A Promise Is A Promise

The Effect Of Commitment Devices On Computer Security Intentions

Alisa Frik
International Computer Science
Institute (ICSI)
University of California, Berkeley
afrik@icsi.berkeley.edu

Nathan Malkin
University of California, Berkeley
nmalkin@cs.berkeley.edu

Marian Harbach
ICSI
mharbach@icsi.berkeley.edu

Eyal Peer
Hebrew University of Jerusalem
eyal.peer@mail.huji.ac.il

Serge Egelman
ICSI
University of California, Berkeley
egelman@cs.berkeley.edu

ABSTRACT

Commitment devices are a technique from behavioral economics that have been shown to mitigate the effects of present bias—the tendency to discount future risks and gains in favor of immediate gratifications. In this paper, we explore the feasibility of using commitment devices to nudge users towards complying with varying online security mitigations. Using two online experiments, with over 1,000 participants total, we offered participants the option to be reminded or to schedule security tasks in the future. We find that both reminders and commitment nudges can increase users’ intentions to install security updates and enable two-factor authentication, but not to configure automatic backups. Using qualitative data, we gain insights into the reasons for postponement and how to improve future nudges. We posit that current nudges may not live up to their full potential, as the timing options offered to users may be too rigid.

CCS CONCEPTS

• **Security and privacy** → Usability in security and privacy; • **Human-centered computing** → Empirical studies in HCI;

KEYWORDS

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

ACM Reference Format:

Alisa Frik, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. 2019. A Promise Is A Promise: The Effect Of Commitment Devices On Computer Security Intentions. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300834>

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. That is, a “human-in-the-loop” would not be necessary [10]. 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 [31]. Recent qualitative studies have attempted to explain why these security precautions are often resisted by end-users [e.g., 48]. 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, leaving these systems vulnerable to attack. This effect has previously been observed for smartphone locking [14], as well as applying software updates [e.g., 21, 36–38, 58, 59, 62].

While the secondary nature of security mitigations is relatively well documented and a variety of approaches have been explored to overcome this (see, e.g., [22], for a survey

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300834>

of the usable security literature), few have examined ways to actually solve this problem using theories from behavioral economics. 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.¹ Thus, in the interim, new methods of nudging users towards engaging with security are needed. In this paper, we present the results of two online studies that explore nudges 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 [34, 46] and therefore to prefer immediate gratification over delayed utility. Acquisti has shown this to impact privacy decision making [1].

Recent research on decision making has identified techniques for overcoming present bias [43], one of which is the use of *commitment devices* [9]. 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. Commitment devices have been shown to be effective at changing behaviors, such as curbing procrastination, saving more for retirement, and donating to charity [5, 8, 54].

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., pre-committing 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 work we apply the principles of behavioral economics to test the effectiveness of reminders and commitment nudges at improving users' intentions to engage in security behaviors. We draw the attention of the usable security community to the need for reducing users' procrastination, in addition to increasing the overall compliance, as delayed security actions increase vulnerability. We performed two experiments with over 1,000 participants to examine the circumstances under which commitment nudges induce a

behavioral intent to improve security behaviors. 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 [31, 48]: applying system updates, enabling two-factor authentication, and configuring automatic backups. Study 1 shows that a commitment nudge (scheduling) can reduce the intention to ignore the request to enable automatic updates by 12%, and a reminder can reduce such intentions by 57%, for users who do not have automatic updates enabled. Study 2 shows that adding a reminder option reduced the willingness to ignore security updates by more than 40% for both Windows and Mac users. Adding an option to commit to installing the update in the future reduced stated ignore rates by about 25% for both Windows and Mac users. While reminders and commitments were not effective at promoting the use of automatic backup tools, they increased the willingness to enroll in two-factor authentication (2FA): reminders halved the intention to ignore (both 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 BACKGROUND AND RELATED WORK

Recent work in computer security has examined ways to improve user compliance with computer security mechanisms through better comprehension and usability of notices and controls [18, 19, 25, 50, 52]; advice on strong password composition [16, 51, 57]; use of memory-augmentation tools, such as password managers [24, 30]; and deployment of behavioral nudges [2, 4, 60]. Yet, even when users understand the importance of good security behaviors, they still do not always act accordingly [31, 48].

