Thesis Proposal

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

Peter Story

December 22, 2020

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

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

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. 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 and overcome obstacles. In my first chapter, I describe my study of nudges designed to encourage adoption of secure mobile payment systems. I tested both implementation intention and protection motivation theory-based nudges, and 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 browsing-related privacy tools, which I describe in my second chapter. Based on these findings, I propose a study of implementation intention nudges designed to help people adopt Tor Browser, which will serve as my final chapter. In this study, I will test nudges based on PMT, action planning implementation intentions, and coping planning implementation intentions. The design of these nudges will incorporate the recommendations from my second chapter study.

Contents

Chapter 1

Introduction

Most Americans express a desire for digital security and privacy [13, 32, 40, 45]. Americans feel a lack of control over their data, and express interest in consumer tools to protect their personal information [13]. However, the limited adoption of security and privacy tools appears inconsistent with these preferences [40, 62]. 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) [5], bounded rationality (e.g., unawareness of effective protections from threats) [51], and various cognitive and behavioral biases (e.g., optimistically underestimating the chance of being affected by a threat [28]). Thankfully, they also suggest a potential solution to these challenges, in the form of nudging interventions [57]. Nudges can take many forms [11, 20, 29], but what nudges have in common is that they should help people make decisions that align with their stated preferences [3].

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 [14, 21]. The effectiveness of implementation intentions has been demonstrated in many other contexts [27, 34, 35, 36, 39, 41, 42], 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 implementation intentions (II). The experiment included three treatment conditions: a PMT-only condition, a PMT+II condition, and a control condition. My results showed that participants in the PMT-only and PMT+II treatment groups were 2.4x ($p = 0.020$) 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 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 II 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. My results show 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. I will incorporate these design recommendations into the nudges in my Chapter 5 study. Furthermore, our choice of which tool to nudge people towards was informed by my descriptive data on current levels of tool adoption and on relative interest in preventing different privacy threats.

Based on these findings, I propose a study of implementation intention nudges designed to help people adopt Tor Browser, a privacy-enhancing technology. In this study, I will test nudges based on PMT, action planning implementation intentions, and coping planning implementation intentions. The design of our nudges will incorporate the recommendations from our earlier browsing survey. I describe the design of this study in Chapter 5.

In summary, I propose that:

> *Nudging interventions can correct people's misconceptions and increase real-world adoption of security and privacy tools. By quantifying and comparing the effect size of nudges based on implementation intentions and protection motivation theory, we can inform their use in the field of computer security and privacy.*

Chapter 2

Related Work

Researchers in the field of usable security and privacy study ways to protect people from digital threats. In recent years, security and privacy researchers have begun testing various *nudging* techniques to guide people towards more secure and privacy protective behavior (Section 2.1). 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 2.2) and the related *protection motivation theory* (Section 2.3).

2.1 Nudges for Security and Privacy

A large body of research examines ways to help people protect their security and privacy. For example, researchers have studied how to improve privacy notices [24], how to guide people towards choices that fit their preferences [30], and how a lack of usability can inhibit adoption of security tools [25, 60]. This varied research is unified by the acknowledgment that people have limited cognitive resources and suffer from behavioral biases.

Inspired by work in psychology and behavioral economics [57], researchers are increasingly studying how *nudges* can improve design for security and privacy [4]. Nudges are design elements that help people overcome their cognitive and behavioral biases in order to make decisions which align with their stated preferences. For example, Almuhimedi et al. used nudges to mitigate the information asymmetry between users and the behaviors of apps on their devices [10, 12]. Their nudges were successful at encouraging users to reassess and restrict permissions settings. Frik et at. tested using nudges to overcome present bias [18]. 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 [7, 8]. In both studies, the videos successfully motivated many participants to enable secure lock screens. However, nudges are not always effective [4], which shows the value of empirical research like ours.

2.2 Implementation Intentions

An implementation intention is a concrete plan to achieve some goal, which is triggered by situational factors like time or location [22]. 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" [48]. 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%) [41], in encouraging people to exercise for 20 minutes a week (from 39% to 91%) [36], and in getting people to take actions in many other contexts [34, 35, 39, 42].

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 [48]. The literature suggests that implementation intentions work by helping people recognize opportunities for action [1, 59] and by helping people perform the action automatically when the opportunity does arise [15, 26]. Also, implementation intentions are most likely to be effective when a person has a strong commitment to both their plan [48] and to the goal which motivates the plan [49]. Implementation intentions offer the greatest benefit when the goal to be achieved is challenging [23].

