

GUIDELINES FOR ETHICAL NUDGING IN PASSWORD AUTHENTICATION

Karen Renaud* and Verena Zimmermann†

* *Division of Cyber Security, Abertay University, Scotland. k.renaud@abertay.ac.uk*

† *Technische Universität Darmstadt, Germany. zimmermann@psychologie.tu-darmstadt.de*

Abstract: Nudging has been adopted by many disciplines in the last decade in order to achieve behavioural change. Information security is no exception. A number of attempts have been made to nudge end-users towards stronger passwords. Here we report on our deployment of an *enriched nudge* displayed to participants on the system enrolment page, when a password has to be chosen. The enriched nudge was successful in that participants chose significantly longer and stronger passwords. One thing that struck us as we designed and tested this nudge was that we were unable to find any nudge-specific ethical guidelines to inform our experimentation in this context. This led us to reflect on the ethical implications of nudge testing, specifically in the password authentication context. We mined the nudge literature and derived a number of core principles of ethical nudging. We tailored these to the password authentication context, and then show how they can be applied by assessing the ethics of our own nudge. We conclude with a set of preliminary guidelines derived from our study to inform other researchers planning to deploy nudge-related techniques in this context.

Key words: nudge, ethics, autonomy

1. INTRODUCTION

The password is intended to be a secret shared exclusively between the password owner and the system the password controls access to. Passwords can leak for a number of reasons. They could be guessed, for example, especially if weak passwords are chosen, as users tend to do [1]. Leaked passwords permit unauthorised access to sensitive personal or organisational data in a way that is sometimes very hard to detect.

It is standard practice for organisations to offset this risk by expiring passwords on a regular basis [2, 3]. The rationale is that this curtails ongoing use of leaked passwords, and also reveals dormant accounts, thereby improving overall system security.

End users seldom contemplate password replacement with enthusiasm. They might well react to password expiry by choosing weak passwords. This, then, appears to justify password expiry, leaving organisations with no choice but to continue enforcing regular password expiry. The end result is deterioration of system security, with ever weaker passwords [4, 5]. The eventual result is weaker, not stronger, password defences.

Some have suggested a middle ground, where the expiry requirement is directly proportional to the strength of the password chosen by the user [6, 7, 8]. Under this scheme, weak passwords would expire much more quickly than strong passwords. This scheme rewards end users for strong passwords by allowing them to use the password for longer, thereby amortising the effort they put into formulating and remembering it. It also satisfies the organisation's need for measures that protect their systems.

We trialled this scheme, using an *enriched nudge* to ensure that the end users were aware of the variable expiry period, and to ensure that they were aware of being “nudged” towards stronger passwords. We wanted them to be aware of the extended expiry time scheme. Our trial proved efficacious: participants chose significantly longer and stronger passwords.

In carrying out the research we were somewhat perturbed by the fact that we were not able to find nudge-specific ethical guidelines to inform our experimentation in this context. By “ethics”, we mean “*Moral principles that govern a person's behaviour or the conducting of an activity*” (OED). In this context, that refers to the way researchers should conduct nudge-related research.

We followed the generic BPS guidelines [9] to obtain ethical approval, as required by our institution, but we were left somewhat dissatisfied that we did not have more nuanced guidelines to apply. In this paper, we report on the deployment of our nudge, and we then consider how we could inform subsequent nudge trials by deriving a set of ethical guidelines specifically for this context.

The following section discusses the password expiry issue, presenting the *raison d'être* for the practice, and the consequences thereof. Then the enriched nudge, and the experiment carried out to assess its impact on end users, is presented (Section 4.) and the results reported (Section 5.). In Section 6. we report on a review we did of the ethical nudge literature, and conclude with a set of guidelines to inform and guide researchers in the password authentication discipline. Section 7. makes some recommendations about future directions for research. Section 8. concludes.

2. PASSWORD EXPIRY

It is common for organisations to require their employees to change passwords regularly,* and indeed this is often considered to be “good practice” [10, 11]. The reason for mandating password expiry is the belief that it improves security [12, 13].

Yet leading academics [4, 5, 14, 15], journalists [16], and standards bodies, such as NIST [17], are urging system administrators to rethink their traditional password expiry practices because it effectively weakens passwords in the long run and is not the cure-all many believe it to be.

2.1 The Burden

Passwords, in and of themselves, can impose an unacceptable burden on computer users [18]. Password expiry exacerbates this. Password expiration has two immediate consequences in terms of human behaviour, given the fact that users cannot amortise the effort involved in memorising a strong password over a long period of time.

The **first** consequence is that people are likely to drift towards ever weaker passwords with each successive change, perhaps incrementing a digit at the end of the password or appending the month and year, merely to offset the expense of memorising a new password each month [4, 5].

The **second** consequence is that they are more likely to record the password either on paper or digitally [19,20,21], because they dread forgetting it [22, 23].

Both of these consequences weaken the mechanism [24], the opposite of what password expiry is intended to achieve. It also undeniably undermines the user experience and often prevents users from accessing their accounts.

An expired password requires immediate action, perhaps when users have other urgent goals to satisfy. If they change the password in haste, merely to gain access to their account when there is some urgency, they are likely to forget the hastily-formulated password. They then have to go through the pain of password replacement, and for some systems this process is more arduous than for others. This can result in a loss of productivity and also financial cost related to help desk resources.

It is unsurprising that this leads to the use of weak passwords. People wish to avoid the pain and inconvenience of a replacement and act to prevent such an occurrence.

*<https://www.symantec.com/connect/articles/simplest-security-guide-better-password-practices>
<http://hitachi-id.com/documents/password-management-best-practices.php>
<https://technet.microsoft.com/en-us/library/ff741764.aspx>

2.2 The Expense

Password expiration is expensive for end-users, but also for organisations. Password forgetting is likely to lead to an extra number of help desk calls.

If each person calls only two times more a year, for whatever reason, consider what the cost would be. Figure 1 shows a compilation of data collected from Gartner, Forrester, and the META Group by Osper [25] using this tool: <http://www.mandylionlabs.com/PRCCalc/PRCCalc.htm>. The table represents a possible cost scenario if a company were to experience just two additional help desk calls per person and year, on average.

The assumption here is that the extra calls can be related to anything, but that 30% relate to passwords, regardless of the number of calls. The numbers make it clear that any number of extra help desk calls have very real financial implications.

	BEFORE	+2 Calls AFTER
Avg. # of Call per Year ¹	11	13
% Related to Password Reset ²	30%	30%
Total # Password Resets	3.3	3.9
Average Cost per Reset ³	\$ 25.00	\$ 25.00
Total Yearly Cost per Reset	\$ 82.50	\$ 97.50
Total # of Users	3,000	3,000
Total Yearly Cost	\$ 247,500	\$ 292,500
Worst Case Scenario - Lost Productivity		\$ (45,000)

Sources: ¹ Gartner Research ² Forrester Research ³ META Group

Figure 1: Cost scenario of just two additional password related helpdesk calls a year, on average [25]

This expenditure and undermined user experience might be warranted if expiry does indeed reduce the probability and duration of unauthorised access. This is only the case where the person who obtains the password plans to carry out long-term forays into the account. The reality is that most password thieves will carry out their nefarious activities as soon as they gain access to the password [4]. Password expiry, in these cases, is the equivalent of shutting the stable door after the horse has bolted. There are other ways of providing superior protection, and these are considered next.

2.3 The Threat

The specific threat that password expiry addresses is password leakage: it limits the period during which the leaked password can be used. Passwords can be leaked either deliberately or inadvertently (Table 1). Password expiry is not intended to address the former. If someone wishes to share a password he/she will simply share

the newly-chosen password when the previous password expires.

Category	Action
(1) Deliberate	Owner shares password [26]
(2) Password Choice	Weak Password Guessed [27]
	Encrypted Weak Password stolen and easily decrypted [28]
	Strong Password Observed During Entry [29] Record of Strong Password Discovered [19]
(3) Technical Failure	Theft during transmission or storage or due to a website bug [30, 31]
(4) Deception	Phishing [32]
	Vishing [33]

Table 1: Password Leakage Categories

Leakage is only a major issue if the credential owner is unaware of the leakage, as and when it happens. If he or she becomes aware of it, he/she can act quickly to change the password and thereby curtail access and prevent further damage.

If the password owner does not discover the leak, the thief can keep using the password as long as there are no detectable side effects. So, for example, if a Phisher bot gets hold of a password it will probably be used pretty much immediately. The legitimate credential owner is likely to observe the side effects of such usage, and will act to curtail access.

Sometimes credentials can be used merely to snoop, which might be the case if an email password has been stolen by an ex-partner, for example. In some cases credentials are used to leapfrog into the organisation's systems and infrastructures. In this case, the side effects will not necessarily be evident in the password owner's account. If there are no immediately-observable side effects, the thief can keep using the leaked password unhindered.

