

The Nudge Puzzle: Matching Nudge Interventions to Cybersecurity Decisions

VERENA ZIMMERMANN*, Technische Universität Darmstadt, Germany

KAREN RENAUD, University of Strathclyde, Scotland, Rhodes University, South Africa

Nudging is a promising approach, in terms of influencing people to make advisable choices in a range of domains, including cybersecurity. However, the processes underlying the concept, the nudge's effectiveness in different contexts, and in the long term, are still poorly understood. Our research thus first reviewed the nudge concept and differentiated it from other interventions before applying it to the cybersecurity area. We then carried out an empirical study to assess the effectiveness of three different nudge-related interventions on four types of cybersecurity-specific decisions. Our study demonstrated that the combination of a simple nudge and information provision, termed a "hybrid nudge", was at least as, and in some decision contexts even more effective in encouraging secure choices as the simple nudge on its own. This indicates that the inclusion of information when deploying a nudge, thereby increasing the intervention's transparency, does not necessarily diminish its effectiveness.

A follow-up study explored the educational and long-term impact of our tested nudge interventions to encourage secure choices. The results indicate that the impact of the initial nudges, of all kinds, did not endure. We conclude by discussing our findings and their implications for research and practice.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; *HCI theory, concepts and models*; • **Security and privacy** → Social aspects of security and privacy;

Additional Key Words and Phrases: Nudging; Security; Privacy; Decision Making, Feedback; Information

ACM Reference Format:

Verena Zimmermann and Karen Renaud. 2020. The Nudge Puzzle: Matching Nudge Interventions to Cybersecurity Decisions. *ACM Trans. Comput.-Hum. Interact.* 1, 1, Article 1 (January 2020), 43 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

We are confronted, daily, with the need to make a plethora of decisions, each of which is influenced both by the dimensions of the decision itself and by the context of the decision i.e. the '*choice architecture*'. The Nobel prize winner, Richard Thaler, together with Cass Sunstein, introduced the world to nudges in 2008 [83]. Nudges are effectively ways of tweaking the choice architecture to influence people's choices. Some countries' governments, including the USA, the UK and Australia [16, 29, 56, 82], have established units to study and deploy nudge-related interventions to improve the welfare of their citizens.

The nudge concept originated from the field of behavioral economics and has been applied in a variety of contexts. It has gained prominence in contexts such as health [45], energy consumption [5, 69] and road safety [83]. Nudges have also been deployed in the digital world, referred to as "digital nudging" [93]. This has become

Authors' addresses: Verena Zimmermann, zimmermann@psychologie.tu-darmstadt.de, Technische Universität Darmstadt, Darmstadt, Germany; Karen Renaud, University of Strathclyde, Glasgow, Scotland, Rhodes University, Grahamstown, South Africa, karen.renaud@strath.ac.uk.

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.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1073-0516/2020/1-ART1 \$15.00

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

increasingly important as the boundaries between the digital and physical worlds blur due to the widespread diffusion of smart technologies.

An important application area of digital nudging is the domain of human-centred cybersecurity. In some ways, cybersecurity decisions are very similar to other kinds of decisions. The information people have, and the biases they are subject to, influence all their decisions. On the other hand, cybersecurity decisions have some distinguishing features. Security is a relatively intangible concept and often invisible to users in the digital world. For example, the appearance of a website does not necessarily align with its security and privacy features. Even a security breach might not be immediately visible or experiential. For example, the link between the unauthorized sharing of one's email address by one service provider and the later receipt of spam mails, might never be revealed. Furthermore, security is often not the user's primary aim. People usually engage in a security ceremony because they are required to do so, not because it is their primary goal. For example, someone wants to connect to a WiFi to check their email while shopping (their primary aim). To do this, they have to choose a WiFi to connect to, and security might not be uppermost in their mind. Nudging can make the security and privacy dimensions of the decision more salient.

The field of human-centred cybersecurity aims to support people in behaving more securely [30, 73, 87], or in adopting measures to preserve their privacy while online [24]. For example, one cybersecurity-related study trialled a number of nudges to identify the one that would encourage stronger passwords [73]. A privacy-related nudge attempted to persuade people to choose the most secure WiFi to connect to [85].

To qualify as a nudge, an intervention should not forbid or significantly alter the economic incentives of the pre-nudge options [83]. Yet the original definition was perhaps not precise enough to delineate exactly what counts as a nudge [37, 75]. For example, if a web page displays password strength requirements, does that count as a nudge? What about nagging people into installing software updates? This kind of ambiguity prompted researchers to develop alternative and more precise definitions [21, 38, 49, 55, 72, 75], in an attempt to bring more clarity to the domain, but their definitions also differ from each other.

The experimental results across the digital nudge domain have been somewhat mixed. While some interventions led to positive behavioral change, others did not. A review of nudging in HCI, for example, found that about a third of the studied nudges did not lead to a significant effect. Moreover, the authors did not uncover an obvious relationship between the applied nudge mechanism and its effectiveness [22]. Even more puzzling is the fact that particular nudges work well in one context but do not exert influence in others. An example is that of visual password strength prompts that worked in some contexts [87] but did not prompt the choice of stronger passwords in others [71, 89].

These examples do not prove that nudges in general, and cybersecurity nudges in particular, are ineffective or unreliable. What they *do* do is to highlight the strong influence exerted by the decision context, the nudge design, and their interaction. The potential interactions between the nudge and the choice architecture are not yet well enough understood and require more evidence from empirical research [22, 28]. Understanding *what* counts as a nudge and *how* nudges exert their influence is important, in terms of informing deployment decisions, and also to facilitate discussions about their ethical implications. The latter includes aspects such as their transparency, long-term and/or side effects (see [38] and [72] for ethical nudge considerations). Moreover, understanding the mechanisms behind nudges might save nudge designers from engaging in unsuccessful and expensive trials before identifying an effective nudge. Guidance to inform effective and responsibly designed nudges would be helpful.

Related work, to date, identifies at least four research areas requiring further investigation to bring us closer to understanding the nudge concept and to inform effective cybersecurity-related nudge design (Figure 1):

1. What counts as a nudge?
2. How do nudges exert their influence?

3. Which nudges should be deployed in different contexts? Context is a complex and multidimensional concept. Here, we focus on the nature of the cybersecurity-related decision as the contextual factor of interest.

4. Does the nudge influence subsequent decisions in the same general choice architecture, taking place in the absence of the nudge?




Fig. 1. The dimensions studied: the Choice Architecture with (1) the Nudge Interventions, (2) the Targeted Information Processes, (3) the Context within the Decision Types, and (4) the Varying Decisions

To support cybersecurity researchers and nudge designers, this research targets these questions by analyzing the mechanisms behind different types of cybersecurity-related nudge interventions and their impacts on various kinds of security-related decisions.

In an extensive two-part study, four different security decisions representing various types of decisions were studied as one contextual factor. These included password creation, choice of a public Wifi, smartphone encryption, and choice of a cloud service provider. In the main study, the effects of three kinds of nudge interventions were analysed in terms of their impact on the four decision types. We differentiated between simple nudges, information provision, and a combination of the two labelled a 'hybrid nudge'. In a follow-up study about two weeks later, in which the intervention was absent, the durability of the previous nudge interventions' influence was tested.

We found that the combination of a nudge and educational information provision, the 'hybrid nudge,' was at least as, or even more effective in encouraging secure user choices than a simple nudge or information provision on its own. This was true across all analysed decision contexts. Our findings indicate that enhancing nudge transparency, by providing explanatory information, does not diminish the power of the nudge and is also ethically more palatable. However, the follow-up study revealed limited durability of all the tested nudge interventions' impact, in terms of their influence on subsequent security-related decisions in the absence of the intervention.

Contributions: *First*, we clarify the nudge concept to arrive at a shared understanding to help us to distinguish different types of interventions from each other based on the human information processes they target.

Second, we analyze the impact of different cybersecurity-related interventions that are designed based on the differentiation resulting from the first contribution on different kinds of representative decisions to measure their individual and combined impact on security-related decisions.

Third, we distinguish different dimensions of security decisions (frequency and complexity) to explore the interplay between the nudge intervention and the type of the decision, as one contextual factor.

Fourth, we explore the durability of the impact of different nudge types by conducting a follow-up study requiring people to make the same decisions in a nudge-free choice architecture.

Structure: We commence with a related work section in Section 2 to address the four questions in more detail, before explaining how this research addresses each of these to derive a more holistic overview of the 'nudge' in the cybersecurity domain. We then proceed, in Section 3, to clarify the nudge-related interventions as applied in this study, and outline the decision dimensions that were used to represent different choice architectures. The

empirical study design is detailed next (Section 4), followed by the results (Section 5) and discussion in Section 6, which includes some guidelines to guide cybersecurity researchers in deploying nudges. We consider the ethical aspects of our nudges in Section 7 and the limitations of this study in Section 8. Section 9 discusses and reflects on our findings and their implications for research and practice in nudge-related research and deployment.

2 RELATED RESEARCH: NUDGING

This section explores the questions outlined in the introduction by summarizing the related work and providing relevant background information. Each subsection ends with a statement on how this research addresses each question. The final subsection considers related research into the use of nudges in cybersecurity and privacy.

2.1 Addressing the Four Questions

1. *What counts as a nudge?*

Thaler and Sunstein defined a nudge as “*any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be cheap and easy to avoid.*” [83, p.6]. Later definitions and extensions by nudge researchers highlight the fact that option-specific economic incentives should be avoided, and also that all options should be equal in terms of cost (e.g., time, effort, or social sanction) [40]. Another core element is the role of automatic cognitive processes in human decision-making and consideration of how nudges exploit these predictably to influence behaviors [21, 37, 40]. Automatic cognitive processes can be described as intuitive, effortless, fast, and unconscious [38, 61]. Examples include cognitive biases such as the hindsight bias (tendency to believe to have known outcomes beforehand), heuristics such as the availability heuristic (tendency to overestimate the likelihood of events easily available in the memory), or other learned processes such as routines.

In general, the term ‘nudging’ has been applied to a wide variety of interventions, and a number of subsequently formulated definitions and classifications extend the original one proposed by Thaler and Sunstein. Additional, related concepts such as ‘sludge’ or ‘code’ [21] have also been introduced. Some researchers suggest that the definition of a nudge might not be sharp enough to separate nudges from related interventions such as incentives or feedback mechanisms [37, 75]. Furthermore, Marchiori *et al.*, referring to decades of psychological research, conclude that nudging is not a “new” research field but “*a clever application of knowledge on behavior change and decision-making, that is now finding its way into policy-making and consumer welfare*” [49, p.3] and argue that many interventions in psychological studies could retrospectively be labelled as ‘nudges’. This kind of ambiguity is likely to contribute to the existing confusion about whether or not a tested intervention actually counts as a nudge.

For example, consider the blacklisting of weak passwords, which could be considered to nudge people towards stronger passwords. This, however, does not satisfy Thaler and Sunstein’s definition, which does not allow the removal of any pre-nudge option. What about making people pay more for software that gives them control over software updates [9]? This scheme makes automatically-updating software the cheaper option. This, too, is not a nudge because nudges ought not to introduce economic differences between pre-nudge options.

To clarify the meaning of the nudge term, we first consolidate the different definitions of nudges touched upon above, and then distinguish those from related intervention types.

In summary, the following criteria apply to an intervention that can be termed a ‘nudge’:

- **Predictability:** Nudges should influence nudgees in a predictable way and towards a predicted outcome.
- **Automatic cognitive processes:** Nudges exploit automatic cognitive processes such as well known biases and heuristics.

- **Equality of costs:** No choice should be more costly financially or economically, or in terms of time, effort, or social sanction.
- **Preservation of choices:** The nudge should not remove or ban any pre-nudge choice.

The concept of nudges, as envisioned by Thaler and Sunstein, is intended to be used “for good”, that is, to facilitate “better” decision making and behaviors. Examples are choices leading to better health, wiser financial decisions, or more secure behaviors. They emphasize this by signing copies of their book with the words: “nudge for good” (as reported by Hansen and Jespersen [38]). Even though the nudge designer might be well intended, it might sometimes be difficult to discern whether an intervention is beneficial for *all nudgees* with idiosyncratic goals and needs, or for decisions where there is no unanimity about what the best choice actually is.

Even so, there are clear cases where the nudge designer might deploy nudges to benefit him or herself or their employer. An example would be an organization deploying nudges to prompt nudgees to buy the most expensive enterprise-level antivirus software merely to increase their profit margins when a home version is all the individual needs. This kind of influence would be termed ‘sludge’ [41].