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* [14, 54]. For example, a person whose goal is to exercise 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 can be formed by themselves [2, 58], or in addition to action plans [31, 53].

2.3 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 [37], we drew on Protection Motivation Theory (PMT) [33, 46, 47] to motivate participants in our studies. PMT has been widely applied in the medical field [38, 61] and in computer security [6, 9, 17, 50, 52, 55]. 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*) [38, 61]. An intervention drawing on PMT can help people form accurate perceptions of these factors, and may thereby help motivate people to act.

Chapter 3

Completed: Nudges to Increase Adoption of Secure Mobile Payments

3.1 Research Goal

Our primary goal was to determine whether nudges based on implementation intentions (II) 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.

3.2 Methodology

To measure the effect of implementation intentions 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.

Qualitative Interviews

The first part of our study focused on gathering qualitative information on people's thoughts about the threat of card fraud, the use of Pay[1] to protect against card fraud, and people's experiences forming implementation intention plans to use Pay. This portion of our study included three surveys and two semi-structured interviews. We recruited participants from Craigslist and Carnegie Mellon University's participant pool.

Survey #1 served as a screening survey. We employed purposive sampling to select diverse participants who were not already using Pay, but whose phones were compatible with Pay and were likely to have opportunities to use Pay. We invited 20 qualifying participants to participate in a semi-structured interview (Interview #1). The interview started with a discussion of prior experiences with card fraud, card information theft, and prior experiences with Pay. Next, the interviewer administered our PMT and implementation intentions nudges.

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

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 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 implementation intention plans and whether they found the plans to be helpful.

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 used thematic coding to analyze transcripts of our interviews and our survey's open-text responses. Our findings informed the design of our controlled experiment, allowing us to refine our interventions, correct common misconceptions, and to understand some of the limitations of our approach.

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 study consisted of three surveys hosted on Qualtrics using recruitment from Prolific. 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, so we used Survey #1 to screen for those participants.

Shortly after completing Survey #1, participants were invited to Survey #2, in which they were randomly assigned to one of 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 and the mitigation of using Apple Pay in order to motivate them to use Apple Pay. In our PMT+II group, we presented participants with the motivational intervention of the PMT group in addition to an opportunity to form an implementation intention. This opportunity took the form of a template we designed to help participants plan where they could use Apple Pay. We did not test an implementation intention intervention without a PMT intervention because the literature suggests that implementation intentions are only effective when participants are motivated [22].

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.

We analyzed our results by conducting three chi-square tests of independence to compare the use of Apple Pay between our three treatment groups. We also performed exploratory statistical analyses to measure the effects of our interventions on participants' attitudes.

3.3 Results

The results of our qualitative interviews helped us improve the design of our nudges. We uncovered several misconceptions, such as participants wondering whether Pay would interfere with card rewards or whether it cost something to use. We created a list of “Frequently Asked Questions” to address these and other misconceptions. Despite our written instructions, we encountered some participants who still had questions about how to use Pay. To make the instructions clearer, we added a video demonstrating how to Pay. After being exposed to our nudges, 35% of our participants used Pay at least once during the remainder of our study. These results encouraged us to run our controlled experiment.

In our controlled experiment, we compared participants’ use of Apple Pay between our PMT-only, PMT+II, and control groups. We found that our nudges did make participants more likely to use Apple Pay. We found a large difference between our PMT+II treatment and the control group (3.9x more likely to use Apple Pay, $p < 0.001$), and a medium difference between our PMT-only treatment and the control group (2.4x more likely, $p = 0.020$). Although the difference between the PMT+II and PMT-only treatments was not statistically significant (1.7x more likely, $p = 0.085$), the difference we observed suggests that adding an implementation intention to our PMT-only treatment increased its efficacy by a small amount.

We also found evidence that our nudges affected participants’ attitudes about Apple Pay. Most notably, our PMT nudge made participants more likely to agree that Apple Pay would protect them from card fraud. In the control group, only 37% of participants agreed that Apple Pay would protect them, whereas in both treatment groups over 84% agreed. We also found that our those in our treatment groups expressed greater intentions to use Apple Pay in the week ahead, and that those in our implementation intention treatment expressed even greater intentions than those in the PMT-only group. However, we observed that in many cases, intention to use Apple Pay did not translate into actual use of Apple Pay. Whereas more than half of those who expressed a strong intention to use Apple Pay did so, those who expressed weaker intentions were much less likely to use Apple Pay. This shows the importance of asking participants about both their intentions and their actual behavior, as we did in our study.

