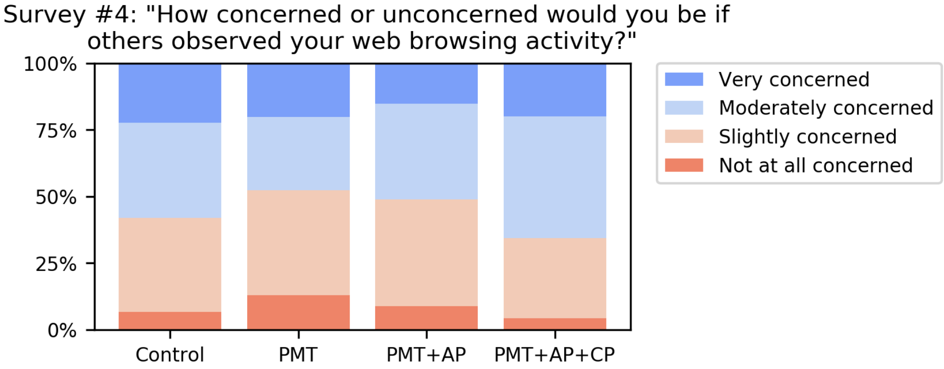

Figure 5.18: At the end of the experiment, we remeasured perceptions of threat severity. Although the Kruskal-Wallis test was significant, none of the post-hoc tests were significant at $\alpha = 0.05$.

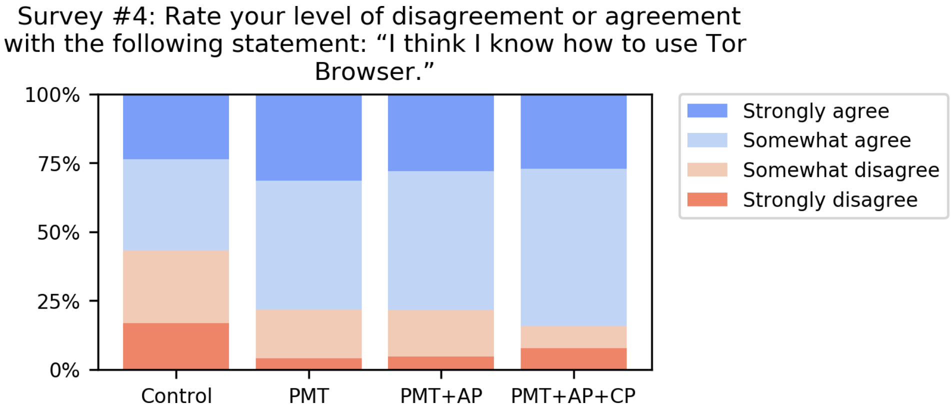

Figure 5.19: At the end of the experiment, we remeasured self-reported knowledge of how to use Tor Browser. Our PMT nudge increased self-reported knowledge of how to use Tor Browser, and this persisted to the end of our experiment. Post-hoc tests: **Control vs PMT, p = 0.003**; **Control vs PMT+AP, p = 0.008**; **Control vs PMT+AP+CP, p = 0.002**; PMT vs PMT+AP, $p = 1.000$; PMT vs PMT+AP+CP, $p = 1.000$; PMT+AP vs PMT+AP+CP, $p = 1.000$.

see if I would like using it." Participants also commonly mentioned the Tor Browser's privacy protections. For example, P90 used Tor Browser "because I did not want my browsing to affect my history or be available to my ISP." Similarly, P62 explained that they used Tor Browser "to keep my browsing of adult oriented websites private. If nothing I am doing is illegal [then] the government can keep their nose out of it." Participants also cited our study as an influence on their behavior. For example, P76 installed Tor Browser because "I was interested in trying it, particularly after reading the information provided across the first few surveys on this topic." Also, P42 explained that they wanted to test their coping plan, writing that they used Tor Browser "To check and see if your tip for faster loading speeds by clicking on lock and clicking on 'New Circuit for this Site' works. It was a definite help with faster loading speeds."

Participants gave different reasons for not installing or using Tor Browser as well. Most commonly, participants explained that they did not need Tor Browser. For example, participants gave answers such as "I don't need it" (P15), "I have no use for it" (P100), and "I did not need any extra internet privacy" (P149). This is notable, because we intentionally recruited participants who we thought would be highly motivated to use Tor Browser – recruitment required that participants were recent users of either private browsing or a VPN, and "very interested" in preventing at least one privacy threat we described. It was also common for participants to cite busyness or forgetfulness as reasons for not using Tor Browser. For example, P120 explained that "I forgot to be honest, it's been a busy week."

5.3.4 What Activities Do People Use Tor Browser For?

In Survey #2, we gave the 265 participants in our PMT+AP treatment group the opportunity to form action plans to use Tor Browser. In their action plans, we invited participants to list privacy-sensitive activities they might perform using Tor Browser. Of these 265 participants, 231 wrote at least one activity in the supplied plan template. In total, participants wrote 598 activities. We coded these activities to identify common themes. Our codebook is shown in Table C.2 in the appendix.

First, we note that in 192 cases, participants indicated that they preferred not to disclose details of activities. In these cases, there is no way to determine whether or not the participant actually had an activity in mind. If the participant did have an activity in mind, there is no way to determine the type of activity. This must be considered when interpreting the prevalence of the other themes we discovered, since participants may be less likely to disclose particularly sensitive browsing activities. The next most common theme was online shopping, which applied to 55 activities. For example, P497 wrote that they planned to use Tor Browser for "browsing Amazon to prevent following advertisements." This appears consistent with Mani et al.'s finding that traffic to www.amazon.com accounts for a large percent of Tor network traffic [103]. Other common activities included those related to finance, the news, Not Safe For Work content (e.g., pornography), and medical topics. Watching videos and accessing YouTube were also commonly described.

When we followed up over the course of the experiment, overall we found that participants reported performing 407 of the 598 activities they described (68.1%). Furthermore, of the 407 activities they performed, participants reported using Tor Browser for 180 of these activities (44.2%). Table C.2 includes a breakdown of which types of activities participants performed and which they used Tor Browser for. Tor Browser was used in at least some cases for nearly all the types of activities participants described.

5.3.5 What Challenges Are Encountered When Trying to Use Tor Browser?

In Survey #3, we asked the 244 participants who reported having ever used or tried to use Tor Browser whether they had encountered any challenges doing so. Participants could indicate that they had not encountered any challenges, select a predefined challenge (i.e., "Websites were extremely slow" or "Websites did not work"), or describe an "Other" challenge. The two predefined challenges were identified by Gallagher et al. [51], and we also observed them in our own pilot study. As shown in Figure 5.20, the majority of participants reported encountering some form of challenge. Additionally, 41 participants described an "Other" challenge. We coded these responses to identify common themes. Our codebook is shown in Table C.3 in the appendix. The challenges of websites being slow or not working were common in both participants' multiple choice selections (Figure 5.20) and in their free text responses (Table C.3). This suggests that our coping plan templates (Figure 5.9) did address participants' greatest challenges. Our findings are also consistent with prior work [51].

5.3.6 Did Participants Form and Follow Coping Plans?

Of the 138 participants who reported encountering challenges, 44 were in our PMT+AP+CP treatment, and so were offered the opportunity to form a coping plan to address their greatest challenge. 26 completed the "Websites were extremely slow" template, in which we explained how to use the "New Circuit" button. Of these participants, 13 reported clicking the "New Circuit" button when we asked the following week. Two participants completed the "Websites did not work" template. One participant planned to use "Old Reddit" to access Reddit, and

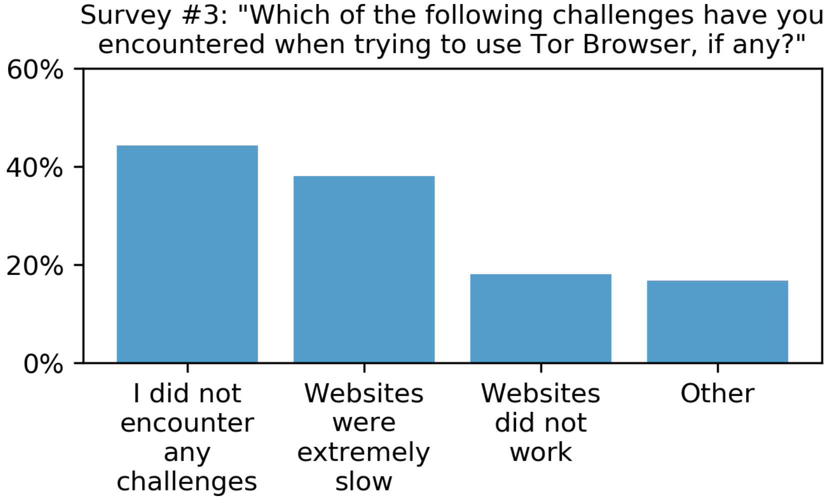

Figure 5.20: Challenges encountered by participants trying to use Tor Browser. Note that participants could select multiple challenges. We coded participants' explanations of their "Other" challenges

the other planned to use Facebook's onion service to access Facebook. When we followed up one week later, the first participant reported successfully using "Old Reddit" in Tor Browser, while the other participant reported not accessing Facebook in the previous week. Finally, 12 participants completed the "Other challenges" template. Participants supplied a very diverse set of responses, such as adjusting their mindset, conducting additional research, and employing external tools (e.g., a third-party password manager). We did not think it would be helpful to code such a small number of diverse responses, so we have simply included them all in Table C.4 in the appendix. Among these responses, only four participants reported following their plans to overcome "Other challenges."

Although we did not directly measure data on the efficacy of participants' coping plans for overcoming their challenges, we did collect free text responses about perceived helpfulness of the plans. Several participants confirmed that their coping plans helped them overcome the challenges they encountered. For example, P33 wrote: "Using the regular reddit web address didn't work but old reddit did." Also, P42 wrote: "Clicking on 'New Circuit for this Site' works. It was a definite help with faster loading speeds."

5.3.7 What Factors Are Associated with Using Tor Browser?

Our pre-registered hypothesis tests showed the effect of our treatments on adoption of Tor Browser (Section 5.3.1). However, we were interested in whether other factors might also influence adoption. Thus, we trained a logistic regression model containing our treatments, demographic factors, and perceptions of Tor Browser. Our model's outcome variable was usage of Tor Browser in either Survey #3 or Survey #4. Our model contains the 20 explanatory variables shown in Table 5.4. Note that for gender, "Female" is the baseline; for income, "Less than $10,000" is the baseline; for employment, "Working (paid employee)" is the baseline; for education, high school or less is the baseline; for living situation, living alone is the baseline; and for treatment, the control group is the baseline. Also, we included an interaction effect between the "encountered a challenge" factor and the PMT+AP+CP treatment, because our pre-registered hypothesis tests found evidence of this interaction. We encoded our Likert scale questions as binary variables (e.g., "Strongly disagree" and "Somewhat disagree" as 0, "Somewhat agree" and "Strongly agree" as 1). We excluded 18 participants who supplied incomplete demographics information, leaving us with 519 participants to train our model. The Hosmer and Lemeshow goodness of fit test did not find evidence of poor model fit ($p = 0.631$). We did not find evidence of multicollinearity, as all VIFs were less than 10. Our model explains approximately 37.8% of the variance in Tor Browser usage (Cox and Snell $R^2 = 0.378$).

Our model suggests that the most influential predictor of Tor Browser use is intention to use Tor Browser; participants who indicated intention to use Tor Browser in the coming week were 21x more likely to use Tor Browser than those who didn't. The model also suggests that when other factors are controlled for, our treatments make participants *less likely* to adopt Tor Browser than the control condition. However, our earlier tests show that our treatments increased intention to use and actual use of Tor Browser (Figure 5.16 and Section 5.3.1). The proper interpretation is that those who intended to use Tor Browser despite being in the control group were even more likely to use it than those we nudged who then expressed intentions to use Tor Browser.

We have several results which are not significant at $\alpha = 0.05$, but approach significance ($\alpha < 0.10$). First, the model suggests that those in the PMT+AP+CP treatment who encountered a challenge (i.e., who were given opportunities to form coping plans) were 3.3x more likely to use Tor Browser than those in the control group who did not encounter challenges. Next, the model suggests that those who are self-employed may be less likely to adopt Tor Browser. Finally, the model suggests that non-females and those who installed Tor Browser prior to the study may be more likely to use it.

5.3.8 Will Participants Become Long-term Users of Tor Browser?

We found that our PMT and coping planning nudges increased use of Tor Browser in Surveys #3 and #4, respectively (Section 5.3.1). However, our exploratory analyses suggest that some of our nudges' effects on participants' perceptions fade over time (Section 5.3.2). In Survey #5, we collected Tor Browser usage data three weeks after Survey #4, so this data may reveal whether the effects of our treatments persist over time. Table 5.5 summarizes use of Tor Browser across our study.

We reran our hypothesis tests on our Survey #5 data, and the results are shown in Table 5.6. Weeks after our interventions, the difference between the Control and PMT conditions remains statistically significant ($p = 0.011$), with an effect size similar to what we observed in Survey #3 (Table 5.1). This suggests that our PMT intervention contributes to long-term adoption of Tor Browser. However, we no longer find our coping planning intervention to significantly increase use of Tor Browser. Although our coping planning intervention temporarily increased adoption of Tor Browser in Survey #4 (Table 5.1), we do not have evidence that it increases long-term use of Tor Browser. One possibility is that between Surveys #3 and #4, these participants used Tor Browser to test their coping plans; after testing their coping plans, they may not have continued using Tor Browser at higher rates. It is possible their coping plans benefited them in ways that are not reflected in these numbers (e.g., using Tor Browser with the same frequency, but Tor

Variable	**p-value**	e^{β}
Age	0.593	0.991
Non-female	*0.068*	1.723
Income: \$10,000 - \$19,999	0.293	0.363
Income: \$20,000 - \$39,999	0.573	1.616
Income: \$40,000 - \$59,999	0.892	0.891
Income: \$60,000 - \$79,999	0.330	0.431
Income: \$80,000 - \$99,999	0.348	0.440
Income: \$100,000 or more	0.645	0.673
Employment: Self-employed	*0.067*	0.384
Employment: Student	0.688	1.220
Employment: Not employed	0.346	0.641
Employment: Retired	0.180	3.483
Education: College or associate degree	0.441	1.303
Education: Graduate degree	0.721	0.859
Computer-related background	0.350	1.283
Living with: Domestic partner	0.338	1.370
Living with: Children	0.449	1.272
Living with: Parents	0.237	0.617
Living with: Other family	0.105	2.043
Living with: Roommates	0.761	1.204
Previously heard of Tor Browser	0.731	1.118
Previously used Tor Browser	0.256	1.490
Knows other users of Tor Browser	0.543	1.220
Installed prior to the study	*0.094*	1.978
S1: Perception of privacy control	0.242	1.348
S2: Perception of threat severity	0.110	1.599
S2: Perception of threat susceptibility	0.322	1.384
S2: Perception of response efficacy	0.307	1.649
S2: Perception of self-efficacy	0.330	1.544
S2: Knowledge of how to use Tor Browser	0.239	1.725
S2: Intention to use Tor Browser	**<0.001**	20.666
S3: Encountered a challenge	0.516	1.249
Treatment: PMT	**0.015**	0.343
Treatment: PMT+AP	**0.013**	0.327
Treatment: PMT+AP+CP	**0.010**	0.284
Treatment: PMT+AP+CP x Challenge	*0.051*	3.298
Constant	**0.001**	0.015

Table 5.4: The variables in our logistic regression model for predicting use of Tor Browser in either Survey #3 or Survey #4 (i.e., the duration of our experiment). 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. p-values significant at $\alpha = 0.10$ are italicized.

	Use of Tor Browser			
Treatment	**In S3**	**In S4**	**In S5**	**Overall**
Control	14.7%	24.3%	15.4%	28.7%
PMT	26.4%	29.1%	27.3%	43.6%
PMT+AP	29.8%	33.0%	32.2%	43.5%
PMT+AP+CP		41.5%	29.2%	49.2%
	Use of Tor Browser, by those who encountered challenges			
Treatment	**In S3**	**In S4**	**In S5**	**Overall**
Control	24.2%	33.3%	15.2%	42.4%
PMT	50.0%	42.9%	35.7%	60.7%
PMT+AP	56.5%	42.9%	47.6%	71.4%
PMT+AP+CP		68.3%	41.5%	78.0%

Table 5.5: Use of Tor Browser across our study. Note that this table only includes the 491 participants who completed Survey #5, our long-term follow-up survey. Of these participants, 123 reported encountering challenges using Tor Browser. In Surveys #3 and #4, we asked participants whether they had used Tor Browser since the previous survey. Since Survey #5 was sent three weeks after Survey #4, in Survey #5 we instead asked whether participants had used Tor Browser in the past week.

Comparison	**Use of Tor Browser**	**Odds Ratio**	**p-value**
Control vs PMT	S5: 15.4% vs 27.3%	2.05	**0.011**
PMT vs PMT+AP	S5: 27.3% vs 32.2%	1.26	0.211
PMT+AP vs PMT+AP+CP	S5: 32.2% vs 29.2%	0.87	0.691
Comparison, for those who encountered challenges			
PMT+AP vs PMT+AP+CP	S5: 47.6% vs 41.5%	0.78	0.678

Table 5.6: One-tailed tests of two independent proportions, run on our Survey #5 data. Results significant at $\alpha = 0.05$ are bolded.

Browser being more pleasant to use).

5.4 Limitations

Our recruitment method and qualification criteria limit the generalizability of our findings (Section 5.2.2). For example, our results would likely differ if we recruited from countries where access to Tor Browser is restricted, or if we recruited less tech-savvy participants [138].

A limitation of our study design is that we rely on self-reported use of Tor Browser, making us reliant on participants' honesty and memory. We mitigated this limitation by reassuring participants that it was optional to use Tor Browser. Also, in most cases, we only required participants to recall their behavior in the past week. We considered an alternative design in which we would monitor participants' behavior using an instrumented Tor Browser. However, awareness of our observation might alter participants' behavior, and browser instrumentation might not capture use of Tor Browser across multiple devices.

Dropout in our study was higher than in other studies we have conducted, but we have no evidence to suggest that this negatively impacted our results. Of the 689 people we invited to participate in our experiment, 77.9% completed our entire experimental protocol (i.e., Survey #2, Survey #3, and Survey #4). We lost 6.4%, 9.6%, and 7.9% of participants between Surveys #1 and #2, Surveys #2 and #3, and Surveys #3 and #4, respectively. Of the 537 participants who completed our experiment, all but two requested an invitation to Survey #5, our optional long-term follow-up survey. Of the participants invited to Survey #5, 91.8% completed Survey #5. Our dropout rates may be partly due to our longitudinal study design, which employed multiple surveys over multiple weeks. It may also be partly due to bugs in the Prolific platform which we encountered while running our study [130, 131] which may have interfered with participation. A Pearson Chi-Square test did not find any evidence of dropout differing between our treatment groups ($p = 0.649$). A Pearson Chi-Square test did not find evidence of Survey #5 completion differing by use of Tor Browser during the experiment ($p = 0.372$).

Another limitation is that our instructions for using Tor Browser were based on a conservative threat model. For example, we recommended that participants not log into online accounts in Tor Browser to avoid deanonymizing themselves. However, it may not be necessary to take this precaution if one is only concerned about protecting one's privacy from one's ISP. We decided against more detailed instructions explaining these nuances, since our intuition was that it might either overwhelm participants or cause them to misunderstand the extent of Tor Browser's protections.

Finally, we identified two instances where improvements to our surveys might make our results clearer. First, we saw that our action planning nudge appeared to negate the increase in perceptions of threat susceptibility from our PMT nudge (Figure 5.11). Perhaps our participants' plans to use Tor Browser made them feel more protected against online observation, since they anticipated using it. But since we were interested in motivation to adopt Tor Browser, we wanted to measure participants' perceptions of threat susceptibility when they *were not* using Tor Browser. In the future, we would use an alternative phrasing to remove this ambiguity (e.g., "If you do not use Tor Browser, what do you think is the likelihood of others observing your web browsing activity?"). Second, we did not ask about perceptions of privacy control in Survey #2;

Comparison	**Use of Apple Pay**	**Odds Ratio**	**p-value**
Control vs PMT	T1: 8.7% vs 18.3%	2.34	**0.010**
PMT vs PMT+AP	T1: 18.3% vs 27.2%	1.67	**0.042**

Comparison	**Use of Tor Browser**	**Odds Ratio**	**p-value**
Control vs PMT	T1: 14.9% vs 24.2%	1.83	**0.026**
PMT vs PMT+AP	T1: 24.2% vs 29.8%	1.33	0.125
PMT+AP vs PMT+AP+CP	T2: 34.4% vs 40.0%	1.27	0.173
Comparison, for those who encountered challenges			
PMT+AP vs PMT+AP+CP	T2: 42.3% vs 65.9%	2.64	**0.027**

Table 5.7: The effect of our interventions on use of Apple Pay and on use of Tor Browser. "T1" and "T2" indicate that the tool usage data was collected one week and two weeks after the experiment began, respectively. For odds ratios, 1.5, 2, and 3 are the conventional thresholds for small, medium, and large effect sizes, respectively [163]. p-values significant at $\alpha = 0.05$ are bolded.

since we do not see differences in Survey #4, it is unclear whether our nudges ever had an effect on these perceptions. In the future, we would asked about privacy control in Survey #2 as well.