2.4 Alternatives

Although there is no empirical evidence that password expiry does indeed result in more secure systems [14], it remains "good practice" in industry. It might limit the damage that can be carried out by undiscovered password leakage. What it definitely does is make those responsible for security feel that they are "doing something" to counteract insecurities caused by poor password practice. That being so, it is unlikely that organisations will discontinue this practice without compelling evidence that doing so will not weaken the security of their systems.

There are other measures that could serve a similar purpose to password expiry. Some of these alternatives perform

better than expiry without the accompanying human cognitive cost. For example, the following techniques are examples of what could be used:

- **Notifications:** Some systems display an informative "last login" message whenever a person accesses their account to reveal illicit activity. This does serve to reveal password leakage if two pre-conditions are met: (1) the person notices the display and realises that he or she was not responsible for the access, (2) the legitimate user accesses the system frequently. Neither of these is, unfortunately, a given. Indeed, it is the infrequently-used account credentials that are most prized by hackers [34].
- **Alerts:** Other systems send an email to the legitimate user whenever anyone accesses the account. If the owner did not access the account, he or she is alerted and can act to terminate the hacker's access. This works unless the person accesses the system frequently, in which case it could become merely an annoyance and the emails would be ignored.
- **Multi-Password:** Some systems require an additional password whenever an action could have consequences. The hacker would then need to steal both passwords in order to carry out his/her nefarious activities, which makes things a bit harder but requires the user to memorize two passwords which increases effort and the possibility of forgetting a password.
- **Multi-Factor:** Some systems use two-factor authentication, requiring a token or biometric in addition to the password. This bolsters the mechanism and renders a single leakage less damaging. Still, tokens are costly and need to be carried around to be available upon log-in. Using biometrics as a second authentication mechanism is promising, but often also requires costly devices or sensors and has implications on privacy. Furthermore, several circumstances (e.g. age, injuries, hand cream, contact lenses) can prevent users from authenticating with biometrics [35]. The most popular fall-back mechanism then still is the password that is associated with the same problems as the multi-password approach.
- **Multi-Channel:** Some systems utilise a separate channel to authorise side-effect actions. So, for example, a hacker could steal a password and be able to log into someone's bank account to see their details. If he or she attempts to transfer money out of the account, a message is sent to the legitimate owner via another channel, perhaps a registered mobile phone, to authorise the transfer. In this way the hacker's purpose is revealed. This mechanism is often used for accounts people really care about, such as their bank accounts. While increasing the security

without having to remember further passwords, this mechanism requires the user to have access to both channels at the same time which isn't always a given, e.g. when paying with a credit card in a foreign country without having access to a mobile network so that they can receive SMS messages.

- **One-Time Passwords:** Some organisations issue their users with a bespoke device that generates one-time passwords. These expire immediately and leakage is rendered a non-threat. However, the inconvenience of carrying the device around, and the expense thereof, probably limits its applicability.
- **Expire Intelligently:** Schneider [36] suggests expiring passwords only when anomalous behaviour is detected on a user's account.

Table 2 summarises this discussion.

Alternative	Advantage	Disadvantage
Notifica-tions	Cheap	Habituation [37]; Only effective for frequently-used accounts
Alerts	Cheap	Habituation unless alerts only signal side-effect actions [38]; Habituation to receiving alerts could be exploited by Phishers
Multi-Pass-Word	Two Passwords have to be leaked; Improved Security	Expense: increases cognitive effort and doubles the number of password resets
Multi-Factor	Improved Security [39]	Expense; Not Scalable [40] Requires hardware on client that limits applicability; Biometrics have privacy implications
Multi-Channel	Improved Security [41]	Expense; Delays on some channels (eg. SMS texts) [42]
One-Time Password Device	High Security [43]	Expense and Inconvenience
Expire In-telligently	Reduces Expiry Burden	Still experimental (needs to be tested empirically)

Table 2: Password Expiry Alternatives

2.5 Status Quo Rationale

Many companies may consider these alternatives too complex or expensive, especially since they themselves

would have to carry the implementation cost. Enforcing password expiry pushes the cost onto the end user and seems cheaper, at first glance. The fact that users migrate towards weaker passwords and thereby compromise the security of their accounts might be a consequence they feel is a reasonable trade-off. They might argue that it is the weak password holder, him or herself, who has to face the consequences. They are probably not considering the possibility that access to a hacked account can be used to leapfrog onto others, or to other organisational systems, using zero-day exploits.

2.6 Proposed Experiment

Our proposal was to implement a strength-dependent password expiry scheme, to determine whether this would encourage stronger password choice. In order to ensure that users were aware of this scheme, we deployed an *enriched nudge*.

Before we introduce the scheme we first review other uses of nudging reported in the research literature.

3. RELATED WORK

Nudging has enjoyed a great deal of media hype over the last few years. The UK [44] and USA [45] Governments, among others, have established units to investigate their use in the public sector. Public health is following suit [46].

The nudging technique [47] manipulates the *choice architecture* (the user interface, in this study's context) to induce people to take the wiser course of action. Nudging is a behavioural economics technique. Other fields also report techniques to change behaviours inexpensively [48, 49].

Not everyone considers nudging a worthwhile endeavour [50], but they have been successful in a range of contexts [51, 52, 53]. For example, small changes in the text of letters sent to citizens made a difference to tax payment rates [53].

What about the information security context? A security-related nudge study [54] successfully nudged users towards a more secure WiFi by using colour and menu order. This finding was confirmed by [55]. Privacy researchers have deployed nudges with some success [56, 57, 58] making people more aware of privacy invasions. People acted upon their new awareness, a strong result.

Password authentication nudge studies have not yet been as successful in delivering change [59, 60, 61, 6]. The password strength meter is also a nudge providing strength feedback, either post-entry or dynamically [62, 63, 64, 60]. Only Vance *et al.* [64] and Ur *et al.* [65] reported a positive result with these meters. The study by Ur *et al.* was an initial scoping study, using Mechanical Turk. As a next step other researchers have tested nudges in

the wild and reported that the meters did not improve password strength, unless users perceived the account to be important. If people *do not* attribute value, then it is understandable that the password meter makes no difference to their choice.

Because password choice is such an important issue in the field of information security it was considered worthwhile to carry out a study to trial a previously-untested nudge in order to identify something that would indeed prove efficacious.

The study described here is part of a long-term project into the deployment of behavioural science techniques in password authentication contexts. The aim was to test the impact of authentication nudges *in the wild*, thus in a natural and realistic setting. Prior to the study reported here all the nudges we trialled were unsuccessful [66].

4. ENRICHED NUDGE

The aim of the investigation was to develop an intervention that was powerful enough to induce people to create stronger passwords. We implemented a strength-dependent password expiry scheme. To ensure that end users understood the scheme, we trialled an *enriched nudge* that comprised a three-pronged approach:

1. the first a user interface tweak (*a nudge*),
2. the second, the mainstay of economic theory: utility (*an incentive*),
3. the third (*a reminder*) at every system login to make users aware of the password expiration date.

4.1 The Theory

The *nudge* was the user interface element that communicated the scheme, the *incentive* was the utility, and the *reminder*, after each successful authentication, that ensured they were warned that their passwords were about to expire.

The idea of offering an *incentive* to encourage actions is based on the concept of utility. The fundamental idea behind neo-classical economics is that people maximise “utility” when they make choices [67]. They weigh up the benefits and costs of each choice option and choose the option that is “best” for them personally. Such an internal utility calculation is possible, and rational, if the information about the choices is complete. If the information is imperfect, on the other hand, Kelman [68] explains that fully rational choice becomes impossible. Hence this manipulation removes all uncertainty: participants were told exactly what the consequences of each choice was. It was unambiguously displayed as they typed in their password.

The idea of a *reminder* is based on the fact that people forget and are easily distracted, especially when a number of other information sources demand their attention [69, 70]. The use of a prominent notification was used to offset this tendency by displaying information about the remaining lifetime of the password every time participants logged into the system.

In effect, the idea was that participants choosing stronger passwords would have to change their passwords less frequently than those who chose weaker passwords (*the incentive*). The *nudge* would make this prominent as and when they formulated a password. The *incentive* would give them a reason (concrete utility) to choose a stronger password. The *reminder* ensured that they were reminded, frequently, about an upcoming expiry date.

4.2 The Implementation

The study was conducted with the help of a self-developed university web application that students could use to look up timetables, coursework information and coursework grades. The application was only accessible from within the campus network and with a valid student ID to prevent outsiders from attempting to use the system. Log in was possible with an alphanumeric password. No password policy or other password requirements were enforced.

