Better Late(r) than Never: Increasing Cyber-Security Compliance by Reducing Present Bias

Alisa Frik, Serge Egelman, Marian Harbach,[1] Nathan Malkin,[2] Eyal Peer.[3]

Abstract

Despite recent advances in increasing computer security by eliminating human involvement and error, there are still situations in which humans must manually perform computer security tasks, such as enabling automatic updates, rebooting machines to apply some of those updates, or enrolling in two-factor authentication. We argue that present bias—the tendency to discount future risks and gains in favor of immediate gratifications—could be the root cause explaining why many users fail to take such actions. Thus, we systematically explore the application of commitment devices, a technique from behavioral economics, to mitigate the effects of present bias on the adoption of end-user security measures. Offering users the option to be reminded or to schedule such tasks in the future could be effective in increasing their proclivity. While some current systems have begun incorporating such commitment nudges into software update messaging, we are unaware of rigorous scientific research that demonstrates how effective these techniques are, how they may be improved, and how they may be applied to other security behaviors. Using two online experiments, with over 1,000 participants total, we find that both reminders and commitment nudges can be effective at reducing the intentions to ignore the request to enable automatic updates (Study 1), and to install security updates and enable two-factor authentication, but not to configure automatic backups (Study 2). We also find that intentions of Mac OS users are generally more affected by these nudges.

Key Words: Usable security, behavioral economics, nudges, decision-making, commitment devices.

1 Introduction

In an ideal world, users would not need to do anything to stay safe and secure online, because systems would automatically protect them from threats such as malware, phishing, and account hijacking. That is, a "human-in-the-loop" would not be necessary [17]. While we are making significant gains towards this goal—many software updates are now applied automatically, email threat detection to recognize scams has vastly improved, and browser and device fingerprinting is used to discover potentially compromised accounts—that is not yet the world in which we live. As a result, users are still expected to perform certain security actions manually. Some of the most prevalent and critical actions are applying system updates and using two-factor authentication [39]. Recent qualitative studies have attempted to explain why these very effective security precautions are often resisted by end-users [e.g., 61]. One of the common findings is that many users are generally unopposed to taking these security measures. Yet, because they are asked at inopportune times, they decline in the moment and then later forget to revisit those decisions. This effect has previously been observed for smartphone locking [21] as well as applying software updates [e.g., 78, 75, 74, 29, 48, 47, 49].

The above examples illustrate that often, users are not necessarily resistant to performing security tasks in the general sense, but instead are opposed to doing them at the moment in which they are first asked. While the secondary nature of security mitigations is relatively well documented and a considerable variety of approaches have been explored to overcome this (see, e.g., [30], for a survey of the usable security literature), few have examined ways to actually solve this problem using theories from behavioral economics and without simply hiding security tasks from the user. Due to the limits of current technology, there is still (and probably will always be) a minimum set of security actions that users must perform themselves.[4] Thus, in the interim, new methods of nudging users towards engaging with security are needed.

In this paper, we therefore present the results of two online studies, aiming to explore options to reduce the effects of primary task interference. Researchers in psychology and behavioral economics have observed that people opt to delay long-term benefits in favor of short-term gains. This phenomenon, called *present bias*, indicates an individual's tendency to discount future outcomes in favor of present values [59, 43] and therefore to prefer immediate gratification over delayed utility. We posit that important security decisions are put off by users when being interrupted, because their decision making suffers from present bias. In the digital realm, Acquisti [1] has shown this to repeatedly impact privacy decision-

[4] In this paper, we make no attempt to define what this minimum set of security actions might be.

making: users succumb to hyperbolic discounting, in that they place greater weight on the immediate gratification of a social media posting, while devaluing long-term privacy costs.

Recent research on decision making has identified techniques for overcoming present bias [56], one of which is the use of *commitment devices* [15]. A commitment device is a mechanism that allows the "present self" to commit to a future action, so that the "future self" is more likely to follow through when the time comes. For example, not wanting to go to the gym today, Alice creates an appointment with a personal trainer for a specific date in the future. While that appointment could be canceled (an example of a "soft commitment"), she is more likely to follow through now that the appointment has been made. As an example of a "hard commitment," Alice could pay a non-refundable registration fee to enter a race in the future, which would motivate her to get into shape prior to that race. Other types of commitment devices involve rewards and punishments: Alice gives a check for $100 to a friend to hold in escrow; if she quits smoking by her target date, the check is destroyed, but if not, the friend mails it to a disagreeable charity. These types of commitment devices have been shown to be effective at changing behaviors such as curbing procrastination, saving more for retirement, and donating to charity [9, 69, 14].

Similar commitment nudges have started appearing in desktop software: both the most recent versions of Windows and Mac OS allow users to schedule system updates to be applied in the future (i.e., precommitting to a time of installation). However, we are unaware of any rigorously controlled experiments to measure the effects of these interventions, systematically improve them, and apply these principles to other security behaviors.

In this paper, we therefore report the results of two exploratory online experiments with more than 1,000 participants total. As a first step in this research area, we examine the circumstances under which commitment nudges induce a behavioral intent to improve security behaviors and how such nudges compare to other existing security decision-making interfaces. For the purposes of our research, we have identified a set of security actions that experts currently agree are important for end-users to perform [39, 61]: applying system updates, enabling two-factor authentication, and configuring automatic backups. Study 1 shows that, a commitment nudge (scheduling) can reduce the intentions to ignore the request to enable automatic updates by 15%, and a reminder can reduce such intentions by about 70%, for users who do not have automatic updates enabled. Study 2 shows that adding a reminder option reduced the willingness to ignore manual security update installation by about 50% for Windows and by about 75% for Mac users. Adding an option to commit to installing the update in the future reduced stated ignore rates by about 30% for Windows users and by about 45% for Mac users. While reminders and commitments were not effective at promoting the use of automatic backup tools, they were effective for promoting two-factor authentication (2FA): reminders halved the intentions to ignore (similarly for Windows and Mac users), and commitments showed a similar effect on Windows users, but no effect on Mac users' willingness to enable 2FA.

2 Related Work

Recent work in computer security has examined ways to improve user compliance with computer security mechanisms. For example, through better comprehension and usability of notices and controls [67, 25, 26, 33, 64]; advice on strong password composition [73, 65, 23]; use of memory-augmentation tools, such as password managers [38, 32]; and deployment of behavioral nudges [7, 76, 2]. Yet, even when users understand the importance of good security behaviors, they still do not always act accordingly [39, 61].

For example, applying software updates is one of the most common security practices users are regularly asked to perform; when promptly installed, they minimize attack surfaces [41, 52]. However, in practice people often avoid, delay, or skip updating their software [71, 74, 48, 49]. Research shows that users often have very rational reasons for declining to perform this important security activity: it may be related to confusion and annoyance [24]; fear of unanticipated user interface changes and satisfaction with the current software version [75, 48, 47]; concerns about reputation, resources usage, bugs, and disruption [74, 48, 47]; misunderstanding or underestimation of the benefits, threats, and consequences, overestimation of skills or time required for the update, and lack of such skills or time [74, 48, 47]. These seemingly rational reasons aside, there are many users who are fully aware of the importance of software updates and other security protective measures, who are hesitant to implement them [54]. Therefore, notifications alone, even when designed well, and when noticed and understood, are not always enough to trigger the desired behavioral change. Although the theories of planned behavior [5], reasoned action [27, 6], and protection motivation [63, 45] predict strong influence of attitudes on behavioral change, such con-

sequentialistic approaches do not always hold in real life. The misalignment between attitudes and actual conduct has been widely documented in privacy research as the "privacy paradox" [66, 11, 16, 3, 10, 53, 4]: people claim to value privacy, but then appear to not act accordingly. Research similarly suggests the existence of a "security paradox": people report high computer security concerns and state that they want to remain secure [72], yet are often resistant to performing the necessary actions [39, 79, 58]. Others have suggested that this is because security is seen as a secondary task, often interfering with a primary task [19, 20, 80]. We hold that this is an example of present bias.

2.1 Present Bias Theory

A useful framework to explore the observed gap between intentions and behaviors, and specifically why people fail to execute their plans, can be found in the research on time inconsistency, a phenomenon that has been studied by economists and psychologists for decades (for a review, see [46]). People are assumed to be present biased: they prefer immediate gratification over delayed utility, and therefore, they discount future outcomes in favor of present values [59, 43]. Due to this time preference, individuals tend to anticipate rewards and delay costs, and as a consequence, they procrastinate the activities that require salient costs and expedite the activities that presume salient benefits [55]. This presents a difficulty for inter-temporal choices, when costs and benefits happen at different moments in time. A classic example of this is saving for retirement, wherein individuals face the cost of not consuming a portion of income today in order to receive it in the future. As the utility from consuming this portion of income "today" always exceeds its "tomorrow" utility, the perfectly present-biased person ends up saving nothing. A similar situation frequently occurs when making security decisions: taking a security action often interrupts the workflow (cost) to protect against a future danger (benefit).