5.5 Comparison to Mobile Payments Study

Next, we will compare the results of our study of nudges for secure mobile payments (Chapter 3) to our study of nudges for Tor Browser (Chapter 5). Table 5.7 shows the effects of our nudges on use of Apple Pay and Tor Browser. Note that in our secure mobile payments study we used more conservative statistical analyses than we did in our study of secure mobile payments[2]. To make the p-values directly comparable, we reanalyzed our secure mobile payments study data using the same statistical techniques we used in our Tor Browser study (i.e., one-tailed tests of two independent proportions), and these results are shown in Table 5.7.

In both studies, PMT nudges increased use of the protective technologies in the week following our interventions, though the effect size was greater in our study of mobile payments. We saw evidence of a small effect from our action planning nudge in our study of mobile payments ($p = 0.042$), but only a non-significant effect was present in our study of Tor Browser ($p = 0.125$). However, we did see an increase in use of Tor Browser from our coping planning nudge. This underscores the importance of selecting a nudge which targets the underlying barriers to adoption: PMT can help people form accurate perceptions of tools, action planning can help people remember to use the tools, and coping planning can help people overcome challenges with using the tools. However, the efficacy of these nudges is an empirical question,

[2] In our mobile payments study we used chi-square tests of independence, which are equivalent to two-tailed tests of two independent proportions. We also used the Holm-Bonferroni method when interpreting p-values to control Type I error. These results are shown in Table 3.1

since it depends on the magnitude of the underlying barriers and on the potency of the nudges themselves.

We also measured participants' perceptions of the tools (Sections 3.4.2 and 5.3.2). In both studies, we found that our PMT-based nudges increased perceptions of the tools' response efficacy. We also found that our PMT-based nudges increased intentions to use each tool, and our data suggest that our action planning nudges may further increase intentions to use. There were also some differences between our studies of mobile payments and Tor Browser. Only in our study of Tor Browser did our PMT intervention increase perceptions of self-efficacy and threat susceptibility. However, our action planning nudge appeared to negate the increase in perceptions of threat susceptibility from our PMT nudge. Our participants' future plans to use Tor Browser may have made them report feeling more protected against online observation. Also, in our study of Tor Browser our PMT nudge *reduced* perceptions of threat severity. This might be because our descriptions of privacy threats did not emphasize the most severe possibilities (Figure 5.2), and perhaps participants' fears in the abstract were greater than those pertaining to the threats we described.

5.6 Discussion and Future Work

Our results suggest that there are opportunities to increase adoption of Tor Browser using nudging techniques, particularly those based on protection motivation theory (PMT). Certainly, not everyone is interested in using Tor Browser (Section 5.3.3 and Figure 5.16). However, our nudging techniques show that many people are willing to give it a try (Table 5.5), and that our PMT-based nudge can encourage a significant percent to continue using Tor Browser in the long term (Section 5.3.8). We also tested nudges based on action and coping planning implementation intentions. Although we did not find evidence of these plans further increasing long-term adoption of Tor Browser (Section 5.3.8), those who were given the opportunity to form coping plans were more likely to use Tor Browser in the subsequent week (Section 5.3.1).

Future work should investigate whether our nudges have effects beyond simply increasing tool usage. First, it is worth testing whether our PMT-based nudge also contributes to more effective use of Tor Browser. For example, our instructions reminded participants that Tor Browser's protections are reduced if one logs in to websites. Future work could confirm that our instructions help people use Tor Browser effectively. Second, although our action planning nudge was designed to help people identify opportunities to use Tor Browser, it did not significantly increase the number of participants who reported using Tor Browser in the previous week. An alternative outcome variable we could not measure was consistency of using Tor Browser: does someone always remember to use Tor Browser for a particular privacy-sensitive activity? A future study could measure whether action plans help in this respect. Also, it should be noted that we intentionally recruited participants who had prior experience with private browsing mode and VPNs (Section 5.2). Similar to Tor Browser, private browsing mode can be enabled for privacy-sensitive browsing; perhaps our action plan template was less helpful to those already familiar with private browsing, since they might be accustomed to the usage model encouraged by our action plan. Future work could test whether our action plans are more effective for a more general audience. Third, our study showed that participants frequently encountered challenges using Tor

Browser (Section 5.3.5). In particular, it was common for participants to report extreme slowness or websites not working in Tor Browser. To help participants mitigate these and other challenges, we tested several coping plans (Figures 5.9 and 5.10). Although we did not find evidence of our coping plans leading to long-term increases in use of Tor Browser, we did see evidence of an effect in the week after participants formed their coping plans, perhaps due to participants using Tor Browser to test their coping strategies (Section 5.3.3). Combined with participants' positive feedback about their coping plans, it seems worth testing whether coping plans like ours have effects beyond what we measured in our study. For example, coping plans may help people persevere in using Tor Browser when they encounter challenges, rather than simply switching to a different browser after encountering difficulty. Future work could study how people react to Tor Browser's usability challenges, and whether coping plans have an impact.

Several things should be considered when translating our results to a real-world deployment of nudges. First, our participants knew our nudges were part of a research study. However, it is widely recognized that how people respond to information depends on which entity delivers that information [40, 98]. Our nudges might be more or less effective, depending on how people perceive the entity administering the nudges. Second, we only recruited participants who we thought would be highly motivated to use Tor Browser (Section 5.2). Specifically, we recruited participants who had prior experience with other privacy tools, and who expressed a high level of interest in preventing at least one privacy threat Tor Browser can protect against. Our intuition was that it would be easier to detect the effects of our nudges among these participants in our experiment; perhaps similar targeting should be employed when deploying nudges in the wild. Nudging someone to adopt Tor Browser when it does not meet a need for them, or when they are not sufficiently motivated to overcome challenges associated with using it, may be problematic. Most notably, people have limited time to devote to privacy and security, so engaging with advice which is unlikely to be followed has a cost to the recipient [67]. Determining the best way to target nudges like ours remains a question for future work. As Tor Browser becomes more usable over time, it might make sense to more broadly deploy nudging to encourage adoption.

Multiple stakeholders have a role to play in increasing the usability of Tor Browser. In particular, website operators may benefit from better supporting users of Tor Browser. Our participants shared that they used Tor Browser for many innocuous activities, including shopping, reading the news, and researching medical issues (Section 5.3.4). Many websites are supported by advertising revenue, and although adoption of ad blockers is widespread [161], Tor Browser actively discourages the use of ad blockers [172]. Thus, websites may have a financial incentive to support Tor Browser users. To support users of Tor Browser, website operators should start by testing that their websites work properly in Tor Browser. If their website is hosted using Cloudflare, they can simply enable Cloudflare's Onion Routing [35] feature. Tor Browser's usability may also be improved through technical enhancements to the Tor Browser and the Tor network [170], and by volunteers contributing computing resources to increase the Tor network's capacity [171].

5.7 Conclusions

In the face of widespread privacy concerns, privacy enhancing technologies offer the possibility of returning control to users. Privacy tools like ad blockers are widely adopted, but other tools, like Tor Browser, are far less commonly used. Is this due to some inherent property of Tor Browser (e.g., is it too slow?), or is there certain information which might convince more people to adopt Tor Browser? To address this question, we tested whether three nudges could increase adoption of Tor Browser: a nudge based on protection motivation theory (PMT), an action planning implementation intention nudge, and a coping planning implementation nudge. Our longitudinal field experiment showed that our coping planning nudge increased short-term use of Tor Browser (Section 5.3.1), and our PMT-based nudge increased both short- and long-term use of Tor Browser (Section 5.3.8). Of course, in the future the usability of Tor Browser might be improved in various ways, but our results suggest that a significant percent of people are ready to start using Tor Browser today, and that nudges can help them do so. Simultaneously, it is important to realize that Tor Browser only addresses particular privacy needs. For example, it cannot prevent a social media company from sharing information about one's profile, or an email provider from analyzing one's emails. For these and other challenges, a combination of different privacy enhancing technologies and legal regulations may be appropriate. In cases where technologies can help, we hope our nudging research will prove helpful in increasing their adoption. In particular, PMT-based nudges may help increase awareness and correct misconceptions, coping planning nudges may help users overcome obstacles, and action planning nudges may help users take protective actions in particular contexts.

Chapter 6

Conclusions and Future Work

As part of this thesis, I conducted two experiments in which I used nudges to increase adoption of security and privacy enhancing technologies (Chapters 3 and 5, respectively). I also conducted a demographically-stratified survey of US residents' awareness, adoption, and beliefs about web browsing privacy tools (Chapter 4). My results show that people are interested in using technology to protect themselves from digital threats, and that nudges based on protection motivation theory (PMT), action planning, and coping planning implementation intentions can help them adopt such technologies. These nudges are most applicable when tools provide strong protections from serious threats, and when the tools are highly usable. Nudges based on PMT can motivate people to adopt the tools by helping correct perceptions of the tools and the threats they can protect against. Nudges based on implementation intentions can help people follow through on their intentions to adopt the tools, by helping people remember to use them and to overcome challenges associated with using them. Our findings suggest that, if deployed in the real world, nudges like these could help many people adopt effective protections against security and privacy threats. Increasing adoption of secure mobile payments by the percentages observed in our study could decrease monetary losses and frustration associated with card fraud for millions of people. Increasing adoption of effective web browsing privacy tools might lead to greater freedom of expression and less anxiety about online tracking. Furthermore, greater awareness of the limitations of VPNs might save people money they might otherwise spend on a commercial VPN that may not provide the protections they require.

Ethical Considerations

Based on my research experience, I believe several ethical guidelines should be followed when deploying nudges. First, PMT nudges should be used to help people form accurate perceptions of threats and protective technologies. Although it might be possible to motivate people to a greater extent by exaggerating threat severity, threat susceptibility, or response efficacy, it seems unethical and unnecessary to do so. For example, the UK's Advertising Standards Authority found an advertisement for NordVPN to be misleading because it exaggerated people's susceptibility to information theft when using public WiFi [7]. The goal of ethical nudging should not be to get everyone to act in the same way (e.g., to adopt a certain tool), but rather to help people make decisions that align with their stated preferences [6] (e.g., to adopt a tool if it solves a problem for them). Second, coping planning nudges are most appropriate when usability challenges cannot be easily resolved through technical improvements. Forming and following coping plans takes time, and would sum to a substantial amount of time if multiplied across millions of tool users. If a few developers can resolve a usability issue through technical means, this is clearly more time-efficient. For example, I contributed a fix for a usability issue with Tor Browser which made it difficult to return to a page of search results [157]. It might have been possible to help users with a coping plan, but this usability issue was straightforward to fix. In contrast, in our study we helped participants form coping plans to visit alternative websites if particular websites blocked users of Tor Browser. Fixing this problem is in the hands of site operators and content delivery networks [35], making it a better candidate for coping plans.

Future Work

My thesis demonstrated the potential of PMT and implementation intention nudges, but there remain unanswered questions about these nudges. Answers to these questions would aid large-scale deployment of these nudges.

Iterating on Nudge Design

In our experiments, our nudges included as many potentially relevant factors as possible. For example, our PMT-based nudges addressed threat severity, threat susceptibility, response efficacy, and self-efficacy. This meant that our nudges were quite lengthy. However, only certain perceptions were affected by our nudges, suggesting that not all parts of our nudges may be strictly necessary. Determining the essential elements of our nudges would make them more feasible to deploy in the real world. Future work could attempt to reduce the length of our PMT-based nudges, perhaps by focusing on the perceptions which these nudges most affected. Alternative implementation intention designs could also be tested. For example, instead of enumerating locations you might use Apple Pay at, you could plan to use Apple Pay if you see the wireless payment symbol at checkout.

We increased our experiments' external validity by asking about actual use of each tool, rather than mere intentions to use each tool. Our experiments showed that although intentions to use tools were associated with using them (Tables A.3 and 5.4), in many cases people expressed intentions but didn't take subsequent actions (Figure 3.9). This makes sense, since even someone who is highly motivated to use a tool might fail to use it if unforeseen circumstances arise. Nevertheless, in some cases intentions might be useful as a proxy for behavior. For example, iteratively testing alternative nudge designs using longitudinal experiments like ours would be time-consuming and costly; testing alternative designs using a single survey, including the nudges and measurement of intentions, would be more practical. After finalizing the design of a nudge, it could be tested using a longitudinal experiment, like we used in our studies. Future work should measure the relationship between intention and behavior, as it would be useful to know whether there are some cases in which it is sufficient to simply measure people's intentions.

Messenger Effects

In our experiments, it was clear to participants that the information they read came from researchers. However, in a large-scale deployment of nudges, the source of the nudges would likely be different. For example, a privacy advocacy group might host nudges on their blog. The literature suggests that how people respond to information depends on which entity delivers that information [40, 98]. For example, government agencies [33], celebrities [76], influencers [94], and local community leaders [70] have all encouraged people to get vaccinated for Covid-19, suggesting recognition that certain messengers may be able to reach different people.

A factor which arose in our research was people questioning whether we had their best interests in mind. In interviews for our study of secure mobile payments, several participants questioned whether we were being funded by Apple to promote Apple Pay. We corrected this misconception in our experiment, since a conflict of interest might reduce participants' belief in our claims about Apple Pay's security. This has implications for deployment of such nudges in the real world: nudges promoting Apple Pay might be more effective if they come from entities whose interests align with users. For example, messaging from banks and credit card issuers might be more credible than messaging from Apple, Samsung, or Google; banks and credit card issuers have a clear motivation to reduce card fraud, whereas these tech companies' motivations are less clear. Different entities may also be able to employ particular types of messaging. For example, users might be more convinced to use secure mobile payments if they received higher card rewards for purchases made using more secure payment methods. Future work should measure messenger effects for security and privacy nudges, since nudge efficacy is likely to differ based on the messenger.

Targeting Nudges

We recruited different types of participants for our two nudging experiments. For our security-focused mobile payments experiment we recruited people who had phones compatible with Apple Pay, who had recently made an in-person payment using a credit or debit card, and who had not recently made an in-person payment using Apple Pay. Similarly, for our privacy-focused Tor Browser experiment we recruited people who had devices compatible with Tor Browser, who had recently used a web browser, and who had not recently used Tor Browser; however, we additionally required participants to have recently used either private browsing mode or a VPN, and we required them to express a high level of interest in preventing at least one of the privacy threats we described to them. By requiring use of another privacy-protecting tool and interest in protection against threats Tor Browser can protect against, we sought to recruit participants who might be more likely to adopt Tor Browser, potentially increasing our ability to detect effects from our nudges. Future work could test the effect of our nudges for Tor Browser on a different segment of the population. Next, I will describe the intuitions which informed our use of targeted recruitment for our Tor Browser experiment, which are also worthy of future research.

First, we thought the average person's concern about privacy threats would be less than their concern about card fraud, so for our Tor Browser study we focused on recruiting participants who expressed concern about privacy. Our intuition was informed by our study of perceptions of web browsing privacy tools, which showed that people were more interested in preventing card fraud than in preventing any of the more privacy-focused threats we described (Figure 4.3). According to protection motivation theory, perceptions of threats are a component of motivation to adopt protective actions. Thus, we thought it would be more difficult to convince the average person to adopt Tor Browser than to adopt Apple Pay, so for our study of Tor Browser we recruited participants who we thought would be more motivated.

Second, we thought Tor Browser would meet a need for fewer users than Apple Pay would, so for our Tor Browser study we recruited users of other privacy tools. Apple Pay offers another payment option that works very similarly to paying with a physical card, and people who shop regularly are likely to have opportunities to use it [14]. It is less clear whether everyone who browses the web would benefit from using Tor Browser, since Tor Browser provides the greatest privacy protections for activities that do not require logging in to websites. For example, Tor Browser provides strong privacy protections if you research a medical condition on Wikipedia, but provides little protection if you log in to your email account and send an email about that medical condition. Thus, Tor Browser is most useful for people who regularly perform privacy-sensitive browsing activities that do not require logging in. This relates to perceptions of response efficacy, since people will be more motivated to adopt Tor Browser if it can protect them during their typical browsing activities. We thought that recent users of other privacy tools would have a greater need for Tor Browser, since their use of other tools suggests a need for browsing privacy protections. Although recruiting users of other privacy tools may have made some of our nudges' effects easier to detect, it may have made other effects more difficult to detect. For example, our action plans were designed to help participants remember to use Tor Browser, but we did not find a significant affect on use of Tor Browser. This might be because our plans were less useful to users of private browsing, since the usage modality of Tor Browser is very similar to that of private browsing. Thus, it is possible that people who do not use private browsing regularly might benefit more from our action plans.

Third, our intuition was that Apple Pay is easier to use than Tor Browser, which would influence perceptions of self-efficacy for each tool. Apple Pay is accepted at the majority of retail locations [14], and should be easy to use at locations that support it. In contrast, challenges like website slowness are a common occurence when using Tor Browser [51]. We thought that participants who were more concerned about privacy threats and who had previously adopted other privacy tools would be more motivated to use Tor Browser despite its usability challenges.

These last two intuitions, about the relative usefulness and ease of use of Apple Pay and Tor Browser, are particularly relevant with respect to deploying nudging interventions (cf. [82]). For which tools should we use nudging to encourage adoption? People have limited time to devote to privacy and security, so advice should be targeted to those who will benefit from it [67]. On the one hand, a tool which is useful for many people and is highly usable seems like a good candidate for widespread adoption. We think that Apple Pay is such a tool, so widely deploying nudges to encourage use of Apple Pay seems appropriate. On the other hand, a tool which is not useful for most people, or which has major usability issues, seems like a poor candidate for widespread adoption. We think that in its current state, Tor Browser falls somewhere between these extremes: it is useful for some people, and it has some usability issues. So it seems appropriate to target nudging interventions to those users who would benefit from using Tor Browser in its current state. This was our approach in our experiment, but determining the optimal candidates for nudging remains a question for future work. Our experiments revealed factors associated with adopting protective technologies in response to our nudges, including gender and prior experience with the tools (Tables A.3 and 5.4). Targeting nudges based on a demographic factor like gender raises questions of equity, but the problem is not as simple as relying on other factors, since even non-demographic factors (e.g., prior use of tools) can be associated with demographic factors (Table 4.2). Thus, the ethics of which users may benefit from nudging should also be considered. Furthermore, nudging a person to adopt a tool when they are unlikely to benefit from using it (e.g., due to insufficient usefulness or usability) could have adverse effects: the person might have a negative experience with the tool, making them less likely to try it in the future, even if it is improved in the meantime. Thus, developers of privacy tools should first seek to make their tools as useful and usable as possible before encouraging widespread adoption. As a tool like Tor Browser becomes more useful and usable over time, nudging to encourage adoption can be deployed more broadly.

Other Effects of Our Nudges

In our experiments, we asked participants whether they had used each protective technology in the past week – this served as our outcome variable. A limitation of our design was that we did not have visibility into *how* participants used each tool. Thus, it is possible that our nudges had effects beyond increasing tool usage. In particular, our nudges might help people use tools more correctly and consistently.

In both our experiments, our PMT-based nudge included instructions about how to use each tool. The instructions for Tor Browser were particularly important, since it is possible to use Tor Browser in a way that provides few privacy protections (e.g., if you log into an account). Our study design did not measure how people used Tor Browser, so we did not have visibility into misuse of Tor Browser. However, it seems likely that those who read our instructions would be more likely to use Tor Browser correctly than those who didn't. Future work could investigate the degree to which our instructions affect how people use Tor Browser.

We tested using action plans to help people remember to use secure mobile payments and Tor Browser. It would be interesting to measure the impact of action planning on people's consistency of tool usage. This is important, because a tool's protections can be undermined if the tool is not used consistently. For example, if you usually use Apple Pay at the grocery store but occasionally forget, you may still fall victim to card fraud if the store's point of sale systems have been compromised. Similarly, if you sometimes forget to use Tor Browser when researching a sensitive medical condition, those sensitive searches you conduct outside of Tor Browser will be visibile in your other browser's history. Action plans are intended to help people remember to take action, so they might help in this respect. Measuring consistency of tool use would require a different experimental design, like instrumenting users' devices to see whether they use the tools consistently.