Nudge: An image of a long dachshund (Figure 2) was displayed above the password entry field, both for initial creation and password replacement. The length of the dog, and the reputation of this particular breed for strength, would, it was hoped, communicate a subtle message to the participants: *go long and strong*. Even if they were unfamiliar with the breed they could hardly miss the presence of the nudge. This was calculated to draw attention to the speech bubble emerging from the dog’s mouth, telling them that the *stronger* the password, the *longer* they would be able to keep it.

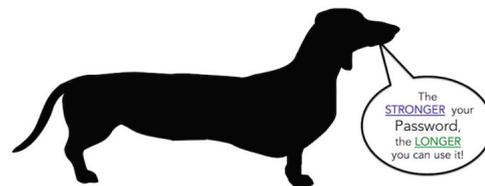


Figure 2: The Nudge

Incentive: The participants were offered extended expiration periods in relation to their password’s strength similar to the suggestion made by [6,7,8]. A survey carried out by Tam *et al.* [71], where participants responded positively to this concept, informed the decision to trial it in the wild. The utility of the password was updated and displayed, below the password field, as they typed in their password. (Figure 3). This communicates a direct benefit related to stronger password choice.

Figure 3: The Incentive

Reminder: A notification made users aware, every time they logged in, when their password would expire. A handy button was provided to facilitate convenient password changes (Figure 4).

Figure 4: The Reminder

Password strength was calculated using the client-based, free and open source JavaScript `zxcvbn.js` [72], a strength estimator that uses pattern matching and minimum entropy calculation. Among other values it delivers a strength score between 0 and 4 that was used in this study and that indicates whether the number of guesses required to break the password is less than 10^2 (score 0), 10^4 (score 1), 10^6 (score 2), 10^8 (score 3), or above (score 4). For example, the password “password” gets a rating of 0, where a password like “BouncyTigger92!” is issued a rating of 4. The script detects 10,000 common passwords, prevalent English words and surnames, as well as common patterns such as dates, repeats (e.g. “aaa”), sequences (e.g. “abcd”), and QWERTY patterns. Calculating strength on the client side ensured no transmission of unhashed passwords to the server. Moreover, the script is used in industry, with the popular Dropbox cloud service being a prominent user [72].

Password length was measured by the number of characters in the password.

4.3 Procedure

When choosing a password for the university web application the participating students were presented *the nudge* and *the incentive* on the registration page of the university web application as described above. Upon every access, the participants were notified of the expiry date by

the reminder. Data was collected between September 2016 and March 2017.

Password recovery was relatively simple: the participants could request a one-time code that was emailed to their registered email address. This allowed them to define a new password and gain access to the system as painlessly as possible.

4.4 Participants

A total of 918 students, the majority of them being Computer Science students, registered to use the system that logged password change events. They logged in 10,317 times during the trial period. Of the 918, 672 opted into the study and the password strength and password length of their passwords were included in the analysis. Unfortunately, due to the requirements of the ethics committee we were not able to collect any further demographics of the students to protect their anonymity.

5. RESULTS

5.1 The Outcome

This particular enriched nudge delivered a positive result: When participants changed their passwords, either before the actual expiration date, or upon request after the expiration date, they chose significantly stronger and longer passwords than the previous one. The analysis is described in more detail below.

5.2 Password Changes

A total of 680 password changes occurred. Of these, 64% were forced changes due to expiration, 36% voluntarily changed their passwords. This could happen because the *reminder* warned them that the password was about to expire. It could also happen because the user decided spontaneously to change their password. The former is more likely. Over the previous year, when users used the same system without password expiration, only eleven voluntary password changes occurred over the six-month period. This confirms Inglesant and Sasse’s finding that very few people voluntarily engage in password changing activities [19].

Most password changes happened in March 2017 (26.47%) when many coursework grades are published, and in January 2017 (24.12%) after the Christmas holidays. Fewest changes happened at the beginning of the term and in October (2.2%) and November (11.47%) 2016.

The password length before a change ranged from 1 to 24 characters with a mean of $\mu = 9.44$ and $\sigma = 2.47$ and a median of $\tilde{x} = 9$. After the change, password length ranged from 1 to 30 characters with a mean of $\mu = 10.42$ and σ

= 3.53 and a median of $\tilde{x} = 10$. A visual inspection of the data revealed deviations from a normal distribution, thus a Wilcoxon signed rank test was conducted. The test was conducted on a significance level of $\alpha = .05$. The increase in length was statistically significant with $W(680) = 87827.00$, $p < .001$. The resulting effect size of $r = .20$ can be interpreted as a small effect. Figure 5 shows the distribution of the password length before and after the change.

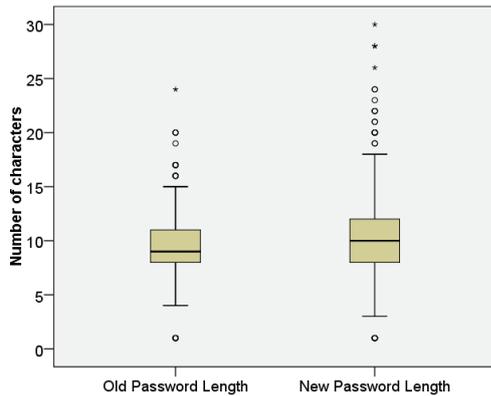


Figure 5: Visualisation of the password length before and after the change (Explanation: o indicates outliers, * indicates extreme values).

We recorded the strength of passwords calculated with the score metric of the strength estimator `zxcvbn.js` both before and after a change. As the scale of the password strength was ordinal, non-parametric Wilcoxon signed rank tests were used to evaluate the differences in password strength. Overall, the strength of changed passwords increased from a median of $\tilde{x} = 1$ to a median of $\tilde{x} = 2$. Figure 6 shows the password strength distribution before and after the change. The increase was statistically significant with $W(680) = 22340.50$, $p < .001$, $r = .27$. In more detail, the effect was significant for both, voluntary changes before expiration ($W(245) = 9751.00$, $p < .001$, $r = .27$) and forced changes after expiration ($W(434) = 29759.00$, $p < .001$, $r = .27$) with a similar effect size. The tests were conducted applying the Bonferroni-Hochberg-procedure for multiple comparisons correction. The effect size of $r=.27$ can be interpreted as a small effect that is close to the threshold of a medium effect ($r = .3$).

5.3 Forgotten Passwords

It was possible to log the number of forgotten password events for 672 of the participants. Of these, 282 forgot their password at least once, that is 41.96%. They forgot passwords between one and twelve times with a mean of $\mu = 2.01$ ($\sigma = 1.52$) and a median of $\tilde{x} = 1$. Looking at the whole group of participants, on average, 0.84 passwords were forgotten per person. Compared to a previous

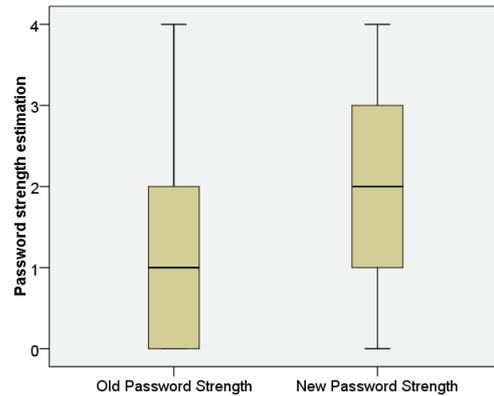


Figure 6: Visualisation of the password strength before and after the change.

nudge study using the same system, in which 737 users participated, the number of people forgetting passwords increased. In the previous study, only 219 of the 737 participants forgot their passwords at least once, that is 29.72%. In that study, on average, 0.49 passwords were forgotten per user.

5.4 Limitations

It is possible that the participants simply reused other strong passwords they already knew [73]. They could also have written down the password, a fairly common response to being confronted with a complex password that people know they are likely to forget. None of these coping mechanisms are easily measured and so we cannot be certain of the extent of their deployment.

The sample in this study consisted of university students that were mainly enrolled in technical courses such as Computer Science. Thus, the sample might be somewhat skewed towards being young and more technically adept than the average user.

The web application allowed us to collect realistic password data of actual users in contrast to more or less artificial data that could be collected with common survey platforms. This, on the one hand, increases the external validity of the study but, on the other hand, prevented us from collecting further demographic information or reasons for people's password choices. This could be valuable data to be collected in future studies.

6. ETHICAL DISCUSSION

This study showed that computer users can indeed be nudged to create stronger passwords. However, the stronger passwords also led to an increase of about 10% in the number of people who forgot passwords, replicating the findings of other studies [74]. Having nudged participants towards stronger passwords, and

having observed the compromised user experience that resulted from our nudge trial, we owed it to our participants to reconsider Sunstein's suggestion that we contemplate whether nudging is indeed warranted in this context [75].