Calo [21] differentiates three different kinds of interventions, one of which is the nudge. The next is a ‘code’, which manipulates the environment to make the undesirable behavior more difficult. Consider, for example, speed bumps that require drivers to slow down if they do not want to damage their cars. The difference from the nudge concept is that a code is not as “cheap and easy to avoid” as a nudge. An oft-cited example of a nudge used to target the same behavior is a traffic sign displaying a sad face if the driver exceeds the speed limit and a happy face if the driver slows down. These could easily be ignored by the driver without undue penalty. Another difference might be the focus of the intervention. While codes aim to *decrease* an undesired behavior, nudges often aim to *increase* the incidence of the desired behavior (though exceptions are possible, see, for example, the differentiation of nudges encouraging or discouraging behavior as proposed by [42]).

The third intervention type proposed by Calo [21] is a ‘notice’, i.e. the provision of information that can take the form of information texts or reminders. Mere information provision is also distinguished from the concept of nudging by other researchers [12, 62]. According to Osman [62], this differentiation is important, because otherwise nearly every intervention could be considered a nudge and the nudge agenda would thus rendered unfalsifiable. Previous studies suggest that mere information provision does not reliably change behaviors [60], perhaps because they do not benefit from the power nudges have by targeting cognitive bias, as suggested by Renaud and Zimmermann [72], or perhaps because of the effort associated with processing the provided information.

Contribution: Before analysing the effects of nudging, this research first establishes a definition of the nudge concept to separate it from related concepts. Building on that definition, this research contributes by examining the effectiveness of different nudges and nudge-related interventions individually and in combination. In doing so, the research targets unresolved questions related to the impact of different interventions aimed at a combination of cognitive processes [28].

2. How do nudges exert their influence?

Nudges activate automatic cognitive processes, such as biases and heuristics, to encourage people to decide in a particular way. Particularly in the area of politics and public policy, this includes guidance provided by the authorities, while preserving the user’s freedom of choice, and has often been linked to the term ‘libertarian paternalism’ [12, 37, 83]. However, this kind of intervention has not been unanimously welcomed and has triggered a discussion around the ethics of nudging and the argument that libertarian paternalism is essentially a contradiction in terms.

One criticism concerns the acknowledgement that nudges essentially manipulate choice by activating *automatic* cognitive processes and nudgees might well be unaware of their influence [65, 95]. In essence, the nudgee might

not actually have the freedom to choose another option than the one they are being nudged towards, which might curtail their freedom of choice. Another criticism concerns the responsibility and power of the ‘choice architect’, i.e. the person or authority deciding on the “best” option for the nudgees. The nudgees’ opinions about the goodness of options might well differ from those of the choice architect. Moreover, choice architects themselves are equally prone to bias and heuristics, and there might also be uncertainty about what the “best” option actually is [3, 21, 94]. Furthermore, choice architects may not only design nudges for good but may, in fact, use their knowledge to manipulate users towards choosing the option that is in the interests of the choice architect rather than the nudgee, perhaps to increase the profit of their organization.

Some researchers thus argue that libertarian paternalism is an oxymoron [54] and that the much-vaunted freedom of choice cannot be assumed when nudges are applied in contexts where rational decision making is known to be deficient.

Supporters of the nudge approach, on the other hand, argue that people cannot avoid being nudged because no decision context is neutral [1, 19, 78]. From their perspective, it would be desirable to actively design choice architectures for the good of the user instead of accepting unanticipated and potentially negative effects created by happenstance. Another argument in favour of nudges is that they can help to facilitate choice by reducing complexity [17, 25].

Hansen and Jespersen [38] developed a framework to encourage the responsible use of the nudge approach by providing a detailed analysis of how nudges exert their influence. Based on psychological Dual Process Theories [44, 53, 61, 77], Hansen and Jespersen differentiate between Type 1 and 2 nudges targeting distinct cognitive processes. While Dual Process Theories differ in their details, most are based on the underlying concept of two different cognitive processes labelled Systems 1 and 2. Basically, System 1 comprises implicit, automatic, fast, and unconscious cognitive processes, and System 2 concerns explicit, controlled, conscious, slow cognitive processes. Even though researchers acknowledge that the two systems might not be completely independent and are likely to be interconnected in a serial or parallel way [47], Type 1 nudges are *primarily* aimed at fast and automatic System 1 processing. Instead, Type 2 nudges *primarily* target reflective System 2 processing via activating System 1 automatic cognitive processing. As an example for a Type 2 nudge, a password meter using color-coding to activate learned color associations targets System 1, but also “attracts reflective attention” [38] by helping people reflect on how to change the color from red to green.

Hansen and Jespersen [38] also characterize nudges in terms of their transparency to the user, which led to a classification of four types of nudges relating to how they exert their influence and how ethically acceptable their use is. In general, the use of transparent System 2 nudges was deemed most acceptable as these allow citizens to “*change their actions and behaviors in a predictable way, while simultaneously leaving them free to choose otherwise - not just as a matter of principle, but also in practice*” [38, p.24]. Other researchers argue for the use of transparent nudge interventions so that citizens are aware of them [19, 27, 57, 72]. Sunstein and Thaler agree with the importance of disclosing the presence of a nudge and advocate that deployers ought to be willing publicly to defend its “goodness” to ensure that ethically acceptable interventions are deployed [81].

The Type 1/2 and opaque/transparent differentiation proposed by Hansen and Jespersen [38] has been widely adopted. For example, recently Caraban *et al.* [22] classified nudge interventions in terms of the cognitive processes they target and their degree of transparency.

Furthermore, a number of fine-grained taxonomies have been developed to describe the mechanisms nudges can be designed to target, e.g., in terms of the effects they produce or the cognitive biases they exploit [22, 29, 42, 83]:

First, the taxonomy by Thaler and Sunstein [83] nominates six distinguishing principles: (1) defaults, (2) expect error, (3) understand mappings, (4) incentives, (5) structure complex choices, and (6) give feedback. Given our previous discussion about nudges being *more* than mere information provision, the latter should provide feedback that includes some kind of nudge: e.g., a smiley communicating goodness/badness.

Second, the British Behavioural Insights Team, also known as Nudge Unit, proposed a framework with the acronym MINDSPACE [29] to describe the cognitive biases and heuristics nudges activate: Messenger, Incentives, Norms, Defaults, Salience, Priming, Affect, Commitment, and Ego.

Third, Caraban *et al.* conducted a rigorous review of nudging in HCI and categorized nudge interventions into 23 distinct mechanisms within six categories leveraging 15 different cognitive biases and heuristics, e.g., invoking feelings of reciprocity by activating the reciprocity bias or evoking feelings of fear and loss by activating the scarcity bias [22].

Further taxonomies of the nudge's mechanism of influence are related to whether nudges *encourage* or *discourage* behavior, and to whether nudges are *externally-* or *self-imposed* [42].

While the taxonomies and research based on Dual Process Theories are promising, in terms of explaining how nudges exert their influence and in terms of selecting or designing nudges, major challenges for future research remain. Dolan *et al.* [28], for example, propose to analyse the joined-up and combined effects of different nudges across the dual processing model of the brain. Other unresolved questions include the extent to which nudges can indeed be transparent without losing their efficacy [49].

Contribution: This research contributes by shining light on the question of how nudges exert their influence by systematically analysing the effects of specific nudge interventions on different security decisions. The interventions were designed to target different cognitive processes based on the Dual Process Theory and included different degrees of transparency. The research also addresses the question of the extent to which nudges can indeed be transparent without losing their efficacy, as called for by Marchiori *et al.* [49].

3. Which nudges should be deployed in different contexts?

Metters and Grinter argue that for security-related technologies to be usable and useful there is a need to match the design to the task (choice architecture) and the context of use [76]. The same is true for nudges: Johnson *et al.* [43], Caraban *et al.* [22] and Brown [20] argue that nudges are not a “one-size-fits-all” solution, but that the effectiveness of nudges depends on the tailoring, in terms of the individual characteristics of the decision-maker, their goals, and the decision context. An indicator of the importance of analysing the target group, task and decision context are nudges that have worked successfully in one context but did not in others. A prime example is that of visual password strength meters [71, 87, 89]. Potential contextual factors that might have influenced the outcome might be that one study included information on how to improve password security [87] while another did not [71]. One study found positive effects when using dynamic information [89] while another condition in the same study using static information was ineffective. Finally, the studies also differed in the samples studied: real users with actual accounts [71] as compared to Mechanical Turk users creating passwords for hypothetical accounts [87].

Lindhout and Reniers [48] propose six steps for designing a nudge: (1) assess the situation at hand, (2) focus on individual behavior, (3) select a nudge type, (4) design, construct, and pre-test the nudge, (5) implement the nudge, and (6) evaluate the nudge. The fact that the situation, i.e. the context, is mentioned first acknowledges the power and importance of the context within which the nudge exerts its influence.

In terms of the decision context, Wansink [92] points out that the effectiveness of nudges might be limited because they attempt to impact *complex* decisions, such as the choice of a smart home device, such as voice-activated assistants. Such decisions are influenced by a large range of factors, e.g., their functionality, price, or privacy concerns, and the nudge itself is only influencing one of these. Thaler and Sunstein's conceptualization of nudging reflects a *simple* choice between more or less equivalent options. This raises the question of whether a nudge targeted at automatic processing will indeed be equal to the task of influencing complex decisions such as choosing to take actions to secure mobile devices [70], or the best ways of preserving privacy [39].

Therefore, this research analyses the type of decision, in terms of complexity and frequency, as an important contextual factor, and a starting point for analysing the impact of context on the effectiveness of nudging.

Contribution: As a step towards exploring the influence of context on the effectiveness of nudge interventions, this research considers different types of decisions: that is, *simple vs. complex*, and *frequent vs. infrequent* [70].

4. How Durable is the Nudge's Influence?

There is relatively little evidence related to how “here and now” nudging transfers to future decisions, how it encourages people to break bad habits, or to habituate to the advisable behavior over time [28, 49, 64]. This is confirmed by Caraban *et al.* [22]. The HCI and cybersecurity nudge fields, being less mature than the wider nudge literature, offer little evidence indicating the durability of a nudge's influence.

Furthermore, studies into the long-term effects of nudging from other areas sometimes report ambiguous results. For example, large-scale studies in which households received regular reports on their energy consumption showed that the effect of decreasing energy consumption was maintained even after several months [10], but also that the effect declined in the months after receiving a report and increased upon receipt of the next letter [4]. In the case of prompts encouraging stair use, the number of people using the stairs declined after the removal of the prompts. Even though differences were still significant after 12 weeks, a downward trend led the authors to conclude that an eventual regression to the baseline values would occur [14].

There is thus a need for further research to determine whether nudges can be applied as a useful intervention with long-term influence on behaviors, and this is particularly true in the cybersecurity domain, where the need to behave securely is critical.

Contribution: To explore the consequences of nudging in cybersecurity, we tested whether the effect of the different nudge interventions transferred to future decisions, in which the nudge intervention was absent in a follow-up study.

2.2 Nudge-Related Research in Cybersecurity & Privacy

Nudges and nudge-related interventions have been applied and tested in a variety of privacy- and security-related decisions.

A number of studies tested interventions designed to nudge users towards privacy-friendly apps or privacy-friendly permission settings, aiming to help users manage the disclosure of their personal information on smartphones. For example, Choe *et al.* [24] used framing to increase users' awareness of the level of privacy-invasiveness of an app. Harbach *et al.* [39] used personalized messages to reveal the potential consequences of granting permissions to apps. For example, the intervention would randomly show personal images from the smartphone's storage to make the user aware of the content the app could access if the chosen permission is granted. Balebako *et al.* [11] and Almuhimedi *et al.* [6] attempted to increase people's awareness of the risks of location disclosure. They developed tools that included privacy nudges to help users manage with whom, or with which app, they shared their location. Almuhimedi *et al.* [6], for example, sent users notifications such as “Your location has been shared 3472 times with organization Facebook and Groupon in the past 10 days” and enabled users to change their settings.

In the security area, many nudge attempts have been concerned with supporting users in creating secure passwords. The most frequently used interventions are password meters that provide users with feedback on password strength and apply nudges within the design of the instruction or the feedback. For example, Ur *et al.* varied the design of the feedback bar [88] and provided textual feedback based on the user's actual input [87], Vance *et al.* [89] used fear appeals to increase the users' motivation, and Dupuis and Khan [30] as well as Ohyama and Kanaoka [59] made use of social influence to increase password strength. Renaud and Zimmermann deployed a number of images to encourage users to choose stronger passwords, e.g., a pair of watching eyes to

activate social norms or an image of a long sausage dog to make the association between password length and strength more salient [73]. Apart from text-based passwords, von Zezschwitz *et al.* [90] even applied nudging to graphical passwords. Using background images and animations during pattern creation encouraged users to create more diverse and unpredictable Android unlock patterns.

Other security-focused interventions included a nudge based on color and ordering to encourage people to choose the most secure WiFi option [85] and images of physical “firewall” metaphors, such as a brick wall, to encourage the use of protective security measures [67].

Apart from researchers, organisations also made use of nudges in a range of settings and for different purposes. For example, Flickr.com shows the image of all people who will be able to see a posted photo to increase privacy awareness (as described by [11]) and Facebook used the image of a dinosaur that popped up to make people aware of the fact they had not updated their privacy settings (as described by [2]).