In our study of Tor Browser, we tested using coping plans to help people overcome obstacles to using Tor Browser. Our experiment showed evidence that people tested their coping plans in the week after forming them, but we did not find evidence of coping plans contributing to a long-term increase in use of Tor Browser. However, note that we measured the effect of our interventions on use of Tor Browser in the past week, but we didn't have visibility into cases where a person might start using Tor Browser, but switch to their regular browser if they encountered difficulty (e.g., a website blocking Tor users). It is possible that those in the coping planning treatment used Tor Browser at similar rates to those in the PMT and PMT+AP groups, but that their coping plans helped them persist in using Tor Browser when they encountered difficulties. A different experimental design could directly measure how people react to the challenges associated with using Tor Browser, and the degree to which coping plans can help people overcome those challenges. If coping plans can help people overcome challenges to using Tor Browser, they might help people use Tor Browser more consistently.

Integrating Privacy and Security Tools

It is unclear how to best integrate privacy and security tools alongside other technologies. For example, the Brave web browser includes a "Private Window with Tor" mode [26], NordVPN includes optional ad blocking and antivirus functionality [114], and Norton offered a "Privacy Manager" which integrated ad and tracker blocking, a VPN, and a privacy-friendly search engine [115, 117]. Bundling such functionalities might bring the underlying technologies to more people, but it might also lead to confusion about the protections provided. For example, Brave's "Private Window with Tor" uses the Tor network, but does not include Tor Browser's fingerprinting protections. Although Brave warns users that "if your personal safety depends on remaining anonymous, use the Tor Browser instead," users might still assume that Brave's Tor mode provides more protections than it actually does. More research is needed into when and how to integrate protective technologies.

Final Thoughts

In closing, my research shows the promise of behavioral nudging to increase real-world adoption of security and privacy tools, particularly nudges based on protection motivation theory and implementation intentions. I showed the generalizability of my approach by testing nudges in two different contexts, for Apple Pay and Tor Browser. I also outlined best practices, which may help practitioners design nudges for other tools. Finally, I described future work that will contribute to even more effective nudges. I hope nudges will empower people to protect themselves from security and privacy threats. I will close by restating my thesis:

> *Nudging interventions can motivate people to adopt security and privacy tools, and can help people start using those tools in the real world. By quantifying and comparing the effect of nudges based on implementation intentions and protection motivation theory, we inform their use in the field of computer security and privacy.*

Appendix A

Nudges to Increase Adoption of Secure Mobile Payments

Table A.1: Final Codebook With Code Frequencies

Code	**Description**	**Interview #1 (n=20)**	**Interview #2 (n=10)**	**Survey #2 (n=20)**	**Survey #3 (n=10)**	**Overall (n=20)**
accidental_ activation	Accidentally activating Pay (e.g., by double-tapping the home button, proximity of NFC devices, etc.).	1	3	1	0	4
additional_ research	Performing additional research about Pay (e.g., asking others for their opinion about it, doing Google searches, etc.).	5	3	1	0	7

Continued on the next page

Code	**Description**	**Interview #1 (n=20)**	**Interview #2 (n=10)**	**Survey #2 (n=20)**	**Survey #3 (n=10)**	**Overall (n=20)**
curiosity_ availability	Wondering which places will accept Pay.	10	0	0	0	10
curiosity_ information_ theft	Wondering how their card information was stolen or could be stolen, how fraud occurred, why a data breach occurred, etc.	9	1	0	0	9
curiosity_ reviewing_ transactions	Wondering whether they will still be able to review their past transactions if they start using Pay.	1	0	0	0	1
curiosity_ technology	Wondering about specific technologies behind Pay (e.g., how NFC works, how the cryptography works, etc.), what cards can be added, how to activate it, how it works, its business model, etc.	16	4	3	1	16
experience_ card_fraud	People's own (or others') experiences with card fraud. Any fraudulent purchase made to a card is card fraud.	19	1	0	0	19
experience_ card_ information_ theft	People's own (or others') experiences with card info theft.	10	0	0	0	10

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
experience_card_theft	People's own (or others') experiences with their physical card being stolen.	5	0	0	0	5
experience_no_card_fraud	People having no experiences of their own (or others') to recount about card fraud.	1	0	0	0	1
experience_no_card_information_theft	People having no experiences of their own (or others') to recount about card info theft.	5	0	0	0	5
experience_other	Security-related experiences that don't fit into the other codes.	2	1	0	0	3
experience_unsure	People saying they are unsure whether their card information has been stolen or whether they have been the victim of fraud.	4	0	0	0	4
implementation_intention_clarified_understanding	Forming the implementation intention clarified the person's understanding of Pay (e.g., realizing it won't work at gas stations).	4	0	0	0	4

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
implementation_intention_forgotten	The participant not being able to remember their plan.	0	4	0	0	4
implementation_intention_helpful	Participants' reasons why the implementation intention would or did help them remember to set up or use Pay.	10	6	1	0	13
implementation_intention_remembered	The participant remembering their plan.	0	8	1	0	8
implementation_intention_unhelpful	Participants' reasons why the implementation intention would not help them remember to use Pay, why it is hard to form a plan, etc.	12	7	0	0	15
influenced_positive_self_report	The participant saying that the interview made them more likely to use or set up Pay.	10	1	0	0	10
misconception_affiliation	Thinking that we are working for or being funded by a company behind one of the technologies we're discussing (e.g., are you guys working for Google?).	4	0	0	0	4

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
misconception_always_resolved	Thinking that fraudulent purchases will always be resolved (e.g., they will always get their money back).	2	0	0	0	2
misconception_cost	Thinking or wondering if Pay costs something to use.	3	0	0	0	3
misconception_opening_app	Thinking that using Pay requires opening the Pay app by tapping on its icon.	1	0	0	0	1
misconception_other	Other misconceptions.	2	2	0	0	4
misconception_required	Thinking that using Pay or following the plan is a required part of the study.	1	0	0	0	1
misconception_rewards	Thinking that they won't get rewards, points, or cash back if they use Pay.	4	0	0	0	4
misconception_screen_scan	Thinking that Pay works by scanning the user's phone or watch screen, rather than by using NFC.	4	0	0	0	4
misconception_square_pos	Thinking that Pay only works at Square POSs, or that Pay is the software running on those Square POS iPads. It is not a misconception that Pay works at most Square POSs.	2	0	0	0	2

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
mitigation_description_helpful	Participants' reasons why the description of Pay or the instructions for how to set up and use Pay are helpful to them.	6	2	1	0	6
protection_action_RFID_wallet	Using an RFID-blocking wallet to protect your card information.	1	0	0	0	1
protection_action_account_access	Protect access to your account (e.g., password, 2FA).	6	2	1	0	8
protection_action_avoid_disclosure	Avoid giving information to others, whether prompted or not; avoiding falling for phishing, etc.	5	0	0	0	5
protection_action_avoid_merchant	Avoid transactions at untrusted merchants, only use trusted merchants, etc.	4	0	0	0	4
protection_action_avoid_online	Avoid making purchases online, avoid putting card information online, etc.	3	0	0	0	3
protection_action_certification_logo	Looking for certification logos (e.g., Trustee, Verisign, McAfee), browser plugin indicators (e.g., Web of Trust), TLS certificates, or any other symbols that attest to security in some way.	4	0	0	0	4

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
protection_action_corporate_resolution	Reporting fraudulent purchases to the card issuer, getting a new card, etc.	19	0	0	0	19
protection_action_data_retention	Preventing a card from being saved on a website either in whole or in part (e.g., not allowing the CVC to be saved).	2	0	0	0	2
protection_action_law_enforcement	Reporting card fraud or theft to law enforcement.	1	0	0	0	1
protection_action_monitor_statements	Looking for unauthorized transactions on card statements.	8	0	0	0	8
protection_action_monitoring_service	Lifelock, credit monitoring, etc.	4	2	2	1	4
protection_action_network	Using a secure network connection (e.g., home Wi-Fi, a VPN when on public Wi-Fi, etc.), avoiding insecure networks, avoiding public computers, etc.	2	0	0	0	2

Continued on the next page

Code	**Description**	**Interview #1 (n=20)**	**Interview #2 (n=10)**	**Survey #2 (n=20)**	**Survey #3 (n=10)**	**Overall (n=20)**
protection_action_other	Other actions people take to protect themselves from card info theft and fraud.	7	1	0	0	8
protection_action_physical_awareness	Looking for card skimmers, hiding PIN, putting things in a place so they won't be stolen, paying close attention to what a shopkeeper does, checking receipts, etc.	10	1	0	1	11
protection_action_use_cash	Using cash.	6	0	0	0	6
protection_action_use_chip	Using the chip in their card (as opposed to the magnetic stripe).	1	0	0	0	1
protection_action_use_credit	Using a credit card, since getting refunded is easier, etc.	6	1	1	0	7
protection_action_use_debit	Using a debit card or using a debit card as a credit card.	1	0	0	0	1
protection_action_use_other_payment_service	Using PayPal, Venmo, or another payment service other than Pay.	2	1	1	1	4
protection_action_use_pay	Using Apple Pay, Google Pay, or Samsung Pay (coded only when brought up prior to us suggesting that they use Pay).	2	0	0	0	2

Continued on the next page

Code	**Description**	**Interview #1 (n=20)**	**Interview #2 (n=10)**	**Survey #2 (n=20)**	**Survey #3 (n=10)**	**Overall (n=20)**
reasons_for_not_setting_up	People's reasons for not setting up Pay.	14	0	9	0	16
reasons_for_not_using	People's reasons why they don't want to or did not use Pay.	7	1	4	6	12
response_efficacy_security_convinced	Reasons why participants are convinced that Pay will protect them.	12	4	2	1	14
response_efficacy_security_unconvinced	Reasons why participants think Pay will not protect them.	8	2	1	0	8
response_efficacy_security_unsure	Reasons why participants are unsure whether Pay will protect them.	12	4	0	0	14
self_efficacy_negative_battery	Using Pay requires a charged phone.	1	0	0	0	1
self_efficacy_negative_learning	Using or setting up Pay requires practice, learning, or attention to detail.	7	1	1	0	7
self_efficacy_negative_limited_availability	Not all places accept Pay. It may be unclear whether a given place accepts it.	7	7	2	0	12
self_efficacy_negative_limited_card_compatibility	Not all cards can be added to Pay.	1	2	0	0	3
self_efficacy_negative_opportunities	Not going shopping, not having any money, etc., and so not having opportunities to use Pay.	4	4	5	3	11

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
self_efficacy_negative_other	Other challenges to using Pay, negative experiences using it, and things that make using it more difficult.	7	5	2	2	9
self_efficacy_negative_overspending	The convenience of Pay makes the participant more inclined to wastefully or accidentally spend money.	4	0	0	0	4
self_efficacy_negative_payment_failure	Pay payments not going through or taking too long/timing out.	2	6	2	2	6
self_efficacy_negative_remembering	Difficulty remembering to use Pay.	5	3	2	1	9
self_efficacy_negative_setup	Difficulty or irritation setting up Pay.	9	2	6	0	14
self_efficacy_negative_time	It taking too long or a long time to use Pay.	1	4	1	2	6
self_efficacy_other	Other comments about Pay usability, that are neither positive nor negative.	2	2	0	0	4
self_efficacy_positive_easy_to_use	Fast, simple, convenient, etc. to make transactions.	14	6	4	1	15
self_efficacy_positive_extensive_availability	Many or enough places accept Pay.	4	1	0	0	4

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
self_efficacy_positive_initiative	People taking the initiative to determine whether Pay is accepted (e.g., asking if Pay is accepted, or attempting to use it if they're unsure). Not coded if people said they didn't take the initiative.	0	1	0	0	1
self_efficacy_positive_no_wallet	If you use Pay, you won't have to carry your wallet, carry your cards, or pull out your cards or wallet.	7	5	2	3	10
self_efficacy_positive_novelty	Using or setting up Pay due to curiosity, wanting to see if it works.	9	2	5	1	10
self_efficacy_positive_only_option	Being more likely to use or using Pay because it's an option if you forget another payment method, another payment method doesn't work, you don't have your cards with you, etc. Also includes making Pay more accessible than cards (e.g., by burying cards in your purse and leaving phone on top).	7	5	1	0	10

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
self_efficacy_positive_opportunities	Going shopping, etc., and so having opportunities to use Pay. Includes inferred opportunities (e.g., if someone says they used Pay, that implies they had opportunities).	15	10	5	3	18
self_efficacy_positive_other	Other non-security perks to using Pay, positive experiences using it, good things about Pay, etc.	4	1	0	0	5
self_efficacy_positive_other_reminders	Other things reminding people to use Pay. Not including the setup instructions or implementation intention plan template we offer users. Not including it being the only option.	1	6	1	0	7
self_efficacy_positive_setup	Positive things said about the setup process (easy, etc.).	14	4	1	0	17
self_efficacy_practice	Wanting to practice (or actually practicing) using Pay in a low-pressure situation (e.g., a vending machine, a self-checkout, etc.).	0	1	1	0	2

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
threat_severity_card_type	The severity of fraud would depend on what type of card was affected by the fraud (e.g., fraud on credit vs debit card).	2	0	0	0	2
threat_severity_high_concern_gets_worse	When fraud or information theft occurs, this might be a precursor to something worse (e.g., a worse hack, more lost money, etc.).	10	0	0	0	10
threat_severity_high_concern_hassle	Resolving the situation would be time-consuming, stressful, irritating, etc.	10	0	0	0	10
threat_severity_high_concern_lost_money	Being concerned about losing money, either from purchases not being refunded, or not being refunded for overdraft or other fees.	6	0	0	0	6
threat_severity_high_concern_other	Other reasons why people perceive the severity to be higher.	7	0	0	0	7
threat_severity_high_concern_violation	People feel violated, helpless, angry, etc. when they suffer from card fraud or information theft.	3	0	0	0	3

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
threat_ severity_ low_concern_ other	Other reasons why people perceive the severity to be lower.	2	0	0	0	2
threat_ severity_ low_concern_ resolution	It would be possible to resolve the situation.	11	0	0	0	11
threat_ severity_ other	Other things that impact perceptions of threat severity.	1	0	0	0	1
threat_ severity_ purchase_size	The severity of fraud would depend on the size of the fraudulent purchase which was made.	6	0	0	0	6
threat_ susceptibility_ comparison	Participants comparing the relative likelihood of one type of card (information) theft/fraud to another type of event. For example, it being more likely for debit information to be stolen than credit information.	11	0	0	0	11

Continued on the next page

Code	Description	Interview #1 (n=20)	Interview #2 (n=10)	Survey #2 (n=20)	Survey #3 (n=10)	Overall (n=20)
threat_susceptibility_high_likelihood	Reasons participants perceive the likelihood of encountering the threat to be higher.	12	0	0	0	12
threat_susceptibility_low_likelihood	Reasons participants perceive the likelihood of encountering the threat to be lower or non-existent (e.g., it's never happened to me before, it's never going to happen, etc.).	9	0	0	0	9
threat_susceptibility_other	Other things that impact perceptions of threat susceptibility. Also includes participants expressing that they are unsure about their threat susceptibility.	2	0	0	0	2

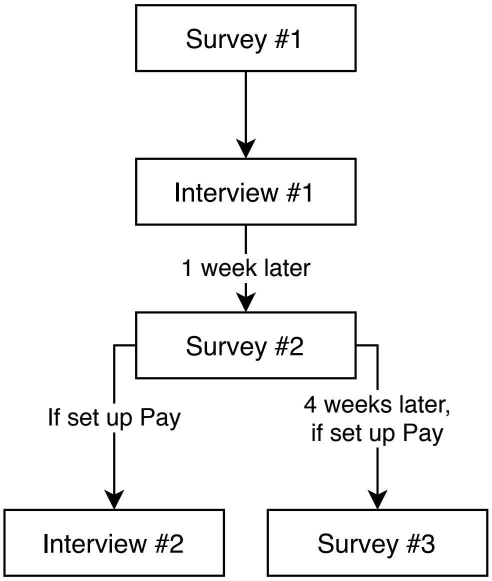

Figure A.1: The process of administering surveys and interviews in the qualitative portion of our study.

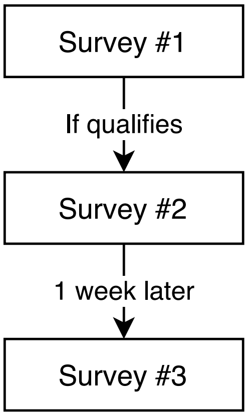

Figure A.2: Our controlled experiment contained three online surveys.

Demographic	**Values**	
Age	Minimum	18
	Median	32
	Mean	34.7
	Maximum	71
Gender	Female	58%
	Male	41%
	Other	1%
Employment	Working	74%
	Student	11%
	Not employed	10%
	Other	6%
Education	High school or less	18%
	College or associate	56%
	Graduate degree	18%
	Professional degree	4%
	Other	3%
Worked or Studied in a Computer-related Field	Yes	25%
	No	75%
Household Income	Median	$60,000 to $79,999

Table A.2: Demographics for the 411 participants who completed all parts of our controlled experiment.

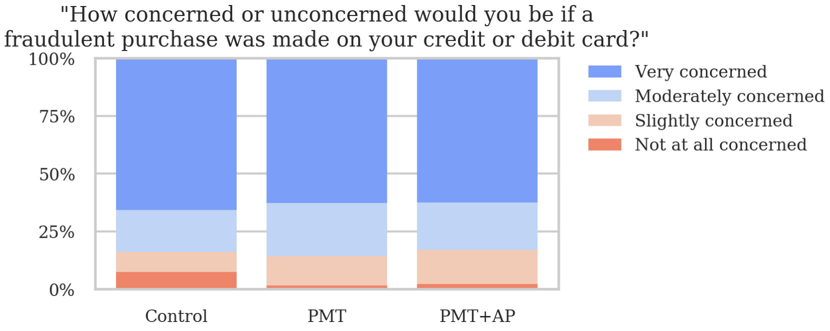

Figure A.3: We did not find statistically significant evidence that our treatments affected perception of threat severity ($p = 0.932$).

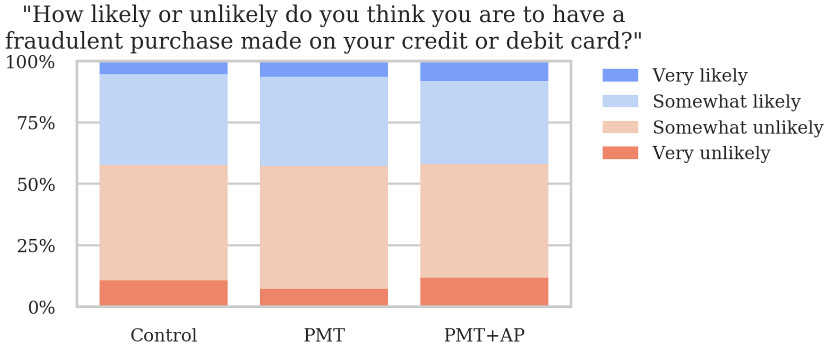

Figure A.4: We did not find statistically significant evidence that our treatments affected perception of threat susceptibility ($p = 0.881$).

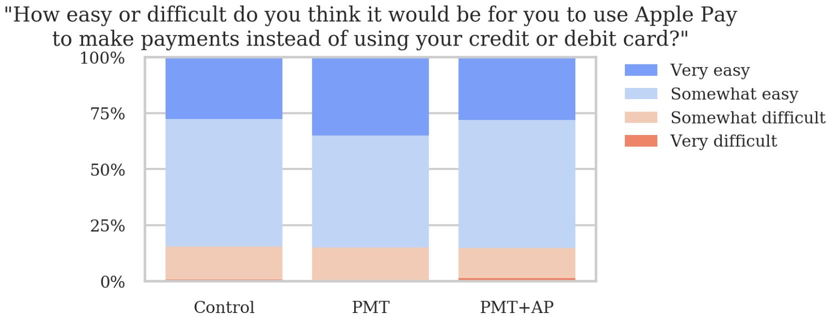

Figure A.5: We did not find statistically significant evidence that our treatments affected perception of self-efficacy ($p = 0.523$).

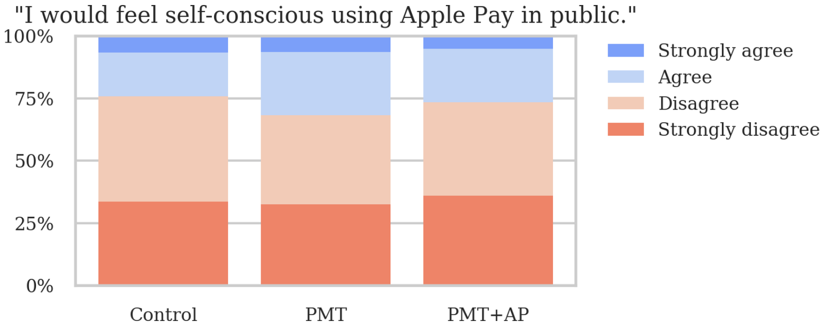

Figure A.6: We did not find statistically significant evidence that our treatments affected self-consciousness ($p = 0.628$).