6.1 Requirements from the Nudge Literature

We ground this discussion by considering the definition of nudging as provided by Thaler and Sunstein [76]: "*Nudges are ways of influencing choice without limiting the choice set or making alternatives appreciably more costly in terms of time, trouble, social sanctions, and so forth*". Two requirements emerge from this definition: (1) the existing available choices must be retained, and (2) the possible choices should be more or less equivalent, in terms of "cost" to the decision maker. Although not included in their definition, Thaler and Sunstein also explain that nudges should be used "for good". Indeed, Hansen and Jespersen [77] report that Thaler signs each copy of their book, *Nudge*, with the words "Nudge for good". Hence a third requirement is revealed: (3) the need for nudgers to ensure that the choice they are nudging people towards is actually beneficial *as judged by the nudgees themselves*.

What about how nudges actually exercise their influence? Nudges, according to Nys and Engelen [78], exploit predictable cognitive biases and heuristics in order to influence people towards wiser actions. This is also emphasised by Saghai [79], who explains that nudges trigger people's automatic cognitive processes. This targeting of automatic processes is reminiscent of the dual processing model proposed by Kahneman and Tversky [80]. They explain that processing happens at System 1 and System 2 levels. System 1 is the automatic and quick way of processing, and System 2 the reflective and more time-consuming kind of thinking. Humans prefer to operate on the System 1 level, and only engage System 2 when they have to because it is cognitively expensive.

System 1 nudges are processed automatically and they, and their impact, might not be transparent to nudgees. Such nudge transparency is certainly something that concerns many researchers [78, 81, 82].

An example of a System 1 nudge is to use smaller plates in canteens to reduce calorie intake [83]. People do not notice the change, and respond by eating less, but are probably unaware of their response to the smaller plate.

Nys and Engelen [78] also argue for transparency of nudges as a pre-requisite for their ethical deployment. They believe that people should be aware of the presence of the nudge. Indeed, Thaler and Sunstein [76] explain that adherence to Rawls' Publicity principle [84], i.e. full disclosure of the presence of the nudge and willingness to defend its "goodness", make it ethically sound [85].

It could reasonably be considered devious to nudge people in a way that they are not fully aware of. On the other

hand, if we restrict ourselves to System 2 nudges, we lose a large arsenal of tools that can be used to improve user behaviours. It might be better to require experimenters to justify their intention to deploy System 1 nudges rather than forbidding them altogether.

Hence the fourth requirement is: (4) nudges should be transparent to nudgees, *unless* the experimenter or nudger is able to make a compelling argument for its opacity.

Finally, Thaler and Sunstein [76] also argue that nudge designers should be able to specify which particular behavioural bias their nudge is mitigating against. This requires people to have thought about the design of the nudge before-hand, so that its impact will be *predictable*. This is our fifth requirement.

Based on the nudge literature, to qualify as an ethical nudge, a mechanism needs to meet the following requirements:

- N1. The original set of options should be retained: none should be removed [76, 78].
- N2. The choices offered to nudgees should be more or less equivalent in terms of cost (effort, time etc.) [77, 76].
- N3. It should nudge "for good", *as judged by the nudgee* [76, 86].
- N4. The nudge mechanism should be transparent [85, 78], unless the experimenter or nudger is able to make a compelling argument for its opacity.
The rationale for this is that nudgers thereby respect the autonomy of the decision maker. Autonomy means that people retain the ability to construct their own goals and values, and are free to decide, plan and act in order to satisfy their goals and in accordance with their values [87]. If nudges are not seen and perceived by the nudgee, they can not reflect on their actions and decide to act in accordance with their own goals and values.
On the other hand, humans sometimes act emotionally on the spur of the moment [88, 89, 90]. A well-designed System 1 nudge could ameliorate this tendency, thus helping the person to act as they would if they were not in a 'hot state' [91].
- N5. The nudge designer must be able to explain which particular behavioural bias the nudge mitigates against; to be willing to justify their choice architecture manipulation, and predict its effects.

6.2 Judging the Enriched Nudge

Does our enriched nudge satisfy these requirements? The first requirement is satisfied because participants were technically free to choose weak passwords if they wanted

to. Our enriched nudge was unmistakable and obvious to participants. It did not attempt to influence people without their knowledge. It thus also satisfies the fourth condition, of transparency.

The fifth requirement justification was that we were attempting to offset the least effort principle [92] (taking the easiest course of action without thinking about the long-term consequences). We did this by making the consequences of a weak password salient, and ensuring that they were frequently reminded of the consequence or benefit of their choices every time they logged in.

The second and third requirements warrant closer scrutiny.

(Requirement 2) Option Equivalence

First let us consider the second requirement, that available options be equivalent in terms of cost. As a proviso, we need to acknowledge that password strength is a continuum, ranging from extremely weak, such as '123' to very strong, such as 'If what's done is done, t'were well it were done Quickly!!!'.

To simplify this discussion we shall refer to 'weakest' passwords as passwords on the left end of the continuum and 'stronger' passwords denoting the better passwords, those whose strength tends towards the right of the weakest passwords on the continuum. We acknowledge the fact that this is an over-simplification of the issue but feel that it allows us to present a helpful delineation of password costs for the purposes of considering the cost of a password.

It seems as if the option we were nudging people towards was the more expensive one: stronger passwords. To tease apart this initial assumption, let us consider three main password costs: (1) *time and effort to memorise* [93], (2) *time and effort to enter the password* [94] and (3) *replacement cost* [19].

At first glance, memorising a weak password appears to be less costly than memorising a strong password i.e. $m^{weakest} < m^{stronger}$ where $m = \text{memory effort}$. The situation is not that straightforward, unfortunately. Certainly the weakest password "password" is far easier to remember than the stronger "h6@g2D". On the other hand, long (and therefore stronger) passwords such as passphrases are more memorable than short complex passwords [95]. A meaningful passphrase, such as "Blue water ocean?", is far more memorable than a complex nonsense password such as "h6@g2D". There is thus no linear correlation between the strength of a stronger password and the cost of memorisation.

Next, consider the entry cost. Stronger passwords, both long and complex, take longer to enter than the weakest ones. The *long* because there are more keys to type and the *complex* because they require the use of different parts of the keyboard [29]. Hence $e^{stronger} > e^{weaker}$ where

$e = \text{password entry effort}$. On the other hand, it is likely that typing speeds improve the more often a person enters a particular password. If someone chooses to reuse a well-known stronger password from another system, the entry cost would not necessarily be significantly higher despite being stronger than the weakest password. Moreover, entry is only costly if the password is entered manually. If a browser add-on password manager is used, or the browser remembers the password, this cost is either non-existent or trivial. Once again, there is no linear correlation between weak and stronger passwords in this respect.

On a superficial level, if we discount the complexity of quantifying m and e correctly, we offered participants to a choice between Option 1 [weak password cost = $\sum_{i=1}^n c$] and Option 2 [stronger password cost = $0|c$] where $c = \text{replacement effort}$. c is composed of $\{p,n\}$, $p = \text{cost of engaging with system's process to recover forgotten password}$ and $n = \text{cognitive cost of coming up with a new password}$. If the person deliberately changes the password, then $p=0$, but if they are locked out, they have engage with a password reset process, making $p > 0$.

Yet m , e and c are not orthogonal. Consider the following options that could be chosen by someone being asked to come up with a password in a real-life setting. Figure 7 presents the different options.

Option A: Reuse a stronger password from another system. The person chooses the reduced replacement effort, but benefits from the password requiring no effort to memorise ($m = 0$) and reduced entry effort because they have practised entering it: $e^{reused\ stronger} < e^{new\ stronger}$.

Option B: Choose a weak password. In this case m and e are minimised, but not 0. This is preferred to the costs related to a stronger password, and this preference is not affected by the consequent replacement cost of $\sum_{i=1}^n c$.

Option C: Choose a stronger password and memorise it. This is the opposite of option B. Here the person chooses to minimise the replacement cost while accepting that m and e are more costly than they would be for a weaker password.

Option D: Choose a stronger password and record it, perhaps by making a written record. Here $m = 0$ and $e^{stronger} > e^{weaker}$ but replacement cost is minimised.

Option E: Choose a stronger password and allow the browser to store it. Here $m = 0$ and $e = 0$, which explains why this option is so popular.

Option F: Use a browser-installed password manager to generate a stronger password, store it and populate

the password entry field when required. The cost to the user is the same as Option E, but far more secure.

Option G: Use a password manager on a Smartphone to manage stronger passwords. This removes the need to memorise or record it, but the password still has to be entered manually. Here $m = 0$ and $e^{stronger} > e^{weaker}$ and replacement cost is minimised.