Further examples of information security and privacy nudges are provided in the review conducted by Caraban *et al.* [22] that concerned HCI-related nudging. As described in the background section, they identified 23 different mechanisms leveraging 15 different biases and heuristics that have already been applied by HCI researchers. The 23 mechanisms were clustered into the following six categories: (1) facilitate, (2) confront, (3) deceive, (4) social influence, (5) reinforce, and (6) fear. Facilitating nudges decrease the effort associated with the favorable choice, examples are default options or opt-out policies. Confronting nudges instill doubt in terms of an unwanted action by, e.g., holding a Facebook post for ten seconds before publishing it [91]. Deception nudges might make use of placebos or deceptive visualizations while social influence nudges make use of social norms and social comparisons to encourage a certain choice. Consider, for example, the password meter using peer feedback to increase password strength [30]. A reinforcing nudge could be a just-in-time prompt such as a notification shown as people are selecting privacy settings on their smartphones. Finally, fear-inducing nudges, such as fear appeals [70], aim to motivate users to avoid an unwanted consequence if the insecure option is chosen, e.g., the hacking of an online account is more likely if an insecure password is chosen.

Contribution: Applying nudges to encourage secure and privacy-friendly choices and behaviors in the digital world is a relatively new application of nudges. Initial findings were promising, but some studies delivered mixed or conflicting results. Our study builds on findings from this application domain but extends previous studies by systematically dissecting interventions to explore the effect of each intervention component individually, and in combination. This research analyses four relevant security decisions in one study thereby decreasing potential differences in the samples and times that are present when comparing different studies. A better understanding of the nudge concept, the ethics involved, the processes underlying nudging, and the influence of the context on the effectiveness of nudge interventions will hopefully help researchers to design effective, ethical and user-respecting nudges.

3 RELEVANT CONCEPTS

3.1 Defining Nudge Interventions

Our review of the nudge-related literature led us to differentiate between three different forms of interventions in this research study:

(1) Simple Nudge: Nudging has often been defined as an intervention primarily targeting automatic processing (i.e. System 1), e.g., by making use of known cognitive biases and heuristics, to make users choose a certain option [83]. Therefore, depending on the transparency of the intervention nudges might either be unaware of the nudge itself, the influence it exerts, or the reasons behind the intervention. This form of intervention is referred to as a simple nudge in this research. Hansen and Jespersen [38] explain that while nudging always affects System 1, the automatic mind, it does not necessarily involve System 2, the reflective mind. The effect is thus supposed to be non-educational. As soon as the nudge is taken away, nudges may not be able to maintain

the behavioral change if they remained unaware of the intervention, its aim or the underlying reasons for the choice they made. However, certain nudges, called Type 2 nudges by Hansen and Jespersen [38], aim to influence System 2 via System 1, e.g., by attracting reflective attention. The concept of the simple nudge, as shown in Figure 2, thus partially reaches into System 2. An example of a simple cybersecurity nudge is Von Zezschwitz *et al.*'s [90] use of background images to encourage secure choices.




Fig. 2. Classification of the interventions used in the study.

(2) Information: Information primarily targets active cognitive involvement and reflective reasoning, i.e. System 2 processing. It can take various forms of educational elements such as information texts, explanations, reminders, or textual feedback. Examples include the general terms and conditions provided to users or nutrition facts on food. Because the user is actively engaged in the decision process and aware of the reasons for their decisions, information nudges are supposed to be educational in essence. It is similar to Calo's 'notice' [21]. Examples of cybersecurity information include a list of password requirements provided next to the password entry field or privacy statements.

(3) Hybrid Nudge: The third type of intervention is a hybrid nudge, a combination of a simple nudge and information provision. Due to its nature, hybrid nudges are supposed to target automatic (System 1) as well as reflective reasoning (System 2). They rely on the cognitive processes involved in nudging, but also on the provision of information which ensures that the intervention and the reasons for encouraging a certain choice are noticed by the nudgee. Similar to information provision, the intervention fosters active nudgee involvement in the decision process. Because the nudgee actively engages in the decision process, the effect is supposed to be educational and the impact might therefore transfer to related or future decisions. An example of a hybrid cybersecurity nudge is Renaud and Zimmermann's [73] combination of a simple nudge and information provision to persuade people to choose stronger passwords.

The considerations above led to the following hypotheses:

Hypothesis H1: *Hybrid nudges, i.e. the combination of a simple nudge and information provision, are most effective in encouraging secure choices, as compared to no intervention, a simple nudge or information provision on their own.*

Simple nudges might not influence future decisions where the choice architecture does not include the intervention, because nudgees might not have learned why the option they are being nudged towards is deemed 'better' or 'more secure'. This lack of justification is also discussed by [22] in explaining why some effects do not persist when nudge interventions are absent. This leads us to the second hypothesis:

Hypothesis H2: *Information and hybrid nudges are more effective in helping people to choose the secure option in subsequent decisions where no nudge intervention is present.*

Apart from analysing the hypotheses, we were also interested in the nudgee's perspective. This has seldom been reported by previous nudge studies. We wanted to understand what nudgees perceived their reasons for a particular choice to be, and what role they felt the different nudge interventions played in their decision-making processes. By analysing these perceptions, we hoped to gain insights into the transparency of interventions and

learn about the effects of these interventions on the nudgee, in addition to merely observing and recording the choices they made.

3.2 Contextual factor decision type

To target the question about which kinds of nudges should be applied in different contexts, we chose to manipulate the type of decision users are facing as one important contextual factor. This aspect was chosen as it concerns every single decision within and outside the cybersecurity context, and also because it can be controlled within a study as compared to contextual factors embodied within the person or in third parties. Yet, we are aware that numerous contextual factors might interact. The type of decision can thus be viewed as a starting point for analysing the impact of context on the efficacy of nudging.

We classified decisions as follows:

Complexity:

(1) Simple Decision: A simple decision constitutes a more or less equivalent one-faceted choice between two or more options: an A/B-decision, e.g., to install a security update or not.

(2) Complex Decision: A complex decision constitutes a choice between multiple, non-equivalent options; a multi-faceted choice. The options differ on a range of factors, such as functionality, cost, time or effort. An example might be the choice of antivirus software.

Frequency:

(3) Infrequent Decisions: These are decisions that people have to make rarely, e.g., deciding whether to allow your health provider to share your health records with medical researchers.

(4) Frequent Decisions: These are decisions that are repeatedly or regularly made, e.g., deciding whether to click on a link in an email message or not.

4 METHODOLOGY

Our study aimed to explore the influence of different nudge interventions on different types of security-related decisions. The following sections detail how the decisions and nudge interventions were chosen before describing each decision and the study procedure in more detail.

As an overview, Table 1 provides a summary of the four decisions and their related experimental conditions.

4.1 Selection of Nudges and Decisions

As the focus of this research is on cybersecurity, we only considered security-related decisions that are taken within digital environments. To identify relevant decisions, we studied related work in terms of the decisions previously analysed. After classifying identified decisions in terms of frequency and complexity we selected one exemplary decision for each of the four combinations, i.e. 'simple vs. complex' and 'infrequent vs. frequent'. While the decisions studied in this research are not exhaustive, they serve as representatives and a starting point for exploring the influence of context, in this case, the type of decision, on the effectiveness of nudging.

For each decision, we followed the six design steps suggested by Lindhout and Reniers [48]. These commence with (1) an analysis of the situation or decision and (2) of the individual behavior within the situation. Based on the analysis' outcome (3) a suitable nudge type should be chosen, and then (4) be designed and pre-tested. For the nudge design, previous work and proven examples should be considered. Finally, the nudge should be (5) implemented and (6) evaluated. We thus analysed each decision context and related work before choosing and designing nudges matching the different decisions. Previously unvalidated materials, as well as the final nudge designs, were evaluated in pilot studies. The evaluation of the nudges and their effectiveness constitutes the core of this research.

	INFREQUENT				FREQUENT			
COMPLEX	Choice of Cloud Service				Password Creation			
	Control	Simple Nudge	Information	Hybrid Nudge	Control	Simple Nudge	Information	Hybrid Nudge
	Textual description of services	„Most popular“ banner above secure option	Summary table of differences	Combination of Information and Simple Nudge	Generation of password not previously used	Bar that changes colour and fills with increasing strength	(Dynamic) Information on what makes a strong password	Combination of Information and Simple Nudge
SIMPLE	Encryption of Smart Phone				Choice of Public WiFi			
	Control	Simple Nudge	Information	Hybrid Nudge	Control	Simple Nudge	Information	Hybrid Nudge
	Yes/No Decision	Default option “Yes”	Yes/No Decision + information about encryption	Combination of Information and Simple Nudge	Choice of a network sorted by strength of connection	Choice of a network sorted by security of connection	Choice of a network marked as secure or insecure	Combination of Information and Simple Nudges

Table 1. Overview of the decisions and conditions tested in the study.

By following the six steps, we aimed to reduce the chances of ineffective results based on an inadequate match and also to ensure relevance for the cybersecurity and HCI communities. The decisions, as well as the interventions applied in each decision context, differ so that comparisons of effects across decisions cannot be traced exclusively to the intervention or decision type. We use the classification of decision types *first*, to increase the awareness of different types of decisions and the importance of matching the nudge to its context, and *second*, to compare the results to derive hypotheses for future research.

In each of the decisions, we tested the same kinds of nudge interventions. The tested interventions were: (1) a control condition, (2) a simple nudge condition, (3) an information provision condition and (4) a hybrid nudge condition, with the intervention combining a simple nudge and information provision. The study design was thus a 2x2x4 factorial design. The decisions were varied within-subjects while the kind of intervention was varied between participants. The following sections describe these factors, the study procedure and the ethical considerations of nudging. Examples of the mock-ups are provided for each decision context, with a complete set being provided in the Supplementary Material.

4.1.1 Pilot Study. The texts, instructions and symbols used in the study were iteratively developed with a number of evaluators. These pilot studies are described in the respective decision condition.

4.1.2 Choice of Public WiFi. An example of a *simple and frequent* decision is the choice of a public WiFi to connect to, e.g., at an airport or coffee shop. From the user’s perspective, the choice between WiFi ‘A’ or ‘B’ is equal, in terms of cost and effort, both fulfilling the same need. The only difference between the WiFi options is that one WiFi option encrypts communications while another does not. This nudge design was informed by Turland *et al.*’s [85] use of color-coding and WiFi network positioning designed to nudge people towards choosing a secure WiFi.

The design of the symbol used for indicating a secure vs. an insecure network was selected based on a pilot study with 18 users who were asked for their understanding of five alternatives. The options were developed based on the literature related to security indicators: Felt *et al.* [33] found that symbols to indicate a secure connection are not universally understood by users and suggest that an indicator of an insecure state be included so that the user has a click target to provide information about the security state of the connection. Furthermore, Turland *et al.* [85] found that part of the decision-making process involved an assessment of the lock symbol.

Their open responses revealed associations such as “locked out” rather than secure/insecure. We thus conducted the pilot study to evaluate the perceived meaning of several security indicators to allow users to understand *why* the network was secure/insecure as an educational measure and combined the final solutions with text (“secure” and “insecure”) to prevent ambiguity. Images of all evaluated variants in the pilot study can be found in the Supplementary Material.

Figure 3 shows the four final conditions of the WiFi choice condition. Figure 4 shows the options and information available when people clicked on a WiFi in any of the conditions.

WiFi Airport Guest	WiFi	Airport Guest WiFi	WiFi	WiFi Airport Guest	unsecured	WiFi	Airport Guest WiFi	secured	WiFi
Airport WiFi Guest	WiFi	Guest Airport WiFi	WiFi	Airport WiFi Guest	unsecured	WiFi	Guest Airport WiFi	secured	WiFi
Airport Guest WiFi	WiFi	WiFi Airport Guest	WiFi	Airport Guest WiFi	secured	WiFi	WiFi Airport Guest	unsecured	WiFi
WiFi Guest Airport	WiFi	Airport WiFi Guest	WiFi	WiFi Guest Airport	unsecured	WiFi	Airport WiFi Guest	unsecured	WiFi
Guest Airport WiFi	WiFi	WiFi Guest Airport	WiFi	Guest Airport WiFi	secured	WiFi	WiFi Guest Airport	unsecured	WiFi
Control		Simple Nudge		Information		Hybrid Nudge			

Fig. 3. Conditions of the Public WiFi choice.



Fig. 4. Options and information available when clicking on a WiFi in the list.

- **Control:** The list of WiFis was sorted by strength of the connection. An information button was provided to allow people to find out more about the connection details of the network when they clicked on it. Besides logging the choice, we also recorded whether participants clicked on the information button.
- **Simple Nudge:** The simple nudge was based on the positioning heuristic [22], that is, people’s tendency to pick the first option of a list. Thus, to increase the number of secure choices, the list of WiFis was sorted

from most to least secure, adapted from [85], ensuring that the secure option always appeared at the top of the list.