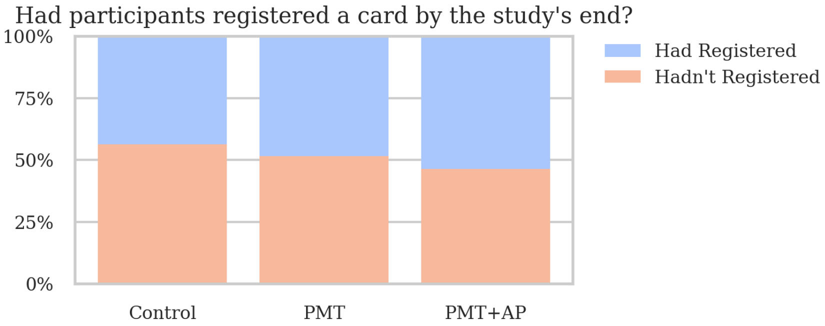

Figure A.7: We did not find statistically significant evidence that our treatments affected whether participants would have a card registered in Apple Pay by the end of our study ($p = 0.237$).

Variable	$\boldsymbol{\beta}$	$\boldsymbol{e^{\beta}}$	**p-value**
CS_background	-1.438	0.237	**0.001**
experienced_fraud	-0.799	0.450	**0.024**
threat_severity	-0.353	0.703	0.414
knows_users	-0.041	0.960	0.911
age	0.010	1.010	0.527
SA6	0.042	1.043	0.259
response_efficacy	0.146	1.157	0.783
self_conscious	0.281	1.324	0.451
usefulness	0.368	1.445	0.509
self_efficacy	0.377	1.458	0.501
threat_susceptibility	0.383	1.466	0.296
own_watch	0.483	1.621	0.198
Face_ID	0.745	2.106	**0.022**
non-female_gender	0.870	2.387	**0.009**
prior_use	1.295	3.653	<**0.001**
intention	1.804	6.077	<**0.001**
treatment			0.297
PMT	0.390	1.477	0.394
PMT+AP	0.698	2.010	0.123
Intercept	-4.856	0.008	<**0.001**

Table A.3: Our logistic regression model for predicting use of Apple Pay by those who completed Survey #1, #2, and #3 ($n = 411$). e^{β} indicates the change in odds of using Apple Pay when the variable is true. p-values significant at $\alpha = 0.05$ are bolded. Cox & Snell $R^2 = 0.238$.

Variable	$\boldsymbol{\beta}$	$\boldsymbol{e^{\beta}}$	**p-value**
CS_background	-1.635	0.195	0.059
experienced_fraud	-1.519	0.219	**0.039**
response_efficacy	-1.326	0.266	0.245
messaged_info	-0.607	0.545	0.451
usefulness	-0.598	0.550	0.598
non-female_gender	-0.563	0.569	0.378
own_watch	-0.243	0.784	0.728
threat_severity	-0.118	0.889	0.904
SA6	-0.011	0.989	0.870
age	0.008	1.009	0.752
checked_intention	0.165	1.179	0.923
knows_users	0.463	1.589	0.461
Face_ID	0.902	2.464	0.207
messaged_plan	0.936	2.549	0.245
self_efficacy	1.032	2.805	0.267
prior_use	1.377	3.964	**0.028**
threat_susceptibility	1.458	4.297	0.058
self_conscious	1.639	5.148	**0.020**
visited_location	3.414	30.374	0.052
intention	20.768	1045582764.370	0.997
Intercept	-24.824	0.000	0.997

Table A.4: Our logistic regression model for predicting whether those who received our action planning implementation intention treatment used Apple Pay ($n = 136$). e^{β} indicates the change in odds of using Apple Pay when the variable is true. p-values significant at $\alpha = 0.05$ are bolded. Cox & Snell $R^2 = 0.385$.

A.1 Qualitative Interviews Materials

A.1.1 Survey #1

Researchers at Carnegie Mellon University are conducting a study to understand people's use of smartphones, credit cards, and debit cards to make payments.

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 preliminary survey. If you are eligible, the preliminary survey will take about 10 minutes to complete. Only some of the participants who take this survey will be invited to participate in subsequent interviews and follow-up surveys. Participants will not be compensated for completing this survey: participants will only be compensated if they are selected to participate in subsequent parts of this study.

Do you live in the United States of America?
(Yes, No)

Do you speak English?
(Yes, No)

What is your age in years?
___

Are you able to visit Carnegie Mellon University's campus for an interview?
(Yes, No)

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)

Do you use a smartphone?
(Yes, No, I don't know)

In which country did you purchase your smartphone?
(The United States, Other: ___, I don't know)

What kind of smartphone do you have? If you have multiple phones, answer based on the phone you use the most.
(iPhone, Samsung phone, Other Android phone, Other: ___, I don't know)

[Here we show the iPhone-specific text, but users saw text appropriate to the type of phone they selected.]
What **model of iPhone** do you have? For example, iPhone 4S, iPhone 5, etc. You can check your phone's model by opening the "Settings" app, going to "General", then "About". Your phone's "Model Name" should be listed on the "About" page.
______

What **version of iOS** is running on your phone? For example, 7.9, 10.3, etc. You can check your phone's iOS software version by opening the "Settings" app, going to "General", then "About". Your phone's "Software Version" should be listed on the "About" page.
______

Do you have an Apple Watch?
(Yes, No)

We would like to understand how you make payments at brick and mortar stores, restaurants, or other physical locations.

Do you have a **credit card**?
(Yes, No)

Do you have a **debit card**?
(Yes, No)

Please select **all options** which accurately complete the following statement: "**Sometime in the past**, I have made in-person payments in physical locations..."
... using cash
... using my credit card
... using my debit card
... using Apple Pay. Apple Pay allows you to make payments using your smartphone.
... using Google Pay. Google Pay allows you to make payments using your smartphone.
... using Samsung Pay. Samsung Pay allows you to make payments using your smartphone.

Please select **all options** which accurately complete the following statement: "**In the past month**, I have made in-person payments in physical locations..."
... using cash
... using my credit card
... using my debit card
... using Apple Pay. Apple Pay allows you to make payments using your smartphone.
... using Google Pay. Google Pay allows you to make payments using your smartphone.
... using Samsung Pay. Samsung Pay allows you to make payments using your smartphone.

Has your **credit or debit card information** ever been stolen?
(Yes, No, I don't know)

How **concerned or unconcerned** would you be if your credit or debit card information was stolen in the future?
(Not at all concerned, Slightly concerned, Moderately concerned, Very concerned)

How **likely or unlikely** do you think you are to have your credit or debit card information stolen in the future?
(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

Has a **fraudulent purchase** ever been made on your credit or debit card?
(Yes, No, I don't know)

How **concerned or unconcerned** would you be if a fraudulent purchase was made on your credit or debit card in the future?
(Not at all concerned, Slightly concerned, Moderately concerned, Very concerned)

How **likely or unlikely** do you think you are to have a fraudulent purchase made on your credit or debit card in the future?
(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

How did you find this study?
(CBDR Participation Pool, Craigslist, Other: ___)

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)

Have you ever lived outside the United States for more than 1 month?
(Yes, No, Prefer not to answer)

Where outside the United States have you lived the longest?

______

If you are eligible for participation in this study, we may email you with an invitation to participate in the study. Because we have a limited number of interview slots available, we may not be able to interview all eligible candidates.

Name:

______

Email address:

______

A.1.2 Interview #1 Script

Hello XXX, my name is YYY [and my assistant's name is ZZZ]. Thank you for agreeing to participate in Interview #1. [I will be asking most of the questions, and ZZZ will be taking notes.] [I am/We are] very interested in your thoughts about credit cards, debits cards, and smartphones. This interview will be recorded, but the audio will not be shared with the public. Your responses will be kept anonymous, but quotes from your responses may be shared with the public.

Prior to completing Survey #1, you expressed your consent to participate in this study. However, the interview is completely voluntary, and you are free to end it at any time. The interview will take up to an hour. Is it alright if I start the audio recording now?

Great! I will start the audio recording now.

Alright, let's get started! Remember that there are no right or wrong answers to any of my questions.

Could you explain how you typically pay when you make a purchase in a physical location, like a brick and mortar store or restaurant?

In the survey, you also indicated that you used [a credit card][a debit card][credit and debit cards] to make purchases.

If has credit and debit: Is there a reason why you use one card instead of another?

If fraudulent purchase: In the survey, you wrote that a fraudulent purchase had been made on your credit or debit card. What happened? [Was it your credit or debit card?] [What did you do?] [How do you think it happened?]

If no fraudulent purchase: In the survey, you wrote that a fraudulent purchase had not been made on your credit or debit card. Do you know anyone who has had a fraudulent purchase made on their credit or debit card? What happened?

If don't know: In the survey, you wrote that you weren't sure if a fraudulent purchase had been made on your credit or debit card. What did you mean by that? [What did you do?]

If card info was stolen: In the survey, you wrote that your credit or debit card information had been stolen before. What happened? [Was it your credit or debit card?] [What did you do?] [How do you think it happened?]

If card info wasn't stolen: In the survey, you wrote that your credit or debit card information had not been stolen before. Do you know anyone who has had their credit or debit card information stolen? What happened?

If don't know: In the survey, you wrote that you weren't sure if your credit or debit card information had been stolen. What did you mean by that? [What did you do?]

[I think most people carry their smartphones all the time, but this is a sanity check.]

What kind of smartphone do you use?

Do you carry your smartphone with you every day?

Are there any times when you do go out without your smartphone?

If Apple Watch: In the survey, you indicated that you have an Apple Watch. Do you wear it every day?

If they have an iPhone: [Pay] = [Apple Pay]

If they have a non-Samsung Android phone: [Pay] = [Google Pay]

If they have a Samsung phone:

If they previously used Google Pay and Samsung Pay:

In the survey, you said that you had previously used Google Pay and Samsung Pay, but haven't used either to pay in a physical location in the last month.

If you were going to use one of them again, which would you use? [Why?] [If you don't have a preference, that's okay, too.]

If they previously used Google Pay xor Samsung pay:

In the survey, you said that you had previously used [Google Pay][Samsung Pay], but haven't used it to pay in a physical location in the last month. Your phone is also compatible with [Samsung Pay][Google Pay], which can also be used to make payments through your phone.

If you were going to use Google Pay or Samsung Pay in the future, which would you use? [Why?] [If you don't have a preference, that's okay, too.]

If they haven't previously used Google Pay or Samsung pay:

In the survey, you indicated that you hadn't used either Google Pay or Samsung Pay to pay in a physical location before. Google Pay and Samsung Pay are both mobile payments systems that allow you to make payments in stores through your phone. Your phone is compatible with both Google Pay and Samsung Pay.

If you were going to start using one, which would you choose? [Why?] [If you don't know enough to choose, that's okay, too.]

If Samsung Pay: [Pay] = [Samsung Pay]

If Google Pay: [Pay] = [Google Pay]

Else: [Pay] = [Samsung Pay]

In that case, let's focus on [Pay] for the rest of the interview.

If they previously used [Pay], but haven't used it recently:

Omit if asked above: In the survey, you said that you had previously used [Pay], but haven't used it to pay in a physical location in the last month.

Tell me about your experiences using [Pay]. [When did you first use it? For how long did you use it? Was your experience using [Pay] good or bad?]

Is there a reason why you haven't used [Pay] recently?

If they have never used [Pay]:

Omit if asked above: In the survey, you indicated that you hadn't used [Pay] to pay in a physical location before. [Pay] is a mobile payments system that allows you to make payments in stores through your phone [Apple watch: or watch].

Had you heard of [Pay] before taking the survey?

If yes: How did you hear about [Pay]? Have you set it up on your phone [or watch]?

If yes: Have you tried using [Pay] before? Is there any reason why you haven't used it to make a payment before?

If no: Is there any reason why you haven't set it up?

There have been many big hacks where credit and debit card information was stolen from retailers. For example, Target was hacked in 2013, Home Depot was hacked in 2014, and Saks Fifth Avenue was hacked last year. Information about millions of cards was stolen in these hacks. If criminals get your credit or debit card information, they might use that information to make fraudulent purchases. If you notice fraudulent purchases on your credit card, you can probably get refunded. But if the purchases are made on your debit card, you might not be able to get your money back. In any case, you would need to get a replacement card with a new number, which would be inconvenient.

How concerned or unconcerned would you be if a fraudulent purchase was made on your credit or debit card [again]? Why?

[Concern Likert] On this scale, which option best reflects your answer?

How likely or unlikely do you think you are to have a fraudulent purchase made on your credit or debit card [again]? Why?

[Likelihood Likert] On this scale, which option best reflects your answer?

How concerned or unconcerned would you be if your credit or debit card information was stolen [again]? Why?

[Concern Likert] On this scale, which option best reflects your answer?

How likely or unlikely do you think you are to have your credit or debit card information stolen [again]? Why?

[Likelihood Likert] On this scale, which option best reflects your answer?

Do you know of anything you can do to prevent your credit or debit card information from being stolen? [Have you done anything to protect your card information?]

Thankfully, there are steps you can take to prevent your card information from being stolen and to protect yourself from card fraud. One of the best things you can do is to start using [Pay]. Instead of paying by swiping or inserting your card, you can make payments through your phone [or watch], which adds an extra layer of security. Payments made with [Pay] will still be charged to your credit or debit card, but because the payments go through [Pay], your card number is not shown to or recorded by retailers. This means that your card number cannot be stolen from transactions made with [Pay]. If your phone [or watch] is stolen, the thief will not be able to make payments because [Pay] is protected by your [Apple: fingerprint/Face ID and lock screen PIN][Other: lock screen]. Although no system is perfectly secure, security experts generally agree that [Pay] is more secure than paying with credit or debit cards. [Pay] takes just a few minutes to set up, and is widely accepted. Apple Pay: As of this year, Apple Pay is accepted in 65% of retail locations in the United States. For example, Giant Eagle, ALDI, Dunkin' Donuts, and CVS all accept Apple Pay. Google Pay: Google Pay is accepted at millions of locations. For example, Giant Eagle, ALDI, Dunkin' Donuts, and CVS all accept Google Pay. Samsung Pay: Samsung Pay is accepted at most retail locations in the United States.

These instructions show you how to set up [Pay] on your phone and how to make payments in stores.

If Apple Watch: Since you wear an Apple Watch, you might also be interested in the instructions for using Apple Pay on your watch. Using your watch might be more convenient than using your phone, and it's just as secure.

Please take a minute to review the instructions. If you want to set up [Pay], feel free to try it right now. If you run into any trouble, I would be happy to help you set it up. However, you do not have to set up [Pay] if you do not want to.

[Pass the handout to the participant]

[If they make a phone call: pause the recording to avoid recording their card number, SSN, or other sensitive information]

[Note whether they simply read the instructions, or tried to set up Pay. Ask if it's unclear.]

[Ask or observe what the participant had to do to verify their card (e.g., whether they had to call their bank, open the bank's app, etc.)]

[After pausing for at least 30 seconds, or however long it takes them to start setting up Pay]

[Remember to resume the recording, if it was paused]

After reviewing the instructions, do you have any questions about [Pay]?

If they simply reviewed the instructions:
How easy or difficult do you think it would be for you to set up [Pay]? Why?
[Difficulty Likert] On this scale, which option best reflects your answer?
Do you plan to try to set up [Pay] later, or would you rather not? Why?

If they tried to set up Pay:
Were you able to complete the setup of [Pay]?
If yes: How easy or difficult was it for you to set up [Pay]? Why?
[Difficulty Likert] On this scale, which option best reflects your answer?
If no: How easy or difficult was it for you try to set up [Pay]? Why?
[Difficulty Likert] On this scale, which option best reflects your answer?
Do you plan to try to set up [Pay] later, or would you rather not? Why?

How easy or difficult do you think it would be for you to use [Pay] to make payments instead of using your credit or debit card? Why?
[Difficulty Likert] On this scale, which option best reflects your answer?

[Agreement Likert] On this scale, please rate your level of disagreement or agreement with the following statement:
“If I were to start using [Pay], I would be less likely to have my card information stolen.”
[And why do you choose that option?]

[Interest Likert] And on this scale, could you show me how interested or uninterested you are in using [Pay]? Why?

[Based on the person’s stated level of interest and why they feel that way, I may skip the entire implementation intention section below.]

[To determine which handout to give the person. If they are ambivalent:
If they set up Apple Pay: handout corresponding to where they set it up
If they wear the Apple Watch all the time: watch handout
Else: iPhone handout]

Apple Watch: If you were going to start using Apple Pay, do you think you would be more likely to pay with your phone or with your watch? [Why?]

If you plan to use [Pay] in order to protect your credit or debit card information, one challenge might be simply remembering to use [Pay]. Forming a simple, concrete plan to use [Pay] can help you remember. If you like, you can use the plan template I have written on this handout.

[Hand the appropriate handout to the person]

Please take a minute to read through the plan. If you want to use [Pay] in the coming week, I encourage you to fill out the plan, since it may help you remember to use [Pay]. However, you do not have to fill out or follow the plan.

[Note the number of locations the person wrote and which boxes they checked]
[Number of locations written: ___]
[Number of boxes checked: ___ / 3]
[Final box checked? ___]

You are welcome to keep the plan and the instructions for using [Pay].

Do you want to use [Pay] in the coming week?

If they did not fill out the plan:
Is there a reason why you didn’t fill out the plan?

If they did fill out the plan:
In the coming week, how likely or unlikely are you to visit at least one of the locations you listed? Why?
[Likelihood Likert] On this scale, which option best reflects your answer?

How likely or unlikely are you to try to use [Pay] at these locations? Why?
[Likelihood Likert] On this scale, which option best reflects your answer?

Do you think this plan will or will not help you remember to use [Pay]? Why?

Before we conclude the interview, do you have any other thoughts or questions?

Thank you for participating in this interview! In about one week, I will send you a short follow-up survey. After you complete that survey, I will email you a $15 Amazon e-Gift Card.

A.1.3 Survey #2

This survey is Survey #2 in the study "Use of Smartphones, Credit Cards, and Debit Cards" that you previously gave your consent to participate in. It will take about 10 minutes to complete this survey. If you complete this survey, we will email you a $15 Amazon e-Gift Card for your participation in our study.

Please answer the following questions about your experiences since our interview. There are no right or wrong answers to any of these questions, so please answer honestly.

You did not set up $PAY during our interview. Did you set up $PAY after the interview?
(Yes, No)

Please write a few sentences explaining why you [set up][did not set up] $PAY.

When did you set up $PAY?
(Today, Yesterday, A few days ago, Right after the interview)

Since our interview, have you **tried to use** $PAY to make a payment in a physical location?
(Yes, No)

Please write a few sentences explaining why you [tried][did not try] to use $PAY.

Since our interview, have you **successfully used** $PAY to make a payment in a physical location?
(Yes, No)

Please write a few sentences describing your experience [using][trying to use] $PAY.

Since our interview, have you done anything else to protect your credit or debit card information from being stolen, or to protect yourself from credit or debit card fraud?
(Yes, No)

Please write a few sentences explaining what other steps you have taken to protect yourself from card information theft or card fraud.

——

Are you interested in meeting for an additional 30 minute interview? If you participate in this interview, you will be compensated with an additional $15 Amazon e-Gift Card.
(Yes, No)

A.1.4 Interview #2 Script

Hello XXX, my name is YYY and my assistant's name is ZZZ. Thank you for coming to Interview #2. This interview is focused on your experiences since Interview #1. I will be asking most of the questions, and ZZZ will be taking notes. This interview will be recorded, but the audio will not be shared with the public. Your responses will be kept anonymous, but quotes from your responses may be shared with the public.

Prior to completing Survey #1, you expressed your consent to participate in this study. However, the interview is completely voluntary, and you are free to end it at any time. The interview will take roughly 30 minutes. Is it alright if I start the audio recording now?

Great! I will start the audio recording now.

Alright, let's get started! Remember that there are no right or wrong answers to any of my questions.

In interview #1, we discussed [Pay].

If setup after Interview #1: During our last interview, you did not [setup][complete the setup of] [Pay].

When did you set up [Pay]? What reminded you to set up [Pay]?

What did you have to do to set up [Pay]? [Did you have to call your bank?]

Did you try to use [Pay] since our last interview?