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 [32, 41]. However, in practice people often avoid, delay, or skip updating their software [37, 38, 55, 58]. Users often have very rational reasons for declining to perform this important security activity [17, 36, 37, 58]. Therefore, notifications alone, even when designed well, and when noticed and understood, are not always enough to trigger the desired behavioral change. Research on the so-called “security paradox” suggests that people report high computer security concerns and state that they want to remain secure [56], yet are often resistant to performing the necessary actions [31, 45, 63]. Some studies argue that this is because security is seen as a secondary task, often interfering with a primary task [12, 13, 64]. We hold that this is an example of present bias.

¹We do not attempt to define what this set of security actions might be.

Present Bias and Commitment Devices

People generally prefer immediate gratification over delayed utility, and therefore discount future outcomes [34, 46]. As a result, individuals often anticipate rewards and delay costs, delaying the activities that require salient costs and expediting the activities that presume salient benefits [42]. This complicates inter-temporal choices, when costs and benefits happen at different moments in time. Saving for retirement is a classic example: 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 income “today” always exceeds its “tomorrow” utility, the perfectly present-biased person saves nothing. Similar situations occur when making security decisions: security actions often interrupt workflows (cost) to protect against future dangers (benefit).

Commitment devices are mechanisms for overcoming present bias [43], which represent sophisticated attempts at self-control by limiting access (e.g., buying smaller packages of sweets), increasing sunk costs (e.g., purchasing an annual gym membership), or setting up clear promises (e.g., college savings accounts). In a broad sense, a commitment device is “an arrangement entered into by an individual with the aim of helping fulfill a plan for future behavior that would otherwise be difficult” [9]. Various forms of commitment devices [9, 33] were shown effective in triggering the behavior change against procrastination, for instance, in charitable giving [8], savings [54], and adherence to deadlines [5].

Present Bias in the Security Domain

Present bias in the security domain is related to the dominance of the user’s primary task over security protective tasks. That is, security is almost never a primary task [12, 13, 64]; 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 [48]. Economically speaking, the value of current needs exceeds the value of future needs.

In a study of why smartphone users do not securely lock their device screens (e.g., with a PIN), several participants indicated a desire to do so, but were asked at inconvenient times, so declined in the moment, thereby leaving their devices in insecure states [14]. Similarly, when examining why 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 [21].

Figure 1: Windows 10 scheduling nudge that allows the user to install now, choose a future time, or be asked again.

Some present bias can be eliminated by automating security tasks, thereby taking them out of users’ hands: automatic software updates increase installation rates and improve computers’ attack immunity [23, 40]. However, forced updates exogenously transposition the order of user tasks, preventing the user from continuing a primary activity. This naturally leads to confusion, irritation, and dissatisfaction [37, 38, 59], such as the waves of indignation that follow automatic updates in Windows [27]. Apart from disrupting users’ workflows at potentially critical moments, automatic updates may undermine user trust in the long term [36, 49]. For example, one-third of the participants in a study by Mathur and Chetty [36] had disabled automatic updates, and these users were more likely to have had past negative experiences. Moreover, keeping the user out of the loop removes control and leads to further divergence of mental models [61, 62].

Many systems now give users an opportunity to delay the task until it is more convenient. For example, Mac OS gives users the option to be reminded to apply software updates in the future. However, these reminders may also be seen as annoyances and therefore may not always be effective in reducing present bias. An alternative approach is to schedule the software update for a certain time in the future. Similar to the results by Ariely and Wertenbroch [5], scheduling could play the role of a commitment device and lead to fewer delays and higher compliance rates. Recently, Windows introduced such a scheduling feature (Figure 1).

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.

3 STUDY 1: ENABLING AUTOMATIC UPDATES

The goal of our first study was to understand how present bias affects people’s security behavior intentions. We hypothesized that when given the option to either act in the moment or not at all, many users will likely choose the latter

(a) Windows

(b) Mac OS

Figure 2: Commitment devices in Windows and Mac OS allow the user to be reminded in the future.

due to present bias. However, when given an option to reconsider the decision in the future, we hypothesized that fewer users will outright decline the suggested behavior, instead opting to revisit the decision at some point in the future.

For this first exploratory study, we decided to focus on a single security behavior: enabling automatic updates. This choice was inspired by the fact that recent versions of Windows (Figure 2a) and Mac OS (Figure 2b) present users with similar options; however, their empirical effectiveness is not known to the scientific community.

To answer this question, we sampled 300 participants (42.3% females, mean age = 34.16, SD = 11.15) from Prolific Academic,² an online research participant recruitment platform.³ 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.” Next, we asked them: “Assuming you don’t have automatic updates enabled on your computer currently, what would you do next?”