- **Information Provision:** The list of WiFi networks was sorted by strength of the connection. Security indicators were displayed next to each network as information.
- **Hybrid Nudge:** The list of WiFis was sorted by security (simple nudge). In addition, the security indicators from the information condition were displayed next to the network name.

4.1.3 Choice of Cloud Service. This is an example of a *complex and infrequent* choice involving consideration of numerous influential factors. Choosing a cloud service provider might depend on a range of aspects, including the price of the service, the offered storage, whether and how the data in the cloud is secured, or whether the service is used by friends or colleagues. Once the decision is made, the person will probably stick to the chosen service because a subsequent change is extremely expensive in terms of time and effort. For this research, we created three fictional cloud service providers with textual descriptions of the various functions and criteria that differed in one aspect for each option to model the multi-dimensionality of the decision, while still controlling for confounding influences. One criterion was security in line with the focus of this research. The idea was to provide three options with different advantages and disadvantages while being relatively balanced overall in line with the nudge concept requiring equivalent choices.

In two pilot studies including 18 and eight evaluators, initial descriptions of the cloud service providers were evaluated in terms of comparability, understanding, and users' choices. Participants were also asked to select criteria from different areas including functionality, usability, and security, that they deemed important when selecting a cloud service provider. All criteria provided to the participants were chosen to be scalable (e.g., storage space) instead of being binary (e.g., offering a desktop app or encryption, or not). This was done to allow for differences between the services while not rendering the options too imbalanced to apply a nudge. For example, offering or not offering encryption might be an exclusion criterion for participants so that other factors are no longer taken into consideration. Next, the evaluators were asked to assign values to the criteria they deemed relevant and to adapt these values to different service providers so that one performed better than the other, but, at the same time, the difference did not exclude the lower performing service from the users' decision. The most frequently chosen criteria and values were then incorporated into the descriptions of the cloud services. Figure 5 shows the hybrid nudge condition which contains the textual description of the service providers as well as the information of the other conditions.

The final conditions were designed in the following way:

- **Control:** Texts about three fictional cloud services with various features were presented in a randomized position. Each service performed best in one feature: security, the number of installations, or data storage. All other features were equal.
- **Simple Nudge:** We used a popularity nudge taking the form of a banner with the text "most popular" above the most secure service. The nudge uses descriptive (what do the majority of users do?) and normative social influences (what do others prefer?).
- **Information Provision:** The information contained in the texts was aggregated and structured in a table to allow participants' to make a quick assessment of the differences, and to make visible which service performs best in terms of each aspect.
- **Hybrid Nudge:** This condition combined the simple nudge with the restructuring of the relevant differences to support an informed decision (see Figure 5).

4.1.4 Smartphone Encryption. This is a *simple and infrequent decision* offered by Android phones with Android 5 and below, as well as some phones upgraded to Android 6 or 7 [15]. The decision to encrypt smartphone storage is probably made only once per phone and is thus not frequent. It is also a rather simple decision because the use

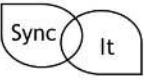






	 <p>The cloud-service provider Sync It and its servers are based in the EU. The available desktop client allows for easy access to your data from up to four different devices. A free account comes with 4GB data storage that can be extended by upgrading your account for a small fee. Sync It offers access control and end-to-end encryption. Since the start of the company in 2013 the number of people already using the service continually increased. You can share links to documents or folders and jointly edit these with other people. The service includes regular security updates and allows recovery of accidentally deleted data or attacked accounts for up to 30 days.</p>	 <p>Cloudy is a cloud-service provider with servers in the EU that started in 2013. It offers 4GB data storage for free, but can be upgraded by paying a small fee. A desktop client is available allowing you to access your data wherever you go from two different devices. The provider offers end-to-end encryption, regular security updates and access control features. Accidentally deleted data or attacked accounts can be recovered for up to 120 days. Sharing and jointly editing documents and folders with your friends is possible. Since the foundation of the company in 2013 a continuously increasing amount of people is using the service.</p>	 <p>„Lift up“ your data to the cloud and access it with the desktop client on two different devices from anywhere in the world. If your data has been accidentally deleted or your account attacked, your data can be recovered for up to 30 days. The cloud-provider’s servers are based in the EU. The company was founded in 2013 and offers free accounts with 6GB data storage. The data storage can be extended by paying a small fee. Lift up comes with access control and end-to-end encryption of your files. The number of users has steadily increased throughout the last years. Folders and documents can be shared and jointly edited with others.</p>
	free	free	free
	4 GB storage	4 GB storage	6 GB storage
	Recovery of hacked or deleted data for 30 days	Recovery of hacked or deleted data for 120 days	Recovery of hacked or deleted data for 30 days
	Installation & synchronization on 4 devices	Installation & synchronization on 2 devices	Installation & synchronization on 2 devices
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Fig. 5. Hybrid Nudge condition of the Cloud Service choice.

of encryption is preferable from a security perspective with few potential disadvantages: In particular, the initial setup takes about an hour but can be scheduled for later, and full-disk encryption requires a lock screen on the smartphone to decrypt the data once the phone is powered on [66]. However, recent studies found that more than two thirds of phone owners already use a lock screen [23] with an increasing trend [8] so that it might well be becoming a default course of action.

The layout and provided information in the mock-up were similar to that of an actual Android phone that offers optional full-disk encryption, e.g., Android 4.4. KitKat. Please note that, similar to the actual Android process, before confirming the encryption decision, all participants were presented an information page that also included the downsides of encryption as detailed above (see screenshot in the Supplementary Material). This was carried out to allow users to make a realistic decision knowing about the advantages and disadvantages. In addition to logging the decision, we recorded whether the information changed their minds about encrypting.

In the study, the participants were asked to imagine they were setting up a new phone and to decide on a number of settings as follows:

- **Control:** The participants were asked to select smartphone settings such as font size, brightness, and the background color to create the impression that they were setting up a new smartphone. One of the settings was security-related, namely the choice to encrypt the smartphone, or not. The decision was presented as a ‘Yes/No’ decision using checkboxes. If the participants selected “No” they were directly forwarded to the next setting page. If they clicked “Yes” they were shown an additional page with information concerning the smartphone encryption process, similar to the information provided by Android. They could either select “Encrypt”, which forwarded them to the next setting page or “Back” which returned them to the decision page.

- **Simple Nudge:** This nudge utilized a default setting, that has delivered robust outcomes in related work [28, 83]. The “Yes” checkbox was pre-selected, but participants could change the selection.
- **Information Provision:** Brief information was displayed similar to that provided by Android to inform the participants about the security benefits of encryption.
- **Hybrid Nudge:** This condition combined the simple (default) nudge with the information text about encryption.

Figure 6 shows the hybrid nudge condition as this contains the other conditions’ information as well.



Fig. 6. Hybrid nudge condition of the Smartphone Encryption choice.

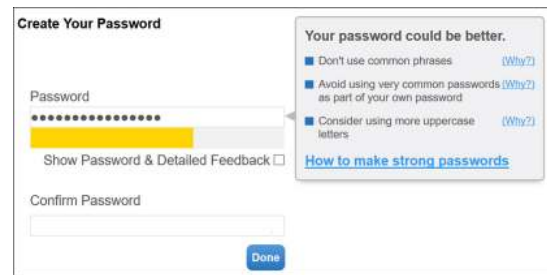


Fig. 7. Hybrid password creation nudge based on Ur *et al.* [87].

4.1.5 Password Creation. This is an example of a *complex and frequent* decision. People have to create passwords for new accounts, when they forget passwords, or when changes are forced. It is a decision with an immense variety of choices. Password creation is influenced by the type of account and data, time to log in, experience with passwords, and habit. Furthermore, complex passwords are generally more expensive in terms of memorability, effort and time it takes to type in the password.

The four conditions were designed in the following way and are depicted graphically in Figure 7:

- **Control:** This condition displayed a password entry field with the instruction to create a new, not previously used password. No password requirements were enforced to permit variance.
- **Simple Nudge:** An additional strength bar was displayed, appealing to learned associations by using color-coding (green = good/secure; red = bad/insecure) and providing users with feedback related to strength but not supporting an understanding of what a good password looks like.
- **Information Provision:** Information based on Ur *et al.*'s password meter [87] and the NIST password recommendations [36] were displayed next to the password entry field to inform users about what makes a good password. The information changed according to the participants' input, as implemented by [87].
- **Hybrid Nudge:** This condition combined the simple nudge and the dynamic information text described above to support the understanding *that* the password should be strong and *how* to achieve it (slightly adapted from [87]).

4.2 Procedure

4.2.1 Pilot Study. The complete study was tested with twelve users to assess understanding, language and functioning. The pretest resulted in slight improvements in terms of the formulation of the informed consent text and some instructions.

4.2.2 Main Study. The study was conducted using the online platform Mechanical Turk to reach a large and heterogeneous sample. Even though an artificial testing environment such as Mechanical Turk has certain limitations, we made this choice for several reasons: (a) the service allowed us to include a large sample of English native speakers, which was important as the password meter we used was based on English dictionaries, (b) we were able to conduct follow-up measurements with the same participants, (c) ethical considerations could be satisfied by properly introducing the study as such and debriefing people about nudges, which is more challenging in a field setting, and (d) it afforded comparisons with nudge studies conducted in a similar context.

To provide an overview, the study procedure is visualized in Figure 8.

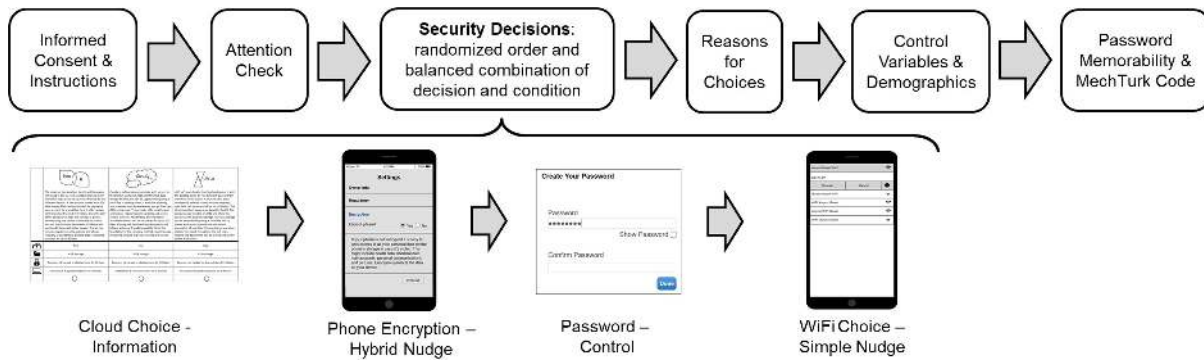


Fig. 8. Graphical depiction of the study procedure and an exemplary combination and order of security decisions.

On the first web page, the participants were provided with study information, researcher contact details and an informed consent form. If the participants agreed to participate in the study, they were presented with an attention check item based on [52] and as described by Egelman and Peer [31] to weed out inattentive participants.

If successful, participants were asked to imagine themselves within the decision scenarios presented on the following pages. The participants were provided with all four security decisions described in Section 3 and asked to make a choice. The sequence of decision types was randomized to balance sequential effects. For each security decision, one of the four interventions was randomly chosen, i.e. control, simple nudge, information, or hybrid nudge. Participants were allocated to each security decision and each type of intervention once but in a randomized order to balance sequential effects and to avoid bias due to individual effects on either the security decisions or the kind of intervention. To allow for this balancing of decision sequences and decision interventions, a ballot-box procedure was used. The ballot box included all possible combinations and was sequentially emptied. When emptied, the procedure started anew. After completion of the four security decisions, we asked the participants to explain their previous choices in an open question to (a) explore whether people mentioned and were aware of the different kinds of interventions, and (b) to explore whether other factors were important in influencing their decisions.

Participants were asked to provide some demographic information such as gender, age group, highest level of education, occupation, technical and security expertise. Their technological affinity was assessed using the Affinity for Technology Interaction (ATI) scale developed by Franke, Attig, and Wessel [34].

Additionally, we asked people about their security knowledge and attitude using the Security Behavior Intentions Scale (SeBIS) [31] and a slightly adapted version of the Human Aspects of Information Security Questionnaire (HAIS-Q) [63]. We used the items of the attitude and behavioral scale and reformulated items to address a general context rather than a work context (e.g., “passwords” instead of “work passwords”, “computers and mobile devices” instead of “computers”). Items dealing with the handling of information on paper and the reporting of colleagues were removed.

To allow us to evaluate the created password as well as the memorability thereof, and to draw a parallel between Ur *et al.*'s [87] and our research, we then asked participants to provide their created password again. They were permitted three attempts. If they failed, they were forwarded to the next page. There, we asked them for their password creation strategy (e.g., create a new password or adapt an existing password) and the method used to remember the password (e.g., remember, write down or store in a password manager). We thereby aimed to ensure that people actually created, rather than copied a password from another account or a password manager.