O'Donoghue et al. [56] present commitment devices as one possible mechanism for overcoming present bias. Commitment devices represent sophisticated attempts at self-control, for example, by limiting access (e.g., buying small packages of sweets, or locking up a mini-bar with alcohol), increasing sunk costs (e.g., purchasing an annual gym subscription), or setting up clear promises (e.g., college savings accounts).

2.1.1 Commitment Devices

Dual-self models describe the inter-temporal choice dilemma as a conflict between short- and long-run selves, where the long-run, or future, self is not able to hold the preferences or execute the plan of the short-run, or present, self [70]. Commitment strategies are often called upon to mitigate the conflict between short- and long-run selves. In a broad sense, Bryan et al. [15, p.1] define commitment devices as "an arrangement entered into by an individual with the aim of helping fulfill a plan for future behavior that would otherwise be difficult."

Commitments can operate through the restriction of a future choice set or setting up a penalty for not fulfilling the goal and a reward for its achievement. The model of Bisin and Hyndman [12] predicts that commitment devices based on gratification are stronger than ones based on penalties, because the cost and probability of not completing the task are lower, when deadline is farther away in time.

Gratification and punishment may be economic (hard commitments) and/or psychological (soft commitments) [15]. An example of a hard commitment device is a Save More Tomorrow (SMarT) plan, which allows employees to increase retirement savings via a pre-commitment to contribute to it as part of their future income [69]. A similar idea was successfully applied to charitable giving as well [14]. In contrast, publicly declaring a goal to complete a task serves as a soft commitment, because failure to achieve it entails primarily psychological consequences of disappointment or shame. These soft commitments of promises and goal-setting are considered the most natural yet effective way to commit [42].

Deadlines are another popular form of commitment device against procrastination. A study by Ariely and Wertenbroch [9] found that students who had a deadline performed better in delivery of the papers than those who did not. Moreover, the group, on whom the deadline was exogenously imposed, performed better than the group who was flexible to choose the deadlines themselves. Instead of hard deadlines, limiting the frequency of when the action can be submitted decreases procrastination and improves the overall success rate, through the increased cost of delay [56].

2.1.2 Present Bias in the Security Domain

The manifestation of present bias in the security domain is related to the common dominance of primary user task over security protective tasks. That is, security is almost never a primary task [19, 20, 80]; people do not sit down at the computer specifically

to "not get phished," "not get infected," or otherwise "do security." Even when users become aware of a potential security hazard, they are likely to see the risks as being in the future. Hence, at the moment of interaction with the computer, current needs are closer in time than the future risks [1], and the aspiration to complete the primary task exceeds the willingness to comply with the security recommendations, which are seen as inconveniences [61]. Economically speaking, the value of current need exceeds the value of future need. As with time, the order of primary and secondary needs is unlikely to change, different user tasks will always take priority over security activities.

We have observed examples of present bias in recent research. While interviewing smartphone users to understand why they do not securely lock their device screens (e.g., with a PIN or fingerprint), several participants indicated a desire to, but were asked at inconvenient times, so declined in the moment, thereby leaving their devices in insecure states [21]. Similarly, when examining why home users disabled automatic updates, several security-conscious users claimed that they wanted to exert more control over their systems. However, they later forgot to follow through with these actions, leaving their systems vulnerable [29].

Some present bias can be eliminated by simply automating security tasks, thereby taking them out of users' hands: automatically applying software updates increases installation rates and improves computers' immunity against attacks [31, 51]. However, forced updates exogenously transposition the order of user tasks, unexpectedly preventing the user from continuing a primary activity. This naturally leads to confusion, irritation, and dissatisfaction [75, 48, 49], such as a wave of indignation that followed the Windows 10 automatic update [35]. Apart from destroying users' workflows at potentially critical moments, automatic updates may undermine user trust in the long term [62, 47]. For example, one-third of the participants in the study of Mathur and Chetty [47] had disabled automatic updates, and these users were more likely to have had past negative experiences with updates. Moreover, keeping the user out of the loop removes control and leads to further divergence of mental models [77, 78].

Instead of automating security tasks, many systems require users' manual intervention but provide an opportunity to revisit the decision later. One example of such solution is Apple's "remind me later" option: when users are prompted with software update messages, they are given an opportunity to postpone the decision for later. Although generally this feature may produce considerable improvement in computer and smartphone security, this approach also precludes interruption of the user's workflow, making some users unhappy with this type of reminder. In response to the query, "Apple remind me later update," Google produces 9 out of 10 results describing how to "disable," "get rid of," "stop," or "turn off" this "annoying" feature.[5] Qualitative data by Fagan et al. [24, p.20] provides further evidence: one user in their study mentioned that "the option to be asked later is the most annoying [...] because it will continually pop up." However, the message continues to pop up because the user continues to delay the action. Therefore, reminders may not always be effective in reducing present bias.

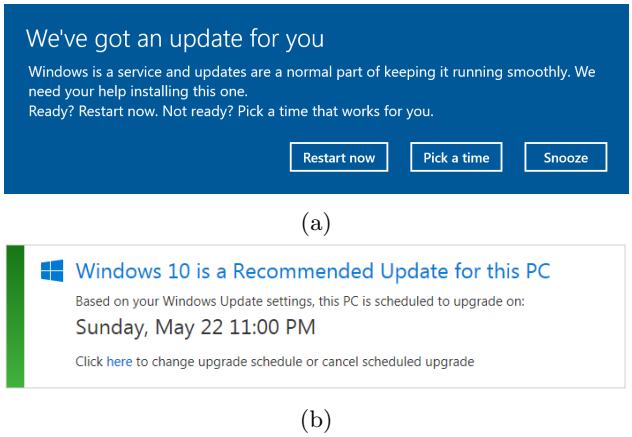

Figure 1: Windows 10 update message with scheduling feature.

An alternative approach to reminding is to schedule the software update for a certain time in the future. Similar to the results by Ariely and Wertenbroch [9], without changing the decision space, scheduling could play the role of a commitment device and lead to lower number of delaying events and higher compliance rates. Recently, Windows introduced such a scheduling feature (Figure 1), along with nudges, also inspired by behavioral science and psychology, such as a default option ("Upgrade now" with highlighted "OK" button), and social pressure ("Over 300 million people have already upgraded").

Research on persuasion profiling theory suggests that persuasive mechanisms constructed for one goal in a certain domain generally applies to other goals within the same domain [40]. Therefore, the methodology of commitment devices discussed for the case of security updates may be extended to other security-enhancing solutions as well, such as use of two-factor authentication or automatic data backups, which we

[5] As of August 23, 2017.

also explore in our study.

Overall, commitment devices have been proven to be a powerful instrument for overcoming present bias by aligning the behavior of the busy "present self" with the intentions of the security-conscious "future self." Despite preliminary adoption in certain security contexts, such as installing system updates, we are unaware that anyone has rigorously evaluated these approaches to gauge their effectiveness and determine ways in which they could be improved. Thus, our research aims to explore various forms of nudges, including commitment device, to evaluate which ones show the most promise in changing user security behaviors and decreasing present bias. In this study we start from testing the nudges that are currently used by the major software companies.

3 Study 1

Based on the prior literature on present bias and the effects of commitment devices in other domains, we decided to conduct an exploratory study to examine the effects of commitment devices on users' willingness to apply automatic software updates. We specifically examined whether users' willingness to enable *automatic* updates increases when they are given an option to make that change in the future, rather than in the moment. That is, our hypothesis was that when given the option to either act in the moment or not at all, many users will likely choose the latter option due to present bias (i.e., not wanting to enable automatic updates because it would interfere with a primary task). However, when given an option to reconsider the decision again in the future, we hypothesized that fewer users will outright decline to enable automatic updates, instead opting to revisit the decision at some point in the future. Recent versions of Windows (Figure 2a) and Mac OS (Figure 2b) present users with similar options. Based on the prior literature, as well as the current Windows dialogue, we also aimed to explore whether an option to be reminded of a pending decision may cause a different response than an option to actually commit to an action even though it will be executed at a later point in time.

To answer this question, we sampled 300 participants (42.3% females, Mean age = 34.16, SD = 11.15) from Prolific Academic, which is an online research participant recruitment platform.[6] We asked participants to imagine that while they were working on their computers, they received the following message: "Enabling Automatic Updates will make sure your operating system is always up-to-date and protected from malicious attacks, viruses or malware."