²<https://www.prolific.ac/>

³For both studies described in this paper, all authors obtained the necessary IRB and ethics approvals from their respective institutions.

The response options in our hypothetical scenarios varied according to three randomly-assigned conditions. In the control group, the options were to either “ignore the message” or “enable automatic updates.” The *Commitment* condition provided a third option, to “set auto-updates to be enabled one week from today,” while the *Reminder* condition provided a third option to be “reminded again in a week” (in contrast to committing 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.

Figure 3: The number of participants who chose to enable auto-updates now, be reminded or commit later, or ignore enabling auto-updates altogether, based on their current use of auto-updates and their randomly-assigned condition.

Results

Figure 3 shows the percentage of participants per condition who chose to ignore the message, reported a willingness to enable auto-updates now or in the future (by committing or setting a reminder). About 53% of participants reported having auto-updates already enabled. Because this question significantly interacted with the condition ($\chi^2(4) = 45.73, p < .001$), we used it to split our analyses.

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

For participants who already had auto-updates enabled, 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 *Commitment* 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 having auto-updates disabled (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 to participants who claimed to already have auto-updates enabled (14.5% chose it compared to 26.2% who chose the commitment option). Among these participants, the reminder option did not reduce the percentage of those who were willing to update now.

Summary. The results of this study show evidence for both present bias and the potential promise of commitment devices. Specifically, our results suggest that committing to enable auto-updates in the future could be an attractive option, sometimes even more so than setting a simple reminder. Overall, 18.9% of all participants in our hypothetical scenario 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. While intentions may not always translate into actual behavior, our goal in this research is studying the relative effectiveness of the nudges; see the Discussion for an in-depth analysis of this question and the advantages and disadvantages of our approach.

4 STUDY 2: UPDATES, BACKUP, AND 2FA

Study 1 provided initial evidence that techniques from behavioral economics have the potential to reduce present bias in security decision making. These results prompted us to formulate additional research questions, in order to dig deeper into the domains and circumstances when the nudges would work best, as well as other factors that may affect their relative effectiveness:

- (1) Can reminder and commitment nudges be effective across various security behavior scenarios?

- (2) Will participants become less likely to ignore the security recommendation if the point in time when the event takes place better fits their schedules?
- (3) Do these effects vary based on participants’ operating systems (i.e., Mac vs. Windows users)?

To answer these questions, we designed several nudges that may improve several end-user security behaviors that experts agree are important [e.g., 31], 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 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 included two additional security behavior scenarios, the timing options were not predefined (participants proposed the time for reminder or future installation using an open-text field), people already engaged in these behaviors were screened out, and we also distinguished between Windows and Mac users.

Method

We deployed a 3 (Control vs. Reminder vs. Commitment) x 3 (Update vs. Backups vs. 2FA) between-subject design. All experimental participants were either Mac OS or Windows users. We asked participants to imagine that after finishing this study, they received a message on their computer screen (Figure 4).⁴ We then 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.”⁵

The response options varied according to three randomly-assigned conditions. In the control group, the options were either to (1) ignore the message or (2) install updates, enable automatic backups, or enable 2FA. In the *Commitment* condition, the options were identical to the control, with a third option to pick a future time to install updates, enable automatic backups, or enable 2FA. In the *Reminder* condition,

⁴In all conditions, full-size screenshots were shown to the participants.

⁵We chose to use this account for the 2FA scenario because we were sure that all subjects in our population have it, and because it contains personal and financial information. For some participants, it may even be their main source of income.

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)
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)

(a) Update scenario, control condition, Mac. **(b) Backups scenario, Commitment condition, Windows.** **(c) 2FA scenario, Reminder condition, Windows.****Figure 4: Example messages shown to Study 2 participants.**

this third option was replaced with a request to be reminded in the future; 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.

Next, we asked respondents to explain, 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 [26] and the Security Behavior Intentions Scale (SeBIS) [15].

Results

We recruited 903 Mac and Windows users from Amazon Mechanical Turk (MTurk),⁶ and randomly assigned them to one of the 9 experimental conditions (Table 1). We told them that the study was 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 had 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 behavior change. The resulting sample included 734 participants (53% female; aged 19–84, mean = 37.78, SD = 12.12).

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. The binary dependent variable represented whether respondents chose "Don't ask me again" (Table 2). We included age, gender, and the 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 were less than 5), to compare the ratio of specific choices in each condition.