On the final page, participants were provided with a code for Mechanical Turk and received information about the planned follow-up study.

Relevant study material, such as additional screenshots of the security decisions and study questions are provided in the Supplementary Material.

4.2.3 Follow-Up Study. Two weeks later, a follow-up study was launched. Only participants who had previously taken part in the main study were eligible to participate in this study.

Participants were first presented with an informed consent sheet. Afterwards, they were required to pass an adapted version of the attention check items used in the main study. We then asked participants to reproduce their previously created password to help us to gauge password retention.

The follow-up study consisted of the same four decision scenarios used before, presented in random order. This time, participants were exposed to the control condition of all four decisions to check whether the impact of any of the interventions they were exposed to in the first study still endured impacting their choice. To avoid bias, due to people remembering the names or icons of the services they were nudged towards in the main study, the cloud services' names and icons were changed, as were the names of the WiFi networks. Apart from that, the descriptions of the services remained unchanged. The instructions used in the password creation task were adapted in that we now asked people to *change* their password for an important online service which we had referred to in the main study. They were asked to create a new password that they had not recently used elsewhere. To control for reuse, participants were not allowed to reuse the password they had previously provided during the main study.

Similar to the main study, participants were asked to explain their choices, and to provide their newly created password and the password creation strategy.

At the conclusion of the study, all participants, those who participated in either the main study or both studies, received a message with detailed information about the concept of nudging and the nudges applied in the study. They were provided with a link to a website providing additional information. Finally, the researchers' contact details were displayed to allow participants to ask questions or raise concerns.

4.3 Sample

The sample in the main study consisted of 450 participants, of whom 264 identified as male and 180 as female. The remaining six identified as 'other' or did not provide an answer. All participants lived in the United States and were aged 18 and over. Detailed demographics are provided in Table 2.

The follow-up study sample consisted of a subset of $N=330$ of the participants that had taken part in the main study and who returned voluntarily. Participants were recruited using Mechanical Turk as a platform and

Measure	N	%
Age (in years)		
18-29	165	36.67
30-39	162	36.00
40-49	65	14.44
50-59	38	8.44
>60	17	3.78
No answer	3	0.67
Education		
Finished High School	113	25.11
Associate Degree	69	15.33
Bachelor's Degree	196	43.78
Master's Degree	48	10.67
PhD or similar	5	1.11
Other/ No answer	18	4.00
Occupation*		
Employee/Civil Servant	311	69.11
Self-Employed	73	16.22
IT-related occupation	40	8.89
Unemployed/Seeking Employment	21	4.67
In School/University	20	4.44
Retired	8	1.78
Other/No Answer	14	2.67

Table 2. Description of the sample in terms of age, education, and occupation. *Multiple answers were possible.

compensated with \$2.50 for the main study. People also taking part in the follow-up study were awarded \$1 plus an additional \$1 bonus to increase the number of returns.

4.4 Ethical Considerations

The study was carried out in accordance with the ethics checklist provided by our university's ethics committee and guidelines for ethical psychological research [58]. The participants were recruited using the online platform Amazon Mechanical Turk designed for this very purpose. Furthermore, the study was implemented in SoSciSurvey [46] that stores data in the EU in accordance with strict EU data protection laws. Participation was voluntary and participants could withdraw at any time during the experiment without negative consequences or penalties. In line with EU data protection laws, the participants were informed about study details such as the purpose of the data collection and the way their data would be handled. Contact details were provided to facilitate asking questions or expressing concerns. The data was anonymized and analyzed on an aggregated level. In line with data economy and to enhance anonymity, only a few relevant demographics were collected, e.g., age ranges were gathered instead of exact ages.

The participants' compensation was equivalent to \$10 per hour, exceeding the USA's minimum wage. In terms of nudge-specific ethical issues, we applied ethical guidelines for nudging in IT security and privacy [72] that were rooted in well-established guidelines for ethical psychological research as suggested by the American

Psychological Association [7] and the British Psychological Society [84]. The derivation of ethical guidelines followed a similar approach to McMillan *et al.*'s [51], who analysed and categorized ethical guidelines for large-scale mobile HCI research. After the completion of the follow-up study, all participants (including those who did not return) received a message containing further information on nudging and an example of how the concept was deployed in the study they participated in.

5 RESULTS

The following section describes the findings of the main and the follow-up study for each of the analyzed decisions. Furthermore, some interesting findings from the security attitude and behavior variables are reported.

For the analysis of both studies, all participants were excluded that (a) did not pass the attention check test, (b) did not complete all four decisions, or (c) completed the survey in less than four minutes (indicating that it was unlikely that instructions and items were read thoroughly).

All tests were conducted on a significance level of $p \leq .05$. Multiple tests were accounted for by comparing the p -values with the corrected significance level calculated with the Benjamini-Hochberg procedure [13]. This procedure dynamically calculates an individual, reduced significance level for each test with the smallest p -value being compared to the strictest significance level. If the p -values displayed in the tables in the following sections exceeded the new significance level, the test was deemed non-significant. In this case, it was marked as such in the related table.

For each of the four decisions, hypothesis **H1**, concerning the assumed effectiveness of hybrid nudges, was tested by comparing the distributions of “secure” vs. “insecure” decisions of each experimental group in the main study.

Hypothesis **H2**, assuming that educational effects from the initial information or hybrid nudge interventions transferred to future decisions, was analysed by comparing the distributions of “secure” vs. “insecure” decisions of the participants in the follow-up study that had been assigned to the different experimental conditions in the main study. Further, differences in the choices of the participants that took part in both studies were analysed using paired tests for the experimental conditions.

Furthermore, the participants' perceived reasons for their choices, and the potential involvement of the nudge interventions, were analysed by coding open answers that the participants provided after completing all four decisions. The answers were coded for the experimental condition the participants belonged to in the main study, the reasons they provided for their choice, and mention of the simple nudge and/or information intervention, where appropriate. Each of the four decisions was analyzed separately following an inductive, open coding approach [50]: Researcher 1 reviewed the first quarter of the responses for all four decision contexts and developed four initial categorical systems. Mayring [50] suggests a revision of the categorical system after reviewing 10-50% of the material. This was done by having another researcher, Researcher 2, apply the categorical system to the same quarter of responses. The codings of Researchers 1 and 2 were checked for inconsistencies by calculating the inter-rater agreement and resulted in slight adaptations of the categorical systems following a discussion. This step was followed by the analysis of the complete material with the revised categorical system by Researcher 1. After the completion of the process, another quarter of the responses was cross-coded by a third researcher, Researcher 3, to again calculate inter-rater agreement as a quality check of the process. The different quarter was chosen to check whether the categorical system also applied well to the material not previously cross-coded. The third researcher was included to test whether the categorical system was understandable for a researcher that had not been involved until that point. According to Rössler [74] the amount of material critically influences the effort associated with calculating the intercoder reliability, yet, the number of test codings has to be sufficiently large. Rössler [74] suggests a minimum of 30 to 50 codings per category (such as “reason for choice” in this research) which is fulfilled given that a quarter of the material equals more than 100 participants' responses for

the main and the follow-up study each. Thus, inter-rater agreement was calculated twice: The first time to reveal areas for improvement in the categorical system, and the second time as a quality check by a person that had not been involved in the process before. In the area of product development, this is often referred to as formative and summative evaluation. The inter-rater agreement for each decision is reported in the respective section. Remaining ambiguities were solved during a discussion between the researchers. The complete codebooks, category descriptions and examples can be found in the Supplementary Material.

5.1 Choice of Public WiFi

5.1.1 Main Study. For the analysis, the participants' choices were clustered in insecure (WiFi 1, 2 and 4 in the list shown for the control condition) and secure choices (WiFi 3 and 5 in the list shown for the control condition). We recorded only 21 cases in which participants clicked on the information button to see additional information on the network type and security: $N = 7$ in the control condition, $N = 5$ in the simple nudge condition, $N = 3$ in the information condition, and $N = 6$ in the hybrid nudge condition.

Overall, a χ^2 test revealed that there were significant differences between the frequency distributions across all four conditions, $\chi^2(3) = 151.16$, $p < .001$, Cramér's $V = .58$. The effect size Cramér's V (and ϕ for 2x2 tables respectively) is indicative of a medium-sized effect [26].

To follow up that finding and test hypothesis H1, the frequencies from the control condition when no intervention was present were compared with the observed values in each experimental condition using one-sided χ^2 goodness-of-fit tests [18]. To do so, the frequency distribution of the control group was applied to the number of people in each of the experimental groups to avoid bias from unequal sample sizes.

The tests revealed that in all experimental conditions the participants chose the secure options more often (see Figure 9 and Table 3). Moreover, the differences between the individual simple nudge and information provision and the combination, the hybrid nudge condition, were significant based on the outcome of the χ^2 goodness-of-fit tests. All p -values were smaller than the Benjamini-Hochberg corrected significance levels.




Fig. 9. Number of participants selecting a secure vs. insecure Public Wifi option in the four experimental conditions of the main study.

	Main Study							Follow-Up Study	
	Secure	Insecure	Comparison	χ^2	df	p	Φ	Secure	Insecure
Control	25	86	-	-	-	-	-	30	52
Simple Nudge	88	17	Control	226.01	1	<.001	1*	22	52
Information	94	28	Control	207.87	1	<.001	.92	31	62
Hybrid Nudge	102	10	Control	301.60	1	<.001	1*	28	50
			Simple Nudge	4.35	1	.018	.14		
			Information	12.45	1	<.001	.24		

Table 3. Descriptive values and χ^2 test results of the WiFi Choice condition, ϕ = Effect Size, p = asymptotic p-value, *As the value of the test statistic exceeded the number of participants and resulted in an effect size larger 1, the value was given as an approximation.

5.1.2 Follow-Up Study. Similar to the main study, the decisions of the 330 participants that took part in both studies were clustered as secure or insecure WiFi choices, and then sorted according to the experimental group they belonged to in the main study. Table 3 shows the absolute numbers of people choosing a secure vs. insecure WiFi network for each group.

To test H2, a χ^2 test compared the decisions of the participants that belonged to different groups in the main study. The test showed no significant deviations from the values that could be expected if the experimental group and the decision were independent, $\chi^2(3) = .989$, $p = .804$, Cramér's $V = .06$. Thus, no further pairwise comparisons were calculated in this regard.

Related-samples McNemar tests conducted to compare the participants' decisions in the main study with the decisions in the follow-up study revealed that participants that had been in the simple nudge condition ($n = 74$, asymptotic $p < .001$), information condition ($n = 93$, asymptotic $p < .001$) and hybrid nudge condition ($n = 78$, asymptotic $p < .001$) chose a secure network less often than in the main study. Only for the control condition was no significant difference found, $n = 82$, $p = .093$. The results are graphically depicted in Figure 12.

5.1.3 Qualitative Analysis. The inter-rater agreement of the two researchers that independently coded about a quarter of the responses with the initial codebook was 78.52%. After coding all responses with the refined codebook that can be found in the Supplementary Material, the inter-rater agreement on a different quarter of the responses was 83.51%.

The most common reasons for choosing a public WiFi provided by 449 participants in the main, and 326 participants in the follow-up study, referred to the immediately visible information displayed on the decision screen, that is signal strength (main study $n = 188$, follow-up study $n = 144$), security (main study $n = 179$, follow-up study $n = 56$), and the position of the WiFi in the list (main study $n = 79$, follow-up study $n = 53$). Apart from that, the WiFi name (main study $n = 55$, follow-up study $n = 69$), i.e. the word order or its sound, and the general "appearance" of the WiFi (main study $n = 48$, follow-up study $n = 55$) were mentioned several times.

An exploratory analysis of differences between groups in the main study revealed that participants in the information and hybrid nudge condition, in which security information was displayed on the decision screen, more often referred to security and privacy as a reason for their choice ($n = 83$ and $n = 82$) as compared to the control and simple nudge groups ($n = 8$ and $n = 6$). Two participants in the information group said: "because it was a guest and secured connection since I care about preserving my privacy" and "I chose the secured connection because I don't trust unsecured networks in public places and only use them if I have absolutely no other choice and even then only for things that are non-sensitive."

Instead, the control and simple nudge group mentioned the position of the WiFi in the list more often ($n = 30$ and $n = 34$) as compared to the other two groups ($n = 7$ and $n = 8$). Exemplary statements are:

“There really wasn’t any reason besides it being listed first.” (Control Group, Main Study)

“The first option for WiFi is usually the strongest signal, therefore I chose that one.” (Simple Nudge Group, Main Study)

However, the position nudge itself, i.e. the intervention that the secure network was put on top, was rarely mentioned explicitly, neither by the simple nudge or the hybrid nudge group.