(a) Windows

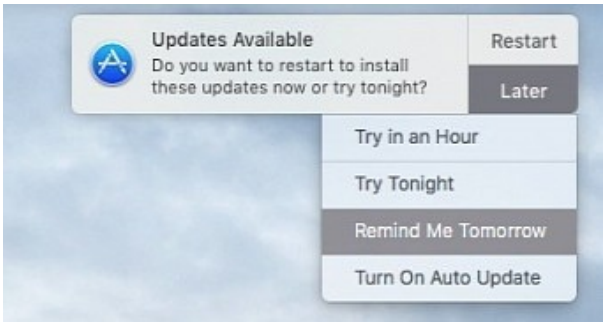

(b) Mac OS

Figure 2: Commitment devices to nudge users to apply software updates in Windows and Mac OS: the user is given the option to be asked again at some point in the future.

Next, we asked them the following question: "Assuming you don't have automatic updates enabled on your computer currently, what would you do next?" The response options varied according to three randomly assigned conditions and were presented as a survey-style question. In the control group, the options were to either ignore the message or enable automatic updates. In the *Commit* condition, the options were identical to the *Control*, with the exception of a third option to set auto-updates to be enabled "one week from today." In the *Reminder* condition, the third option was replaced with a request to be reminded again in a week (rather than committing to applying the change in a week). We also asked participants whether they currently had auto-updates enabled on their computers, in order to control for their prior experiences and attitudes towards automatic updates.

3.1 Results

Figure 3 shows the percentage of participants in each condition who chose to ignore the message, reported

[6] https://www.prolific.ac/

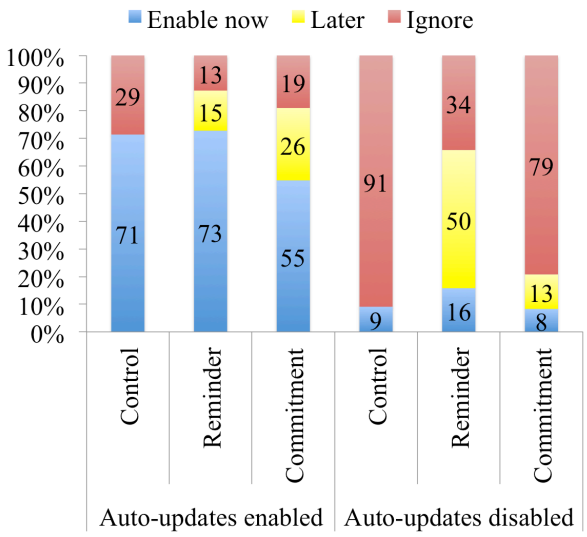

Figure 3: Response distribution in Study One (Auto-Updates scenario).

willingness to enable auto-updates now, or pledge to do so in the future (by either committing or setting a reminder). About 53% of participants reported to have auto-updates already enabled, and this question showed a significant interaction with the condition ($\chi^2(4) = 45.73, p < .001$). Thus, we split our analyses between participants who had vs. did not have auto-updates already enabled.

Commitment Condition As can be seen, among those who did not have auto-updates enabled, giving users the option to commit to enabling auto-updates in the future reduced the willingness to ignore from 90.9% to 79.2%, as 12.5% of participants in the *Commit* condition expressed the intention to enable auto-updates in a week ($\chi^2(4) = 39.85, p < .001$). This reduction did not significantly change the proportion of participants, who said they would enable auto-updates right now.

For participants who already have auto-updates enabled on their computers, we found that 26.2% expressed the intention to commit, when that option was given, and it reduced rates of "Ignore" choices from 28.6% to 19%. This time, however, there was a significant reduction in the share of participants willing to enable auto-updates now, from 71.4% in the control condition to 54.8% in the *Commit* condition ($\chi^2(4) = 16.72, p = .002$).

Reminder Condition The reminder option was, not surprisingly, more attractive than the commitment option among those who reported not having auto-updates already enabled (50% chose it compared to 12.5% who chose the commitment option). Among these participants, the reminder option actually increased the percentage of those who were willing to update now to 15.8% (compared to 9.1% in the control group). However, it was less attractive among participants who claimed to already have auto-updates enabled (14.5% chose it compared to 26.2% who chose the commitment option). Amongst these participants, the reminder option did not reduce the percentage of those who were willing to update now.

Summary The results of this study suggest that committing to enable auto-updates in the future (in one week) could be an attractive option, sometimes even more so than setting a simple reminder. Overall, 18.9% of all participants who were given that option opted for it, and among those who already had auto-updates enabled (but imagined they did not), it was slightly more popular than setting a reminder. Importantly, introduction of the commitment option reduced the intention to ignore the message for both groups, and it reduced the rates of those willing to enable auto-updates now only for those who reported already having auto-updates enabled.

The results of Study 1 demonstrate that one technique for mitigating present bias, commitment devices, show potential for increasing users' compliance with software update notifications. In this experiment, when presented with commitment option, the number of participants who expressed intention to completely reject the security task decreased. This encouraged us to run the second study to better inform future designs of these nudges.

4 Study 2

The most salient question is whether the translation of a stated behavioral intent to actual behavior will be any different between the reminder and commitment conditions. Before we can address this central research question in our future studies, we decided to first refine our nudge design, as studying behavior instead of behavioral intent will require a more complex and resource-intensive study design.

Building upon the results of Study 1, in Study 2 we aim at answering three key questions: (1) Will participants become less likely to ignore the security recommendation, if the point in time when the event takes place better fits their schedules? (2) Can reminder and commitment nudges be effective across various security behavior scenarios? And (3) do users of the two major operating systems – Windows and

Mac OS – react differently to such nudges?

To address this goal, we designed several nudges that can be applied towards addressing several end-user security behaviors that experts agree are important [e.g., 39], but require user action because they cannot yet be completely automated: applying operating system security updates, enrolling in two-factor authentication (2FA), and configuring automatic backups. We performed a hypothetical online experiment, similar to Study 1, to evaluate how our commitment and reminder nudges impact participants' stated willingness to comply with the requested security actions. In contrast to Study 1, Study 2 considers manual updates instead of automatic, it includes two additional security behavior scenarios, the timing options are not predefined but open-ended (participants propose the time for reminder or future installation themselves), people already engaged in certain security behaviors (e.g., automatic backups or 2FA) are screened out, and we also distinguish between the groups of Windows and Mac users.

4.1 Experimental design

We deployed the 3 (Control *vs.* Reminder *vs.* Commitment) x 3 (Update *vs.* Backups *vs.* 2FA) between-subject design. We sampled participants among users of both major operating systems — Mac OS and Windows — to compare the effects. We asked participants to imagine that after finishing this study, they received a message on the screen of their personal computer (the Updates and Backups scenarios) or in the browser (the 2FA scenario). These scenarios were chosen to increase the realism of the role playing task.

Next, we showed the participants a screenshot of the message (Figure 4)[7] and asked: "Among the following, what option would you click in response to this message in a real situation?" In the *Update* scenario, the message said: "A security update is available. Installing this update will protect your computer from attackers." In the *Backups* scenario, the message said: "The automatic backup tool is available. It will provide 50 GB of free virtual storage space and protect from data loss due to malicious software." In the *2FA* scenario, the message said: "Two-step verification for your Amazon Mechanical Turk account is available. It will add an extra layer of security because no one will be able to access your account if the password alone is compromised." We chose to use Mturk account for the 2FA scenario because we are sure that all subjects in our population have it, and because it contains personal and financial information (for some Mturk workers it is even a main source of income), therefore, Mturk account is subject to serious security, privacy, and financial loss risks.

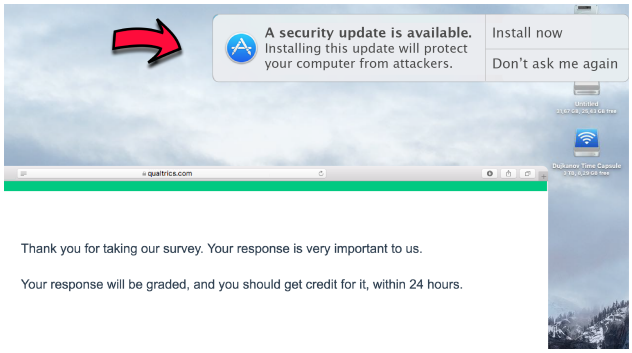

(a) Update scenario, Control condition, Mac.

(b) Backups scenario, Commitment condition, Windows.

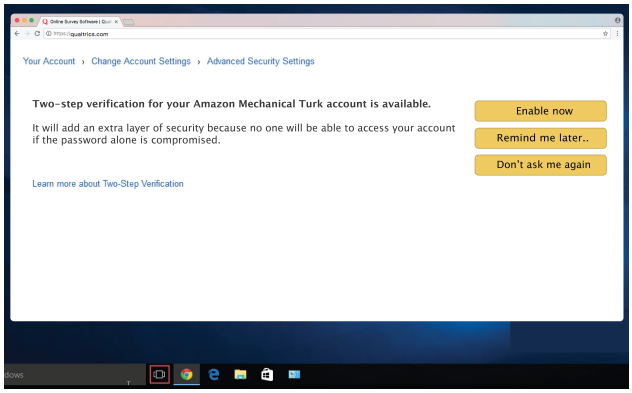

(c) 2FA scenario, Reminder condition, Windows.