Update scenario. Regression coefficients (Table 2) demonstrate that introduction of the nudges, either the reminder or the commitment to acting in the future, significantly reduces the proportion of people willing to outright ignore the message, especially among Mac-using participants. Not surprisingly, people with positive security updating intentions, measured by the SeBIS Updating subscale, are less likely to choose to ignore the update message.

Figure 5 depicts 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 ($\chi^2(2) = 39.71, p < .0005$) and from 60% to 15% among Mac users ($\chi^2(2) = 57.61, p < .0005$), compared to the control. In

⁶We limited participation to subjects in the United States who had completed at least 500 tasks with an approval rate of at least 95%.

Table 2: Logit regression on respondents' choices for "Don't ask me again" (0 - no, 1 - yes).

	Update	2FA	Backups
Control × Windows		(baseline)	
Control × Mac	-1.434**	-0.146	-0.202
Reminder × Windows	-2.430***	-1.774***	-0.647
Reminder × Mac	-3.819***	-1.530**	-1.944**
Commitment × Windows	-1.673***	-1.628**	1.019
Commitment × Mac	-2.552***	-0.400	-1.103
Age	0.0323*	0.00961	0.00264
Female	-0.00696	-0.130	-0.346
SeBIS Updating	-1.350***		
SeBIS Password Generation		-0.571**	
SeBIS Proactive Awareness			1.340***
Back up manually			-1.219***
Constant	0.580	0.489	1.398
N	297	234	192
χ^2	87.13	37.30	50.55

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 5: Response distribution in the Update scenario.

the presence of the "Pick a time" option, these proportions decreased to 57% for Windows users ($\chi^2(2) = 21.68, p = .0062$) and 34% for Mac users ($\chi^2(2) = 29.32, p = .0101$).

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. Additionally, they are more favorable to the reminders than Windows users: they chose the "Remind me later" option more often than Windows users ($\chi^2(2) = 7.54; p = .0063$).

Thus, both the Reminder and Commitment nudges were effective at increasing *willingness* to install updates. While a stronger effect was observed in the Reminder condition, the

Figure 6: Response distribution in the Backups scenario. **Figure 7: Response distribution in the 2FA scenario.**

differences between the Reminder and Commitment conditions are unclear for behavior change: 100% of the users who *commit* to the behavior will ultimately perform it, whereas a non-zero number of those being reminded may ask to be reminded *ad infinitum*—effectively never performing the behavior—or may ignore and permanently dismiss a future reminder.

Backup scenario. The regression model (Backups in Table 2) revealed that the Reminder nudge is effective in decreasing the willingness to dismiss the message about the automatic backup tool for Mac users. However, no other significant effects were discovered (Figure 6). Therefore, we conclude that the examined commitment nudges are not effective at increasing willingness to configure automatic backups.

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 2). 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 ($p = .529$; $\chi^2(2) = 8.28$; $Pr = .016$). 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 matches our hypothesis as, unlike the other two scenarios, this scenario was platform-independent: it discussed only an online account, and the message itself was presented in a browser window.

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 outright dismissing it). Reminders demonstrated a greater potential for 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.

Decreasing Immediate Action

In most conditions, we noticed a “side effect” of the nudges: although they decreased the overall negative response rate (i.e., the proportion of “Don’t ask me again” choices), they also drove the proportion of “now” choices down. That is, the option to delay was chosen by some participants who otherwise would have chosen to act immediately. As we mentioned in Section 4, 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*” (P639).

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, encouraging users’ choices may be a double-edged sword: while the additional options resulted in decreased tendency to dismiss messages in all scenarios, it is clear that the option to delay encourages some to procrastinate, when they would otherwise act immediately. That said, while we cannot know how many who choose reminders will ultimately act in the future upon being reminded, we do know that all who choose commitments will eventually act. Thus, while we took a conservative approach in our analysis by using the rates of outright dismissals as the dependent variables, an alternative approach may be to examine the rates at which people choose to either commit immediately or at some point in the future (combining the “now” and “later” choices in the *Commitment* condition, because both of these groups are ultimately choosing to act).

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 the “install now” option is easier and more attractive than the delaying option) could be one way to mitigate the negative impact of nudges on the immediate action options. Additionally, longitudinal studies are needed to examine what proportion of “remind me later” users ultimately decide to perform the recommended security action.

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 use an open-ended text field to specify when they would like to receive a reminder or commit to performing the action. When being able to select a time that fits (Study 2), instead of choosing from a predefined set of options (Study 1), more participants chose the reminder and commitment options. However, there are some confounds, as the two studies were performed at different times, on different samples, using slightly varying protocols.