This study was published at SOUPS 2020 [56].

Chapter 4

Completed: Measuring Adoption of and Beliefs About Web Browsing Tools

4.1 Research Goal

Our research goal was to understand people's awareness, adoption, and understanding of web browsing-related privacy tools. We choose this goal to inform the design of nudges based on coping planning, which we propose to study in Chapter 5.

4.2 Methodology

We gathered our 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 [44]. We collected data in August 2020 from 500 participants.

Our survey contained four main parts. First, we asked participants general questions about their perceptions of online privacy. Second, we asked participants specific questions about the tools we studied, such as whether they had heard of or used each tool. We asked questions about private browsing, VPNs, Tor Browser, ad blockers, and antivirus software. Third, we asked participants how effective they thought each tool would be at preventing different scenarios from happening. Each scenario was introduced as a question of the form: "When you browse the web, how effective are the tools below at preventing *advertisers from seeing the websites you visit*?" We asked each participant about six scenarios, which were randomly assigned from twelve total scenarios. Our scenarios included seeking protection from hackers, advertisers, the government, and other entities. Finally, we asked participants demographic questions, such as about their education and device usage patterns.

To better understand participants' misconceptions about the tools, we asked participants to explain their answer for one randomly selected tool for each scenario. We used thematic coding to analyze these free-text responses. Our other findings are based on statistical analyses and descriptive statistics.

4.3 Results

Some of the tools we studied are already widely used, such as antivirus software and ad blockers, which about half of participants reported using the same day they took our survey. However, other tools are used less frequently, particularly Tor Browser, which less than 1% of participants reported using the day they took our survey. Overall, we observed that 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. However, as revealed by our subsequent analyses, many participants have misconceptions about the protections offered by the tools they are using.

We observed large numbers of unsure and incorrect responses across tools and scenarios. Participants answered fewer than half of our assessment questions correctly for all but one scenario. For all tools, participants answered fewer than 60% of the associated questions correctly.

We were interested in whether participants' level of experience with each tool was associated with their ability to answer questions about each tool correctly. Our statistical tests show that greater levels of experience are typically associated with answering more questions correctly. 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. Our results suggest that participants who have more experience with tools are more willing to definitively answer questions about the tools. However, greater experience is not necessarily associated with a more accurate understanding of the tools' protections.

Our analysis of participants' free-text responses revealed a number of themes. Here, we list just four:

- Across all tools and scenarios, participants cited true aspects of tool functionality when explaining their incorrect answers. Participants' responses show that they 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.
- Participants frequently expressed resignation, writing that nothing could be done to protect against an entity, or that the entity's resources were overwhelmingly powerful. This theme was present even in scenarios where effective tools are available.
- Participants also expressed overconfidence in tools' protections, writing that tools provided total protection or anonymity even though they do not.
- 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.

We submitted this study for publication at PETS 2021.

Chapter 5

Proposed: Nudges to Increase Adoption of Tor Browser

5.1 Research Goal

My primary goal is to determine whether nudges based on action planning implementation intentions (APII), coping planning implementation intentions (CPII), and protection motivation theory (PMT) can increase real-world adoption of anonymity systems like Tor Browser. In particular, I am interested in comparing the relative effectiveness of these nudges.

This research will also allow me to study three additional goals. First, I will measure the effect of these nudging interventions on participants' attitudes about Tor Browser. Second, I will compare the mechanism of this study's interventions to those from my earlier study of secure mobile payments (Chapter 3). Third, I will identify and quantify obstacles to widespread adoption of Tor Browser.

5.2 Proposed Methodology

My proposed study design involves encouraging the adoption of Tor Browser. I chose to focus on Tor Browser because both action planning and coping planning nudges are applicable to successfully adopting Tor Browser. In order to use Tor Browser to protect one's privacy, Tor Browser should only be used for particular purposes. Thus, users must remember to activate Tor Browser before they perform particular privacy-sensitive operations (e.g., researching a medical condition). An action plan may help people remember to use Tor Browser. For example, "If I want to research my mental illness, then I will launch and use Tor Browser." But even if someone remembers to open Tor Browser, the literature identifies a number of obstacles they may encounter, including high latency and websites blocking Tor traffic [19]. However, these obstacles can often be overcome if a user is persistent, and coping planning has the potential to increase resilience in the face of obstacles. For example: "If I encounter excessive slowness when using Tor Browser, then I will use the broom icon to restart Tor Browser."