Figure 4: Examples of the messages shown to the survey respondents.

The response options varied according to three randomly-assigned conditions. In the control group, the options were either to ignore the message or to install updates, enable automatic backups, or 2FA. In the *Commitment* condition, the options were identical to the *Control*, with the exception of a third option to pick a time to install updates, enable automatic backups, or enable 2FA in the future. In the *Reminder* condition, the third option was replaced with a request to be *reminded again* in the future

[7] We resized the images in the paper for the sake of space economy. In all conditions full-size screenshots were shown to the participants.

(rather than committing to *applying* the change). If participants chose the third option, we asked them to specify, in an open-ended manner, when they would prefer to receive a reminder or to apply the change.

Then, we asked respondents to explain, again in an open-ended manner, why they selected each option, and what circumstances would make them more likely to choose a different one. Finally, we surveyed participants' basic demographic information, and responses to a computer expertise scale [34] and the Security Behavior Intentions Scale (SeBIS) [22].

4.2 Results

We recruited 903 respondents among Mac and Windows users on Amazon Mechanical Turk (MTurk),[8] and randomly assigned them to one of the 9 experimental conditions (Table 1). We told participants that the study is about basic computer use preferences to not prime them to think about computer security specifically or induce self-selection bias. We screened out 108 participants who performed automatic computer backups and respondents who have Amazon two-step verification enabled. However, we did not exclude participants who reported backing up their computers manually, because they may do it irregularly and therefore could also benefit from the behavior change. The resulting sample includes 734 participants (53% females; age between 19 and 84, Mean = 37.78, SD = 12.12) (Table 2).

To estimate main treatment effects, for each scenario, we ran a logit regression with the participants' responses to the computer message as the dependent variable and two-way interactions between conditions and operating systems. We used two binary variable as dependent variables to represent respondents' choices of "Install/enable now" (Appendix, Table 4) and "Don't ask me again" (Appendix, Table 3). We also included age, gender, and corresponding SeBIS subscales as control independent variables. Additionally, we ran tests of proportions, and χ^2 test or Fisher's exact test (if numbers of observations in some cells was less than 5), to compare the ratio of specific choices in each condition.

4.2.1 Update scenario

Regression coefficients (Table 3) demonstrate that introduction of the nudges, either the reminder about available computer update or possibility to pick a time to install it in the future, significantly reduces the proportion of people willing to dismiss the message compared to the Control group, especially among participants who use Mac computers. The reminder nudge shows a larger effect size, and therefore has higher potential in improving the eventual installation rate than the commitment nudge (though it is unclear, of course, how many people will actually perform the behavior in the future, upon being reminded). Not surprisingly, people with positive security updating intentions, measured by the SeBIS Updating subscale, tend to choose positive update installation options more often. Age is negatively correlated with the willingness to install updates. Therefore, nudges are expected to be especially useful in increasing computer security amongst the older population and could be specifically targeted to them.

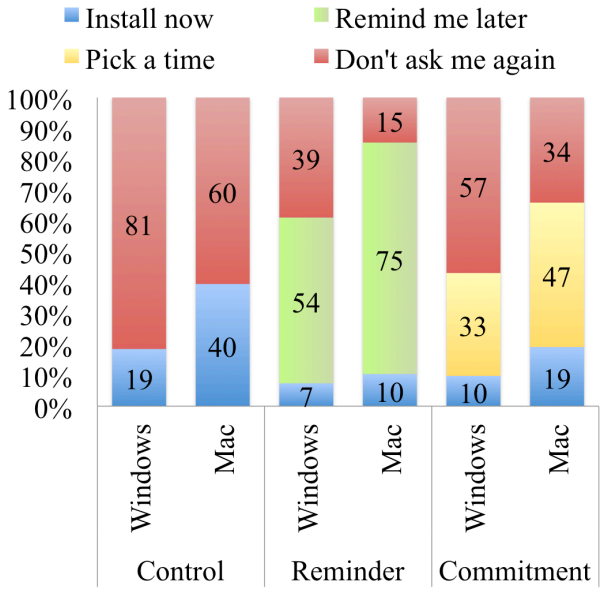

Figure 5: Response distribution in the Update scenario.

Figure 5 illustrates the distribution of responses in the Update scenario. In the presence of the "Remind me later" option, the share of participants who chose "Don't ask me again," dropped from 81% to 39% among Windows users ($p = .000; \chi^2(2) = 39.71, Pr = .000$) and from 60% to 15% among Mac users ($p = .000; \chi^2(2) = 57.61, Pr = .000$), compared to the Control. In the presence of the "Pick a time" option, these proportions decreased to 57% for Windows users ($p = .0062; \chi^2(2) = 21.68, Pr = .000$) and 34% for Mac users ($p = .0101; \chi^2(2) = 29.32, Pr = .000$).

Although regression coefficients (Table 4) do not reveal this effect, according to the results of test of proportions nudges reduced the share of "Install now"

[8] `www.mturk.com`. Subjects could participate if they lived in the United States, had previously completed at least 500 tasks, and had an approval rate of at least 95% on MTurk.

Table 1: Number of observations in experimental conditions.

Scenario	Condition (Windows — Mac)			Total
	Control	Reminder	Commitment	
Updates	102 (54 — 48)	102 (54 — 48)	98 (51 — 47)	302 (159 — 143)
Backups	68 (34 — 34)	60 (32 — 28)	67 (36 — 31)	195 (102 — 93)
- Do Not Back Up	37 (20 — 17)	23 (14 — 9)	33 (22 — 11)	93 (56 — 37)
- Manually Back Up	31 (14 — 17)	37 (18 — 19)	34 (14 — 20)	102 (46 — 56)
2FA	79 (37 — 42)	80 (42 — 38)	78 (38 — 40)	237 (117 — 120)
Total	249 (125 — 124)	242 (128 — 114)	243 (125 — 118)	734 (378 — 356)

Table 2: Respondents' demographic information.

Demographic	Participants
Age	
18-24	14.85%
25-34	33.65%
35-44	26.98%
45-54	12.40%
55-64	9.40%
65+	2.72%
Gender	
Male	45.46%
Female	52.86%
Other	0.54%
Prefer not to answer	0.95%
Education	
High school diploma or less	6.32%
Some college but no degree	19.23%
Associate's degree	11.81%
Bachelor's degree	42.72%
Master's degree	17.58%
Doctoral degree	2.34%
Number of participants	734

responses among Mac users ($p = .001$ in the Reminder condition, and $p = .029$ in the Commitment condition), but not Windows users ($p = .086$ in Reminder, and $p = .202$ in Commitment conditions). One explanation is that in the Control condition people had to choose essentially between two options – positive and negative. In the treatment groups, they had one negative choice (to dismiss the message forever) and two positive choices (either installing immediately or some time in the future). Therefore, naturally, the decisions of people, who are generally supportive of the proposed security practice, split between two positive categories based on their more nuanced preferences. Moreover, when the reminder or scheduling option was not offered, some participants could have chosen the "Install now" option because they may have anticipated a possibility to forget about it later. In the presence of such options, choosing "Remind me later" or "Pick a time" better corresponded to the preferences of these participants. We discuss this observation further and corroborate it with quotations from respondents' open-ended answers in Section 5.

A comparison of Windows and Mac users' responses reveals that the latter are less likely to dismiss the message about available updates in all conditions. Specifically, Mac users choose the "Don't ask me again" option less often than Windows users in the Control ($p = .0186$), Reminder ($p = .0256$), and Commitment ($p = .0235$) conditions. Mac users are also more likely to choose "Install now" in the Control condition (Appendix, Table 4). Additionally, they are more favorable to the reminders than Windows users: they chose the "Remind me later" option more often than Windows users ($p = .006; \chi^2(2) = 7.54; Pr = .023$). One potential explanation is that Windows displays notifications more frequently than Mac OS. For instance, in the past 10 years, Mac OS was releasing, on average, about 6 major security updates per year [8], compared to at least 12 Windows security updates, which happen every month [50]. Moreover, while Windows often offers an option to dismiss as a response to their messages (e.g., "Dismiss", "Cancel", "X", "Do not notify me again" buttons (Figure 2a)), Apple usually does not provide such an option in their messages. Therefore, Mac users may simply believe that the option to dismiss in our study is synonymous with asking them again tomorrow (Figure 2b).

In comparison to Study 1, the rate of users who choose to ignore the message in Commitment condition appears lower (even though Study 1 was about enabling automatic updates and Study 2 was about installing a specific update). This indicates that the point in time for which a commitment is demanded needs to be selected with care in order for such nudges to be effective. We will explore this aspect in more

detail in future work.