However, anecdotally, inspection of the open-ended responses 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, for instance:

- “When I’m done using my PC for the day”
- “Next time I log in”

Among the responses that specified a concrete time, the most popular response in all scenarios was “tomorrow.” The reasons for this may be endogenous or a learned preference: they may be accustomed to offering “tomorrow” as a procrastination tactic in a wide variety of scenarios. The second most popular suggested delay was in the “evening,” which is presumably when users expect to be finished with the

computer for the day. While participants in the Updates and Backups scenarios did not suggest delaying the actions for more than one week, several participants proposed to delay 2FA enrollment for up to 3 months.

These findings have several practical implications for future nudge designs. First, when people are asked to find a specific slot to schedule an action or reminder, they may simply have no suitable time in mind. As a result, showing them too many time-related options may confuse and annoy them, leading to a decision to dismiss the message altogether. A conditional—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. For example, the system might offer the option to apply the update “after the computer has not been used for at least one hour.”

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 can provide insights into how to best design these intuitive options and account for contextual factors [e.g., 7, 11, 20, 28, 29, 35, 44, 47, 53], which is a subject for future work.

Reasons for Delaying

Based on the open-ended responses, which were independently coded by two researchers, 83% of participants in the Update scenario chose to delay the action because of inconvenient timing. In the 2FA and Backup scenarios, this rate was lower (62% and 56%, respectively), which partially explains why our nudges were most effective in the Updates scenario: because we designed them to mitigate inter-temporal biases, they worked in the scenario where present bias was the most salient reason to delay. Other reasons, not related to timing—such as lack of trust in the notification, reduced awareness of security risks and benefits of suggested mitigations, expected computer performance deterioration, limited resources (e.g., disk space), and annoyance—appear to be strong predictors of the decision to completely *dismiss* the notification rather than to delay action.

Moreover, 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 to the average user and therefore elicit low willingness to enable them for reasons not related to timing. For instance, a perceived lack of

importance of local files, and therefore expressing low interest in protecting them, may also lead to users’ unwillingness to install backup tools (e.g., “*I don’t have anything important enough on my personal computer to back up,*” [P461]).

Additionally, a few participants’ comments hint at negative prior experiences with spam-like messaging that prevented them from complying with the request for enabling automatic backups. One participant reported that the dialogue seemed “*almost like an advertisement. I’d have to scan my PC [for] viruses afterwards,*” (P329).

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. For instance, our findings suggest that commitment nudges have potential for increasing compliance, especially when users are looking for a more convenient time to do it or need additional time to decide. They appear to be less effective when users have more fundamental doubts about the security recommendations; in these cases, other mitigations, beyond commitment devices, may be more appropriate. For instance, more information facilitating decision-making could reduce the perceived lack of necessity. Better communication of security risks and benefits of protection behaviors may also be more effective in certain cases.

Mac and Windows Users Behave Differently

Finally, we noticed that under most conditions, Windows users were more likely to choose to dismiss the security recommendation messages than Mac users. As speculated in Section 4, Windows seems to show more notification dialogues and may hence cause greater levels of habituation and/or fatigue with regard 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.

Our qualitative data also supports the idea that Windows users trust Microsoft less than Mac users trust Apple. For instance, one participant said that he believes that “*malware ... doesn’t often happen on a Mac,*” (P393). Another participant believed that usually Apple products are not “*getting attacked by viruses. This is why I own an Apple,*” (P279). Future work should explore whether different messaging strategies are needed to target users of different platforms.

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 actual behavior that was the consequence of our nudges. However, a number of factors suggest that behavior would be likewise impacted. First, prior research has shown that intentions are a precursor to behavior change [3, 39], with especially strong predictive

power for immediate and non-contingent actions [6], like the “now” and commitment choices in our study. Second, 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. Third, while we acknowledge that intentions may be overestimated with respect to actual behavioral rates, in this paper we focus on the comparison of *relative* effectiveness of the nudges. Therefore, although 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. Finally, we would like to emphasize the advantage (and even, in our view, important prerequisite) of running *hypothetical* studies in the early stages of designing and testing a large variety of messages in a safe environment. Despite the positive intention to improve 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.

In this paper we demonstrated through a controlled internally valid study that user intentions can in fact be swayed. We are now in the process of designing a follow-up longitudinal and externally valid field study to examine in more realistic settings whether this leads to actual behavior change. Because participants recruited via online crowdsourcing platforms may not be fully representative, in future work we will engage a wider population of computer users to ensure the generalizability of our results.