To study my APII, CPII, and PMT nudges I propose a longitudinal, randomized experiment. The experiment will contain four treatment groups: Control, PMT, PMT+APII, and

PMT+APII+CPII. Treatments will be administered twice, with one week between the first and second administrations; this is because the literature suggests that effective coping planning relies on having prior experience with challenges [54]. My outcome variable will be whether participants used Tor Browser in the week after the intervention. I will also measure perceptions of Tor Browser and reasons for abandoning or not adopting Tor Browser.

The literature suggests that implementation intentions are most effective when people are highly motivated [49]. To increase the likelihood that my APII and CPII nudges will have a measurable effect, I propose recruiting participants who I expect to become highly motivated to use Tor Browser after reading the information in my PMT-based nudge. My Chapter 4 study gives insight into how to accomplish this. First, my study showed that a sizable percent of people express a high level of interest in achieving privacy protections. For example, 43% of participants indicated that they were "Very interested" in "preventing advertisers from seeing the websites [they visited]." Second, my study revealed that many people have already adopted privacy-enhancing tools like private browsing, but that they overestimate the protections offered by those tools. For example, over 25% of participants thought that private browsing was "Very effective" at "preventing advertisers from seeing the websites [they visited]," which is simply incorrect. Many participants' responses suggest they are interested in and willing to take action to protect their privacy, but that their behavior is misdirected. Therefore, I suggest recruiting participants who express a high level of interest in protection from scenarios which Tor Browser can protect against, and who are actively using privacy-enhancing technologies which offer inadequate protection against those scenarios. I expect that revealing the limitations of the tools they have already adopted and describing the benefits of using Tor Browser will highly motivate them to try using Tor Browser to achieve their privacy goals.

5.3 Expected Results

Based on the literature and the results of my Chapter 3 study, I anticipate that I will observe the following differences in Tor Browser usage between the four treatment groups:

$$\text{Control} < \text{PMT} \leq \text{PMT+APII} \leq \text{PMT+APII+CPII}$$

That is, usage will be lowest in the Control group, and highest in the PMT+APII+CPII group. My ability to detect statistically significant differences will depend on effect sizes and sample sizes. I expect that the PMT nudge alone will have a medium effect, and that the APII and CPII nudges will add small effects.

Next, I anticipate that my PMT nudge will work by increasing perceptions of response efficacy (e.g., belief that Tor Browser is effective at preventing observation), threat susceptibility (e.g., realizing that existing precautions, like using Private Browsing, are ineffective), and self-efficacy (e.g., believing they are personally able to use Tor Browser). This differs from my payment system nudge, which did not affect perceptions of threat susceptibility and self-efficacy.

Finally, I expect that inability to load pages in Tor Browser will be the most cited reason for abandoning it, since I expect that the highly motivated users I recruit will be willing to tolerate increased latency.

5.4 Contingency Plans

As with any empirical study, there is a possibility of negative results. The only problem with negative results in this context is that they might seem obvious in retrospect; since adoption of Tor Browser is low and associated usability issues are well-known, it may seem unsurprising if participants in my study do not adopt it. For that reason, my research contribution will be much more interesting if my interventions do increase adoption of Tor Browser. Below, I have outlined steps I will take to increase my chances of finding positive results.