In summary, we conclude that both Reminder and Commitment nudges are effective in increasing the willingness to install updates, though the Reminder has a stronger effect than the Commitment. While these findings are true for both user groups, the impact is stronger on Mac users. Of course, the differences between the Reminder and Commitment conditions are unclear for behavior change: 100% of the users who commit to the behavior will ultimately perform the behavior, whereas a non-zero number of those being reminded may ask to be reminded *ad infinitum*—effectively never performing the behavior—or may outright dismiss a future reminder. We also found that older users tend to dismiss software update messages more often than younger users. Therefore, nudges targeted to this population could be especially useful in increasing their security.

4.2.2 Backups scenario

The first regression model (Backups(1) in Table 3) revealed that the Reminder nudge is effective in decreasing the willingness to dismiss the message about automatic backup tool for Mac users. Participants, who already back up manually are also less likely to dismiss the message, than those, who currently do not do backups. Among Windows users, Commitment nudge demonstrates negative effect on the willingness to choose "Enable now" option (Backups(1) in Table 4). We ran two additional regressions, separately for the participants, who do not back up their computers at all, and those who back them up manually (Backups(2) and (3) in Table 3 and Table 4). For simplicity, we further separately analyze these two subgroups.

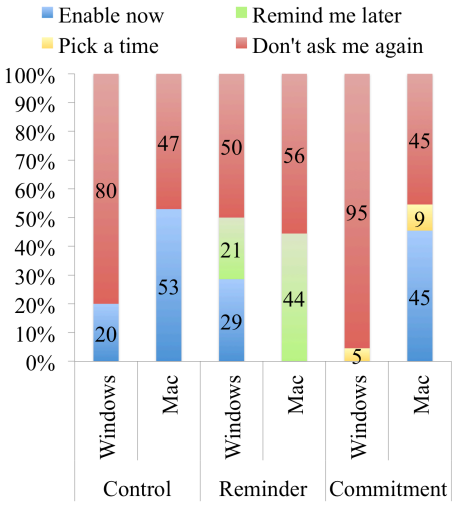

(a) Users not backing up their computers.

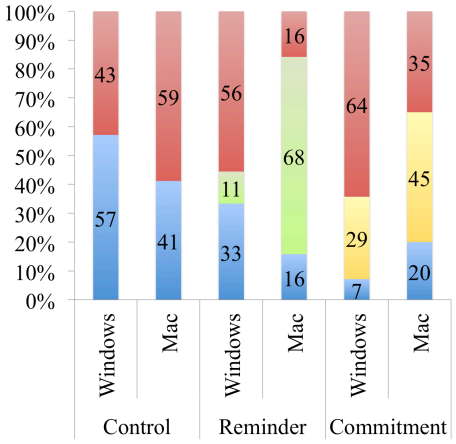

(b) Users who manually backup.

Figure 6: Backups scenario response distribution.

Currently do not back up We found that none of the nudges was significantly effective in driving the dismissing choices down to the respondents, who do not back up their computers (Backups(2) in Table 3). Figure 6a illustrates the distribution of responses. Even assuming all users who chose a Reminder or Commitment option eventually enable the automatic backups in real life, the overall compliance rate would not significantly outperform the baseline. More importantly, introduction of the Reminder to Mac users ($p = .0069$; Fisher's exact: $Pr = .001$) and Commitment nudge to Windows users (p= .0274; Fisher's exact: $Pr = .043$) dropped their willingness to immediately enable automatic backups to zero.[9] Ninety five percent of the Windows users among our respondents chose "Don't ask me again" in the Commitment condition, without giving it a second chance.

Generally, among respondents, who do not back up their computers, Windows users chose to dismiss the message about an available automatic backup tool more often than Mac users in the Control ($p = .0365$; Fisher's exact: $Pr = .047$) and Commitment ($p = .0009$; Fisher's exact: $Pr = .001$) conditions.

Currently manually back up For users who currently back up their computers manually, regression coefficients did not reveal significant treatment effects on the willingness to dismiss the message (Backups(3) in Table 3). Introduction of the Reminder did not significantly change the proportion of "Enable now" choices for either user groups (for Windows $p = .178$,

[9]This explains omitted regression coefficients for these effects in Backups(2) model in Table 4 of the Appendix).

for Mac $p = .09$). However, introduction of the Commitment nudge significantly decreased the proportion of "Enable now" choices for Windows users (Backups(3) in Table 4; test of proportions: $p = .005$), but not for Mac users ($p = .16$).

Comparing the responses of Mac and Windows users overall in this subgroup, we found that Mac users are more favorable to "Remind me later" ($p = .0113$) and "Pick a time" ($p = .0109$) options than Windows users. In the Reminder condition, they also tend to dismiss the message less often than Windows users ($p = .0004$; Fisher's exact: $Pr = .001$).

Therefore, we conclude that the examined nudges are not effective in increasing the willingness to install an automatic backup tool, and might even decrease it.

4.2.3 Two-factor authentication scenario

The regression model shows that the Reminder is an effective nudge in decreasing the willingness to dismiss the request to enroll in 2FA for users of both major operating systems (with a larger effect on Windows users), while the Commitment nudge only has a significant impact on Windows users' intentions (Table 3). As can be seen in Figure 7, a reminder reduced the ignore rates from 70 to 33 percent for Windows users ($p = .000; \chi^2(2) = 19.85; Pr = .000$), and from 64 to 34 percent for Mac users ($p = .007; \chi^2(2) = 25.67; Pr = .000$). A commitment nudge was also effective among Windows users, reducing ignore rates to 29 percent ($p = .000; \chi^2(2) = 16.2; Pr = .000$), but not so among Mac users, of which 58 percent still chose to ignore the message ($p = .529; \chi^2(2) = 8.28; Pr = .016$). Introduction of the Reminder and Commitment nudges did not significantly change the proportion of "Enable now" choices for either user groups (Table 4).[10] Positive intentions to generate secure passwords in general (as measured by the SeBIS Password subscale) are correlated with the reported tendency to add an extra layer of security to one's online account.

2FA is the only scenario in our experiment in which there was no significant difference between the choices of Windows and Mac users in the Control and Reminder conditions. This may be evidence of the efficacy of the role playing scenarios, as this was the only condition where the message was displayed in a browser window instead of the desktop of the operating system.

[10]Reminder condition: for Windows $p = .434$, for Mac $p = .073$. Commitment condition: for Windows $p = .084$, for Mac $p = .292$.

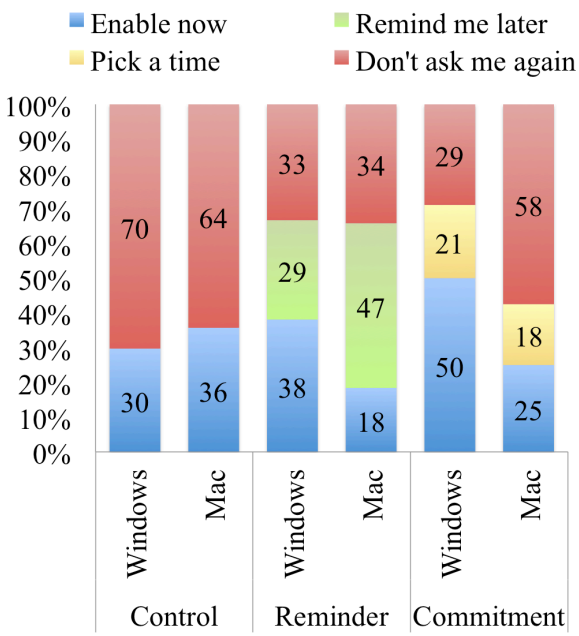

Figure 7: Response distribution in the 2FA scenario.

We conclude that reminders have potential to effectively increase 2FA adoption among users of all major operating systems, while the Commitment nudge is expected to be effective in improving the intention to enroll in 2FA among Windows users.

5 Discussion

Our online experiments revealed that offering users an opportunity to delay or schedule a security action for a future time often increases their stated willingness to accept the proposed security behavior (or rather, decreases the likelihood of them dismissing it outright). Reminders demonstrated a greater potential in improving security intentions than commitment devices. However, their effects differ across user groups and security behaviors. In this section, we offer interpretation and implications for our findings and discuss viable ways for improving nudges and commitment devices as means for mitigation of present bias in the security domain.

5.1 Decreasing Immediate Action

In most conditions, we noticed a "side effect" of the nudges – although they decreased the overall negative response rate (i.e., share of "Don't ask me again" choices), they also drove the proportion of "now" choices down. As we mentioned in Section 4.2.1, introduction of the second positive option generated

a split between people with positive security intentions into those willing to act immediately and those who prefer to defer the security action until later. Anecdotally, open-ended responses appear to support this explanation: several participants in the control condition noted that their choices of "now" options were driven by the fear of forgetting about it later. Similarly, other participants noted that they *"would consider in the future but in this moment [...] most likely wouldn't take the time to set it up right now."* For example, P37 in the Update-Control condition wrote: *"I would forget about doing it at a later time if I didn't choose to do it right away."*

Therefore, restricting users' choices may be a double-edged sword: it is possible that the presence of an option to delay may make those who would otherwise act in the moment to procrastinate. In our future work, we will try to address this issue. For example, increasing the "behavioral cost" of procrastination (i.e. manipulating the choice architecture so that "Install now" option is easier and more attractive than delaying option) could be one way to mitigate the negative impact of nudges on the immediate action options.