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.

6 CONCLUSIONS

We performed an online study to test the effectiveness of reminders and commitment nudges in improving users’ intentions to engage in security behaviors by reducing the present bias effect. As a first step in exploring this application of behavioral economics to the computer security domain, we found that both the Reminder and Commitment nudges have the potential to increase willingness to engage in beneficial computer security behaviors by up to 85%. However, at the same time, introducing nudging options decreased the fraction of users who reported a willingness to take immediate action. In sum, our results suggest that offering people the opportunity to take action at a later time may increase compliance with security mitigations, but people will also procrastinate, when given the opportunity.

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 willingness to enable automatic backups. Furthermore, we posit that current nudging dialogues may not live up to their full potential, if the timing options offered to users are too rigid.

ACKNOWLEDGMENTS

This work was made possible by the U.S. National Science Foundation through grants CNS-1528070 and CNS-1817249, the U.S.–Israel Binational Science Foundation through grants 2014626 and 2017751, the Center of Long-Term Cybersecurity (CLTC) at U.C. Berkeley, as well as feedback from Arunesh Mathur and Refjohürs Lykkewe.

REFERENCES

- [1] Alessandro Acquisti. 2004. Privacy in Electronic Commerce and the Economics of Immediate Gratification. In *Proceedings of the ACM Electronic Commerce Conference (EC '04)*. ACM Press, New York, NY, 21–29.
- [2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for privacy and security: Understanding and assisting users’ choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 44.
- [3] Icek Ajzen. 1991. The theory of planned behavior. *Organizational behavior and human decision processes* 50, 2 (1991), 179–211.
- [4] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. 2014. *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.
- [5] Dan Ariely and Klaus Wertenbroch. 2002. Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological science* 13, 3 (2002), 219–224.
- [6] Richard P Bagozzi. 1992. The self-regulation of attitudes, intentions, and behavior. *Social psychology quarterly* (1992), 178–204.
- [7] Peter Bogunovich and Dario Salvucci. 2011. 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)*. ACM, New York, NY, USA, 3123–3126. <https://doi.org/10.1145/1978942.1979404>
- [8] Anna Breman. 2011. Give more tomorrow: Two field experiments on altruism and intertemporal choice. *Journal of Public Economics* 95, 11 (2011), 1349–1357.
 - [9] Gharad Bryan, Dean Karlan, and Scott Nelson. 2010. Commitment devices. *Annual Review of Economics* 2, 1 (2010), 671–698.
 - [10] Lorrie Faith Cranor. 2008. A Framework for Reasoning about the Human in the Loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*. USENIX Association, Berkeley, CA.
 - [11] Laura Dabbish, Gloria Mark, and Víctor M. González. 2011. 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)*. ACM, New York, NY, USA, 3127–3130. <https://doi.org/10.1145/1978942.1979405>
 - [12] Paul Dourish, E. Grinter, Jessica Delgado de la Flor, and Melissa Joseph. 2004. Security in the Wild: User Strategies for Managing Security As an Everyday, Practical Problem. *Personal Ubiquitous Comput.* 8, 6 (Nov. 2004), 391–401. <https://doi.org/10.1007/s00779-004-0308-5>
 - [13] W. Keith Edwards, Erika Shehan Poole, and Jennifer Stoll. 2008. Security Automation Considered Harmful?. In *Proceedings of the 2007 Workshop on New Security Paradigms (NSPW'07)*. ACM, New York, NY, USA, 33–42. <https://doi.org/10.1145/1600176.1600182>
 - [14] Serge Egelman, Sakshi Jain, Rebecca S. Portnoff, Kerwell Liao, Sunny Consolvo, and David Wagner. 2014. 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)*. ACM, New York, NY, USA.
 - [15] S. Egelman and E. Peer. 2015. 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)*. ACM, New York, NY, USA.
 - [16] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. 2013. 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*. ACM, 2379–2388.
 - [17] Michael Fagan, Mohammad Maifi Hasan Khan, and Ross Buck. 2015. A study of users' experiences and beliefs about software update messages. *Computers in Human Behavior* 51 (2015), 504–519.