A total of 18 people in the information and hybrid nudge conditions mentioned an association of an unsecured network with being free and open to use without having to type in a password and therefore actively chose an unsecured network. For example, two people said that they *“chose the unsecure network as it was most likely to let me on.”* (Information Group, Main Study) and *“it was unsecured so I did not have to type in a password”* (Hybrid Nudge Group, Main Study).

The strength of the network was important across all groups, in the main study as well as in the follow-up study. It was sometimes associated with quick data transfer, popularity and legitimacy of the network, e.g., *“I figure the real airport WiFi would have a strong signal”* (Information Group, Follow-up) or *“I’m hoping that the first network on the list is the strongest, most popular. Less of a chance that I’m connecting to a spoofed network.”* (Control Group, Main Study).

Comparing the main and the follow-up study, in which no security information was directly visible any more, the number of people referring to security dropped from $n = 179$ to $n = 56$ despite the number of people in the follow-up study being reduced to 330 instead of 450 in the main study. Instead, the number of people referring to the “look and feel”, that is the sound of the WiFi name or its appearance increased from $n = 55$ to $n = 69$ and $n = 48$ to $n = 55$. Exemplary quotes are:

“Liked the word airport and wifi to be in that order.” (Hybrid Nudge Group, Follow-up)

“Since the wording was in order, it seemed the most secure and legitimate.” (Control Group, Follow-Up)

5.2 Choice of Cloud Service

5.2.1 Main Study. The description of the cloud services was equal except that each service performed “better” than the others in terms of one decision factor: Cloudy - Security, SyncIt - Installations, and Lift Up - Storage.

A comparison of an even distribution among cloud services and the results of the control condition did not detect significant differences indicating that no option was rated significantly “better” than the others beforehand, $\chi^2(2) = 2.82$, $p = .244$.

Overall, a χ^2 test revealed that there were significant differences in the frequencies with which people chose each service across the four experimental conditions, $\chi^2(6) = 26.61$, $p < .001$, Cramér’s $V = .17$.

The results of the comparisons to analyse H1 are depicted in Table 4. The procedure was similar to that described in Section 6.2.1. The results are graphically depicted in Figure 10 and indicate significant differences from the control group in terms of the participants’ choice of cloud services with the security-focused cloud service being the most frequently chosen in the information, the simple nudge and the hybrid nudge condition. The frequency patterns of the simple nudge and the hybrid nudge conditions do not differ significantly. All p -values, except for the one value that was already larger than .05, were below the individual Benjamini-Hochberg corrected significance levels.

5.2.2 Follow-Up Study. As described in the procedure section, we changed the cloud services’ names and icons in the follow-up study. Apart from that, the descriptions remained the same. Similar to the main study, the sequence of the services in the table was randomized.




Fig. 10. Number of participants selecting each Cloud Service in the four experimental conditions of the main study.

In order to test H2, a χ^2 test comparing the participants' cloud service choices that belonged to different groups in the main study showed slight deviations from the values that could be expected if the experimental group and the decision were independent, $\chi^2(6) = 12.870$, $p = .45$, Cramér's $V = .14$.

Similar to the main study, pairwise χ^2 goodness-of-fit tests were carried out to compare the control group participants' decisions in the main study to those of the participants in the other experimental conditions. The distribution of the previous control group was set as expected values when no intervention takes place. The results that are depicted in Table 4 show that the distributions of the participants that previously belonged to the control group differ significantly from the participants in the other experimental conditions. However, no clear pattern of difference is discernable. Each of the services received higher values by one of the groups while the choice distribution of the participants that previously belonged to the hybrid nudge condition was nearly equal.

To compare the participants' choices in the main study with those in the follow-up study we used a paired-samples McNemar test. To achieve this, the choices had to be binary. The participants' responses were therefore clustered into their choice of the security-focused service or one of the other two services. The test revealed that participants that had been in the simple nudge ($n = 74$, asymptotic $p = .001$), and hybrid nudge conditions ($n = 90$, asymptotic $p < .001$) chose a secure network less often than they did in the main study. Those who had been in the information ($n = 87$, asymptotic $p = .082$) or control conditions ($n = 76$, $p = .839$) demonstrated no significant differences. The results are graphically depicted in Figure 12.

5.2.3 Qualitative Analysis. After applying the initial codebook to about a quarter of the responses, the inter-rater agreement was 87.5%. Subsequently, two categories were refined and one was added leading to the codebook provided in the Supplementary Material. After coding all responses, the inter-rater agreement on a different quarter of the responses was 87.87%.

Overall, the most commonly provided reasons for choosing a cloud service related to the features the services differed on: 33.55% of the 450 people referred to security in terms of the recovery of data (main study $n = 117$, follow-up study $n = 74$), 33.11% to the amount of free storage (main study $n = 95$, follow-up study $n = 90$), and 18.67% to the number of possible installations/devices (main study $n = 57$, follow-up study $n = 45$). Apart from that, 20.67% (main study $n = 62$, follow-up study $n = 47$) said their choice was the best option or suited them best without providing functionality-related reasons, and 11.78% (main study $n = 25$, follow-up study $N = 39$) said that they relied on the "look and feel" of the service. For example, two participants said: "I liked the logo and the name the most" and "It has the catchiest name".

	Main Study								Follow-Up Study						
	Sec	Inst	Stor	Comparison	χ^2	df	p	V	Sec	Inst	Stor	χ^2	df	p	V
Control	38	26	38	-	-	-	-	-	24	16	36				
Simple Nudge	69	14	26	Control	31.95	2	<.001	.38	30	21	23	7.91	2	.02	.23
Information	51	36	30	Control	6.78	2	.017	.17	26	37	24	26.33	2	<.001	.39
Hybrid Nudge	74	19	29	Control	28.60	2	<.001	.34	31	29	30	9.31	2	.01	.23
				Simple Nudge	.85	2	.655	.06				1.44	2	.49	.09
				Information	17.47	2	<.001	.27				2.44	2	.30	.12

Table 4. Descriptive values and χ^2 test results of the Cloud Service Choice condition, Sec = Security, Inst= Installations, Stor= Storage, Cramér's V = Effect Size, p = asymptotic p-value

In terms of the hypotheses, the information table to allow for a quick comparison was never explicitly mentioned by the participants as a reason for their choice or an intervention. However, more participants mentioned security in terms of the recovery of data in the two conditions that saw the table, the information and hybrid nudge condition, as compared to participants in the other two conditions (63.25% vs. 36.75%) in the main study. In contrast to that, the social norm nudge was explicitly mentioned by 74 of the 240 (30.83%) participants that were assigned to the information or hybrid nudge condition as a reason for their choice in the simple nudge and hybrid nudge condition. Sometimes just the popularity argument was mentioned as the reason for choosing the service, e.g., “*I went with the most popular option.*” Beyond that, other participants did not only select the service because of its popularity, but because of the advantages they associated with this attribute:

“The services were fairly similar, but I choose ‘Cloudy’ as it was the most popular. I figure if it’s that popular, then it must be good and reliable. If it was poor quality, then few people would use it.” (Condition Hybrid Nudge)

“I went with the most popular one because it had good features and if it was the most popular it must also mean that it’s reliable.” (Condition Simple Nudge)

People in the main study’s control condition slightly more often referred to other security features that did not differ across services ($N = 13$ vs. $N = 8, 6$, and 3) or the “look and feel” of the service ($N = 13$ vs. $N = 6, 2$, and 3).

In the follow-up study, the only categories where numbers increased despite the reduced number of participants were the “look and feel” of the service and security features that did not differ across accounts. Thus, when the information was not present, many people relied on aspects other than the content or actual differences. For example, a participant in the information condition in the main study chose the service because “*it had a longer day recovery*” while in the follow-up study the same person made the choice because it “*looked the best*”. The difference between the number of people in the information and hybrid nudge conditions that referred to the security argument in terms of the recovery of data as compared to the other two conditions decreased (55.41% vs. 44.59%) in the follow-up study.

5.3 Encryption of Smartphone

5.3.1 Main Study. Based on hypothesis H1, the participants’ choices were clustered in insecure, non-encrypted and secure, encrypted choices. Of the 450 participants, only 16 decided to change their mind in terms of encrypting the phone after reading the information on the confirmation page. That is, the participants first chose to encrypt but then changed their minds. These were relatively evenly distributed across conditions: control $N = 4$, simple nudge $N = 5$, information $N = 3$, and hybrid nudge $N = 4$.

A χ^2 test showed significant differences across the four conditions, $\chi^2(3) = 12.73$, $p = .005$, Cramér's $V = .17$. Follow-up one-sided χ^2 goodness-of-fit tests [18] revealed significant differences between the frequency distribution of the control group and the frequency distribution of the information, simple nudge and hybrid nudge condition. The differences were largest between the control and the hybrid nudge condition as graphically shown in Figure 11. The descriptive and test values are summarized in Table 5. All p -values were below the Benjamini-Hochberg corrected significance levels.




Fig. 11. Number of participants selecting the secure, encrypted option vs. insecure, non-encrypted option in the four experimental conditions of the main study.

	Main Study							Follow-Up Study	
	Secure	Insecure	Comparison	χ^2	df	p	Φ	Secure	Insecure
Control	89	35	-	-	-	-	-	70	22
Simple Nudge	103	20	Control	8.69	1	.002	.19	67	17
Information	85	19	Control	5.09	1	.012	.16	59	20
Hybrid Nudge	89	10	Control	16.05	1	<.001	.28	58	14
			Simple Nudge	2.76	1	.048	.12		
			Information	4.42	1	.018	.15		

Table 5. Descriptive values and χ^2 test results of the Smartphone Encryption Choice condition, ϕ = Effect Size, p = asymptotic p -value

5.3.2 Follow-Up Study. The choices of the 330 people that participated in the follow-up study were clustered into insecure, non-encrypted and secure, encrypted choices.

A χ^2 test comparing the decisions of the participants that belonged to different groups in the main study showed no significant deviations from the values expected if the experimental group and the decision were

independent, $\chi^2(3) = 1.097$, $p = .778$, Cramér's $V = .06$. Thus, no further pairwise comparisons between groups were calculated.

Related-samples two-sided McNemar tests, comparing the participants' decisions in the main study with the decisions in the follow-up study, revealed that participants that had been in the hybrid nudge condition ($n = 72$, $p = .012$) chose to encrypt less often as compared to the main study. For the control condition ($n = 92$, $p = .210$), the simple nudge condition ($n = 84$, $p = .289$), and the information condition ($n = 79$, $p = .092$) no significant differences were found. Figure 12 illustrates the findings.




Fig. 12. Comparison of the percent of secure choice in the main vs. the follow-up study separated by experimental conditions.

5.3.3 Qualitative Analysis. When the initial codebook was applied to a quarter of the responses the inter-rater agreement was 75.86%. Applying the refined codebook (see Supplementary Material) to another quarter of responses led to an increase of the inter-rater agreement to 83.93%.

The major reasons for encrypting the phone were security and privacy ($n = 273$ Main Study, $n = 196$ Follow-up Study):

"I wanted to encrypt my phone to protect my data." (Hybrid Nudge Group, Main Study)

"I care about my privacy and want to avoid any chances of information being stolen." (Control Group, Follow-up Study)

Other reasons for encrypting the phone included the encryption being of general importance ($n = 32$ main Study, $n = 18$ follow-up study), encrypting out of habit ($n = 15$ main study, $n = 12$ follow-up study), and encrypting because of missing disadvantages ($n = 30$ main study, $n = 25$ follow-up study). One example for each would be:

"I think that encryption is important and is almost a necessity." (Control Group, Main Study)

"I usually encrypt my devices, so I decided to do it with this one." (Simple Nudge Group, Main Study)

"I always favor encrypting over not encrypting, simply because it just feels like it's a safe move to do it regardless of the situation, I don't see cause to NOT do it so I do do it" (Control Group, Follow-up Study)

Some people that decided to encrypt referred to the information that included the need for a passcode when encrypting and mentioned the irreversibility of the process. However, while this made 16 people change their mind (see above), the large majority decided to proceed anyway, e.g.:

"I am concerned with privacy, so I would prefer to have my data encrypted. I always lock my phone, so having to unlock it doesn't bother me. Also, I have my phone's data backed up, so I'm not worried about losing anything." (Simple Nudge Group)

Reasons for not encrypting the phone included perceived disadvantages ($n = 21$ main study, $n = 16$ follow-up study), lack of knowledge in terms of encryption ($n = 22$ main study, $n = 13$ follow-up study), and a perceived lack of need to do so ($n = 24$ main study, $n = 23$ follow-up study):

"Encrypting the whole phone seems like a hassle." (Control Group, Follow-up Study)

"I actually don't know what it means to encrypt a phone, so I chose no on this one assuming I could go back and change it at a later time once I understand it more." (Simple Nudge Group, Main Study)

"I chose not to encrypt my phone because I felt like it wasn't needed in order for me to stay secure." (Simple Nudge Group, Main Study)

In terms of the hypotheses, five people referred to the information text. However, the simple nudge, i.e. the pre-selected default yes-option, was never explicitly mentioned. People in the control and simple nudge group slightly more often decided not to encrypt (22.76%) as compared to the information and hybrid nudge group (10.38%). No other obvious patterns or large differences became apparent when comparing the qualitative responses across groups.