If yes:

These instructions show how to review the transactions you made with [Pay]. Please take a minute to review the transactions you made since our interview. [Hands handout]

What were your experiences trying to use [Pay]?

Where did using [Pay] work the best? What happened?

Where was using [Pay] the most difficult? What happened?

Are there any other experiences you'd like to share?

If no:

Did you have any opportunities to use [Pay]?

Did you visit any stores, restaurants, or other locations where you thought Pay might be accepted?

Why did you not end up using [Pay]?

Did you use [Pay] [if yes: more or] less than you thought you would?
Did anything about [Pay] surprise you?
Did you encounter any challenges trying to use [Pay]?
Do you plan to use [Pay] in the future? Why?
What is your overall impression of [Pay]?

During our last interview, we discussed making a concrete plan to help you remember to use [Pay].

If filled out in interview: You filled out the plan template during our last interview. Do you remember what the plan was?

If not filled out in interview: You did not fill out the plan template during our last interview. Did you fill out the plan after the interview? Do you remember what the plan was?

If filled out at some point: Part of the plan was listing three stores or restaurants you thought you might visit. Do you remember what stores or restaurants you listed?

If yes: Did you visit any of those locations? Did you try to use [Pay] there? Did you try to use [Pay] at any other locations?

Did you find the plan to be helpful or not helpful? [Did the plan help you remember to use [Pay]?] [Was the plan more or less helpful than you thought it would be?] [Do you think you would have remembered to use Pay if you hadn't made the plan? Why?] [Can you think of any other strategies to help you remember to use Pay?]

Did you do anything else to help you remember to use [Pay]?

Potentially ask for clarification about free-text responses to survey.

Before we conclude the interview, do you have any other thoughts or questions?

Thank you for participating in this interview! In the next couple days, I will email you a $15 Amazon e-Gift Card.

A.1.5 Survey #3

This survey is Survey #3 in the study "Use of Smartphones, Credit Cards, and Debit Cards" that you previously gave your consent to participate in. It will take about 10 minutes to complete this survey. If you complete this survey, we will email you a $5 Amazon e-Gift Card.

Please answer the following questions about your experiences in the past week. There are no right or wrong answers to any of these questions, so please answer honestly.

In the past week, did you **try to use** $PAY to make a payment in a physical location?
(Yes, No)

Please write a few sentences explaining why you [tried][did not try] to use $PAY.
______

In the past week, did you **successfully use** $PAY to make a payment in a physical location?
(Yes, No)

Please write a few sentences describing your experience [using][trying to use] $PAY.

______

How likely or unlikely are you to use $PAY in the future?
(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

In the past week, have you done anything else to protect your credit or debit card information from being stolen, or to protect yourself from credit or debit card fraud?
(Yes, No)

Please write a few sentences explaining what other steps you have taken to protect yourself from card information theft or card fraud. ______

A.2 Controlled Experiment Materials

A.2.1 Survey #1

Researchers at Carnegie Mellon University are conducting a study to understand people's use of Apple services.

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 #1. If you are eligible, Survey #1 will take about 5 minutes to complete. Only some of the participants who take Survey #1 will be invited to participate in two follow-up surveys (Surveys #2 and #3).

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

What operating system (OS) does your primary mobile phone have?
(iOS (iPhone), Other, I don't know)

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 Survey #1.

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)

In which country did you purchase your iPhone?
(United States, Other country ___, I don't know)

What model of iPhone do you have? For example, iPhone 4S, iPhone 5, etc. You can check your phone's model by opening the "Settings" app, going to "General", then "About". Your phone's "Model Name" should be listed on the "About" page.
(Original iPhone, iPhone 3G, ..., iPhone 11 (or 11 Pro or 11 Pro Max))

What version of iOS is running on your phone? For example, 7.9, 10.3, etc. You can check your phone's iOS software version by opening the "Settings" app, going to "General", then "About". Your phone's "Software Version" should be listed on the "About" page.
______

Do you own an Apple Watch?
(Yes, No)

Please select **all options** which accurately complete the following statement: "**Sometime in the past**, I have made in-person payments in physical locations..."
... using cash.
... using my credit card.
... using my debit card.
... using Apple Pay. Apple Pay allows you to make payments using your iPhone.

Please select **all options** which accurately complete the following statement: "**In the past week**, I have made in-person payments in physical locations..."
... using cash.
... using my credit card.
... using my debit card.
... using Apple Pay. Apple Pay allows you to make payments using your iPhone.

A.2.2 Survey #2

Researchers at Carnegie Mellon University are conducting a study to understand people's use of Apple services.

This survey is Survey #2 in the "Apple Services Study" that you previously gave your consent to participate in. It will take up to 30 minutes to complete this survey. If you complete **both** Survey #2 and Survey #3 **within 3 days of each survey invitation**, you will be compensated $7 total. We will invite you to Survey #3 one week after you complete Survey #2.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**.

[Control Group]
Apple Pay allows you to make payments in stores using your iPhone. Payments made with Apple Pay are charged to credit or debit cards that have been registered in Apple Pay.

[PMT and PMT+AP Groups]
[See Figure 3.1]
[See Figure 3.2]
[See Figure 3.3]

[PMT+AP Group]
[See Figure 3.4]
Please explain why you did not fill out the plan.
______

How **concerned or unconcerned** would you be if a fraudulent purchase was made on your credit or debit card?
(Not at all concerned, Slightly concerned, Moderately concerned, Very concerned)

How **likely or unlikely** do you think you are to have a fraudulent purchase made on your credit or debit card?
(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

How **easy or difficult** do you think it would be for you to use Apple Pay to make payments instead of using your credit or debit card?
(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 Apple Pay regularly, I would be **less likely** to be a victim of card fraud.”
(Strongly disagree, Disagree, Agree, Strongly agree)

How **useful or not useful** do you think Apple Pay would be for making payments?
(Not at all useful, Slightly useful, Moderately useful, Very useful)

Rate your level of **disagreement or agreement** with the following statement: “I would feel self-conscious using Apple Pay in public.”
(Strongly disagree, Disagree, Agree, Strongly agree)

Do you know anyone who uses Apple Pay?
(Yes, No, I’m not sure)

Do you have a credit or debit card registered in Apple Pay?

(Yes, No, I don't know)

When did you register a card in Apple Pay?
(Prior to taking this survey, While taking this survey)

Please explain why you do not know whether you have a credit or debit card registered in Apple Pay.
______

Rate your level of disagreement or agreement with the following statement: "I intend **to register** a credit or debit card in Apple Pay in the next week." (Strongly disagree, Disagree, Agree, Strongly agree)

Rate your level of disagreement or agreement with the following statement: "I intend **to use** Apple Pay in the next week."
(Strongly disagree, Disagree, Agree, Strongly agree)

What is your overall opinion of Apple Pay? (Please write a few sentences)
______

This is a link to the information about Apple Pay that we showed you earlier:
Apple Pay Setup, Use, and FAQ
Would you like us to send you a message on Prolific containing this link?
(Yes, No)

This is a link to your plan for using Apple Pay:
My Plan for Using Apple Pay
Would you like us to send you a message on Prolific containing this link?
(Yes, No)

Has a fraudulent purchase ever been made on your credit or debit card?
(Yes, No, I don't know)

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)

Each statement below describes how a person might feel about the use of security measures. Examples of security measures are laptop or tablet passwords, spam email reporting tools, software updates, secure web browsers, fingerprint ID, and anti-virus software.

Please indicate the degree to which you agree or disagree with each statement. In each case, make your choice in terms of how you feel **right now**, not what you have felt in the past or would like to feel.

There are no wrong answers.
(Strongly disagree, Somewhat disagree, Neither disagree nor agree, Somewhat agree, Strongly agree)

I seek out opportunities to learn about security measures that are relevant to me.

I am extremely motivated to take all the steps needed to keep my online data and accounts safe.

Generally, I diligently follow a routine about security practices.

I often am interested in articles about security threats.

I always pay attention to experts' advice about the steps I need to take to keep my online data and accounts safe.

I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.

A.2.3 Survey #3

Researchers at Carnegie Mellon University are conducting a study to understand people's use of Apple services.

This survey is Survey #3 in the "Apple Services Study" that you previously gave your consent to participate in. It will take up to 5 minutes to complete this survey. If you complete this survey **within 3 days of the survey invitation**, you will be compensated $7 total for completing Survey #2 and Survey #3.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**.

In Survey #2, you indicated that you [did not have][did not know whether you had] a credit or debit card registered in Apple Pay.

Since taking Survey #2 on $DATE, **have you registered a credit or debit card in Apple Pay**?
(Yes, No)

Please explain why you did not register a credit or debit card in Apple Pay.
______

Rate your level of disagreement or agreement with the following statement: “I intend **to register** a credit or debit card in Apple Pay in the next week.”
(Strongly disagree, Disagree, Agree, Strongly agree)

Since completing Survey #2 on $DATE, have you made an in-person payment in a physical location using Apple Pay?
(Yes, No, I don’t know)

Since completing Survey #2 on $DATE, how many payments have you made **with Apple Pay** in physical locations?
___

Please explain why you [used][did not use][do not know whether you used] Apple Pay.
______

Did you use Apple Pay in a location where you had previously paid with a credit or debit card?
(Yes, No, I don’t know)

[PMT+AP Group, if they wrote at least one location]
In Survey #2, you made a plan to use Apple Pay.
Since completing Survey #2 on $DATE, which of the locations in your plan, if any, **have you visited**?
($LOCATION_1, $LOCATION_2, $LOCATION_3)

Please select all options which accurately complete the following statement: “**Since completing Survey #2 on $DATE**, I have made in-person payments at **$LOCATION_N**...”
...using cash
...using my credit card
...using my debit card
...using Apple Pay. Apple Pay allows you to make payments using your iPhone.
...using another payment method. Please specify: ___

How **concerned or unconcerned** would you be if a fraudulent purchase was made on your credit or debit card?
(Not at all concerned, Slightly concerned, Moderately concerned, Very concerned)

How **likely or unlikely** do you think you are to have a fraudulent purchase made on your credit or debit card?
(Very unlikely, Somewhat unlikely, Somewhat likely, Very likely)

How **easy or difficult** do you think it would be for you to use Apple Pay to make payments instead of using your credit or debit card?
(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 Apple Pay regularly, I would be **less likely** to be a victim of card fraud.”
(Strongly disagree, Disagree, Agree, Strongly agree)

How **useful or not useful** do you think Apple Pay would be for making payments?
(Not at all useful, Slightly useful, Moderately useful, Very useful)

Rate your level of **disagreement or agreement** with the following statement: “I would feel self-conscious using Apple Pay in public.”
(Strongly disagree, Disagree, Agree, Strongly agree)

Rate your level of disagreement or agreement with the following statement: “I intend **to use** Apple Pay in the next week.”
(Strongly disagree, Disagree, Agree, Strongly agree)

Appendix B

Measuring Adoption of and Beliefs About Web Browsing Privacy Tools

B.1 Bundled Products and Incorrect Answers

As we describe in our limitations section (§ 4.2.2), 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 [114]. 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 [118]. 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 tracking 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.

B.2 Survey Materials

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 Carnegie Mellon University 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)

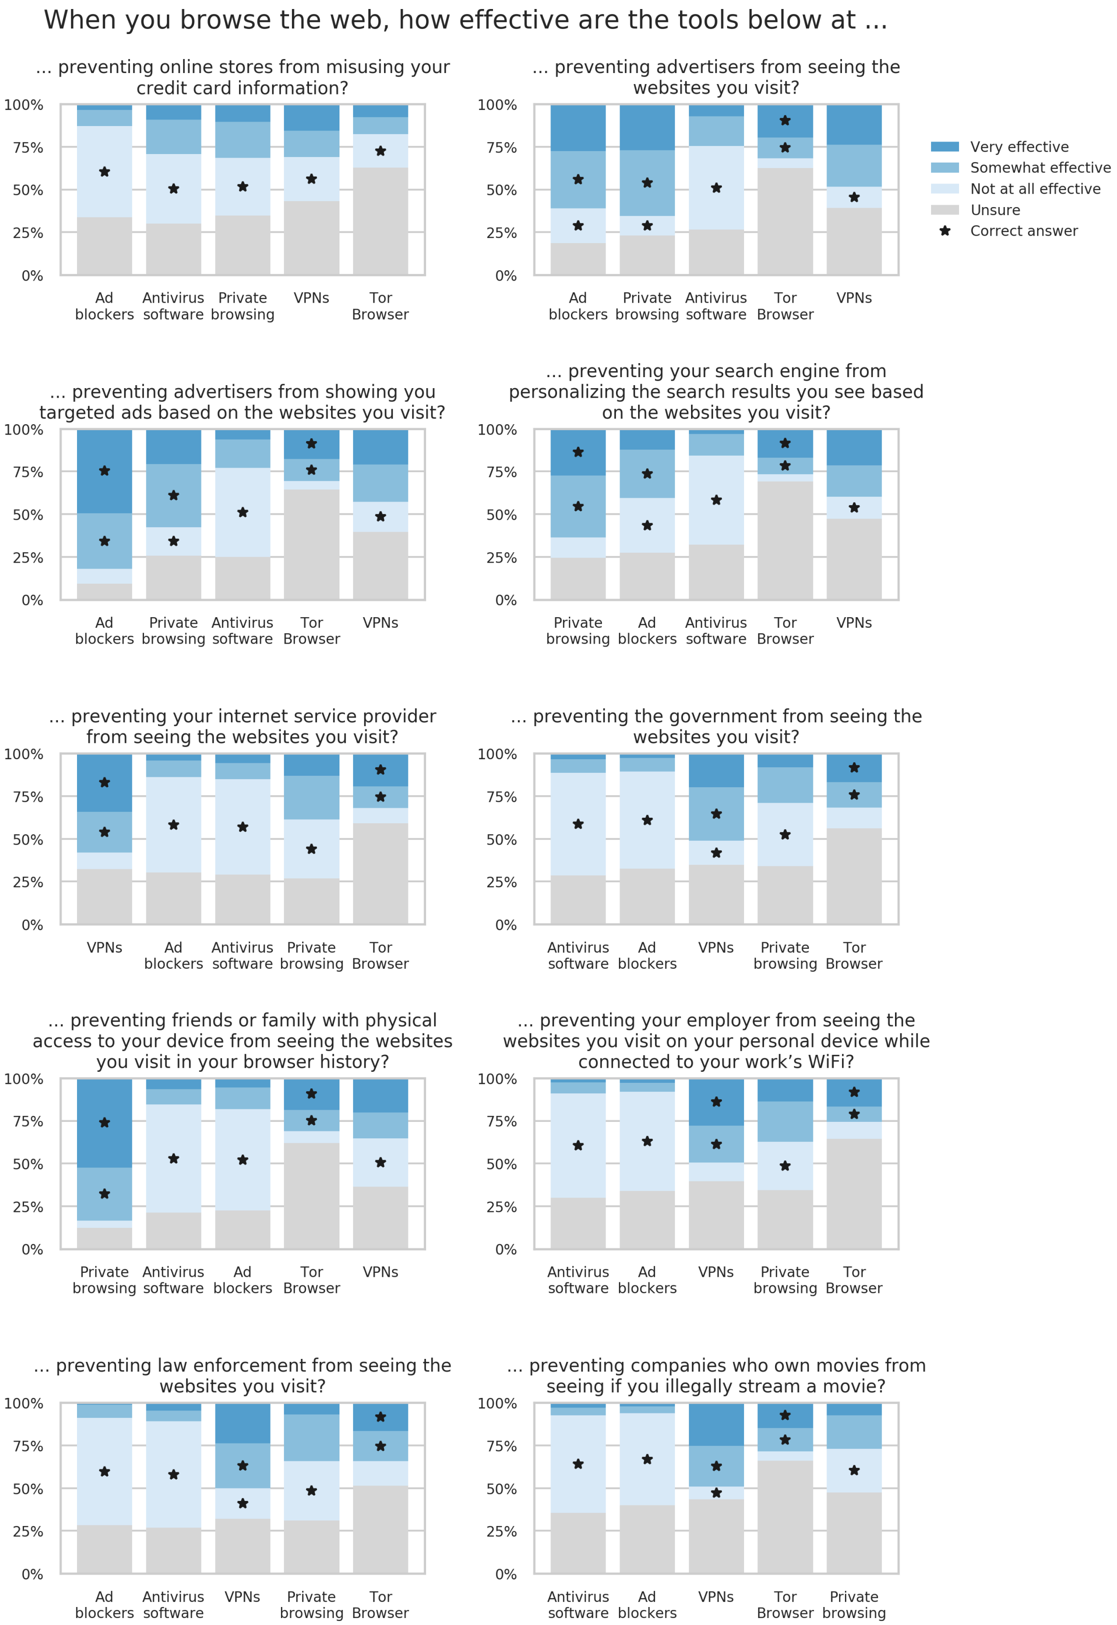

Figure B.1: Responses consistent with our threat model are indicated with a star. Tools are sorted by the percent of correct responses.

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 B.1: 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 [132]. We include data from the US Census Bureau for comparison [30].

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 B.2: 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 B.3 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
		Continued on the next page

Second Pass Code	**Description**	**Number of Occurrences**
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

Continued on the next page

Second Pass Code	**Description**	**Number of Occurrences**
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 [39, 114]. 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 [77, 115, 116, 118]: we count these extra features as true functions.	262

Table B.3: 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.

Appendix C

Nudges to Increase Adoption of Tor Browser

Use/Install Code	Description	Number of Occurrences
NOVELTY	Wanting to test out Tor Browser, compare it to other browsers, etc.	36
PRIVACY	Installing/using Tor Browser for its privacy protections. We count “not seeing ads” as a privacy issue (i.e., intrusion upon seclusion). We count listing topics which would commonly be considered privacy-sensitive (e.g., medical).	34
STUDY	Installing/using Tor Browser explicitly because of the study (e.g., if they think we asked them to use it, that it is required, they made a promise to do so, or they explicitly state that their plan or the study information is influencing them).	21
SECURITY	Installing/using Tor Browser for security protections. We count listing topics which would commonly be considered security-sensitive (e.g., logging into your bank account).	10

Continued on the next page

Use/Install Code	**Description**	**Number of Occurrences**
VAGUE_POSITIVE	Installing/using Tor Browser for vague positive reasons. If their other answers remove the ambiguity, it's okay to use the other answers to inform the choice of a different code.	9
CONTENT	For accessing inaccessible content (e.g., viewing country-specific content, paywalls, censored content, piracy, them accessing dark web sites, etc.).	6
HIGH_SELF_ EFFICACY	Thinking it would be easy to install/use.	4
GOOD_REVIEWS	Hearing/reading good things about Tor Browser (aside from those in the survey itself), or knowing others who use it. Not applicable if you heard about it before, but don't specify whether you heard good things or not.	3
JOB	For one's work or school.	1
Not Use/Install Code		
NOT_NEEDED	Not needing Tor Browser, whether stated generally, or for a particular reason (e.g., I don't need that level of protection, my needs are already met by another tool, etc.).	80
NOT_INSTALLED	Not using Tor Browser because it's not installed.	53
BUSY	Not having time to install/use Tor Browser (e.g., general busyness, vacation, being away from devices, etc.).	32
FORGOT	Forgetting to install/use Tor Browser.	30
LOW_SELF_ EFFICACY	Thinking it would be difficult, inconvenient, etc. to install/use. More vague than the explicit difficulties mentioned for USABILITY.	19
USABILITY	Usability challenges, such as Tor Browser being slow, websites not working, not functioning or opening (e.g., due to antivirus software), etc.	16
		Continued on the next page

Not Use/Install Code	**Description**	**Number of Occurrences**
SAFETY_DOUBTS	Doubting that Tor Browser is safe to install/use.	15
RESEARCH	Needing to do more research before installing/using Tor Browser.	13
VAGUE_NEGATIVE	Not installing/using Tor Browser for vague negative reasons. For example, writing just "I'm not interested," or "I'm lazy". If their other answers remove the ambiguity, it's okay to use the other answers to inform the choice of a different code.	13
DEVICE	Device-related limitations discouraging installation/use of Tor Browser (e.g., workplace prohibitions on installation, lack of disk space, a slow computer, etc.).	7
LOW_RESPONSE_EFFICACY	Doubting that Tor Browser is effective at protecting one's privacy, or doubting that anything can be done to protect one's privacy.	5
LOGINS	Tor Browser not being useful for activities that require logging in.	4
BAD_REVIEWS	Hearing/reading bad things about Tor Browser from sources that are clearly other than the survey itself, or finding it suspicious that they've heard nothing about it before.	4

Table C.1: Throughout our study, we asked participants whether they had installed or used Tor Browser, and their reasons for either doing so or not doing so. We collected multiple responses from all 537 participants who completed our experiment. We stopped coding after reaching code saturation; in total, we coded 558 free text responses from 150 randomly selected participants. Note that codes are not mutually exclusive, and that we count each code at most once per participant.