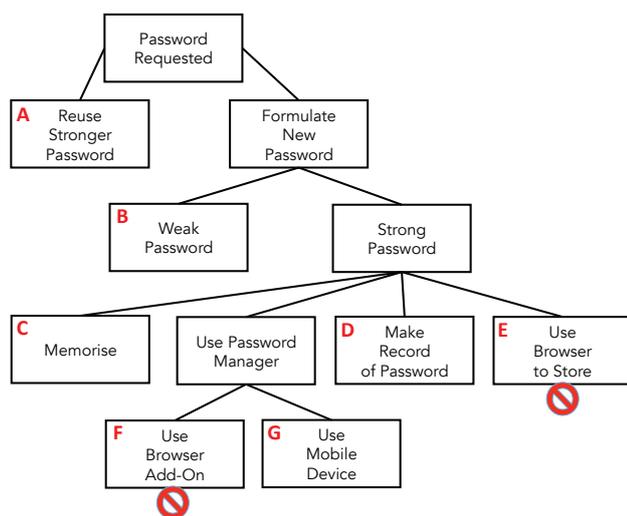


Figure 7: Password Choice Options & Costs

In our experiment, participants were not permitted to install browser add-ons nor to let the browser remember their passwords on lab machines. This means options E and F were not available. Students are permitted to connect their own laptops to eduroam in the lab, but the student system is only accessible from the University network, not eduroam, so that they could not use their own laptop browser to access the system. Hence the password choice options open to our participants were A, B, C, D and G.

It should also be mentioned, at this stage, that c (the replacement cost) was minimised. Participants simply requested a one-time code be sent to their email address, and then used that code to effect a password reset. It would be interesting to run another longitudinal experiment to ascertain whether a more arduous process would encourage even greater adoption of stronger passwords that are more durable.

An empirical investigation into the exact nature of the trade-offs people make between these costs should be the subject of a future study. Without such an investigation, we cannot argue that we either satisfied or violated the second ethical requirement.

(Requirement 3) Nudge for Good

The third requirement, to nudge “for good”, is equally hard to judge. On the one hand, it is undeniable that weak

passwords compromise system security. A serious hacking attack seldom affects only one computer user. Hackers will generally use one compromised account as a stepping stone to other accounts or other software systems on the same infrastructure. Because one person’s poor password choice thus potentially has an impact on others we might well feel justified in nudging individuals towards stronger passwords.

On the other hand, the fact that we deployed nudges implicitly suggests that we thought the participants were either unaware of the need for strong passwords or not willing or able to transfer their intention of creating strong passwords into behaviour. The potential reasons for that are numerous and might even be reasonable from the user’s point of view given his/her decision context and knowledge.

For example, users might underestimate the risk of an attack, shy away from the cost of stronger passwords, perceive the data as not having enough value to warrant a stronger password, or have a perception of a strong password that does not match actual password security.

Still, whatever the reason for choosing weaker passwords, the underlying assumption is that computer users are unable to choose a password to match the actual risk associated with the value of the asset and the vulnerabilities represented by a particular password. Deploying nudges might therefore be considered to demonstrate a lack of respect for the participants; a paternalistic intervention that does not respect human autonomy [96].

Hall and Carter [97] admit that nudges infringe autonomy but argue that it is justified to offset the nudges used by others because nudges are intended “for good” [77]. Moreover, Gordijn and Ten Have [98] claim that autonomy has not proved the “cure-all” for all ethical issues in society. Brooks [99] also acknowledges the autonomy issue but argues that nudges are inescapable as no choice architecture can be completely neutral. Thus, from his point of view the question is not whether to nudge, but how to do it in an ethical way. He argues for better mechanisms for obtaining informed consent and for nudge transparency. Finally, Norman [100] argues that it is not autonomy, *per se*, that is important, but rather that people can indeed retain some measure of control, and that they understand and have knowledge of what they are able to control.

With respect to our experiment, we did indeed obtain informed consent and participants could opt out of the study by checking a box during enrolment and all our interventions were fully visible. Our participants thus retained full control.

Still, this third requirement creates the most difficulty and brings us back to the initial question: “were we justified in deploying the nudge at all?” Would our participants have

considered the choice we nudged them towards as being *for their own benefit*? If not, should we not allow end users to decide on how strong to make their password, without interference? Were we nudging “for good”, which Thaler and Sunstein consider core to the ethical nudge [101]?

The real difficulty we have with nudging in this domain is that nudges are required to align with the end user’s judgement of personal benefit, or Thaler’s “goodness”. The push to stronger passwords is generally good for the organisation, for the current software system, for the other users. Is it good for the individual, and should we be concerned about that? Conly [102] argues that when our actions can have a negative impact on others we can have no expectation of autonomy. As detailed previously, poor password choice can indeed have an impact on others.

This is a complex issue, with researchers arguing both for and against the humans having the right to individual autonomy [94, 103] and whether nudging violates this or not [96, 76, 102].

We therefore reserve judgement with respect to whether our enriched nudge justifiably infringed end-user autonomy, or not.

6.3 Authentication Requirements

Whereas nudging in other contexts has focused on simplifying and easing behaviours [53], that is not always an option in authentication, especially when it comes to password authentication.

Authentication is a context that operates under different constraints from other contexts. For example, in many contexts the aim is to simplify the process the human needs to engage with [104]. Some of the most common techniques in achieving this is to maximise feedback on actions [105], and to ensure that people can recover from errors [106, 107].

Now consider the authentication context. The system is interacting with someone who is required to prove his or her identity. An error could be a signal of an intrusion attempt, and the system cannot help the person to recover from his or her error. Feedback is minimal in this context, too, for the same lack of trust in the person interacting with the system.

How about nudging in authentication? The option we are nudging people towards, in the context of stronger passwords, must be warranted. We are trying to influence users towards expending more effort, and to do so every time they use the system. Hence:

A1. The nudge designer *must* be able to argue that nudges that encourage such efforts are indeed justified.

A2. Nudge proposers should have to make an argument for deploying nudging at all, purely based on the value of the resource being protected by the password, *to the nudgee him or herself*. User effort, with respect to passwords, is not free as many developers seem to believe.

A3. Those deploying nudges have to monitor the impact of the nudge by checking for the number of users being locked out, the impact of the nudge on password strength and the satisfaction of the users. If there are unanticipated and negative side effects the nudge can be disengaged.

We also need to remove one of the previously-listed principles. It is not possible for the user’s options to be equivalent with a password authentication nudge. The pure nudge requirement for retention is not feasible and N2 will therefore be dropped.

6.4 Ethical Password Authentication Guidelines

The information security field does indeed consider efforts to persuade people to choose stronger passwords worthwhile. We do not suggest that researchers have deployed any unethical nudges in the information security context, but there is undoubtedly scope for unethical use of nudges in this context. This is especially true if nudgers do not have applicable guidelines to inform their deployment.

If we concede that nudges, as a technique, are indeed acceptable in password authentication, we should ensure that such nudges, in an organisational and real-time setting, exhibit a number of characteristics.

We believe it would be helpful to encode the principles we have derived from the literature as a list, to inform and assist researchers working in this area.

The first over-arching directive is that organisations should prioritise the use of technical measures to scaffold and bolster the security of the system as much as possible before focusing their efforts on changing end-user behaviour.

Having exhausted technical measures and having judged that end users need to use stronger passwords, it should be ensured that password nudges are:

N1: retentive: End users must still be able to vary password strength to resist the influence of the nudge. It might be necessary to mandate a particular minimum password strength, and then conceivably use the nudge to encourage passwords that exceed the minimum. However, the weaker options should still be available so that the nudge respects their autonomy and agency.

N3: respectful : Choice architecture manipulations should respect user autonomy. Users should never feel that they have no agency when interacting with an authentication system. Otherwise the system risks triggering a reactance response. If the nudge can make a coherent argument for deploying an opaque System 1 nudge, and is able to satisfy the ethics review board that this is essential, then it is crucial for nudges to be apprised of the manipulation once the experiment is over.

They should be told that: (1) they were nudged, (2) why this was done, (3) what the nudge was intended to achieve, and (4) what the impact of the nudge was over all participants.

N4: transparent : Nudges ought to be fully aware of the nudge and the influence it is attempting to exert. If this would negate the influence of the nudge, it is necessary for the nudge designer to argue convincingly that this is the case, and for the Ethics Review Board to be persuaded of the need for opacity. For example, if the experimenter is concerned about the impact of demand characteristics [108], that could be a reason to make the nudge opaque.