To avoid exhausting my budget on a study with negative results, I will begin by running a pilot study. In this pilot study, I will test the PMT+APII+CPII intervention with crowd workers, since the literature suggests this intervention will be most effective. If a high percent of pilot study participants adopt Tor Browser (e.g., 30%), this will be an indication that our intervention affected behavior. But if a low percent of pilot study participants adopt Tor Browser (e.g., less than 10%), this will be a sign that the effect of our intervention is small, and would be hard to detect, even with a much larger group of participants. In this situation, I would first attempt to increase the effect size of my intervention by requiring that participants install Tor Browser to qualify for the intervention. This would eliminate participants unwilling to install Tor Browser and would make it easier for qualifying participants to use Tor Browser, since they would already have it installed. This would negatively impact ecological validity, since in the real world people are not paid to install Tor Browser, but would still allow me to measure the effect of implementation intentions on helping participants remember to use Tor Browser. My findings would be applicable to a real-world situation where Tor Browser comes pre-installed in users' browsers (e.g., as in Brave browser [16]). After making this change to recruitment, I would pilot the study with another set of crowd workers. If the percent of participants who adopt Tor Browser is still very low, my final contingency plan would be targeting my interventions to a highly motivated segment of the population, such as journalists. In this case, recruitment would be more challenging, so I likely wouldn't have enough data to perform my original experiment. I would still gather data on the challenges encountered, quantify the reasons for lack of adoption, and gather anecdotal data supporting the importance of the information in my action and coping plans. In addition, even with a small group of highly motivated participants, I could still perform a statistical test to compare these participants to crowd workers. I would likely show a statistically significant difference in the level of adoption between the two groups, and I would identify components of motivation which differ between the two groups.

5.5 Impact

My expected findings have implications for Tor Browser, for other privacy and security tools, and for the broader nudging research community.

With respect to Tor Browser, I am very likely to identify populations which are ready to adopt Tor Browser. If I find that my interventions successfully encourage many crowd workers to adopt Tor Browser, I will have strong evidence that Tor Browser is suitable for more widespread adoption. This may encourage major browser vendors to integrate Tor into their browsers [16] and to promote this feature. Increasing the public's adoption of anonymity tools could have positive effects on society, such as mitigating the chilling effect observed after the Snowden revelations [43]. Alternatively, if I find that my interventions are only effective among highly motivate participants, this will suggest that Tor Browser is not suitable for widespread adoption, at least in its current form. This could encourage Tor developers to address the barriers to adoption which we identify. Also, in this case I am likely to find that certain highly motivated people, such as journalists, are able to successfully adopt Tor Browser (see Section 5.4). This may encourage journalism-related organizations to promote the adoption of Tor Browser, thereby protecting journalists and their sources.

With regard to other privacy and security tools, comparing my findings from this study to those of my mobile payments study will help generalize the findings from both studies. For example, in my mobile payments study my PMT nudge greatly increased participants' belief in the response efficacy of Apple Pay (i.e., belief that it would protect them from card fraud). In my study of Tor Browser, I expect to see changes in other components of motivation, such as perception of threat susceptibility (e.g., due to participants realizing that private browsing does not provide the protections they assumed). This will emphasize the importance of addressing all components of protection motivation theory, at least in the design phase of an intervention. In addition, another datapoint about the effect size of implementation intentions will help calibrate expectations about implementation intentions in the domain of security and privacy. Bélanger-Gravel et al.'s meta-analysis of studies using implementation intentions to promote physical activity estimates their effect size to be small to medium in that domain [14], but my study of mobile payments provides the only datapoint in the domain of security and privacy. Finally, if I find a sufficiently large effect from my interventions, this will present an opportunity to test different variations of my nudges. For example, I could test a variant of my PMT nudge which omits information which seem less likely to influence behavior (e.g., information pertaining to threat severity). If I find that this information can be omitted without decreasing the effectiveness of the nudge, that would allow this nudge to be shortened, facilitating its deployment in the real world. In summary, my findings will contribute to the deployment of nudges to encourage the adoption of other privacy and security enhancing tools.

With regard to the broader research community, this study will contribute to a greater understanding of interventions based on coping planning. The literature suggests that [14] coping planning is effective, but relatively few studies have directly compared coping planning to other interventions, such as those based on action planning [53]. This study will help fill this gap in the literature, since we will compare PMT+APII to PMT+APII+CPII.

Chapter 6

Timeline

Timeline without contingency plans

Task	Time Needed	Planned Completion Date
PETS Rebuttal	1 week	January 7–11
Design and Implement Study	1 month	February 5
Pilot Test and Iterate on Study Design	1 month	February 26
Collect Data from Prolific	1 month	March 26
Analyze Data, Write Paper, Write Thesis	1 month	April 30
Defend	1 day	May 7
Graduate	1 day	May 20
Finalize Thesis	1 month	May 28

Timeline with contingency plans

Task	Time Needed	Planned Completion Date
PETS Rebuttal	1 week	January 7–11
Design and Implement Study	1 month	February 5
Pilot Test and Iterate on Study Design	1 month	February 26
Collect Data from Prolific	1 month	March 26
Collect Data from Highly Motivated Population	**2 months**	**May 31**
Analyze Data, Write Paper, Write Thesis	1 month	June 30
Defend	1 day	July 2
Finalize Thesis	1 month	July 30
Graduate	1 day	May 2022