5.2 Timing Is Important

In our second study, we intentionally did not impose a fixed delay option in the Reminder and Commitment conditions. Instead, we allowed participants to specify when they would like to receive a reminder or pick a time for future security action in an open-ended manner. As already noted in Section 4.2.1, the fraction of participants choosing the reminder and commitment options appeared to increase between Studies 1 and 2. We hypothesize that this is because when being able to select a time that fits, instead of choosing from a predefined set of options, participants were more easily able to choose these options.

A brief inspection of the open-ended responses also hints at this effect: while many respondents indicate the preferred delay in terms of time (e.g., "tomorrow," "in 1 hour," "at 2am"), roughly a third of participants who chose these options specified *conditions* when it might be more convenient for them (e.g., *"When I'm done using my PC for the day," "Next time I log in"*).

These findings have several practical implications for future nudge designs. First, when people are asked to find a slot to schedule an action or reminder, they may simply have no suitable time in mind. Then, showing them too many time-related options may confuse and annoy them, leading to a decision to dismiss the message altogether. A heuristic-based—as opposed to time-based—option to defer a decision may better address the preferences of this group of users, reduce their negative emotions, and avoid formation of general negative attitudes to these kind of messages. Second, providing more information about how much time the process will take would also help to plan ahead and schedule the activity properly. Alternatively, a non-action option, e.g., simply letting the message hang on unobtrusively or be moved around desktop as a post-it sticker until the user has time to deal with it, could also be a solution.

Prior work on user interruptions and a thorough inspection of our own qualitative data will provide insights into how to best design these intuitive options and account for contextual factors [e.g., 36, 28, 44, 13, 18, 37, 57, 68, 60], which is a subject for future work.

5.3 Nudges Do Not Always Work

We found that the nudges we used were least (and in some cases even negatively) effective in the Backups scenario. One reason for this effect could be that while security updates and 2FA provide explicit protection against security risks, automatic backups may have less straightforward implications for reducing security risks from the ordinary user's perspective, and therefore elicit low willingness to enable them for the reasons not related to timing. For instance, a perceived lack of importance of local files, and therefore low interest in protecting them, may also lead to users' unwillingness to install a back up tool.

Additionally, a few participants' comments hint at negative prior experience with spam-like messaging that prevented them from complying with the request for enabling automatic backups, as the dialogue seemed *"almost like an advertisement. I'd have to scan my PC viruses afterwards."* We thus hypothesize that security behaviors that include installing software or enabling a third-party product cannot be easily nudged, as additional barriers come into play. This too, is an area for future study.

5.4 Mac and Windows Users Behave Differently

Finally, we noticed that in most conditions Windows users were more likely to dismiss security recommendation messages than Mac users. As speculated in Section 4.2.1, Windows seems to show more notification dialogues and may hence cause habituation and/or fatigue with regards to similar requests. More research should be done in understanding this difference, addressing implicit concerns, and customization

of messages to both populations of software users.

There is also a possibility that Windows users trust their software company less Mac users. For instance, one participant said that he believes that *"malware ... doesn't often happen on a Mac."* Another participant believed that usually Apple products are not *"getting attacked by viruses. This is why I own an Apple."* We plan to investigate this potential explanation further.

5.5 Limitations and Future Research

The main limitation of our study is its hypothetical nature, as we measured users' *stated willingness* to engage in certain behaviors, but did not observe consequential, actual behavior. However, in line with prior decision-making research, our results show that commitment nudges *do* change computer security *intentions*, which are a precursor to behavior change [5]. Moreover, to bring our scenarios closer to real life, we presented participants with screenshots of the messages rather than describing the situation in a purely textual form. We acknowledge that intentions may be overestimated with respect to actual behavioral rates, but in this paper we focus on the comparison of *relative* effectiveness of the nudges. Therefore, while in absolute terms the actual compliance rates may differ from the estimated intentions, we believe that the general relative trends observed in our hypothetical study are likely to hold in real life. We plan to further develop and test various nudges in more realistic settings with rich contextual ambiance in the future longitudinal field experiment.

On the other hand, we would like to emphasize the advantage (and even, in our view, important prerequisite) of running *hypothetical* studies on the early stages of designing and testing a large variety of messages in a safe environment. Despite the positive intention to improve the cyber-security behaviors, our findings reveal that mis-targeted or poorly designed nudges not only can be ineffective, but even harmful. For instance, Windows is actively experimenting with A/B-testing of their security messages, manipulating wording, design, and choice architecture. In the real world, poorly targeted nudges may increase users' vulnerability and actual security risks. Therefore, we warn researchers and practitioners to attentively consider nudge design and thoroughly test them in safe environments before the full-scale implementation, or even small pilot field trials.

Regression analysis revealed lower willingness to install updates among older respondents in our hypothetical scenarios. Therefore, researchers and practitioners should be especially attentive to the inclusion of a diverse population in their testing to address their concerns and control for the potentially adverse effects of their nudges, especially on sensitive populations, such as children or older users.

We believe that compliance with security recommendations is time-sensitive: for instance, the longer the delay in applying software updates (or enrolling in 2FA, configuring automatic backups, etc.), the longer the devices are vulnerable to attacks and the larger are the potential losses. Therefore, we believe that researchers should not only try to increase the overall engagement with certain security activities (when they cannot be automated), but they should also decrease the time it takes for users to comply. In our analysis we calculate the upper bound of treatment effects on the willingness to perform security task, under assumption of a 100% compliance rate upon receiving the reminder or scheduled prompt. In the real world, users may renege on their pledge (e.g., dismissing the message upon being reminded in the future) or continue to procrastinate indefinitely. Therefore, in future research we will test the number and duration of repeating delays and actual compliance rate in a longitudinal fashion.

Windows currently offers users to install the available updates immediately, "snooze" until later, or schedule the installation for the future ("pick a time"). After testing each nudge separately, we will also consider testing the interaction effect of multiple nudges presented to the user simultaneously, as nudges may have adversary effects.

6 Conclusions

We performed an online study to test the effectiveness of reminders and commitment nudges in improving users' intentions to engage in cyber-security behaviors by reducing the present bias effect. As a first step in exploring this application of behavioral economics to computer security domain, we found that both Reminder and Commitment nudges have the potential to increase willingness to engage in beneficial computer security behaviors up to 85%. However, at the same time, introducing nudging options decreased the fraction of users reported willingness to take immediate action, effectively placing a bet on this stated intention translating to actual behavior at a later point in time. Our results also show that commitment devices may not be equally successful in nudging users towards all security behaviors, as we were unable to establish positive effects for enabling automatic backups. Furthermore, we posit that current nudging dialogues may not live up to their full

potential, as the timing options offered to users may be too rigid.

In future work, we will establish whether or not this gamble will pay off after using our results to improve the security messages designs currently used by major software companies. We will test which nudges not only increase the overall compliance rate, but also reduce the delay period. As an ultimate goal of our research agenda, we will run longitudinal and field studies to test behavioral outcomes of commitment nudges in realistic settings.

References

[1] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Electronic Commerce Conference (EC '04)*, pages 21–29, New York, NY, 2004. ACM Press. `http://www.heinz.cmu.edu/~acquisti/papers/privacy-gratification.pdf`.

[2] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. F. Cranor, S. Komanduri, P. G. Leon, N. Sadeh, F. Schaub, M. Sleeper, et al. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3):44, 2017.

[3] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the facebook. In *International workshop on privacy enhancing technologies*, pages 36–58. Springer, 2006.

[4] A. Acquisti, C. R. Taylor, and L. Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–492, 2016.

[5] I. Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211, 1991.

[6] I. Ajzen and T. Madden. Prediction of goal directed behaviour: Attitudes, intentions and perceived behavioural control. *Journal of Experimental Social Psychology*, 22(5):453–474, 1986.

[7] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. Cranor, and Y. Agarwal. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. Technical Report CMU-ISR-14-116, Carnegie Mellon University, 2014.

[8] Apple Support. Apple security updates. Technical report, Accessed [15 February 2018]: `https://support.apple.com/en-us/HT201222`, 2018.

[9] D. Ariely and K. Wertenbroch. Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological science*, 13(3):219–224, 2002.

[10] S. B. Barnes. A privacy paradox: Social networking in the united states. *First Monday*, 11(9), 2006.