 - [18] Adrienne Porter Felt, Serge Egelman, Matthew Finifter, Devdatta Akhawe, and David Wagner. 2012. How to ask for permission. In *Proceedings of the 7th USENIX conference on Hot Topics in Security (HotSec'12)*. USENIX Association, Berkeley, CA, USA, 7–7. <http://dl.acm.org/citation.cfm?id=2372387.2372394>
 - [19] Adrienne Porter Felt, Robert W. Reeder, Hazim Almuhiemedi, and Sunny Consolvo. 2014. Experimenting at Scale with Google Chrome's SSL Warning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2667–2670. <https://doi.org/10.1145/2556288.2557292>
 - [20] James Fogarty, Jennifer Lai, and Jim Christensen. 2004. Presence Versus Availability: The Design and Evaluation of a Context-aware Communication Client. *International Journal of Human-Computer Studies* 61, 3 (Sept. 2004), 299–317. <https://doi.org/10.1016/j.ijhcs.2003.12.016>
 - [21] Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang. 2016. 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)*. USENIX Association, Denver, CO, 97–111. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/forget>
 - [22] Simson Garfinkel and Heather Richter Lipford. 2014. *Usable Security: History, Themes, and Challenges*. Morgan & Claypool, 124+ pages. <https://doi.org/10.2200/S00594ED1V01Y201408SPT011>
 - [23] Christos Gkantsidis, Thomas Karagiannis, and Milan Vojnović. 2006. Planet scale software updates. *ACM SIGCOMM Computer Communication Review* 36, 4 (2006), 423–434.
 - [24] Eric Grosse and Mayank Upadhyay. 2013. Authentication at scale. *IEEE Security & Privacy* 11, 1 (2013), 15–22.
 - [25] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. 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)*. ACM, New York, NY, USA, 2647–2656. <https://doi.org/10.1145/2556288.2556978>
 - [26] Eszter Hargittai and Yuli Patrick Hsieh. 2012. Succinct survey measures of web-use skills. *Social Science Computer Review* 30, 1 (2012), 95–107.
 - [27] S. Hollister. 2017. Microsoft won't fix the most frustrating thing about Windows. Cnet. <https://www.cnet.com/news/microsoft-windows-10-forced-updates/>.
 - [28] Scott Hudson, James Fogarty, Christopher Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny Lee, and Jie Yang. 2003. 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)*. ACM, New York, NY, USA, 257–264. <https://doi.org/10.1145/642611.642657>
 - [29] Christophe Hurter, Benjamin R. Cowan, Audrey Girouard, and Nathalie Henry Riche. 2012. 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)*. British Computer Society, Swinton, UK, UK, 99–108. <http://dl.acm.org/citation.cfm?id=2377916.2377928>
 - [30] Alexa Huth, Michael Orlando, and Linda Pesante. 2012. Password security, protection, and management. *United States Computer Emergency Readiness Team* (2012).
 - [31] Iulia Ion, Rob Reeder, and Sunny Consolvo. 2015. "...No One Can Hack My Mind": Comparing Expert and Non-Expert Security Practices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 327–346. <https://www.usenix.org/conference/soups2015/proceedings/presentation/ion>
 - [32] Moazzam Khan, Zehui Bi, and John A Copeland. 2012. Software updates as a security metric: Passive identification of update trends and effect on machine infection. In *Military Communication Conference 2012*. IEEE, 1–6.
 - [33] Alexander K Koch and Julia Nafziger. 2011. Self-regulation through Goal Setting. *The Scandinavian Journal of Economics* 113, 1 (2011), 212–227.
 - [34] David Laibson. 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112, 2 (1997), 443–478.
 - [35] Brian Y. Lim, Oliver Brdiczka, and Victoria Bellotti. 2010. 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)*. ACM, New York, NY, USA, 263–272. <https://doi.org/10.1145/1880071.1880115>
 - [36] Arunesh Mathur and Marshini Chetty. 2017. Impact of User Characteristics on Attitudes Towards Automatic Mobile Application Updates. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 175–193. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/mathur>
 - [37] Arunesh Mathur, Josefina Engel, Sonam Sobti, Victoria Chang, and Marshini Chetty. 2016. "They Keep Coming Back Like Zombies": Improving Software Updating Interfaces. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 43–58. <https://www.usenix.org/conference/soups2016/>