The display of a password strength meter meets the transparency requirement, as does the enriched nudge we tested. One could imagine someone using a scary background on the web page subliminally to induce a fear of hacking and thereby attempting to nudge people towards stronger passwords. Such a nudge would not be transparent and therefore questionable as far as ethics is concerned.

N5: defensible : It must be trivial for nudges to contact those deploying the nudge should they have any questions or concerns about it. This is in line with Rawl's Publicity principle [84] and requires nudgers to be able to justify the behavioural biases they are attempting to ameliorate with the nudge.

A1: justified : Strong passwords have a cost associated with them and user cost is not free. Nudgers should be aware of the fact that users may rationally respond to demands for greater strength by deploying less secure practices, such as allowing their browser to remember the password. This weakens the security of the system, while putting a greater burden on the end users. Nudgers should, if possible, ensure that the nudges are apprised of the motivations for the nudge so that they understand why they are being asked to put extra effort into authenticating.

A2: sufficient : Nudging should only be deployed when the asset being protected requires stronger passwords than the *status quo* average password, based on the previous usage of a system. Designers should put some thought into applying a rule such as "*require passwords to be as strong as needed, as matched to the value of the asset, but no stronger*".

A3: monitored : A number of monitoring set points should be defined and adhered to after roll-out to carry out close examination so as to detect unexpected and undesirable side effects. Amelioration or abandonment should be considered seriously if the user experience is being compromised unacceptably.

So, for example, the number of forgotten passwords should be monitored and compared to the usual number, to determine whether the user experience is being impoverished. Other data should also be scrutinised, such as the number of logins, the password strength profile and any complaints from the users.

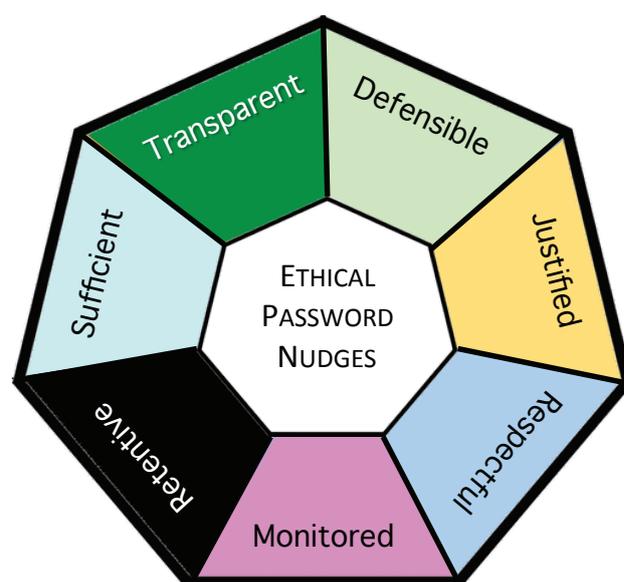


Figure 8: Ethical Password Nudge Characteristics

6.5 Summary

We started off this section by contemplating whether nudging is indeed warranted in this context, as advised by Sunstein [75]. Our investigation required us to consider the ethics of nudging. As we perused the literature five distinct requirements of ethical nudges emerged. We extended this with three authentication-specific ethical nudge principles, dropping one of the previous five principles. In total we arrived at seven ethical principles for nudging in information security.

We assessed the ethics of our enriched nudge using these as metrics and uncovered difficulties related to two areas. The first is the requirement for nudge options to be equivalent, and the second that the nudge be intended for the "good" of the nudgee.

Ajudging the latter seems to come down to the fact that it is considered necessary by the information security community to justify the deployment of these kinds of

techniques when individuals' unwise actions can have serious and undesirable side effects on others. Weak password choice can lead to compromises that impact large numbers of people. Is it acceptable to violate autonomy for the greater good? Conly [102] would say yes, White [96] would disagree.

We conclude the paper with a list of ethical password nudge requirements, intended for those who decide that password nudging is indeed ethical and warranted. We do not yet make a strong argument for, or against, nudging in this context, in particular because of the fact that they could be considered to infringe autonomy.

7. FUTURE WORK

There is scope for further work in a number of directions.

Autonomy

We plan to carry out a more extensive investigation into the meaning of autonomy in this context and the meaning of potential violation that nudges can commit in the information security context.

Change Costs

The change costs, in our experiment, were as low as we could reasonably make them. It would be interesting to run the experiment again with a more expensive replacement process in order to see what impact that would have on password strength.

Password Costs

An investigation into the interplay between the different aspects making up actual password cost would be very insightful and is something worth pursuing.

Generalising the Ethical Guidelines

It would obviously be helpful if we were able to produce a more general set of guidelines to inform other research areas within Information Security and Privacy, and we plan to pursue this goal next.

8. CONCLUSION

In this paper we report on an investigation into the efficacy of an *enriched nudge*, comprising a nudge, an incentive and a reminder, in terms of influencing people towards choosing stronger passwords.

The focus of this paper was on learning lessons from our experiences, and from the nudge literature, in order to derive nudge-specific ethical guidelines. Our purpose was to provide guidance to other researchers experimenting with nudges in authentication, and to ethics review boards having to assess and approve research proposals.

We thus conclude this paper with a set of ethical guidelines for nudging in password authentication and we demonstrate how these can be applied.

9. ACKNOWLEDGMENTS

We obtained ethical approval from the College of Science and Engineering at the University of Glasgow to carry out nudge-related research on the website (Approval #300140006). We wish to thank the support staff in the School of Computing Science for their advice and assistance during the course of this research.

This work was supported by the German Federal Ministry of Education and Research (BMBF) as well as by the Hessen State Ministry for Higher Education, Research and the Arts (HMWK) within CRISP (Center for Research in Security and Privacy).

REFERENCES

- [1] D. Florencio and C. Herley, "A large-scale study of web password habits," in *Proceedings of the 16th international conference on World Wide Web*. ACM, 2007, pp. 657–666.
- [2] E. H. Spafford, "Opus: Preventing weak password choices," *Computers & Security*, vol. 11, no. 3, pp. 273–278, 1992.
- [3] —, "Preventing weak password choices," Computer Science Technical Reports, Tech. Rep. Paper 875, 1991, <http://docs.lib.purdue.edu/cstech/875>.
- [4] S. Chiasson and P. C. Van Oorschot, "Quantifying the security advantage of password expiration policies," *Designs, Codes and Cryptography*, vol. 77, no. 2-3, pp. 401–408, 2015.
- [5] Y. Zhang, F. Monrose, and M. K. Reiter, "The security of modern password expiration: An algorithmic framework and empirical analysis," in *Proceedings of the 17th ACM conference on Computer and Communications Security*. ACM, 2010, pp. 176–186.
- [6] T. Seitz, E. von Zezschwitz, S. Meitner, and H. Hussmann, "Influencing Self-Selected Passwords Through Suggestions and the Decoy Effect," in *EuroUSEC*. Darmstadt: Internet Society, 2016.
- [7] G. R. Walters, "Variable expiration of passwords," USA Patent US 7 200 754 B2, US20040177272, <https://www.google.com/patents/US7200754>.
- [8] R. Childress, I. Goldberg, M. Lechtman, and Y. Medini, "User policy manageable strength-based password aging," USA Patent, Feb. 5, 2013. [Online]. Available: <https://www.google.com/patents/US8370925>