[11] B. Berendt, O. Günther, and S. Spiekermann. Privacy in e-commerce: stated preferences vs. actual behavior. *Communications of the ACM*, 48(4):101–106, 2005.

[12] A. Bisin and K. Hyndman. Procrastination, self-imposed deadlines and other commitment devices. *MPRA Working Paper Nr. 16235*, 2009.

[13] P. Bogunovich and D. Salvucci. The effects of time constraints on user behavior for deferrable interruptions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 3123–3126, New York, NY, USA, 2011. ACM.

[14] A. Breman. Give more tomorrow: Two field experiments on altruism and intertemporal choice. *Journal of Public Economics*, 95(11):1349–1357, 2011.

[15] G. Bryan, D. Karlan, and S. Nelson. Commitment devices. *Annual Review of Economics*, 2(1):671–698, 2010.

[16] R. K. Chellappa and R. G. Sin. Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information technology and management*, 6(2):181–202, 2005.

[17] L. F. Cranor. A framework for reasoning about the human in the loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, Berkeley, CA, 2008. USENIX Association.

[18] L. Dabbish, G. Mark, and V. M. González. Why do i keep interrupting myself?: Environment, habit and self-interruption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI'11, pages 3127–3130, New York, NY, USA, 2011. ACM.

[19] P. Dourish, E. Grinter, J. Delgado de la Flor, and M. Joseph. Security in the wild: User strategies for managing security as an everyday, practical problem. *Personal Ubiquitous Comput.*, 8(6):391–401, Nov. 2004.

[20] W. K. Edwards, E. S. Poole, and J. Stoll. Security automation considered harmful? In *Proceedings of the 2007 Workshop on New Security Paradigms*, NSPW'07, pages 33–42, New York, NY, USA, 2008. ACM.

[21] S. Egelman, S. Jain, R. S. Portnoff, K. Liao, S. Consolvo, and D. Wagner. Are you ready to lock? understanding user motivations for smartphone locking behaviors. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer & Communications Security*, CCS '14, New York, NY, USA, 2014. ACM.

[22] S. Egelman and E. Peer. Scaling the Security Wall: Developing a Security Behavior Intentions Scale (SeBIS). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI'15, New York, NY, USA, 2015. ACM.

[23] S. Egelman, A. Sotirakopoulos, I. Muslukhov, K. Beznosov, and C. Herley. Does my password go up to eleven?: the impact of password meters on password selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2379–2388. ACM, 2013.

[24] M. Fagan, M. M. H. Khan, and R. Buck. A study of users' experiences and beliefs about software update messages. *Computers in Human Behavior*, 51:504–519, 2015.

[25] A. P. Felt, S. Egelman, M. Finifter, D. Akhawe, and D. Wagner. How to ask for permission. In *Proceedings of the 7th USENIX conference on Hot Topics in Security*, HotSec'12, pages 7–7, Berkeley, CA, USA, 2012. USENIX Association.

[26] A. P. Felt, R. W. Reeder, H. Almuhimedi, and S. Consolvo. Experimenting at scale with google chrome's ssl warning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 2667–2670, New York, NY, USA, 2014. ACM.

[27] M. Fishbein. A theory of reasoned action: some applications and implications. *Nebraska Symposium on Motivation*, 27(1):65–116, 2008.

[28] J. Fogarty, J. Lai, and J. Christensen. Presence versus availability: The design and evaluation of a context-aware communication client. *International Journal of Human-Computer Studies*, 61(3):299–317, Sept. 2004.

[29] A. Forget, S. Pearman, J. Thomas, A. Acquisti, N. Christin, L. F. Cranor, S. Egelman, M. Harbach, and R. Telang. Do or do not, there is no try: User engagement may not improve security outcomes. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 97–111, Denver, CO, 2016. USENIX Association.

[30] S. Garfinkel and H. R. Lipford. *Usable Security: History, Themes, and Challenges*. Morgan & Claypool, 2014.

[31] C. Gkantsidis, T. Karagiannis, and M. Vojnović. Planet scale software updates. *ACM SIGCOMM Computer Communication Review*, 36(4):423–434, 2006.

[32] E. Grosse and M. Upadhyay. Authentication at scale. *IEEE Security & Privacy*, 11(1):15–22, 2013.

[33] M. Harbach, M. Hettig, S. Weber, and M. Smith. Using personal examples to improve risk communication for security and privacy decisions. In *Proceedings of the 2014 CHI Conference on Human FActors in Computing Systems*, CHI'14, pages 2647–2656, New York, NY, USA, 2014. ACM.

[34] E. Hargittai and Y. P. Hsieh. Succinct survey measures of web-use skills. *Social Science Computer Review*, 30(1):95–107, 2012.

[35] S. Hollister. Microsoft won't fix the most frustrating thing about windows. Cnet. `https://www.cnet.com/news/microsoft-windows-10-forced-updates/`, 2017.

[36] S. Hudson, J. Fogarty, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee, and J. Yang. Predicting human interruptibility with sensors: A wizard of oz feasibility study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI'03, pages 257–264, New York, NY, USA, 2003. ACM.

[37] C. Hurter, B. R. Cowan, A. Girouard, and N. H. Riche. Active progress bar: Aiding the switch to temporary activities. In *Proceedings of the 26th Annual BCS Interaction Specialist Group Conference on People and Computers*, BCS-HCI'12,

pages 99–108, Swinton, UK, UK, 2012. British Computer Society.

[38] A. Huth, M. Orlando, and L. Pesante. Password security, protection, and management. *United States Computer Emergency Readiness Team*, 2012.

[39] I. Ion, R. Reeder, and S. Consolvo. "...no one can hack my mind": Comparing expert and non-expert security practices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 327–346, Ottawa, 2015. USENIX Association.

[40] M. Kaptein and D. Eckles. Selecting effective means to any end: Futures and ethics of persuasion profiling. In *Persuasive*, pages 82–93. Springer, 2010.

[41] M. Khan, Z. Bi, and J. A. Copeland. Software updates as a security metric: Passive identification of update trends and effect on machine infection. In *Military Communication Conference 2012*, pages 1–6. IEEE, 2012.

[42] A. K. Koch and J. Nafziger. Self-regulation through goal setting. *The Scandinavian Journal of Economics*, 113(1):212–227, 2011.

[43] D. Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.

[44] B. Y. Lim, O. Brdiczka, and V. Bellotti. Show me a good time: Using content to provide activity awareness to collaborators with activityspotter. In *Proceedings of the 16th ACM International Conference on Supporting Group Work*, GROUP'10, pages 263–272, New York, NY, USA, 2010. ACM.

[45] J. E. Maddux and R. 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.

[46] D. Malhotra, G. Loewenstein, and T. O'donoghue. Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2):351–401, 2002.

[47] A. Mathur and M. Chetty. Impact of user characteristics on attitudes towards automatic mobile application updates. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 175–193, Santa Clara, CA, 2017. USENIX Association.

[48] A. Mathur, J. Engel, S. Sobti, V. Chang, and M. Chetty. "They keep coming back like zombies": Improving software updating interfaces. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 43–58, Denver, CO, 2016. USENIX Association.

[49] A. Mathur, N. Malkin, M. Harbach, E. Peer, and S. Egelman. Quantifying users' beliefs about software updates. *arXiv preprint arXiv:1805.04594*, 2018.

[50] Microsoft Security TechCenter. Microsoft security updates bulletins. Technical report, Accessed [15 February 2018]: `https://technet.microsoft.com/en-us/security/bulletins`, 2018.

[51] A. Nappa, R. Johnson, L. Bilge, J. Caballero, and T. Dumitras. The attack of the clones: A study of the impact of shared code on vulnerability patching. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 692–708. IEEE, 2015.

[52] K. Nayak, D. Marino, P. Efstathopoulos, and T. Dumitraş. Some vulnerabilities are different than others. In *International Workshop on Recent Advances in Intrusion Detection*, pages 426–446. Springer, 2014.

[53] P. A. Norberg, D. R. Horne, and D. A. Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.

[54] D. A. Norman. Cognitive engineering. *User centered system design*, 31:61, 1986.

[55] T. O'Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, pages 103–124, 1999.

[56] T. O'Donoghue, M. Rabin, et al. Incentives and self-control. *Econometric Society Monographs*, 42:215, 2006.

[57] T. Okoshi, J. Ramos, H. Nozaki, J. Nakazawa, A. K. Dey, and H. Tokuda. Reducing users' perceived mental effort due to interruptive notifications in multi-device mobile environments. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp'15, pages 475–486, New York, NY, USA, 2015. ACM.

[58] Pew Research Center. Americans and cybersecurity. Technical report, Accessed [11 April 2018]: http://assets.pewresearch.org/wp-content/uploads/sites/14/2017/01/26102016/Americans-and-Cyber-Security-final.pdf, 2017.

[59] E. S. Phelps and R. A. Pollak. On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2):185–199, 1968.