Activities Code	**Description**	**Described**	**Performed**	**Performed Using Tor Browser**
PNTD	Either the literal text “prefer not to disclose,” or something close to it.	192	100	33
SHOPPING	Looking up information about consumer products, regardless of intention to purchase.	55	39	13
FINANCIAL	Looking up information about financial products (e.g., stocks, bitcoin), mortgages, banking, insurance, salaries, applying to jobs, etc.	49	35	8
VAGUE	A vaguely defined activity, such as “using a search engine” or “researching things.”	49	41	25
NEWS	Looking up information about politics, celebrities, current events, document leaks, etc.	40	29	17
NSFW	Pornography or other “Not Safe For Work” content.	36	29	15
MEDICAL	Accessing medical information. Includes personal care and cannabis.	33	21	11

Continued on the next page

Activities Code	**Description**	**Described**	**Performed**	**Performed Using Tor Browser**
VIDEOS	Watching videos, movies, or streaming. We don't assume that all pornography is video-based. Since there is a "community" aspect to YouTube, simply mentioning "YouTube" isn't enough to assume this code applies.	26	18	10
YOUTUBE	Using YouTube.	25	22	12
SNOOPING	Looking up information about non-celebrities (e.g., ex's, friends, background checks) or similar entities (e.g., employers, competitors).	20	11	4
OTHER_ENTERTAINMENT	Websites about hobbies (e.g., emulation, listening to music), reading stories, blogs, etc.	16	10	7
N_A	Not plans or activities. For example, "none". Or "I will install Tor Browser."	16	0	0
MISC	Activities which are well-described but difficult to categorize.	14	11	4
TRAVEL	Travel-related browsing.	12	7	4
OTHER_SOCIAL	Using a generically specified social media website (e.g., "social media," "dating website," "forums," "anonymous messaging").	12	8	4

Continued on the next page

Activities Code	**Description**	**Described**	**Performed**	**Performed Using Tor Browser**
GOOGLE	Using Google search.	11	9	7
LOCAL	Apartment hunting, researching schools, wedding venues, etc.	10	9	9
REDDIT	Using Reddit.	9	8	6
OTHER_NAMED	Using another named website.	9	7	4
PIRACY	Pirating music, software, etc.	7	5	2
WIKI	Using Wikipedia or other wikis. Wikileaks doesn't count, since it isn't actually a wiki.	7	5	4
EMAIL	Accessing email.	7	6	0
AMAZON	Using Amazon.	4	3	1
FACEBOOK	Using Facebook.	4	3	2
DARK_WEB	Accessing the dark web.	2	2	2
TWITTER	Using Twitter.	2	2	1
RELIGION	Accessing religious information.	2	0	0
LEGAL	Accessing legal information.	2	0	0
LINKEDIN	Using LinkedIn.	1	1	0
PINTEREST	Using Pinterest.	1	0	0

Table C.2: In Survey #2, we gave participants in our PMT+AP treatment group the opportunity to plan to use Tor Browser for privacy-sensitive activities. Each participant was given the option to list up to three activities, so in some cases they contributed multiple times to the counts of the same codes. Also, note that codes were not mutually exclusive; for example, it was common for the VIDEOS and YOUTUBE codes to occur together. The "Described" column shows the number of activities with each code described in participants' plans. The "Performed" column shows the number of activities participants reported performing in either Survey #3 or Survey #4. The "Performed Using Tor Browser" column shows the number of activities participants reported performing using Tor Browser in either Survey #3 or Survey #4.

Challenges Code	**Description**	**Number of Occurrences**
BIT_SLOW	Websites were somewhat slow, but not extremely slow.	9
VAGUE	A vaguely defined challenge, or it's unclear whether there was a challenge at all. For example, "It didn't work."	7
SEARCHING	Difficulty finding pages (e.g., poor results from DuckDuckGo).	5
NOT_WORKING	An answer substantially the same as the predefined "Websites did not work" option (e.g., CAPTCHAs). Problems likely originating from the website, rather than the browser.	4
N_A	Clearly not a challenge "encountered when trying to use Tor Browser." For example, "I don't need it," or using the free text fields to explain other responses (e.g., "I made a mistake earlier in the survey").	4
FEATURES	Lacking features (e.g., bookmarks, ad blockers, login persistence, etc.).	4
CONNECTION	Tor Browser taking time or failing to connect to the Tor network.	3
CONFUSED	Expressing confusion about how to use Tor Browser, or what it is for.	2
EXTREMELY_SLOW	An answer substantially the same as the predefined "Websites were extremely slow" option (e.g., that might be mitigated by creating a new circuit).	2
IP_RELATED	IP address-related issues. Don't make inferences when coding (e.g., don't assume that a CAPTCHA is IP-related, unless the participant explicitly makes that connection).	2
LANGUAGE	Pages appearing in the wrong language.	2
LOW_RESPONSE_EFFICACY	The participant doesn't believe Tor Browser can protect their privacy.	1
CONFIGURATION	Configuration being a challenge.	1

Continued on the next page

Challenges Code	**Description**	**Number of Occurrences**
NO_CHALLENGE	An answer substantially the same as the predefined "I did not encounter any challenges" option.	1
COST	Tor Browser costing money.	1
SPACE	Lack of space on one's device.	1
NOT_ALLOWED	Not being allowed to install or use Tor Browser due to company policies, etc.	1

Table C.3: In Survey #3, we asked participants whether they had encountered challenges when trying to use Tor Browser. Some participants indicated that they had encountered a challenge other than those we listed. All participants were given a free text field to explain these challenges, and those in the coping planning treatment were asked to explain further. We coded such responses from 41 different participants, 3 of whom we determined not to have actually encountered a challenge (i.e., their free text responses were only coded with as "N_A"). Note that codes were not mutually exclusive, and if participants gave two free text responses, their responses might have different codes. However, we count each code at most once from each participant.

Challenge	**Coping Plan**	**Reencountered Challenge?**	**Followed Plan?**
I thought it was only for .onion sites and got confused! I was under the impression it was only for accessing hidden sites on the internet, like .onion domains and the silk road as was discussed in the first studies	I will engage with more tutorials and review the previously provided guide on actually using the tor browser.	No	Mostly yes
I was not always able to open websites, even when they had ”are you a human?” checks because they somehow saw me as not a legitimate access request	I think I will just have to open those sites in another browser	No	No
It was very difficult for me to save images from my search. viewing the image or image source only worked half the time. The web was a little slow but nothing bad. It was just frustrating to try and download images	It looks like its a somewhat common issue for android users, which is where I used Tor. An image would only have the ”save image” option half the time.	No	No
I used Tor the only time several years ago. I vaguely recall it being a bit slower, but don’t know that this reflects the current situation.	When I re-download it I’ll keep in mind the security benefits the browser offers and how this outweighs any lag.	No	Mostly no
I tried to use TOR on my phone, even with the workarounds offered it simply didn’t work. I wasn’t able to get it to function on my phone.	I plan to keep researching and see if there is a different method to get TOR on my smartphone. Also I plan to try to download TOR on my laptop just to see what my options for private browsing when I’m not working are.	Yes	Yes

Continued on the next page

Challenge	**Coping Plan**	**Reencountered Challenge?**	**Followed Plan?**
Functionality is limited because privacy protection is based so much on individual sites' policies (e.g. if I go to Google Maps or YouTube) that it doesn't actually help that much	Sorry, to be honest, you cannot get around this. It's not a matter of my individual will.	Yes	No
Migrating bookmarks and other personalization such as passwords was either difficult or not present (which I understand the password portion). It was frustrating using it for any activities that required usage of account or cookie-based websites. This is because the passwords and accounts don't save, for obvious reasons.	The best next step would be to use an independent password manager if I ever want to use Tor Browser again. Services like 1Password exist for a solid reason, so it might not be a bad idea to look into it.	Yes	Mostly no
The browser itself is not stable. When I launched the browser the popup screen to load it got stuck a couple of times. Then when I finally did "search" it took so long I gave up and looked up the information on Google Chrome instead.	I have no idea how to cover come this since it's a tech issue that I don't have control over.	Yes	Mostly no

Continued on the next page

Challenge	Coping Plan	Reencountered Challenge?	Followed Plan?
Sometimes some sites were slow but it was manageable. No challenges, just slower than usual. Videos play back just fine, the initial load time is just slow.	So if it becomes a real big issue, I would look at disconnecting from TOR, and relaunching. Perhaps I could find a faster Peer to connect to that isn't as slow. Worst case, if say watching a video, I could pause it, let it buffer and then proceed.	No	Yes
I don't really like DuckDuckGo so I was trying to use Google but every time I did it was in German. I have a hard time remembering to use it and when I do remember, I am usually not willing to wait for it to load. I like using Google search, but I see why they use DuckDuckGo as the default. The results on DuckDuckGo aren't terrible but I know there are some times when its hard to find what I'm looking for.	I can add the Tor browser to my task bar next to the other browsers so I will remember to use it. I could also leave it open so that it is ready for me to use when I need it.	Yes	Mostly yes

Continued on the next page

Challenge	**Coping Plan**	**Reencountered Challenge?**	**Followed Plan?**
It seemed a little slower than my other browsers but I wouldn’t describe it as ”extremely slow.”	Any additional challenges that I encountered I’d search in DuckDuckGo to learn more about.	Yes	No
It wouldn’t let me install it because it said I was lacking space on my computer.	Buy a new computer? But that would cost a lot of money. I don’t know what I could delete that I don’t need.	Yes	Mostly no

Table C.4: In Survey #3, we gave participants in our PMT+AP+CP treatment group who reported encountering challenges using Tor Browser the opportunity to form coping plans to overcome their challenges. Participants who did not select a listed challenge (i.e., “Websites were extremely slow” or “Websites did not work”) were given an open-ended plan template (Figure 5.10). This table contains these participants’ responses, lightly edited for clarity. In Survey #4, one week later, we checked whether participants reencountered the challenges they described, and whether they followed their plans to overcome the challenges.

Demographic Factor	**Values**	
Age	Minimum	18
	Median	33
	Mean	34.0
	Maximum	76
Gender	Female	32.6%
	Male	65.4%
	Other	2.0%
Employment	Working (paid employee)	65.7%
	Working (self employed)	8.2%
	Student	11.9%
	Not employed	10.8%
	Retired	2.6%
	Prefer not to answer	0.7%
Education	High school or less	22.5%
	College or associate	50.8%
	Graduate degree	21.8%
	Professional degree	3.9%
	Other	0.9%
	Prefer not to answer	0.0%
Worked or studied in a computer-related field	Yes	34.5%
	No	65.5%
Living situation	Domestic partner	50.5%
	Children	29.6%
	Parents	24.2%
	Other family	11.9%
	Unrelated roommates	5.0%
	I live alone	18.2%
	Other	0.2%
	Prefer not to answer	0.6%
Household income	Less than $10,000	3.5%
	$10,000 - $19,999	4.5%
	$20,000 - $39,999	11.5%
	$40,000 - $59,999	17.3%
	$60,000 - $79,999	15.5%
	$80,000 - $99,999	14.5%
	$100,000 or more	31.7%
	Prefer not to answer	1.5%

Table C.5: Demographics of the 537 participants who completed Surveys #1-#4.

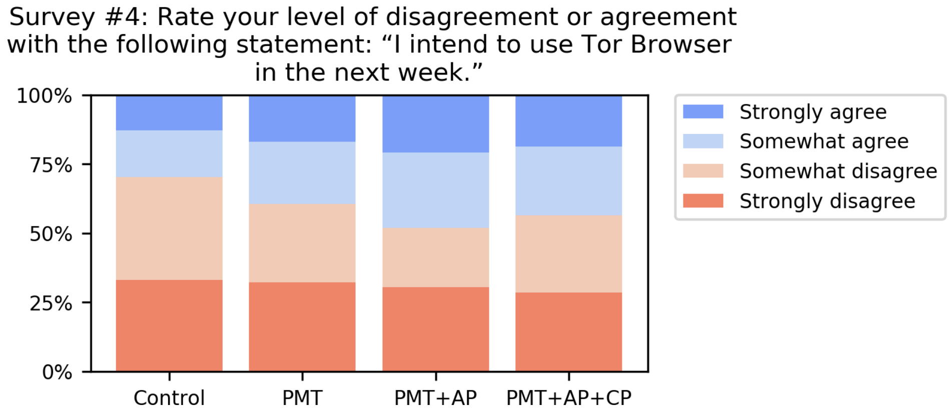

Figure C.1: This question measured intention to use Tor Browser. We did not find statistically significant differences in Survey #4.

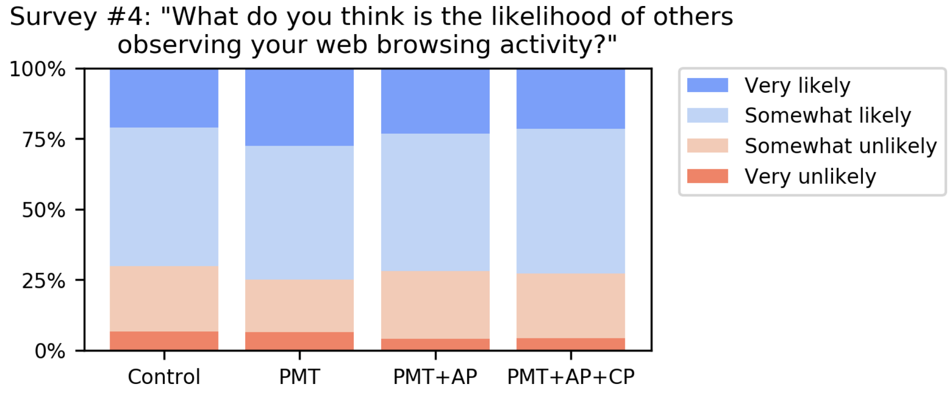

Figure C.2: This question measured perceptions of threat susceptibility. We did not find statistically significant differences in Survey #4.

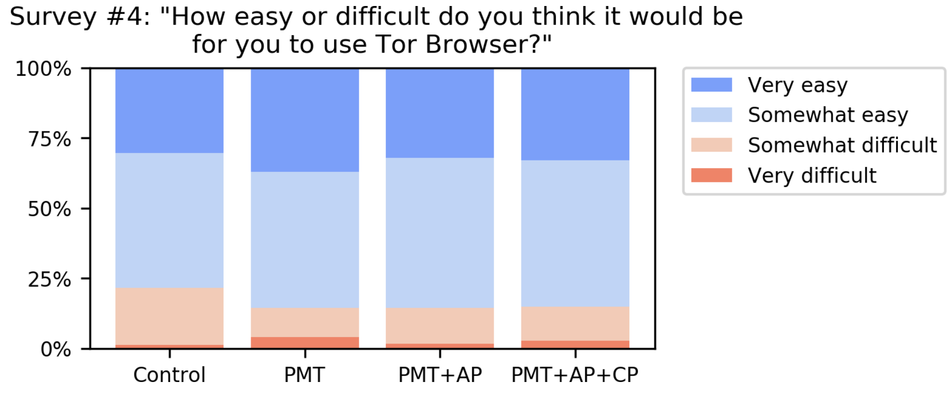

Figure C.3: This question measured perceptions of self-efficacy. We did not find statistically significant differences in Survey #4.

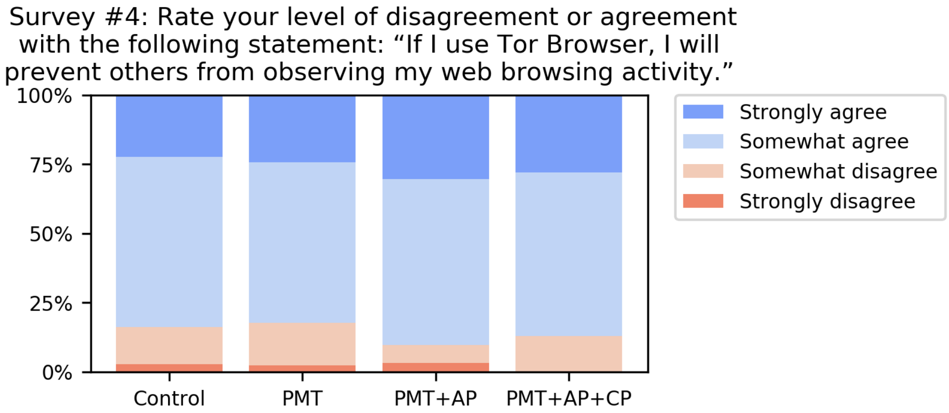

Figure C.4: This question measured perceptions of response efficacy. We did not find statistically significant differences in Survey #4.

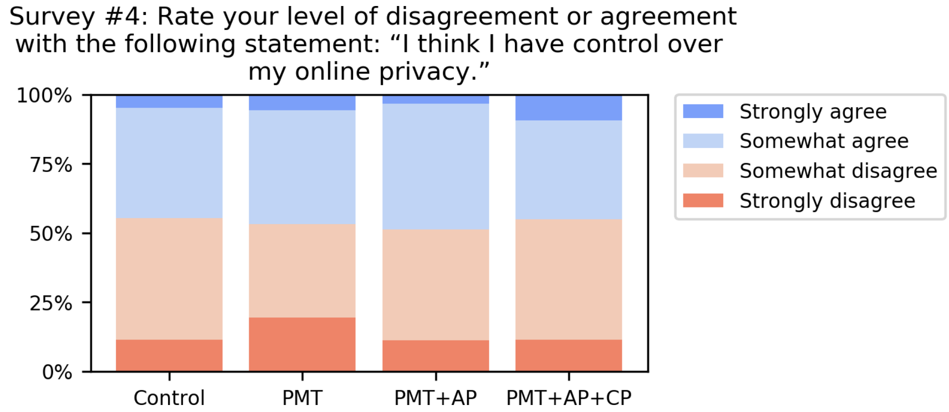

Figure C.5: This question measured perceptions of privacy control. We did not find statistically significant differences in Survey #4.

C.1 Survey Materials

C.1.1 Survey #1

Researchers at Carnegie Mellon University are conducting a research study to understand people's use of web browsing-related tools.

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 #1. If you are eligible, Survey #1 will take about 4 minutes to complete. Only some of the participants who take Survey #1 will be invited to participate in four follow-up surveys (Surveys #2, #3, #4, and #5).

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

Which operating system does your primary personal computer run?
(Windows, macOS, Ubuntu, Other, I don't know)

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 Survey #1.
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)

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

Have you **heard of** private browsing before?
(Yes, No, Unsure)

[If Yes]
Have you **used** private browsing before?

(Yes, No, Unsure)

[If Yes]
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)

VPNs
Have you **heard of** VPNs before?
(Yes, No, Unsure)

[If Yes]
Have you **used** a VPN before?
(Yes, No, Unsure)

[If Yes (i.e., used a VPN before)]
Do you use a VPN **primarily for work purposes**?
(Yes, primarily for work purposes; No, primarily for other purposes; About equally for work and other purposes)

[If Yes (i.e., used a VPN before)]
When did you most recently use a VPN?
(Today, In the past week, In the past month, In the past year, More than a year ago)

Tor Browser
Have you **heard of** Tor Browser before?
(Yes, No, Unsure)

[If Yes]
Have you **used** Tor Browser before?
(Yes, No, Unsure)

[If Yes]
When did you most recently use Tor Browser?
(Today, In the past week, In the past month, In the past year, More than a year ago)

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, 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 a response other than Never was selected for Laptop or Desktop]

In general, are you comfortable installing software on your [laptop or desktop]?
(Yes, No, Unsure)

Rate your level of **disagreement or agreement** with the following statement:
“I think I have control over my online privacy.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How interested or uninterested would you be in **preventing advertisers from seeing the websites you visit**?
(Not at all interested, Slightly interested, Moderately interested, Very interested)

How interested or uninterested would you be in **preventing the websites you visit from seeing what physical location you are browsing from**?
(Not at all interested, Slightly interested, Moderately interested, Very interested)

How interested or uninterested would you be in **preventing your internet service provider from seeing the websites you visit**?
(Not at all interested, Slightly interested, Moderately interested, Very interested)

How interested or uninterested would you be in **preventing the government from seeing the websites you visit**?
(Not at all interested, Slightly interested, Moderately interested, Very interested)

C.1.2 Survey #2

Researchers at Carnegie Mellon University are conducting a research study to understand people’s use of web browsing-related tools.

This survey is Survey #2 in the “Research Study for Internet Users” that you previously gave your consent to participate in. It will take up to 8 minutes to complete this survey. If you complete Survey #2, Survey #3, **and** Survey #4 **within 2 days of each survey invitation**, you will be compensated $3.50 total. We will invite you to each survey one week after you complete the previous survey.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**. All links to external resources are optional: your compensation will not be affected by whether you follow them.