- [9] The British Psychological Society, "Code of human research ethics," 2014, <http://www.bps.org.uk/publications/policy-and-guidelines/research-guidelines-policy-documents/research-guidelines-poli>.
- [10] D. A. Curry, *Unix system security: a guide for users and system administrators*. Addison-Wesley Longman Publishing Co., Inc., 1992.
- [11] SANS Institute, "Password protection policy," <https://www.sans.org/security-resources/policies/general/pdf/password-protection-policy>.
- [12] W. Cheswick, "Rethinking passwords," *Queue*, vol. 10, no. 12, pp. 50:50–50:56, Dec. 2012.
- [13] C. Herley and P. Van Oorschot, "A research agenda acknowledging the persistence of passwords," *IEEE Security & Privacy*, vol. 10, no. 1, pp. 28–36, 2012.
- [14] M. Bishop, "Best practices and worst assumptions," in *Proceedings of the 2005 Colloquium on Information Systems Security Education (CISSE) pp*, 2005, pp. 18–25.
- [15] L. Cranor, "Time to rethink mandatory password changes," 2016, <https://www.ftc.gov/news-events/blogs/techftc/2016/03/time-rethink-mandatory-password-changes>.
- [16] K. Hickey, "Mandatory password changes – not as secure as you think," 2016, <https://gcn.com/articles/2016/06/07/mandatory-password-changes.aspx>.
- [17] P. A. Grassi, J. L. Fenton, E. M. Newton, R. A. Perlner, A. R. Regenscheid, W. E. Burr, J. P. Richer, N. B. Lefkowitz, J. M. Danker, Y.-Y. Choong, K. K. Greene, and M. F. Theofanos, "NIST Special Publication 800-63B. Digital Identity Guidelines Authentication and Lifecycle Management," 2017, <https://pages.nist.gov/800-63-3/>.
- [18] M. A. Sasse, "'Technology Should Be Smarter Than This!': A Vision for Overcoming the Great Authentication Fatigue," in *Workshop on Secure Data Management*. Springer, 2013, pp. 33–36.
- [19] P. G. Inglesant and M. A. Sasse, "The true cost of unusable password policies: password use in the wild," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2010, pp. 383–392.
- [20] W. C. Summers and E. Bosworth, "Password policy: the good, the bad, and the ugly," in *Proceedings of the Winter International Symposium on Information and Communication Technologies*. Trinity College Dublin, 2004, pp. 1–6.
- [21] K. Renaud, "Blaming noncompliance is too convenient: What really causes information breaches?" *IEEE Security & Privacy*, vol. 10, no. 3, pp. 57–63, 2012.
- [22] C. L. Huntley, "A developmental view of system security," *Computer*, vol. 39, no. 1, pp. 113–114, 2006.
- [23] C. Herley, "So long, and no thanks for the externalities: the rational rejection of security advice by users," in *Proceedings of the 2009 workshop on New Security Paradigms Workshop*. Colorado: ACM, 2009, pp. 133–144.
- [24] J. Bonneau, C. Herley, P. C. van Oorschot, and F. Stajano, "Passwords and the evolution of imperfect authentication," *Communications of the ACM*, vol. 58, no. 7, pp. 78–87, 2015.
- [25] J. Osper, "Password expiration policy best practices," 2016, 4 May <https://www.portalguard.com/blog/2016/05/04/password-expiration-policy-best-practices/>.
- [26] K. Renaud, R. Blignaut, and I. Venter, "Smartphone owners need security advice. how can we ensure they get it?" in *International Conference on Information Resources Management (CONF-IRM)*, 2016.
- [27] L. Gong, M. A. Lomas, R. M. Needham, and J. H. Saltzer, "Protecting poorly chosen secrets from guessing attacks," *IEEE Journal on Selected Areas in Communications*, vol. 11, no. 5, pp. 648–656, 1993.
- [28] J. Bonneau, "The science of guessing: analyzing an anonymized corpus of 70 million passwords," in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 538–552.
- [29] F. Tari, A. Ozok, and S. H. Holden, "A comparison of perceived and real shoulder-surfing risks between alphanumeric and graphical passwords," in *Proceedings of the Second Symposium on Usable Privacy and Security*. ACM, 2006, pp. 56–66.
- [30] L. Falk, A. Prakash, and K. Borders, "Analyzing websites for user-visible security design flaws," in *Proceedings of the 4th symposium on Usable Privacy and Security*. ACM, 2008, pp. 117–126.
- [31] C. Yue and H. Wang, "Characterizing insecure javascript practices on the web," in *Proceedings of the 18th international conference on World wide web*. ACM, 2009, pp. 961–970.
- [32] K. Shah, "Phishing: An evolving threat," *International Journal of Students' Research in Technology & Management*, vol. 3, no. 1, pp. 216–222, 2015.

- [33] E. O. Yeboah-Boateng and P. M. Amanor, "Phishing, SMiShing & Vishing: an assessment of threats against mobile devices," *Journal of Emerging Trends in Computing and Information Sciences*, vol. 5, no. 4, pp. 297–307, 2014.
- [34] C. Stoll, *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage*. Gallery Books, 2005.
- [35] N. C. Sickler and S. J. Elliott, "An evaluation of fingerprint image quality across an elderly population vis-a-vis an 18-25 year old population," in *39th Annual 2005 International Carnahan Conference on Security Technology, 2005. CCST'05*. IEEE, 2005, pp. 68–73.
- [36] J. P. Schneider, "Managing password expiry," USA Patent US 8 959 618 B2, Feb 17, 2015, <https://www.google.com/patents/US8959618>.
- [37] C. Bravo-Lillo, S. Komanduri, L. F. Cranor, R. W. Reeder, M. Sleeper, J. Downs, and S. Schechter, "Your attention please: designing security-decision uis to make genuine risks harder to ignore," in *Proceedings of the Ninth Symposium on Usable Privacy and Security*. ACM, 2013, pp. 6–24.
- [38] S. Breznitz, *Cry Wolf: The Psychology of False Alarms*. Psychology Press, 2013.
- [39] D. Glynos, P. Kotzanikolaou, and C. Douligeris, "Preventing impersonation attacks in manet with multi-factor authentication," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2005. WIOPT 2005. Third International Symposium on*. IEEE, 2005, pp. 59–64.
- [40] W. Harwood, "Locking up passwords—for good," *Network Security*, vol. 2016, no. 4, pp. 10–13, 2016.
- [41] C.-Y. Huang, S.-P. Ma, and K.-T. Chen, "Using one-time passwords to prevent password phishing attacks," *Journal of Network and Computer Applications*, vol. 34, no. 4, pp. 1292–1301, 2011.
- [42] M. Al Fairuz and K. Renaud, "Multi-channel, multi-level authentication for more secure ebanking," in *Information Security South Africa*, Johannesburg, South Africa, 2010.
- [43] C. Herley, P. C. Van Oorschot, and A. S. Patrick, "Passwords: If were so smart, why are we still using them?" in *International Conference on Financial Cryptography and Data Security*. Springer, 2009, pp. 230–237.
- [44] The Behavioural Insights Team, "Who we are," 2014, <http://www.behaviouralinsights.co.uk/about-us/> Accessed 19 Sept, 2016.
- [45] J. Holden, "Memorandum to the Heads of Executive Departments and Agencies. Implementation Guidance for Executive Order 13707: Using Behavioral Science Insights to Better Serve the American People," Washington, DC, 2015, sept 15. Executive Office of the President. Office of Science and Technology Policy <https://www.whitehouse.gov/the-press-office/2015/09/15/executive-order-using-behavioral-science-insights-better-serve-american> Accessed 19 September 2016.
- [46] M. Verweij and M. v. d. Hoven, "Nudges in public health: paternalism is paramount," *The American Journal of Bioethics*, vol. 12, no. 2, pp. 16–17, 2012.
- [47] R. H. Thaler, C. R. Sunstein, and T. C. Leonard, "Nudge: Improving decisions about health, wealth, and happiness," *Constitutional Political Economy*, vol. 19, no. 4, pp. 356–360, 2008.
- [48] M. Bateson, L. Callow, J. R. Holmes, M. L. R. Roche, and D. Nettle, "Do images of 'watching eyes' induce behaviour that is more pro-social or more normative? A field experiment on littering," *PloS one*, vol. 8, no. 12, p. e82055, 2013.
- [49] A. Dijksterhuis, J. A. Bargh, and J. Miedema, "Of men and mackerels: Attention, subjective experience, and automatic social behavior," in *The message within: The role of subjective experience in social cognition and behavior*, H. Bless and J. Forgas, Eds. New York: Psychology Press, 2000, ch. 3, pp. 37–51.
- [50] G. Rayner and T. Lang, "Is nudge an effective public health strategy to tackle obesity? No," *British Medical Journal*, vol. 342, 2011.
- [51] A. Oliver, "Is nudge an effective public health strategy to tackle obesity? Yes," *British Medical Journal*, vol. 342, 2011.
- [52] J. K. Turland, "Aiding information security decisions with human factors using quantitative and qualitative techniques," Ph.D. dissertation, Newcastle University, 2016.
- [53] D. Halpern, *Inside the Nudge Unit: How small changes can make a big difference*. London: WH Allen, 2015.
- [54] D. Jeske, L. Coventry, P. Briggs, and A. van Moorsel, "Nudging whom how: It proficiency, impulse control and secure behaviour," in *Personalizing Behavior Change Technologies CHI Workshop*. Toronto: ACM, 27 April 2014.
- [55] I. Yevseyeva, C. Morisset, and A. van Moorsel, "Modeling and analysis of influence power for information security decisions," *Performance Evaluation*, vol. 98, pp. 36–51, 2016.