[60] M. Pielot, B. Cardoso, K. Katevas, J. Serrà, A. Matic, and N. Oliver. Beyond interruptibility: Predicting opportune moments to engage mobile phone users. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(3):91:1–91:25, Sept. 2017.

[61] E. M. Redmiles, S. Kross, and M. L. Mazurek. How I learned to be secure: A census-representative survey of security advice sources and behavior. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS'16, pages 666–677, New York, NY, USA, 2016. ACM.

[62] E. Rescorla. Security holes... who cares? In *USENIX Security Symposium*, pages 75–90. Washington, DC, 2003.

[63] R. W. Rogers. A protection motivation theory of fear appeals and attitude change1. *The journal of psychology*, 91(1):93–114, 1975.

[64] F. Schaub, R. Balebako, and L. F. Cranor. Designing effective privacy notices and controls. *IEEE Internet Computing*, 21(3):70–77, 2017.

[65] R. Shay, P. G. Kelley, S. Komanduri, M. L. Mazurek, B. Ur, T. Vidas, L. Bauer, N. Christin, and L. F. Cranor. Correct horse battery staple: Exploring the usability of system-assigned passphrases. In *Proceedings of the eighth symposium on usable privacy and security*, page 7. ACM, 2012.

[66] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47. ACM, 2001.

[67] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor. Crying wolf: an empirical study of ssl warning effectiveness. In *Proceedings of the 18th USENIX Security Symposium*, SSYM'09, pages 399–416, Berkeley, CA, USA, 2009. USENIX Association.

[68] D. Tasse, A. Ankolekar, and J. Hailpern. Getting users' attention in web apps in likable, minimally annoying ways. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI'16, pages 3324–3334, New York, NY, USA, 2016. ACM.

[69] R. H. Thaler and S. Benartzi. Save more tomorrow^{TM}: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(S1):S164–S187, 2004.

[70] R. H. Thaler and H. M. Shefrin. An economic theory of self-control. *Journal of Political Economy*, 89(2):392–406, 1981.

[71] Y. Tian, B. Liu, W. Dai, B. Ur, P. Tague, and L. F. Cranor. Supporting privacy-conscious app update decisions with user reviews. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 51–61. ACM, 2015.

[72] Unisys. Unisys security index. Technical report, Accessed [11 April 2018]: http://www.unisys.com/unisys-security-index/us, 2017.

[73] B. Ur, P. G. Kelley, S. Komanduri, J. Lee, M. Maass, M. L. Mazurek, T. Passaro, R. Shay, T. Vidas, L. Bauer, et al. How does your password measure up? the effect of strength meters on password creation. In *USENIX Security Symposium*, pages 65–80, 2012.

[74] K. Vaniea and Y. Rashidi. Tales of software updates: The process of updating software. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 3215–3226. ACM, 2016.

[75] K. E. Vaniea, E. Rader, and R. Wash. Betrayed by updates: how negative experiences affect future security. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2671–2674. ACM, 2014.

[76] Y. Wang, P. G. Leon, A. Acquisti, L. F. Cranor, A. Forget, and N. Sadeh. A field trial of privacy nudges for facebook. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2367–2376. ACM, 2014.

[77] R. Wash. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, page 11. ACM, 2010.

[78] R. Wash, E. Rader, K. Vaniea, and M. Rizor. Out of the loop: How automated software updates cause unintended security consequences. In *Symposium on Usable Privacy and Security (SOUPS)*, pages 89–104, 2014.

[79] R. Wash and E. J. Rader. Too much knowledge? security beliefs and protective behaviors among United States Internet users. In *SOUPS*, pages 309–325, 2015.

[80] R. West. The psychology of security. *Communications of the ACM*, 51(4):34–40, Apr. 2008.

Appendix

Table 3: Logit regression on the respondents' choices of "Don't ask me again" (0 - no, 1 - yes).

	Update	2FA	Backups (1)	Backups (2)	Backups (3)
Control × Windows (baseline)					
Control × Mac	-1.434**	-0.146	-0.202	-1.188	0.993
	[-2.40,-0.47]	[-1.13,0.84]	[-1.31,0.90]	[-2.75,0.37]	[-0.74,2.73]
Reminder × Windows	-2.430***	-1.774***	-0.647	-1.622+	0.461
	[-3.39,-1.46]	[-2.77,-0.78]	[-1.76,0.47]	[-3.28,0.04]	[-1.15,2.07]
Reminder × Mac	-3.819***	-1.530**	-1.944**	-1.766+	-1.720+
	[-4.99,-2.65]	[-2.54,-0.52]	[-3.20,-0.68]	[-3.67,0.13]	[-3.59,0.15]
Commitment × Windows	-1.673***	-1.628**	1.019	1.538	1.189
	[-2.64,-0.71]	[-2.65,-0.61]	[-0.22,2.25]	[-0.82,3.90]	[-0.60,2.97]
Commitment × Mac	-2.552***	-0.400	-1.103+	-1.616+	-0.514
	[-3.57,-1.54]	[-1.38,0.58]	[-2.24,0.03]	[-3.34,0.11]	[-2.14,1.12]
Age	0.0323*	0.00961	0.00264	-0.0214	0.0259
	[0.01,0.06]	[-0.01,0.03]	[-0.03,0.03]	[-0.07,0.02]	[-0.02,0.07]
Female	-0.00696	-0.130	-0.346	-0.617	-0.269
	[-0.55,0.53]	[-0.71,0.45]	[-1.04,0.34]	[-1.67,0.44]	[-1.26,0.72]
SeBIS Updating subscale	-1.350***				
	[-2.00,-0.70]				
SeBIS Password Generation subscale		-0.571**			
		[-0.95,-0.19]			
SeBIS Proactive Awareness subscale			1.340***	1.011+	1.360**
			[0.67,2.01]	[-0.03,2.05]	[0.40,2.33]
Back up manually			-1.219***		
			[-1.93,-0.51]		
Constant	0.580	0.489	1.398+	2.829*	-1.312
	[-0.64,1.80]	[-0.75,1.73]	[-0.03,2.83]	[0.62,5.04]	[-3.33,0.71]
Including those, who do not back			Yes	Yes	No
Including those, who back up manually			Yes	No	Yes
N	297	234	192	92	100
χ^2	87.13	37.30	50.55	26.59	24.51

95% confidence intervals in brackets. $^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Logit regression on the respondents' choices of "Install/enable now" (0 - no, 1 - yes).

	Update	2FA	Backups (1)	Backups (2)	Backups (3)
Control × Windows (baseline)					
Control × Mac	1.809***	0.241	0.540	1.578^{+}	-0.640
	[0.78,2.84]	[-0.74,1.23]	[-0.53,1.61]	[-0.00,3.16]	[-2.32,1.04]
Reminder × Windows	-0.785	0.297	-0.0323	0.464	-0.866
	[-2.08,0.51]	[-0.68,1.27]	[-1.13,1.07]	[-1.24,2.17]	[-2.46,0.73]
Reminder × Mac	-0.227	-0.612	-1.238	0	-1.748^{+}
	[-1.45,1.00]	[-1.73,0.50]	[-2.71,0.24]	[0.00,0.00]	[-3.59,0.10]
Commitment × Windows	-0.330	0.676	-2.985**	0	-3.133*
	[-1.56,0.90]	[-0.30,1.65]	[-5.13,-0.84]	[0.00,0.00]	[-5.55,-0.71]
Commitment × Mac	0.494	-0.276	-0.306	1.158	-1.600^{+}
	[-0.61,1.60]	[-1.31,0.75]	[-1.44,0.83]	[-0.54,2.86]	[-3.30,0.10]
Age	-0.0142	-0.00653	0.0264	0.0618*	-0.00330
	[-0.05,0.02]	[-0.03,0.02]	[-0.01,0.06]	[0.01,0.12]	[-0.05,0.05]
Female	-0.606^{+}	-0.190	-0.253	-0.0520	-0.147
	[-1.29,0.08]	[-0.77,0.39]	[-1.02,0.51]	[-1.29,1.19]	[-1.22,0.93]
SeBIS Updating subscale	2.223***				
	[1.27,3.18]				
SeBIS Password Generation subscale		0.489*			
		[0.10,0.88]			
SeBIS Proactive Awareness subscale			-0.891**	-0.269	-1.041*
			[-1.54,-0.24]	[-1.35,0.82]	[-1.95,-0.13]
Back up manually			0.624		
			[-0.14,1.39]		
Constant	-1.210	-0.468	-1.927*	-3.814**	0.505
	[-2.68,0.26]	[-1.71,0.77]	[-3.45,-0.40]	[-6.26,-1.37]	[-1.63,2.64]
Including those, who do not back			Yes	Yes	No
Including those, who back up manually			Yes	No	Yes
N	297	234	192	62	100
χ^2	52.38	16.97	38.24	11.33	20.44

95% confidence intervals in brackets. $^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$