[Control Group]
Tor Browser is an alternative web browser.

[PMT, PMT+AP, and PMT+AP+CP Groups]
[Threat information: Figure 5.2]
[Response information: Figure 5.3]

Please review these materials about Tor Browser.
[Usage and installation instructions: Figure 5.4]
[Optional technical details: Figure 5.5]
[Frequently asked questions: Figure 5.6]
[Common Problems: Figure 5.7]
If you want to use Tor Browser, we encourage you to install it now [169]. It only takes a minute to install Tor Browser. However, you do not have to install Tor Browser if you do not want to: your compensation will not be affected.

[PMT+AP and PMT+AP+CP Groups]
[Action plan: Figure 5.8]
For your convenience, here is a link to the information about Tor Browser that we showed you earlier:
Tor Browser Setup, Use, and FAQ
If you want to use Tor Browser in the coming week, we encourage you to fill out the plan, since it may help you remember to use Tor Browser. However, you do not have to use Tor Browser if you do not want to: your compensation will not be affected. Do you want to continue without writing any activities?
(Yes, I would like to continue without writing any activities)

[PMT, PMT+AP, and PMT+AP+CP Groups]
Thank you for reviewing this information about Tor Browser.

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 Tor Browser.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How **easy or difficult** do you think it would be for you to use Tor Browser?
(Very difficult, Somewhat difficult, Somewhat easy, Very easy)

Rate your level of **disagreement or agreement** with the following statement:
“If I use Tor Browser, I will prevent others from observing my web browsing activity.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Do you know anyone who uses Tor Browser?
(Yes, No, I’m not sure)

Is Tor Browser currently installed on one of your devices?
(Yes, No, I don't know)

[If Yes (i.e., Tor Browser is installed)]
When did you install Tor Browser?
(Prior to taking this survey, While taking this survey)

[If Yes (i.e., Tor Browser is installed)]
Please explain why you installed Tor Browser.
___

[If I don't know (i.e., whether Tor Browser is installed)]
Please explain why you do not know whether you have Tor Browser installed.
___

[If No or I don't know (i.e., whether Tor Browser is installed)]
Rate your level of disagreement or agreement with the following statement:
"I intend **to install** Tor Browser in the next week."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Rate your level of disagreement or agreement with the following statement:
"I intend **to use** Tor Browser in the next week."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

What is your overall opinion of Tor Browser? (Please write a few sentences)
___

[PMT, PMT+AP, and PMT+AP+CP Groups]
This is a link to the information about Tor Browser that we showed you earlier:
Tor Browser Setup, Use, and FAQ
Would you like us to send you a message on Prolific containing this link?
(Yes, No)

[PMT+AP and PMT+AP+CP Groups]
This is a link to your plan for using Tor Browser:
My Plan for Using Tor Browser
Would you like us to send you a message on Prolific containing this link?
(Yes, No)

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)

C.1.3 Survey #3

Researchers at Carnegie Mellon University are conducting a research study to understand people's use of web browsing-related tools.

This survey is Survey #3 in the "Research Study for Internet Users" that you previously gave your consent to participate in. It will take up to 6 minutes to complete this survey. If you complete **both** Survey #3 and Survey #4 **within 2 days of each survey invitation**, you will be compensated $3.50 total. We will invite you to Survey #4 one week after you complete this survey.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**.

[If not installed, or unsure whether installed]
In Survey #2, you indicated that you [did not have][did not know whether you had] Tor Browser installed on any of your devices.

Since completing Survey #2 on $DATE, **have you installed Tor Browser**?
(Yes, No)

[If Yes] Please explain why you installed Tor Browser.
___

[If No] Please explain why you did not install Tor Browser.
___

If you are interested in installing Tor Browser but require technical assistance, you are welcome to message us on Prolific to request help.

Rate your level of disagreement or agreement with the following statement:
“I intend **to install** Tor Browser in the next week.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Since completing Survey #2 on $DATE, **have you used Tor Browser**?
(Yes, No, I don’t know)

Since completing Survey #2 on $DATE, **on which days did you use Tor Browser**, if any?
($DATE, $DATE - 1, $DATE - 2, . . .)

[If Yes (i.e., used Tor Browser)]
Please explain why you used Tor Browser.
___

[If No (i.e., did not use Tor Browser)]
Please explain why you did not use Tor Browser.
___

[If I don’t know (i.e., whether they used Tor Browser)]
Please explain why you do not know whether you used Tor Browser.
___

[PMT+AP Group, if wrote at least one activity]
In Survey #2, you made a plan to protect your privacy when performing privacy-sensitive browsing activities.

Since completing Survey #2 on $DATE, which of the following privacy-sensitive activities **have you performed**, if any?
($ACTIVITY_1, $ACTIVITY_2, $ACTIVITY_3)

[If performed $ACTIVITY_N]
When performing the [first/second/third] activity (“$ACTIVITY_N”), **did you use Tor Browser**?
(Yes, No, I don’t know)

[If No (i.e., did not use Tor Browser) or I don’t know (i.e., whether they used Tor Browser)]
Have you ever **tried to use** Tor Browser?
(Yes, No, I don’t know)

[If used or tried to use Tor Browser]
Which of the following challenges have you encountered when trying to use Tor Browser, if any?
I did not encounter any challenges, Websites were extremely slow, Websites did not work, Other:___)

[If multiple choices were selected]
Which of these challenges was the greatest obstacle to using Tor Browser?
(Websites were extremely slow, Websites did not work, Other: "$OTHER_CHALLENGE")

[If PMT+AP+CP Group, and Websites were extremely slow]
[Figure 5.9]

[If PMT+AP+CP Group, and Websites did not work]
[Figure 5.9]

[If PMT+AP+CP Group, and Other]
[Figure 5.10]

If you want to use Tor Browser in the coming week, we encourage you to fill out the plan, since it may help you overcome challenges associated with using Tor Browser. However, you do not have to fill out or use the plan if you do not want to: your compensation will not be affected. Do you want to continue without filling out the plan?
(Yes, I would like to continue without filling out the plan)

Rate your level of disagreement or agreement with the following statement:
"I intend **to use** Tor Browser in the next week."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

[If PMT+AP+CP Group, and reported a challenge]
This is a link to your plan(s) for using Tor Browser:
My Plan(s) for Using Tor Browser
(Information about your latest plan will appear shortly after you submit this survey)
Would you like us to send you a message on Prolific containing this link?
(Yes, No)

C.1.4 Survey #4

Researchers at Carnegie Mellon University are conducting a research study to understand people's use of web browsing-related tools.

This survey is Survey #4 in the "Research Study for Internet Users" that you previously gave your consent to participate in. It will take up to 3 minutes to complete this survey. If you complete this survey **within 2 days of the survey invitation**, you will be compensated $3.50 total for participating in our study.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**.

[Installation and usage checkup, the same as in Survey #3]

[Action plan checkup, the same as in Survey #3]

[If PMT+AP+CP Group, and made the “Websites were extremely slow” plan]
In Survey #3, you made a plan to click the “New Circuit” button if you encountered extremely slow websites when using Tor Browser.

Since completing Survey #3 on $DATE, did you encounter extremely slow websites when using Tor Browser?
(Yes, No, I don’t know)

[If I don’t know]
Please explain why you do not know whether you encountered extremely slow websites when using Tor Browser.
___

Since completing Survey #3 on $DATE, did you click the “New Circuit” button?
(Yes, No)

[If PMT+AP+CP Group, and made the “Websites did not work” plan]
In Survey #3, you made a plan to use alternative websites if particular websites did not work for you in Tor Browser.

Since completing Survey #3 on $DATE, which of the following websites did you **try to visit with Tor Browser**, if any?
($ORIGINAL_WEBSITE_1,
$ORIGINAL_WEBSITE_2,
$ORIGINAL_WEBSITE_3)

[If $ORIGINAL_WEBSITE_N]
Since completing Survey #3 on $DATE, did **$ORIGINAL_WEBSITE_N** work when you tried to visit it with Tor Browser?
(Yes, every time I tried to visit it; Yes, but only some of the times I tried to visit it; No, it never worked)

Since completing Survey #3 on $DATE, which of the following **alternative websites** did you try to visit with Tor Browser, if any?
($ALTERNATIVE_WEBSITE_1,
$ALTERNATIVE_WEBSITE_2,
$ALTERNATIVE_WEBSITE_3)

[If PMT+AP+CP Group, and made the “Other” plan]
In Survey #3, you described the challenge(s) you encountered when trying to use Tor Browser:
“$CHALLENGE”

Since completing Survey #3 on $DATE, did you encounter the challenge(s)?
(Yes, No, I don’t know)

[If I don’t know]
Please explain why you do not know whether you encountered the challenge(s).
–––

In Survey #3, you described your plan to overcome the challenge(s):
“$PLAN”

Since completing Survey #3 on $DATE, did you follow your plan?
(Yes, Mostly yes, Mostly no, No)

[If formed an action or coping plan]
Were your plans helpful or not helpful? Please explain in a few sentences.
–––

Rate your level of **disagreement or agreement** with the following statement:
“I think I have control over my online privacy.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

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 Tor Browser.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How **easy or difficult** do you think it would be for you to use Tor Browser?
(Very difficult, Somewhat difficult, Somewhat easy, Very easy)

Rate your level of **disagreement or agreement** with the following statement:
“If I use Tor Browser, I will prevent others from observing my web browsing activity.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Rate your level of disagreement or agreement with the following statement:
“I intend **to use** Tor Browser in the next week.”
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Would you like to share any other thoughts about this study or about Tor Browser?

You are eligible to complete a final, optional survey (Survey #5), which would take up to 3 minutes to complete. If you complete Survey #5 **within 7 days of being invited**, you will be compensated **an additional $1**. You would receive your invitation in three weeks.

Your compensation will not be otherwise affected: you will receive $3.50 of compensation shortly after completing this survey (Survey #4).

Would you like to be invited to Survey #5 in three weeks?
(Yes, No)

C.1.5 Survey #5

Researchers at Carnegie Mellon University are conducting a research study to understand people's use of web browsing-related tools.

This survey is Survey #5 in the "Research Study for Internet Users" that you previously gave your consent to participate in. It will take up to 3 minutes to complete this survey. If you complete this survey **within 7 days of being invited**, you will be compensated $1.

There are no right or wrong answers to any of our questions, so please answer honestly. Also, **please take the time to read the information in this survey carefully**.

[Installation checkup, the same as in Survey #3]

[Note that for the use and plan checkups, we only ask about activity in the past week, since multiple weeks had passed since Survey #4. See an example below.]
In the past week, **have you used Tor Browser**?
(Yes, No, I don't know)

[Use checkup, the same as in Survey #3]

[Action plan checkup, the same as in Survey #3]

[Coping plan checkups, the same as in Survey #4]

Rate your level of **disagreement or agreement** with the following statement:
"I think I have control over my online privacy."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

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 Tor Browser."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

How **easy or difficult** do you think it would be for you to use Tor Browser?
(Very difficult, Somewhat difficult, Somewhat easy, Very easy)

Rate your level of **disagreement or agreement** with the following statement:
"If I use Tor Browser, I will prevent others from observing my web browsing activity."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

Rate your level of disagreement or agreement with the following statement:
"I intend **to use** Tor Browser in the next week."
(Strongly disagree, Somewhat disagree, Somewhat agree, Strongly agree)

What is your overall opinion of Tor Browser? Has your opinion changed since the beginning of the study? (Please write a few sentences)
___

Would you like to share any other thoughts about this study or about Tor Browser?
___

Bibliography

[1] Henk Aarts, Ap Dijksterhuis, and Cees Midden. To plan or not to plan? Goal achievement or interrupting the performance of mundane behaviors. *European Journal of Social Psychology*, 29(8):971–979, 1999. 7

[2] Hervé Abdi. Holm's Sequential Bonferroni Procedure. *Encyclopedia of research design*, pages 1–8, 2010. 12, 21

[3] 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. 6, 45, 46

[4] Anja Achtziger, Peter M. Gollwitzer, and Paschal Sheeran. Implementation Intentions and Shielding Goal Striving From Unwanted Thoughts and Feelings. *Personality and Social Psychology Bulletin*, 34(3):381–393, March 2008. 7

[5] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, January 2015. 6

[6] 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. 1, 6, 7, 82

[7] Advertising Standards Authority. ASA Ruling on Tefincom SA t/a NordVPN, May 2019. `https://www.asa.org.uk/rulings/tefincom-sa-a19-547668.html`. 81

[8] George A. Akerlof. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–599, August 1970. 1, 6

[9] Elham Al Qahtani, Mohamed Shehab, and Abrar Aljohani. The Effectiveness of Fear Appeals in Increasing Smartphone Locking Behavior among Saudi Arabians. *SOUPS @ USENIX Security Symposium*, 2018. 7, 8

[10] 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. 7, 8, 49

[11] Nora Alkaldi and Karen Renaud. Why Do People Adopt, or Reject, Smartphone Password Managers? In *Proceedings 1st European Workshop on Usable Security*, Darmstadt,

Germany, 2016. Internet Society. 6

[12] Hazim Almuhimedi. Helping Smartphone Users Manage their Privacy through Nudges. Technical Report CMU-ISR-17-111, December 2017. 6

[13] Hazim Almuhimedi, 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. 1, 6

[14] Apple. Apple Pay coming to Target, Taco Bell and more top US retail locations, Jan 2019. `https://www.apple.com/newsroom/2019/01/apple-pay-coming-to-target-taco-bell-and-more-top-us-retail-locations`. 13, 84

[15] Apple. Apple Pay security and privacy overview, Oct 2019. `https://support.apple.com/en-us/HT203027`. 13, 19

[16] arma. Bittorrent over Tor isn't a good idea, April 2010. `https://blog.torproject.org/bittorrent-over-tor-isnt-good-idea`. 59

[17] Richard A Armstrong. When to use the Bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508, April 2014. 21

[18] 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. 1, 30, 35

[19] Matthew Ball, Roderic Broadhurst, Alexander Niven, and Harshit Trivedi. Data Capture and Analysis of Darknet Markets. *SSRN Electronic Journal*, 2019. 49

[20] Sean Barnum, Michael Gegick, and C.C. Michael. Defense in Depth — CISA, September 2005. `https://us-cert.cisa.gov/bsi/articles/knowledge/principles/defense-in-depth`. 45

[21] Adam Beautement, 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. 5

[22] Ariane Bélanger-Gravel, Gaston Godin, and Steve Amireault. A meta-analytic review of the effect of implementation intentions on physical activity. *Health Psychology Review*, 7(1):23–54, March 2013. 1, 7, 62, 63

[23] 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. 45

[24] Veronika Brandstätter, Angelika Lengfelder, and Peter M Gollwitzer. Implementation Intentions and Efficient Action Initiation. *Journal of personality and social psychology*, 81(5):946, 2001. 7

[25] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, January 2006. 16, 34, 64

[26] Brave. What is a Private Window with Tor Connectivity?, February 2021. https://support.brave.com/hc/en-us/articles/360018121491-What-is-a-Private-Window-with-Tor-Connectivity-. 86

[27] Pamela Briggs, Debbie Jeske, and Lynne Coventry. Behavior change interventions for cybersecurity. In *Behavior Change Research and Theory*, pages 115–136. Elsevier, 2017. 8

[28] Jon Brodkin. Fake tech support scam is trouble for legitimate remote help company, November 2013. https://arstechnica.com/information-technology/2013/11/fake-tech-support-scam-is-trouble-for-legitimate-remote-help-company/. 31

[29] Dalvin Brown. Is streaming video from sketchy websites illegal?, December 2019. https://www.usatoday.com/story/tech/2019/12/16/can-get-arrested-streaming-illicit-movies-its-complicated/2662072001/. 33

[30] United States Census Bureau. Current Population Survey (CPS), 2018. https://www.census.gov/cps/data/cpstablecreator.html. 134

[31] Yinzhi Cao, Song Li, and Erik Wijmans. (Cross-)Browser Fingerprinting via OS and Hardware Level Features. *Network and Distributed System Security Symposium*, March 2017. 31

[32] Kevin Cash. Credit card vs. debit card: Which is safer online?, Sep 2015. https://www.nerdwallet.com/blog/credit-cards/credit-card-vs-debit-card-safer-online-purchases/. 13, 18

[33] CDC. Safety of COVID-19 Vaccines, July 2021. https://www.cdc.gov/coronavirus/2019-ncov/vaccines/safety/safety-of-vaccines.html. 83

[34] Cisco. 2013 Cisco Annual Security Report. January 2013. 31

[35] Cloudflare. Understanding Cloudflare Tor support and Onion Routing, February 2021. https://support.cloudflare.com/hc/en-us/articles/203306930-Understanding-Cloudflare-Tor-support-and-Onion-Routing. 79, 82

[36] 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. 44, 49

[37] Cyrus Farivar and Andrew Blankstein. Feds take down the world's 'largest dark web child porn marketplace', October 2019. https://www.nbcnews.com/news/crime-courts/feds-take-down-world-s-largest-dark-web-child-porn-n1066511. 46

[38] Roger Dingledine. [tor-announce] Tor security advisory: Old Tor Browser Bundles vulnerable, August 2013. 32

[39] Disconnect. Best privacy VPN app for iOS and Mac. Powerful protection with one tap., Sep 2020. https://disconnect.me. 137

[40] P. Dolan, M. Hallsworth, D. Halpern, D. King, R. Metcalfe, and I. Vlaev. Influencing behaviour: The mindspace way. *Journal of Economic Psychology*, 33(1):264–277, February 2012. 79, 83

[41] DuckDuckGo. A Study on Private Browsing: Consumer Usage, Knowledge, and Thoughts. Whitepaper, January 2017. 42

[42] EFF. The Playpen Cases: Frequently Asked Questions, August 2016. https://www.eff.org/pages/playpen-cases-frequently-asked-questions. 49

[43] EFF. Choosing the VPN That's Right for You, March 2019. https://ssd.eff.org/en/module/choosing-vpn-thats-right-you. 32

[44] 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. 49

[45] ExpressVPN. How to Combine a VPN and Tor Browser for Online Anonymity, October 2020. https://www.expressvpn.com/. 47

[46] Cori Faklaris, Laura A Dabbish, and Jason I Hong. A Self-Report Measure of End-User Security Attitudes (SA-6). *SOUPS @ USENIX Security Symposium*, 2019. 25

[47] Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2):175–191, May 2007. 12

[48] Jim Finkle. Web tools help protect human rights activists. *Reuters*, August 2009. 59

[49] Alisa Frik, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. A Promise Is A Promise. In *the 2019 CHI Conference*, pages 1–12, New York, New York, USA, 2019. ACM Press. 7

[50] FTC. Tech Support Scams, October 2018. https://www.ftc.gov/tips-advice/business-center/small-businesses/cybersecurity/tech-support-scams. 31

[51] Kevin Gallagher, Sameer Patil, Brendan Dolan-Gavitt, Damon McCoy, and Nasir Memon. Peeling the Onion's User Experience Layer. *ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, October 2018. 6, 53, 57, 60, 71, 84

[52] 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. 6, 45

[53] Samuel Gibbs. Antivirus software is dead, says security expert at Symantec, May 2014. http://www.theguardian.com/technology/2014/may/06/antivirus-software-fails-catch-attacks-security-expert-symantec. 31

[54] Julia Gideon, Lorrie Faith Cranor, Serge Egelman, and Alessandro Acquisti. Power strips,

prophylactics, and privacy, oh my! *Symposium on Usable Privacy and Security*, page 133, 2006. 1

[55] Vindu Goel and Rachel Abrams. Card Data Stolen From 5 Million Saks and Lord & Taylor Customers, Apr 2018. `https://www.nytimes.com/2018/04/01/technology/saks-lord-taylor-credit-cards.html`. 13

[56] Peter M. Gollwitzer. Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 1999. 1, 7, 10, 11, 15, 61, 62, 63

[57] Peter M Gollwitzer and Veronika Brandstätter. Implementation Intentions and Effective Goal Pursuit. *Journal of personality and social psychology*, 73(1):186, 1997. 7

[58] Dan Goodin. Millions exposed to malvertising that hid attack code in banner pixels, December 2016. `https://arstechnica.com/information-technology/2016/12/millions-exposed-to-malvertising-that-hid-attack-code-in-banner-pixels/`. 31

[59] Dan Goodin. Majority of Android VPNs can't be trusted to make users more secure — Ars Technica, January 2017. `https://arstechnica.com/information-technology/2017/01/majority-of-android-vpns-cant-be-trusted-to-make-users-more-secure/`. 56