- [56] E. K. Choe, J. Jung, B. Lee, and K. Fisher, "Nudging people away from privacy-invasive mobile apps through visual framing," in *IFIP Conference on Human-Computer Interaction*. Springer, 2013, pp. 74–91.
- [57] R. Balebako, P. G. Leon, H. Almuhimedi, P. G. Kelley, J. Mugan, A. Acquisti, L. F. Cranor, and N. Sadeh, "Nudging users towards privacy on mobile devices," in *Proc. CHI 2011 Workshop on Persuasion, Nudge, Influence and Coercion*. ACM, 2011.
- [58] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, "Your location has been shared 5,398 times!: A field study on mobile app privacy nudging," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. New York, NY, USA: ACM, 2015, pp. 787–796.
- [59] M. Ciampa, "A comparison of password feedback mechanisms and their impact on password entropy," *Information Management & Computer Security*, vol. 21, no. 5, pp. 344–359, 2013.
- [60] 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*. Paris: ACM, 2013, pp. 2379–2388.
- [61] B. M. Josiam and J. P. Hobson, "Consumer choice in context: the decoy effect in travel and tourism," *Journal of Travel Research*, vol. 34, no. 1, pp. 45–50, 1995.
- [62] X. de Carné de Carnavalet, "A large-scale evaluation of high-impact password strength meters," Ph.D. dissertation, Concordia University, 2014.
- [63] A. Sotirakopoulos, "Influencing user password choice through peer pressure," Ph.D. dissertation, The University Of British Columbia (Vancouver), 2011.
- [64] A. Vance, D. Eargle, K. Ouimet, and D. Straub, "Enhancing password security through interactive fear appeals: A web-based field experiment," in *2013 46th Hawaii International Conference on System Sciences (HICSS)*. Hawai'i: IEEE, 2013, pp. 2988–2997.
- [65] 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 *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*. Bellevue: USENIX, 2012, pp. 65–80.
- [66] K. Renaud, V. Zimmermann, J. Maguire, and S. Draper, "Lessons learned from evaluating eight password nudges in the wild," in *LASER Workshop. Arlington. 18-19 October, 2017*.
- [67] W. S. Jevons, *The theory of Political Economy*. Macmillan and Company, 1879.
- [68] M. Kelman, "Choice and utility," *Wisconsin Law Review*, p. 769, 1979.
- [69] S. Misra and D. Stokols, "Psychological and health outcomes of perceived information overload," *Environment and Behavior*, vol. 44, no. 6, pp. 737–759, 2012.
- [70] G. Pijpers, *Information overload: A system for better managing everyday data*. John Wiley & Sons, 2010.
- [71] L. Tam, M. Glassman, and M. Vandenwauver, "The psychology of password management: a tradeoff between security and convenience," *Behaviour & Information Technology*, vol. 29, no. 3, pp. 233–244, 2010.
- [72] D. L. Wheeler, "zxcvbn: Low-budget password strength estimation," in *USENIX Conference 2016*. Vancouver: USENIX, August 2016, Dropbox Inc.
- [73] R. Wash, E. Rader, R. Berman, and Z. Wellmer, "Understanding password choices: How frequently entered passwords are re-used across websites," in *Symposium on Usable Privacy and Security (SOUPS)*, 2016.
- [74] R. Shay, S. Komanduri, P. G. Kelley, P. G. Leon, M. L. Mazurek, L. Bauer, N. Christin, and L. F. Cranor, "Encountering stronger password requirements: user attitudes and behaviors," in *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM, 2010, p. 2.
- [75] C. R. Sunstein, "Nudges that fail," *Behavioural Public Policy*, vol. 1, no. 1, pp. 4–25, 2017.
- [76] R. H. Thaler and C. R. Sunstein, *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, 2008.
- [77] P. G. Hansen and A. M. Jespersen, "Nudge and the manipulation of choice: A framework for the responsible use of the nudge approach to behaviour change in public policy," *European Journal of Risk Regulation*, vol. 4, no. 1, pp. 3–28, 2013.
- [78] T. R. Nys and B. Engelen, "Judging nudging: Answering the manipulation objection," *Political Studies*, vol. 65, no. 1, pp. 199–214, 2017.

- [79] Y. Saghai, "Salvaging the concept of nudge," *Journal of Medical Ethics*, vol. 39, no. 8, pp. 487–493, 2013.
- [80] D. Kahneman, *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- [81] T. Haugh, "The ethics of intracorporate behavioral ethics," *California Law Review Online*, April 2017.
- [82] C. R. Sunstein, "Fifty shades of manipulation," 2015, https://dash.harvard.edu/bitstream/handle/1/16149947/manipulation2_18.pdf?sequence=1.
- [83] B. Wansink, "Environmental factors that increase the food intake and consumption volume of unknowing consumers," *Annual Review of Nutrition*, vol. 24, pp. 455–479, 2004.
- [84] J. Rawls, *A Theory of Justice*. Harvard university press, 2009.
- [85] C. R. Sunstein and R. H. Thaler, "Libertarian paternalism is not an oxymoron," *The University of Chicago Law Review*, pp. 1159–1202, 2003.
- [86] C. R. Sunstein, "Nudges Do Not Undermine Human Agency," *Journal of Consumer Policy*, vol. 38, no. 3, pp. 207–210, 2015.
- [87] B. Friedman and H. Nissenbaum, "Software agents and user autonomy," in *Proceedings of the first international conference on Autonomous agents*. ACM, 1997, pp. 466–469.
- [88] J. Benhabib, A. Bisin, and A. Schotter, "Present-bias, quasi-hyperbolic discounting, and fixed costs," *Games and Economic Behavior*, vol. 69, no. 2, pp. 205–223, 2010.
- [89] T. Sharot, A. M. Riccardi, C. M. Raio, and E. A. Phelps, "Neural mechanisms mediating optimism bias," *Nature*, vol. 450, no. 7166, pp. 102–105, 2007.
- [90] E. Castano, V. Yzerbyt, M.-P. Paladino, and S. Sacchi, "I belong, therefore, I exist: Ingroup identification, ingroup entitativity, and ingroup bias," *Personality and Social Psychology Bulletin*, vol. 28, no. 2, pp. 135–143, 2002.
- [91] P. G. Hansen, "The definition of nudge and libertarian paternalism: Does the hand fit the glove?" *European Journal of Risk Regulation*, no. 1, pp. 1–20, 2015.
- [92] G. K. Zipf, *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio Books, 2016.
- [93] E. Stobert and R. Biddle, "The password life cycle: user behaviour in managing passwords," in *Proc. SOUPS*, 2014.
- [94] K. K. Greene, M. A. Gallagher, B. C. Stanton, and P. Y. Lee, "I cant type that! p@ w0rd entry on mobile devices," in *International Conference on Human Aspects of Information Security, Privacy, and Trust*. Springer, 2014, pp. 160–171.
- [95] M. Keith, B. Shao, and P. Steinbart, "A behavioral analysis of passphrase design and effectiveness," *Journal of the Association for Information Systems*, vol. 10, no. 2, p. 2, 2009.
- [96] M. White, *The manipulation of choice: Ethics and libertarian paternalism*. Springer, 2013.
- [97] L. Zhang and W. C. McDowell, "Am i really at risk? determinants of online users' intentions to use strong passwords," *Journal of Internet Commerce*, vol. 8, no. 3-4, pp. 180–197, 2009.
- [98] B. Gordijn and H. Ten Have, "Autonomy, free will and embodiment," *Medicine, Health Care and Philosophy*, vol. 13, no. 4, p. 301302, 2010.
- [99] T. Brooks, "Should we nudge informed consent?" *The American Journal of Bioethics*, vol. 13, no. 6, pp. 22–23, 2013.
- [100] D. A. Norman, "How might people interact with agents," *Communications of the ACM*, vol. 37, no. 7, pp. 68–71, 1994.
- [101] C. R. Sunstein, "The ethics of nudging," *Yale Journal on Regulation*, vol. 32, p. 413, 2015.
- [102] S. Conly, "Against autonomy: justifying coercive paternalism," *Journal of Medical Ethics*, vol. 40, no. 5, pp. 349–349, 2014.
- [103] B. O'Neill, "A message to the illiberal nudge industry: push off," 2000, spiked, 1 November.
- [104] J. Nielsen, *Designing web usability: The practice of simplicity*. New Riders Publishing, 1999.
- [105] K. Renaud and R. Cooper, "Feedback in human-computer interaction-characteristics and recommendations," *South African Computer Journal*, vol. 2000, no. 26, pp. 105–114, 2000.
- [106] K. P. O'Hara and S. J. Payne, "Planning and the user interface: The effects of lockout time and error recovery cost," *International Journal of Human-Computer Studies*, vol. 50, no. 1, pp. 41–59, 1999.
- [107] G. Buchanan, S. Farrant, M. Jones, H. Thimbleby, G. Marsden, and M. Pazzani, "Improving mobile internet usability," in *Proceedings of the 10th international conference on World Wide Web*. ACM, 2001, pp. 673–680.
- [108] A. L. Nichols and J. K. Maner, "The good-subject effect: Investigating participant demand characteristics," *The Journal of General Psychology*, vol. 135, no. 2, pp. 151–166, 2008.