5.4 Password Creation

5.4.1 Main Study. In terms of password creation, the password strength, as a score from 0 to 100, password entropy in bits, and password length as the number of characters in the password were analysed. The password strength score was based on the heuristics and algorithms used by Ur *et al.* [87]. Descriptive values can be found in Tables 6, 7, and 8.

Participants who said that they reused a previous password ($N = 23$ or 5.11%), or used a password manager ($N = 16$ or 3.55%), were excluded from the quantitative password strength analysis. Similar to Ur *et al.*'s [87] study, the reasoning was that the nudge could not influence those who did not create a password, i.e. those who reused another password or used a password manager to create one for them. After excluding these participants, the conditions comprised the following number of participants: Control $n = 101$, Simple Nudge $n = 105$, Information $n = 48$ and Hybrid Nudge $n = 54$. The numbers in the last two conditions were smaller since the two conditions were subdivided into dynamic password information as analysed by Ur *et al.* [87] and static password information. Yet, only the first was of interest for this study.

As the password strength and entropy values were not metric but ordinal, and the length values not normally distributed, non-parametric tests were conducted. Overall, a Kruskal-Wallis test showed that there were significant differences in password strength ($H(3) = 36.67$, $p < .001$), entropy ($H(3) = 40.84$, $p < .001$), and length ($H(3) = 33.80$, $p < .001$) across all four conditions. Following up that finding to analyse H1, Mann-Whitney-U tests were conducted to localize the effects. The results of these are displayed in Tables 6, 7, and 8. Within the tables, the values of the effect size r around .1 indicate a small, values around .3 a medium, and values from .5 a large effect [26]. Further, the medians of the strength, length and entropy measures are graphically depicted in Figure 13. The tables indicate that password strength, length, and entropy were significantly higher in the simple nudge, information, and hybrid nudge condition as compared to the control group. Comparing the single simple nudge and information condition to the combined hybrid nudge condition revealed that the values in the hybrid nudge condition were significantly higher than those in the simple nudge condition. However, the difference to the information condition was smaller and not significant.

5.4.2 Follow-Up Study. As in the main study, $n = 11$ participants who said that they reused an existing password or used a password manager or other tool to generate their password were excluded from the analysis. This




Fig. 13. Medians of the password strength entropy, and length values in the four experimental conditions of the main study.

	PW Strength Main Study									PW Strength Follow-Up Study		
	<i>M</i>	<i>SD</i>	<i>Md</i>	<i>Comparison</i>	<i>Z</i>	<i>df</i>	<i>p</i>	<i>r</i>		<i>M</i>	<i>SD</i>	<i>Md</i>
Control	31.55	25.80	31.69	-	-	-	-	-	-	36.50	23.36	38.20
Simple Nudge	41.06	23.28	39.33	Control	-2.80	1	.005	.20		37.29	23.60	38.06
Information	50.38	25.42	50.96	Control	-3.94	1	<.001	.32		39.38	28.34	39.33
Hybrid Nudge	58.27	27.72	65.25	Control	-5.29	1	<.001	.44		36.63	25.79	32.35
				Simple Nudge	-3.71	1	<.001	.31				
				Information	-1.58	1	.058	.17				

Table 6. Descriptive values and Mann-Whitney-U test results of the password strength values, *M* = Mean, *SD* = Standard deviation, *Md* = Median, *Z* = standardized test statistic, *df* = degrees of freedom, *p* = level of significance, *r* = Effect size

	PW Entropy Main Study									PW Entropy Follow-Up Study		
	<i>M</i>	<i>SD</i>	<i>Md</i>	<i>Comparison</i>	<i>Z</i>	<i>df</i>	<i>p</i>	<i>r</i>		<i>M</i>	<i>SD</i>	<i>Md</i>
Control	54.74	22.82	52.31	-	-	-	-	-	-	58.42	19.96	56.87
Simple Nudge	67.71	23.88	65.39	Control	-4.13	1	<.001	.29		62.72	20.68	56.87
Information	75.54	29.15	66.73	Control	-4.51	1	<.001	.37		61.87	22.97	56.87
Hybrid Nudge	78.20	26.33	72.35	Control	-5.44	1	<.001	.45		60.67	23.90	57.86
				Simple Nudge	-2.59	1	.005	.21				
				Information	-1.27	1	.102	.13				

Table 7. Descriptive values and Mann-Whitney-U test results of the password entropy values, *M* = Mean, *SD* = Standard deviation, *Md* = Median, *Z* = standardized test statistic, *df* = degrees of freedom, *p* = level of significance, *r* = Effect size

resulted in the following distribution: After excluding these participants, the conditions comprised the following number of participants: Control $n = 71$, Simple Nudge $n = 72$, Information $n = 37$ and Hybrid Nudge $n = 44$.

	PW Length Main Study								PW Length Follow-Up Study		
	<i>M</i>	<i>SD</i>	<i>Md</i>	<i>Comparison</i>	<i>Z</i>	<i>df</i>	<i>p</i>	<i>r</i>	<i>M</i>	<i>SD</i>	<i>Md</i>
Control	9.82	3.16	9.00	-	-	-	-	-	10.11	2.93	10.00
Simple Nudge	11.18	3.39	11.00	Control	-3.24	1	.001	.23	10.74	3.07	10.00
Information	12.56	4.61	12.00	Control	-4.01	1	<.001	.33	10.62	3.40	10.00
Hybrid Nudge	12.89	3.79	12.00	Control	-5.07	1	<.001	.42	10.45	3.40	9.50
				Simple Nudge	-2.93	1	.002	.24			
				Information	-1.06	1	.145	.11			

Table 8. Descriptive values and Mann-Whitney-U test results of the password length values, *M* = Mean, *SD* = Standard deviation, *Md* = Median, *Z* = standardized test statistic, *df* = degrees of freedom, *p* = level of significance, *r* = Effect size

A Kruskal-Wallis test to analyse H2 showed that neither the password strength ($H(3) = .311, p = .958$), nor the password entropy ($H(3) = .915, p = .822$), or password length ($H(3) = .955, p = .812$) in the follow-up study differed when data was sorted by the experimental groups the participants belonged to in the main study (see Figure 14). Thus, no further pairwise comparisons were conducted.

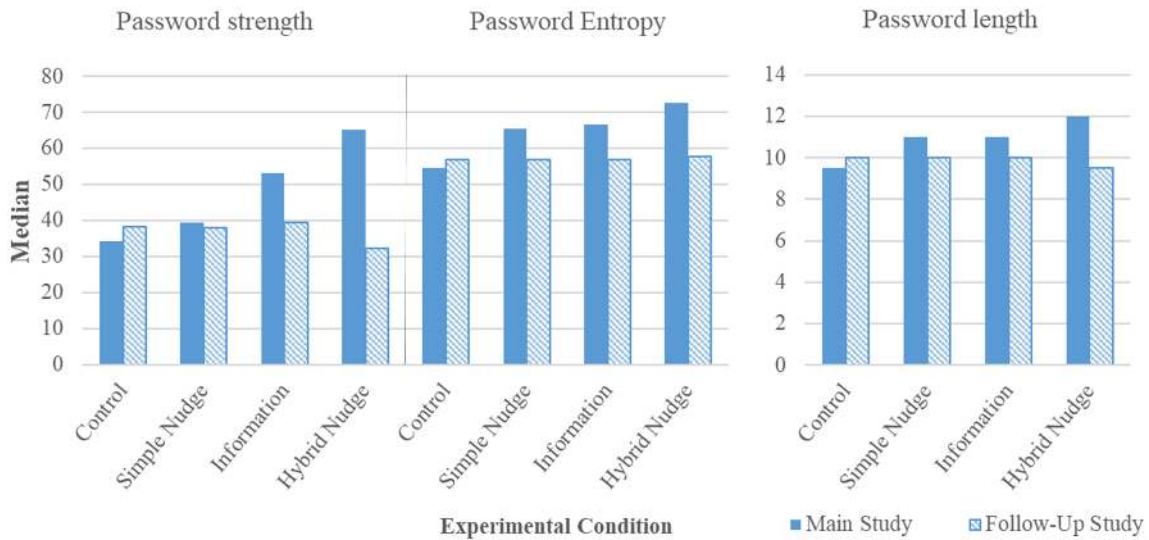


Fig. 14. Comparison of the password strength, entropy and length medians in the main vs. the follow-up study separated by experimental conditions.

To analyse differences between the participants' password choices in the main study and in the follow-up study paired-sample Wilcoxon tests were conducted. They revealed that password strength ($Z = -1.340, p = .180$), entropy ($Z = -.416, p = .677$), and length ($Z = .000, p = 1$) did not differ significantly for the participants that had been assigned to the control condition in the main study. The results were similar for participants that had been previously been in the simple nudge condition: Password strength ($Z = -1.462, p = .144$) and length ($Z = -1.358, p = .174$) did not differ significantly, neither did entropy ($Z = -2.079, p = .038$) given the corrected significance

level of $p = .0375$. Participants who had been assigned to the hybrid nudge condition created significantly weaker passwords in the follow-up study when the nudge was no longer present in terms of strength ($Z = -3.236$, $p = .001$), entropy ($Z = -3.260$, $p = .001$), and length ($Z = -3.132$, $p = .002$). Results for the participants of the information condition were mixed. They created shorter ($Z = -2.445$, $p = .014$) and less complex ($Z = -2.528$, $p = .011$) passwords. However, the overall strength value did not differ significantly ($Z = -1.581$, $p = .114$).

5.4.3 Qualitative Analysis. The qualitative analysis included all participants regardless of password strength and whether they reused a password or made use of a password manager. Following the exploratory nature of the qualitative analysis, we were interested in what made people create a certain password, whether they were aware of the intervention, how it was perceived, and why, in some cases, people did not generate a new password with the help of the intervention. The inter-rater agreement using the initial codebook for a quarter of the responses was 83.44%. After refining some category definitions the final codebook shown in the Supplementary Material was applied to the complete data set. The inter-rater agreement of another quarter increased to 87.78%.

In the open answers most people referred to security ($n = 127$ main study, $n = 69$ follow-up study) and memorability ($n = 160$ main study, $n = 120$ follow-up study) as factors they considered when creating a new password. These were often mentioned in combination, examples include: “*I made a password that was easy enough to remember but hard to guess.*” (Hybrid Nudge Group, Main Study) or “*Something secure yet easy enough for me to remember*” (Control Group, Follow-up Study).

Apart from that, many people provided information on how they created a password, i.e. the character sets they included or the information and strategies the password was based on ($n = 162$ main study, $n = 103$ follow-up study).

“*I just used the first odd phrase that came to mind and added a number to it.*” (Control Group, Main Study)

“*I based it on a phrase with which I’m familiar and then I changed out some of the letters to be numbers or symbols.*” (Information Group, Main Study)

“*It is based on a historical day.*” (Information Group, Follow-up study)

Some people said that they either reused or adapted a previous password or that they followed a strategy that they always apply to create passwords ($n = 34$ main study, $n = 33$ follow-up study), e.g., “*I went with a variation of something I’d used before*” (Information Group, Follow-up Study). A total of 14 people in the main study stated to have used a password manager or generator to create or store the password.

In terms of the hypotheses, 15 people in the main study said they followed the information and suggestions provided by the tool, and only 3 people mentioned the strength bar in their response:

“*The password suggestion gave helpful tips on how to make a more secure password.*” (Hybrid Nudge Group, Main Study)

“*I tried to make security box green, so I made an inordinately long password.*” (Hybrid Nudge Group, Main Study)

Considering the unequal sample sizes, no obvious differences across groups or relevant patterns were visible.

5.4.4 Password Memorability. At the end of the main study, and the beginning of the follow-up study after approximately two weeks, participants were asked to reproduce their password. After the main study, participants were permitted three attempts to enter their password and received feedback whether the password matched. Of the 450 participants, 427 (94.89%) were able to reproduce their password. When asked to re-enter their password at the beginning of the follow-up study, only 83 of 330 (25.15%) participants correctly remembered their passwords.