[60] Dan Goodin. Malvertisers target Mac users with steganographic code stashed in images, January 2019. `https://arstechnica.com/information-technology/2019/01/malvertisers-target-mac-uses-with-stenographic-code-stashed-in-images/`. 31

[61] Google. Google Pay may be preinstalled on your phone, May 2020. `https://support.google.com/pay/answer/7644010`. 17

[62] Google. How payments work, Jan 2020. `https://support.google.com/pay/merchants/answer/6345242?hl=en`. 19

[63] Google. Requests for User Information FAQs - Transparency Report Help Center, 2021. `https://support.google.com/transparencyreport/answer/9713961?hl=en`. 56

[64] Yael Grauer. The impossible task of creating a "Best VPNs" list today — Ars Technica, June 2016. `https://arstechnica.com/information-technology/2016/06/aiming-for-anonymity-ars-assesses-the-state-of-vpns-in-2016/`. 32

[65] Glenn Greenwald and Ewen MacAskill. NSA Prism program taps in to user data of Apple, Google and others. *The Guardian*, June 2013. 32

[66] 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. 6, 42

[67] Cormac Herley. So long, and no thanks for the externalities - the rational rejection of security advice by users. *NSPW*, pages 133–144, 2009. 5, 79, 84

[68] Cormac Herley. More Is Not the Answer. *IEEE Security & Privacy*, 12(1):14–19, 2014. 5

[69] Rae Hodge. Why you should be skeptical about a VPN's no-logs claims, July 2020. https://www.cnet.com/news/why-you-should-be-skeptical-about-a-vpns-no-logs-claims/. 56

[70] Jan Hoffman and Chang W. Lee. 'I Won't Be Used as a Guinea Pig for White People'. *The New York Times*, October 2020. 83

[71] Aaron Holmes. The dark web turns 20 this month — here's how it changed the world - Business Insider, March 2020. https://www.businessinsider.com/dark-web-changed-the-world-black-markets-arab-spring-2020-3. 46

[72] Bryan Horling and Matthew Kulick. Personalized Search for everyone, December 2009. 32

[73] Jun Ho Huh, Saurabh Verma, Swathi Sri V Rayala, Rakesh B Bobba, Konstantin Beznosov, and Hyoungshick Kim. I Don't Use Apple Pay because it's less secure...: perception of security and usability in mobile tap-and-pay. *Proceedings of the Workshop on Usable Security*, 2017. 22, 28

[74] 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. 32

[75] 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. 30

[76] Mike Ives. Celebrities Are Endorsing Covid Vaccines. Does It Help? *The New York Times*, May 2021. 83

[77] Kaspersky. Disk and File Encryption, Sep 2020. https://www.kaspersky.com/enterprise-security/wiki-section/products/encryption. 137

[78] Kaspersky Lab. The main sources of malware infection, November 2018. https://web.archive.org/web/20201127210358/https://support.kaspersky.com/789. 31

[79] Katie Kasunic. How To Use Tor Browser: Everything You MUST Know (2020), August 2020. https://www.vpnmentor.com/blog/tor-browser-work-relate-using-vpn/. 47

[80] Joseph Keller. How to stop Apple Pay from pestering you into signing up, Apr 2018. https://www.imore.com/how-stop-apple-pay-pestering-you. 17

[81] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices - an online study of the nutrition label approach. *CHI*, page 1573, 2010. 5

[82] William R. King and Jun He. A meta-analysis of the technology acceptance model. *Information & Management*, 43(6):740–755, September 2006. 84

[83] Brian Krebs. Tools for a Safer PC — Krebs on Security. 31

[84] Brian Krebs. Post-FCC Privacy Rules, Should You VPN? — Krebs on Security, March 2017. 32

[85] 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. 33

[86] Pierre Laperdrix. Browser Fingerprinting: An Introduction and the Challenges Ahead — Tor Blog, September 2019. `https://blog.torproject.org/browser-fingerprinting-introduction-and-challenges-ahead`. 31

[87] Selena Larson. Infant Social Security numbers are for sale on the dark web, January 2018. `https://money.cnn.com/2018/01/22/technology/infant-data-dark-web-identity-theft/index.html`. 46

[88] 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. 44, 49

[89] Linda Lee, David Fifield, Nathan Malkin, Ganesh Iyer, Serge Egelman, and David Wagner. A Usability Evaluation of Tor Launcher. *Proceedings on Privacy Enhancing Technologies*, 2017(3):257–20, June 2017. 5, 6

[90] Angelika Lengfelder and Peter M Gollwitzer. Reflective and Reflexive Action Control in Patients With Frontal Brain Lesions. *Neuropsychology*, 15(1):80, 2001. 7

[91] Howard Leventhal, Robert Singer, and Susan Jones. Effects of fear and specificity of recommendation upon attitudes and behavior. *American Psychologist*, 2(1):20–29, 1965. 1

[92] Sarah Lichtenstein, Baruch Fischhoff, and Lawrence D Phillips. Calibration of probabilities: The state of the art. In *Decision making and change in human affairs*, pages 275–324. Springer, 1977. 1, 6

[93] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. Follow My Recommendations - A Personalized Privacy Assistant for Mobile App Permissions. *Symposium on Usable Privacy and Security*, 2016. 1, 5

[94] Taylor Lorenz. To Fight Vaccine Lies, Authorities Recruit an 'Influencer Army'. *The New York Times*, August 2021. 83

[95] Aleksandra Luszczynska and Catherine Haynes. Changing Nutrition, Physical Activity and Body Weight among Student Nurses and Midwives: Effects of a Planning Intervention and Self-efficacy Beliefs. *Journal of Health Psychology*, 14(8):1075–1084, November 2009. 7

[96] Ben Luthi. Is Apple Pay Safe?, Apr 2019. `https://creditcards.usnews.com/articles/is-apple-pay-safe`. 13

[97] Monica G. Maceli. Encouraging patron adoption of privacy-protection technologies: Challenges for public libraries. *IFLA Journal*, 44(3):195–202, August 2018. 30

[98] Johanna Catherine Maclean, John Buckell, and Joachim Marti. Information Source and Cigarettes: Experimental Evidence on the Messenger Effect. Technical Report w25632, National Bureau of Economic Research, Cambridge, MA, March 2019. 79, 83

[99] Mary Madden and L. Rainie. *Americans' Attitudes about Privacy, Security and Surveillance*. Pew Research Center, May 2015. 1

[100] 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. 8

[101] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. "What Can't Data Be Used For?": Privacy Expectations about Smart TVs in the U.S. In *Proceedings 3rd European Workshop on Usable Security*, London, England, 2018. Internet Society. 44, 49

[102] Salvatore S. Mangiafico. Kruskal–wallis test, Feb 2020. `https://rcompanion.org/handbook/F_08.html`. 22, 23, 65

[103] Akshaya Mani, T. Wilson-Brown, Rob Jansen, Aaron Johnson, and Micah Sherr. Understanding Tor Usage with Privacy-Preserving Measurement. In *Proceedings of the Internet Measurement Conference*, pages 175–187, Boston MA USA, October 2018. ACM. 71

[104] Jonathan R. Mayer and John C. Mitchell. Third-Party Web Tracking: Policy and Technology. *IEEE Symposium on Security and Privacy*, 2012. 31

[105] 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. 16, 34

[106] 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. *IEEE European Symposium on Security and Privacy (EuroS&P)*, March 2017. 31, 32, 47

[107] Katherine L Milkman, John Beshears, James J Choi, David Laibson, and Brigitte C Madrian. Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences*, 108(26), 2011. 1, 7

[108] Katherine L Milkman, John Beshears, James J Choi, David Laibson, and Brigitte C Madrian. Planning prompts as a means of increasing preventive screening rates. *Preventive Medicine*, 56(1):92–93, January 2013. 1, 7

[109] Sarah Milne, Sheina Orbell, and Paschal Sheeran. Combining motivational and volitional interventions to promote exercise participation: Protection motivation theory and implementation intentions. *British Journal of Health Psychology*, 7(2):163–184, May 2002. 1, 7, 8, 9, 53, 54

[110] Sarah Milne, Paschal Sheeran, and Sheina Orbell. Prediction and Intervention in Health-Related Behavior: A Meta-Analytic Review of Protection Motivation Theory. *Journal of Applied Social Psychology*, 30(1):106–143, January 2000. 8, 10, 11, 13, 56, 57, 60

[111] Chris Morran. House Votes To Allow Internet Service Providers To Sell, Share Your Personal Information, March 2017. `https://www.consumerreports.`

org/consumerist/house-votes-to-allow-internet-service-providers-to-sell-share-your-personal-information/. 56

[112] David W Nickerson and Todd Rogers. Do You Have a Voting Plan? *Psychological Science*, 21(2):194–199, January 2010. 1, 7

[113] 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. 6, 50, 60

[114] NordVPN. Block ads and malicious websites with CyberSec, Sep 2020. https://nordvpn.com/features/cybersec/. 33, 86, 127, 137

[115] Norton. Norton Privacy Manager — How it works, Mar 2019. https://www.youtube.com/watch?v=iKsHl-uzVrU. 33, 86, 137

[116] Norton. Browse the Internet securely with Norton Safe Web, Oct 2020. https://support.norton.com/sp/en/us/home/current/solutions/v19116982. 137

[117] Norton. Norton Privacy Manager, Sep 2020. https://us.norton.com/norton-privacy-manager. 33, 86

[118] Norton. Norton Secure VPN, Sep 2020. https://us.norton.com/products/norton-secure-vpn. 33, 127, 137

[119] Gabriele Oettingen and Gaby Ho. Effective self-regulation of goal attainment. *International journal of educational research*, pages 705–732, 2000. 61, 62, 63

[120] Office of the Privacy Commissioner of Canada. What an IP Address Can Reveal About You, May 2013. https://www.priv.gc.ca/en/opc-actions-and-decisions/research/explore-privacy-research/2013/ip_201305. 56

[121] Kenny Olmstead and Aaron Smith. *Americans and Cybersecurity: Pew Research Center*. January 2017. 1

[122] Sheina Orbell, Sarah Hodgkins, and Paschal Sheeran. Implementation intentions and the theory of planned behavior. *Personality and Social Psychology Review*, 23(9):945–954, 1997. 1, 7

[123] Sheina Orbell and Paschal Sheeran. Motivational and Volitional Processes in Action Initiation: A Field Study of the Role of Implementation Intentions1. *Journal of Applied Social Psychology*, 30(4):780–797, 2000. 1, 7

[124] Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, and Alain Forget. Let's Go in for a Closer Look. In *the 2017 ACM SIGSAC Conference*, pages 295–310, New York, New York, USA, 2017. ACM Press. 20

[125] Sarah Pearman, Shikun Aerin Zhang, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. Why people (don't) use password managers effectively. *USENIX Symposium on Usable Privacy and Security*, August 2019. 6

[126] 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. 16, 33, 55

[127] Nicole Perlroth. Tor Project, a Digital Privacy Group, Reboots With New Board (Published 2016). *The New York Times*, July 2016. 47

[128] Pew. Are Americans Embracing Mobile Payments? pages 1–22, October 2019. 28

[129] Nathaniel Popper. The Tax Sleuth Who Took Down a Drug Lord. *The New York Times*, December 2015. 49, 59

[130] Prolific Support Team. Ineligibility issues, April 2021. https://www.reddit.com/r/ProlificAc/comments/ms114h/ineligibility_issues/. 76

[131] Prolific Support Team. Ineligibility issues: Fix update, April 2021. https://www.reddit.com/r/ProlificAc/comments/n0f1me/ineligibility_issues_fix_update/. 76

[132] Prolific Team. Representative Samples on Prolific, March 2019. https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-Samples-on-Prolific. 29, 134

[133] 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-. 30

[134] ProtonVPN. Why use Tor over VPN, July 2018. 47

[135] Steve Ragan. What you need to know about the Home Depot data breach, Sep 2014. https://www.csoonline.com/article/2604320/what-you-need-to-know-about-the-home-depot-data-breach.html. 13

[136] L. Rainie, S. Kiesler, R. Kang, and Mary Madden. *Anonymity, Privacy, and Security Online. Pew Research Internet Project.* 2013. 30

[137] Lee Rainie. Americans' complicated feelings about social media in an era of privacy concerns. *Pew Research Center*, March 2018. 1

[138] 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. 29, 33, 76

[139] 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. 30

[140] Ronald W Rogers. A Protection Motivation Theory of Fear Appeals and Attitude Change. *The Journal of Psychology*, 91(1):93–114, 1975. 8

[141] Ronald W Rogers and Steven Prentice-Dunn. Protection motivation theory. 1997. 8

[142] Samsung. How secure is Samsung Pay?, Jan 2020. https://www.samsung.com/us/support/answer/ANS00043932/. 19

[143] Samsung. Set up Samsung Pay on your phone, May 2020. https://www.samsung.

com/us/support/answer/ANS00045081/. 17

[144] Choe Sang-Hun. In Reporting on North Korea, Tech Helps Break Through Secrecy. *The New York Times*, July 2017. 59

[145] Paschal Sheeran, Sarah Milne, Thomas L. Webb, and Peter M. Gollwitzer. *Implementation Intentions and Health Behaviour*. 2005. 7, 61, 62, 63

[146] Paschal Sheeran, Thomas L. Webb, and Peter M. Gollwitzer. The Interplay Between Goal Intentions and Implementation Intentions. *Personality and Social Psychology Bulletin*, 31(1):87–98, January 2005. 7, 9, 53, 54

[147] Ruth Shillair, Shelia R. Cotten, Hsin-Yi Sandy Tsai, Saleem Alhabash, Robert LaRose, and Nora J. Rifon. Online safety begins with you and me: Convincing Internet users to protect themselves. *Computers in Human Behavior*, 48(C):199–207, July 2015. 8

[148] Herbert A Simon. *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. J. Wiley and Sons, 1957. 1, 6

[149] Mikko Siponen, M. Adam Mahmood, and Seppo Pahnila. Employees' adherence to information security policies: An exploratory field study. *Information &amp; Management*, 51(2):217–224, March 2014. 8

[150] 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. 42

[151] Falko F. Sniehotta, Urte Scholz, and Ralf Schwarzer. Action plans and coping plans for physical exercise: A longitudinal intervention study in cardiac rehabilitation. *British Journal of Health Psychology*, 11(1):23–37, 2006. 7

[152] Falko F. Sniehotta, Ralf Schwarzer, Urte Scholz, and Benjamin Schüz. Action planning and coping planning for long-term lifestyle change: Theory and assessment. *European Journal of Social Psychology*, 35(4):565–576, 2005. 7, 62, 63

[153] Teodor Sommestad, Henrik Karlzén, and Jonas Hallberg. A Meta-Analysis of Studies on Protection Motivation Theory and Information Security Behavior. *Dewald Roode Information Security Research Workshop*, pages 1–32, May 2014. 8

[154] Ben Stegner. How to Avoid Fake Ads Disguised as Fake Download Links, July 2019. https://www.makeuseof.com/tag/spot-avoid-ads-disguised-download-buttons/. 31

[155] Birgit Steller. *Vorsätze und die Wahrnehmung günstiger Gelegenheiten [Implementation intentions and the detection of good opportunities to act]*. tuduv-Verlag-Ges., 1992. 10, 15

[156] steph. How Has Tor Helped You? Send Us Your Story., February 2019. https://blog.torproject.org/how-has-tor-helped-you-send-us-your-story. 59

[157] Peter Story. Switch from POST to GET for DDG Searches, January 2021. https://gitlab.torproject.org/tpo/applications/tor-browser/-/issues/40287. 82

[158] Peter Story. Tor browser study: Preregistration, Mar 2021. `https://osf.io/bc42h`. 63

[159] Peter Story and Daniel Smullen. Apple services study: Preregistration, Dec 2019. `https://osf.io/k3nrd`. 12

[160] Peter Story, Daniel Smullen, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. From Intent to Action - Nudging Users Towards Secure Mobile Payments. *SOUPS @ USENIX Security Symposium*, 2020. 9, 49

[161] Peter Story, Daniel Smullen, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. Awareness, Adoption, and Misconceptions of Web Privacy Tools. *Proceedings on Privacy Enhancing Technologies*, 2021(3):308–333, July 2021. 29, 79

[162] O. Sukwong, H. S. Kim, and J. C. Hoe. Commercial Antivirus Software Effectiveness - An Empirical Study. *Computer*, 44(3):63–70, 2011. 31

[163] Gail M. Sullivan and Richard Feinn. Using Effect Size—or Why the PValue Is Not Enough. *Journal of Graduate Medical Education*, 4(3):279–282, September 2012. 22, 64, 77

[164] 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. 49

[165] Karyn A. Temple. U.S. Copyright Office Responses To Specific Questions, July 2019. 33

[166] Richard H Thaler and Cass R Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. J. Wiley and Sons, 2008. 1, 6

[167] The Tor Project. Doc/TorPlusVPN – Tor Bug Tracker & Wiki, October 2019. `https://trac.torproject.org/projects/tor/wiki/doc/TorPlusVPN`. 45, 47

[168] The Tor Project. Overview, April 2020. `https://2019.www.torproject.org/about/overview.html.en`. 58

[169] The Tor Project. Download, 2021. `https://www.torproject.org/download/`. 56, 157

[170] The Tor Project. GitLab, 2021. `https://gitlab.torproject.org/tpo/team`. 79

[171] The Tor Project. Relay Operations, 2021. `https://community.torproject.org/relay/`. 79

[172] The Tor Project. Should I install a new add-on or extension in Tor Browser, like AdBlock Plus or uBlock Origin?, 2021. `https://support.torproject.org/tbb/tbb-14/`. 79

[173] Bart Thoolen, Denise Ridder, Jozien Bensing, Kees Gorter, and Guy Rutten. Beyond good intentions: The role of proactive coping in achieving sustained behavioral change in the context of diabetes management. *Psychology & health*, 24:237–54, March 2009. 7

[174] Maciej Tomczak and Ewa Tomczak. The need to report effect size estimates revisited. An overview of some recommended measures of effect size. *Trends in Sport Sciences*, pages 19–25, July 2014. 22, 23, 65

[175] Matt Traudt. VPN + Tor: Not Necessarily a Net Gain - Matt Traudt, November 2016. `https://matt.traudt.xyz/posts/vpn-tor-not-mRikAa4h.html`. 45, 47

[176] Elham Vaziripour, Justin Wu, Mark O'Neill, Daniel Metro, Josh Cockrell, Timothy Moffett, Jordan Whitehead, Nick Bonner, Kent Seamons, and Daniel Zappala. Action Needed! Helping Users Find and Complete the Authentication Ceremony in Signal. *SOUPS @ USENIX Security Symposium*, pages 47–62, August 2018. 6

[177] Gregory Wallace. Target credit card hack: What you need to know, Dec 2013. `https://money.cnn.com/2013/12/22/news/companies/target-credit-card-hack/index.html`. 13

[178] Rick Wash and Emilee J Rader. Influencing mental models of security - a research agenda. *NSPW*, page 57, 2011. 19

[179] Thomas L. Webb and Paschal Sheeran. Identifying good opportunities to act: Implementation intentions and cue discrimination. *European Journal of Social Psychology*, 34(4):407–419, 2004. 7

[180] Zachary Weinberg, Shinyoung Cho, Nicolas Christin, Vyas Sekar, and Phillipa Gill. How to Catch when Proxies Lie. *IMC '18: Internet Measurement Conference*, October 2018. 32

[181] WhatIsMyIPAddress.com. How does geolocation work?, November 2020. `https://whatismyipaddress.com/geolocation`. 32

[182] Alma Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability case study of PGP 5. *USENIX Security Symposium*, August 1999. 5

[183] Wikipedia. Google Personalized Search. *Wikipedia*, April 2020. 32

[184] Wikipedia. Internet geolocation. *Wikipedia*, September 2020. 32

[185] Wikipedia. PRISM (surveillance program). *Wikipedia*, September 2020. 32

[186] Wikipedia. Web tracking. *Wikipedia*, May 2021. 56

[187] 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. 44, 49

[188] Kim Witte and Mike Allen. A meta-analysis of fear appeals: Implications for effective public health campaigns. *Personality and Social Psychology Review*, 27(5):591–615, 2000. 8

[189] Justin Wu, Cyrus Gattrell, Devon Howard, Jake Tyler, Elham Vaziripour, Daniel Zappala, and Kent Seamons. "Something isn't secure, but I'm not sure how that translates into a problem": Promoting autonomy by designing for understanding in Signal. *SOUPS @*

USENIX Security Symposium, 2019. 6

[190] Yuxi Wu, Panya Gupta, Miranda Wei, Yasemin Acar, Sascha Fahl, and Blase Ur. Your Secrets Are Safe. *Companion of the The Web Conference 2018*, 2018. 6

[191] 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. 1, 5, 30, 42, 60