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. 2.2

[2] 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. 2.2

[3] 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

[4] 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. 2.1

[5] 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] 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. 2.3

[7] 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. 2.1

[8] 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. 2.1

[9] 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. 2.3

[10] Hazim Almuhimedi. Helping Smartphone Users Manage their Privacy through Nudges. Technical Report CMU-ISR-17-111, December 2017. 2.1

[11] 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

[12] 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. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015. 2.1

[13] 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

[14] 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, 2.2, 5.5

[15] Veronika Brandstätter, Angelika Lengfelder, and Peter M Gollwitzer. Implementation Intentions and Efficient Action Initiation. *Journal of personality and social psychology*, 81(5):946, 2001. 2.2

[16] Brave. What is a Private Window with Tor Connectivity? https://support.brave.com/hc/en-us/articles/360018121491-What-is-a-Private-Window-with-Tor-Connectivity-, December 2020. 5.4, 5.5

[17] Pamela Briggs, Debbie Jeske, and Lynne Coventry. Behavior change interventions for cybersecurity. In *Behavior Change Research and Theory*, pages 115–136. Elsevier, 2017. 2.3

[18] 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. 2.1

[19] Kevin Gallagher, Sameer Patil, Brendan Dolan-Gavitt, Damon McCoy, and Nasir Memon. Peeling the Onion's User Experience Layer. *the 2018 ACM SIGSAC Conference*, October 2018. 5.2

[20] 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

[21] Peter M. Gollwitzer. Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 1999. 1

[22] Peter M Gollwitzer. Implementation intentions: strong effects of simple plans. *American Psychologist*, 54(7), 1999. 2.2, 3.2

[23] Peter M Gollwitzer and Veronika Brandstätter. Implementation Intentions and Effective Goal Pursuit. *Journal of personality and social psychology*, 73(1):186, 1997. 2.2

[24] 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. 2.1

[25] 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. 2.1

[26] Angelika Lengfelder and Peter M Gollwitzer. Reflective and Reflexive Action Control in Patients With Frontal Brain Lesions. *Neuropsychology*, 15(1):80, 2001. 2.2

[27] 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

[28] 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

[29] 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

[30] 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. 2.1

[31] 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. 2.2

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

[33] 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. 2.3

[34] 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, 2.2

[35] 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, 2.2

[36] 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, 2.2

[37] 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. 2.3

[38] 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. 2.3

[39] David W Nickerson and Todd Rogers. Do You Have a Voting Plan? *Psychological Science*, 21(2):194–199, January 2010. 1, 2.2

[40] Kenny Olmstead and Aaron Smith. *Americans and Cybersecurity: Pew Research Center*. January 2017. 1

[41] 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, 2.2

[42] 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, 2.2

[43] Jonathon W. Penney. Chilling Effects: Online Surveillance and Wikipedia Use. *Berkeley Technology Law Journal*, 31(1):117–182, 2016. 5.5

[44] Prolific Team. Representative Samples on Prolific. https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-Samples-on-Prolific, March 2019. 4.2

[45] Lee Rainie. Americans' complicated feelings about social media in an era of privacy concerns. *Pew Research Center*, March 2018. 1

[46] Ronald W Rogers. A Protection Motivation Theory of Fear Appeals and Attitude Change. *The Journal of Psychology*, 91(1):93–114, 1975. 2.3

[47] Ronald W Rogers and Steven Prentice-Dunn. Protection motivation theory. 1997. 2.3

[48] Paschal Sheeran, Sarah Milne, Thomas L. Webb, and Peter M. Gollwitzer. *Implementation Intentions and Health Behaviour*. 2005. 2.2

[49] 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. 2.2, 5.2

[50] 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. 2.3

[51] 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

[52] 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. 2.3

[53] 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. 2.2, 5.5

[54] 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. 2.2, 5.2

[55] 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. 2.3

[56] 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. 3.3

[57] Richard H Thaler and Cass R Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. J. Wiley and Sons, 2008. 1, 2.1

[58] 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. 2.2

[59] 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. 2.2

[60] Alma Whitten and J D Tygar. Why Johnny can't encrypt: a usability case study of PGP 5. *USENIX Security Symposium*, August 1999. 2.1

[61] 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. 2.3

[62] 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