- technical-sessions/presentation/mathur
- [38] Arunesh Mathur, Nathan Malkin, Marian Harbach, Eyal Peer, and Serge Egelman. 2018. Quantifying Users' Beliefs About Software Updates. *arXiv preprint arXiv:1805.04594* (2018).
 - [39] Daniel E Montano and Danuta Kasprzyk. 2015. Theory of reasoned action, theory of planned behavior, and the integrated behavioral model. *Health behavior: Theory, research and practice* (2015), 95–124.
 - [40] Antonio Nappa, Richard Johnson, Leyla Bilge, Juan Caballero, and Tudor Dumitras. 2015. 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*. IEEE, 692–708.
 - [41] Kartik Nayak, Daniel Marino, Petros Efstathopoulos, and Tudor Dumitras. 2014. Some vulnerabilities are different than others. In *International Workshop on Recent Advances in Intrusion Detection*. Springer, 426–446.
 - [42] Ted O'Donoghue and Matthew Rabin. 1999. Doing it now or later. *American Economic Review* (1999), 103–124.
 - [43] Ted O'Donoghue, Matthew Rabin, et al. 2006. Incentives and self-control. *Econometric Society Monographs* 42 (2006), 215.
 - [44] Tadashi Okoshi, Julian Ramos, Hiroki Nozaki, Jin Nakazawa, Anind K. Dey, and Hideyuki Tokuda. 2015. 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)*. ACM, New York, NY, USA, 475–486. <https://doi.org/10.1145/2750858.2807517>
 - [45] Pew Research Center. 2017. *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>.
 - [46] Edmund S Phelps and Robert A Pollak. 1968. On second-best national saving and game-equilibrium growth. *The Review of Economic Studies* 35, 2 (1968), 185–199.
 - [47] Martin Pielot, Bruno Cardoso, Kleomenis Katevas, Joan Serrà, Aleksandar Matic, and Nuria Oliver. 2017. Beyond Interruptibility: Predicting Opportune Moments to Engage Mobile Phone Users. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 91 (Sept. 2017), 25 pages. <https://doi.org/10.1145/3130956>
 - [48] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2016. 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)*. ACM, New York, NY, USA, 666–677. <https://doi.org/10.1145/2976749.2978307>
 - [49] Eric Rescorla. 2003. Security holes... Who cares?. In *USENIX Security Symposium*. Washington, DC, 75–90.
 - [50] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* 21, 3 (2017), 70–77.
 - [51] Richard Shay, Patrick Gage Kelley, Saranga Komanduri, Michelle L. Mazurek, Blase Ur, Timothy Vidas, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2012. Correct horse battery staple: Exploring the usability of system-assigned passphrases. In *Proceedings of the eighth symposium on usable privacy and security*. ACM, 7.
 - [52] Joshua Sunshine, Serge Egelman, Hazim Almuhammedi, Neha Atri, and Lorrie Faith Cranor. 2009. Crying wolf: an empirical study of SSL warning effectiveness. In *Proceedings of the 18th USENIX Security Symposium (SSYM'09)*. USENIX Association, Berkeley, CA, USA, 399–416. <http://dl.acm.org/citation.cfm?id=1855768.1855793>
 - [53] Dan Tasse, Anupriya Ankolekar, and Joshua Hailpern. 2016. 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)*. ACM, New York, NY, USA, 3324–3334. <https://doi.org/10.1145/2858036.2858174>
 - [54] Richard H Thaler and Shlomo Benartzi. 2004. Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy* 112, S1 (2004), S164–S187.
 - [55] Yuan Tian, Bin Liu, Weisi Dai, Blase Ur, Patrick Tague, and Lorrie Faith Cranor. 2015. 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*. ACM, 51–61.
 - [56] Unisys. 2017. *Unisys security index*. Technical Report. Accessed [11 April 2018]: <http://www.unisys.com/unisys-security-index/us>.
 - [57] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L. Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, et al. 2012. How does your password measure up? The effect of strength meters on password creation.. In *USENIX Security Symposium*. 65–80.
 - [58] Kami Vaniea and Yasmeen Rashidi. 2016. Tales of software updates: The process of updating software. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3215–3226.
 - [59] Kami E Vaniea, Emilee Rader, and Rick Wash. 2014. Betrayed by updates: how negative experiences affect future security. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2671–2674.
 - [60] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2367–2376.
 - [61] Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM, 11.
 - [62] Rick Wash, Emilee Rader, Kami Vaniea, and Michelle Rizer. 2014. Out of the loop: How automated software updates cause unintended security consequences. In *Symposium on Usable Privacy and Security (SOUPS)*. 89–104.
 - [63] Rick Wash and Emilee J Rader. 2015. Too Much Knowledge? Security Beliefs and Protective Behaviors Among United States Internet Users.. In *SOUPS*. 309–325.
 - [64] Ryan West. 2008. The Psychology of Security. *Commun. ACM* 51, 4 (April 2008), 34–40. <https://doi.org/10.1145/1330311.1330320>