Measure	M	Md	SD	Measure	M	Md	SD
HAIS-Q Attitude				HAIS-Q behavior			
Password Management	12.42	13.00	2.64	Password Management	12.82	13.00	2.32
Email Use	12.86	14.00	2.63	Email Use	12.79	14.00	2.60
Internet Use	12.8	14.00	2.37	Internet Use	11.06	11.00	2.62
Social Media Use	12.9	14.00	2.38	Social Media Use	11.39	11.00	2.76
Mobile Devices	12.74	14.00	2.55	Mobile Devices	12.67	14.00	2.68
Information Handling*	4.24	5.00	1.28	Information Handling*	4.34	5.00	1.12
Incident Reporting*	8.31	9.00	1.81	Incident Reporting*	7.84	8.00	1.84

Table 9. Descriptive values of the HAIS-Q sub scales (M = Mean, Md = Median, SD = Standard Deviation) *Note: The sub scales Information Handling and Incident reporting consisted of 1 and 2 items respectively instead of 3.

Measure	M	Md	SD
SEBIS			
Device Securement	4.04	4.25	.94
Password Generation	3.75	3.75	.84
Proactive Awareness	3.70	3.80	.96
Updating	3.87	4.00	.80

Table 10. Descriptive values of the SEBIS sub scales (M = Mean, Md = Median, SD = Standard Deviation)

5.5 Security Attitude and Behavior

The sample's security attitude scale score of the HAIS-Q was $M = 76.26$, $Md = 81.00$ ($SD = 13.22$). As a slightly reduced version of the HAIS-Q was used, the maximum score was 90 points based on a 5-point scale for each of the 18 items as compared to 105 points for the original 21 items. The behavior score of the HAIS-Q was $M = 72.92$, $Md = 76.00$ ($SD = 12.10$). The difference between the sample's attitude and behavior score was significant with $t(448) = 9.92$, $p < .001$. The descriptive values for each of the HAIS-Q's subscales are provided in Table 9.

In terms of the sample's security behavior intention measured with the SEBIS, the sample's mean score was $M = 3.83$, $Md = 3.88$ ($SD = .80$) on a scale ranging from 1 to 5. The results for each sub-area are shown in Table 10.

The technological affinity measured with the ATI scale with scores ranging from 1 to 6 was $M = 4.06$, $Md = 4.00$ ($SD = .91$).

6 DISCUSSION

This section will first provide a brief summary of the main findings referring to the hypotheses derived in the introduction. Next, some relevant findings and their implications will be discussed in more detail before exploring the limitations of this study and the resulting potential for future work.

6.1 Hypotheses 1 & 2

The first hypothesis assumed that the combination of a simple nudge and information provision, a hybrid nudge, would be most effective in encouraging secure choices, as compared to no intervention, a simple nudge or information provision on its own. An initial comparison of the values or frequency distributions across experimental conditions, in each of the four decision contexts, revealed that the type of nudge intervention did indeed influence the participants' choices. Pairwise comparisons revealed that the hybrid nudge was most

effective in encouraging secure choices in the decision contexts ‘Encryption of Smartphone’, and ‘Choice of Public WiFi’.

In terms of password creation, the hybrid nudge was more effective than the simple nudge but failed to reach significance when compared to information provision, even though the descriptive values appeared to indicate that participants in the hybrid nudge condition made slightly more secure choices.

With respect to the cloud service choice, the hybrid nudge was more effective than information provision on its own, but no more effective than the social norm (simple) nudge.

Overall, the results speak in favor of hypothesis 1 and indicate that increasing the transparency of the intervention, by informing people about why one option is deemed more favourable than another, does not diminish the power of the nudge. Indications are that the transparency might even increase its power to influence. The qualitative results support the finding in that people in the pure information and hybrid nudge conditions referred to security as a reason for their choices more often than participants in other conditions. The latter referred to non-functional features, such as option placement or ‘look and feel’.

We have to consider whether the combination of a nudge and information, here referred to as a ‘Hybrid Nudge’, still counts as a nudge intervention as envisioned by its inventors. In this regard, Sunstein [78, p. 207] states that “*there is no opposition between education on the one hand and nudges on the other. Many nudges are educative. Even when they are not, they can complement, and not displace, consumer education.*” For future research, it would be interesting to determine whether the effect still persists if the nudge and information are even more tightly coupled, e.g., by directly referring to the simple nudge in the information. An example related to this research would be “We re-ordered the names of the publicly available WiFis to support you in choosing a high connectivity network that also secures your data transfer and prevents unwanted access by using encryption.”

The second hypothesis assumed that the educational impact of the intervention would endure impacting future decisions. This could not be confirmed. The values and frequency distributions in the follow-up study revealed no significant differences when clustered by the main study experimental condition. Only in terms of the choice of a cloud service were slight differences observed, but without a clear pattern that would speak for or against the hypothesis. This suggests that the influence of the nudge intervention, especially the hybrid nudge, was not as durable as we hoped. They exerted an impact when present, but the desired secure choice did not manifest when the intervention was absent. This might be seen as a slight reflection of the finding that the participants’ HAIS-Q security attitude score was higher than their behavioral score. The finding can be interpreted in the way that the participants care about security and are willing to behave securely when assisted by interventions, but that in the absence of the intervention maintaining that behavior is difficult. The implications of these findings are discussed below.

6.2 Type of Decision and Nudge

6.2.1 Choice of Public WiFi. With regards to public WiFi choice, our results confirm Turland *et al.*’s [86] findings in that the combination of the WiFi position in the list, and some form of information, a text plus a color-coded security indicator in our case, was most effective. However, in contrast to their study, our findings show that positioning secure WiFis at the top of the list was effective on its own. Furthermore, even fewer people chose a secure WiFi when no other information was provided in our control condition, as compared to the Android default that differentiates between “trusted”, “secure” and “open” as analyzed by Turland *et al.* [85]. In Turland *et al.*’s study [86], some participants associated the padlock symbol with being locked out or needing a password. Similarly, the open answers in this study revealed that 4% of our participants interpreted an unsecured network with the network being open and accessible. The low percentage reporting this confusion indicates that our design of the security symbol and removal of the padlock from the signal strength indicator was indeed effective in communicating the security of the network instead of indicating that they had been locked out.

The qualitative findings in our study show that when security information was visible on the decision page, participants often referred to security and it influenced their decisions, thus indicating that they were aware of the information and understood its importance. The positioning nudge seemed to have been successful because many participants mentioned this in explaining their choice across all conditions, but the participants did not seem to realize that a nudge intervention had influenced their decision.

Yet, it is rather depressing that less than 5% of participants clicked on the information button to obtain additional information about the network. This was especially noticeable in the control and simple nudge conditions in which participants were not shown any visible information supporting their decision-making apart from the signal strength indicator. This indicates that relevant information for making a decision should be displayed on the same page and as close to the point the person will be focusing on as possible. Even so, it is clear that clicking on the “i”, which seems trivial and almost effortless, is seen as an additional step that only very few participants felt compelled to take. Another interesting finding is that when no other information was available on the decision page, more people relied on the “look and feel” and the sound of the WiFi’s name. Even though these were designed to sound equal and consisted of the same words, there might have been a preference towards a certain word order and thus some side-effect of the WiFi name. Apart from including salient security information, another implication for security engineers and designers might be an increased awareness that even seemingly trivial design elements, not related to security, should be given consideration as these might influence security decisions.

Reflecting on the information and hybrid nudge condition, it should be acknowledged that the color-coding used for the lock symbols to make them more easily distinguishable in the tiny information symbol could also have “nudged” users. It is possible that the color-coding activated a learned association red-danger/bad and green-safe/good similar to the feedback bar in the password creation task, thereby functioning as a simple nudge. This example demonstrates the difficulty of clearly separating actual nudges from related interventions.

6.2.2 Cloud Service choice. In the case of the multi-faceted, complex cloud service decision, the relatively equal distribution of user choices across accounts in the control condition can be viewed as an indication that the options, as intended, were viewed as relatively equivalent by the participants. The combination of a table facilitating comparison (information) together with a popularity nudge, was as effective in encouraging users to choose the secure option as the popularity nudge on its own. Furthermore, both were more effective than information provision. An explanation could be that the information table, on its own, allowed users to make a quick decision following their wishes and needs without necessarily thinking about security. The popularity nudge changed that: participants seemed to be aware of its presence because it was explicitly mentioned by about a third of participants. They associated a number of different positive aspects with this, such as a recommendation, a high number of users, or pressure for the service provider to maintain a high-quality service.

Reflecting on the design of the cloud service choice, we faced the difficulty of identifying scalable as compared to binary features for security and functionality that the services differed on to ensure relatively equivalent choices. While this was relatively easy for functionality for which the number of installations and data storage were chosen in the pilot study, this was more difficult for security. Prominent security features such as encryption were either given or not given so that we again relied on the pilot study to have participants choose a scalable feature deemed as relevant. From the selection also including the support team response time or the frequency of security update checks the participants chose the recovery time of lost data due to an attack or technical problem. While this might not be the foremost example for a security feature the nudge worked as intended in encouraging users to select the service performing best with regards to this feature.

Finally, we found indications for the same effect that manifested in the choice of a public WiFi. Especially when no other information was salient, more people tended to rely on the “look and feel” of the cloud service provider. This was supported by a larger number of people referring to this as a reason for their decision in the

control condition and the follow-up study during which all participants were assigned to the control condition. This, again, highlights the importance of providing salient security information close to the point where the user makes a decision. Furthermore, it shows how even minimal differences, even those that do not appear to be security-relevant, such as the logo or name, can influence security-related decisions. The example suggests that choice architects should holistically consider all visible design elements and the context within which a decision is made.

6.2.3 Encryption of Smartphone. In terms of smartphone encryption, it is interesting that the majority of participants in the control condition had already decided to encrypt, without any intervention being required. This could be due to personal experience or the media generally displaying encryption as something positive and important. Still, the interventions increased the number of people deciding to encrypt even more, especially so when brief information about the benefits of encryption was paired with the simple nudge, which made encryption the default choice. The fact that only 3.5% of users changed their minds even though the confirmation page also included the possible disadvantages of encryption (e.g., the need for a passcode to be entered each time the phone is powered on) could be explained by a couple of factors: participants might prefer to stick to the decision they had already made, they might consider the disadvantages negligible, especially in an artificial test environment, or they did not read the text. The qualitative data analysis does not provide an overwhelmingly convincing explanation but does show that, of the 35 people who mentioned the confirmation page, 26 decided to encrypt their phones, despite having to unlock the phone at each use, and in the knowledge that they could not reverse the process.

6.2.4 Password Creation. In line with Ur *et al.*'s [87] findings, our study shows that the hybrid nudge is more effective than the colored feedback bar on its own. This supports the assumption that it is not sufficient to indicate *when* the password is secure, but that it helps to explain *why* this is the case. This is especially true as security is a complex topic often invisible to the user, and the impact of adding a certain character on the password's security is not readily understood. With regards to this aspect, security and privacy decisions might differ from other more easily visible or experiential effects, such as understanding that walking the stairs might positively impact fitness.

Similar to Ur *et al.* [87], our results suggest that the combination of the following is most effective in encouraging secure passwords:

- (1) a feedback bar to provide an indication of the current security level (i.e. the current state), to attract attention and encourage reflection,
- (2) information about what makes a good password (i.e. the aim), and
- (3) supporting information on how to get from the current password to the recommended 'good password'.

Furthermore, the descriptive values indicate that it might be more effective than the information provision alone (even though the significance level was not reached). This finding should be followed up in future research with a larger sample. Our sample was reduced because we had to exclude those participants who reused passwords or used password managers. The qualitative responses did not reveal differences in terms of password creation strategies across groups. It seems that the password suggestions and strength bar did not impact password creation strategies in general but helped by aligning the perception of a secure and technically secure password, and increased the extent to which people adopted strategies to create secure passwords.

6.2.5 Type of Decision. One implication of the findings above, assuming that they will transfer to actual environments, would be that choice architects should consider employing hybrid nudges (consisting of a nudge and educational information) rather than automatically reaching for traditional simple nudges. Moreover, the results of the follow-up study suggest that the intervention should always be present because one exposure to a nudge can not be guaranteed to influence future decisions.

Dissipation of the impact of a nudge might not matter as much for infrequently-made decisions. For example, the choice of a cloud service or smartphone encryption are one-off decisions. In this regard, Dolan *et al.* [28] differentiate between nudges that can be considered ‘triggers’ that only exercise a short-term impact, nudges that impact behavior in the long-term, and others that might be self-sustaining. That is, once the choice is made, a certain behavior is self-sustaining, e.g., once the smartphone is encrypted, users are unlikely to do a factory reset to reverse the decision.

The impact of ordering on WiFi choice, however, might not exhibit this quality, with the decision-maker reverting to insecure choices when the nudge is no longer present. However, the long-term effects of a frequent exposure to the nudge in the case of frequent decisions, such as password creation or WiFi choice, is not yet well understood. If the same or a very similar intervention is present to support every future decision, does the efficacy of the intervention endure? Or, will the effect decrease and even lead to habituation or resistance over time? The answer to this question might depend on the type of decision and the bias or heuristic the nudge is exploiting or relying on. Potential influencing factors might be the transparency of the nudge, its intrusiveness and the ease with which it can be ignored.

6.2.6 Type of Nudge. For each of the four decision contexts, a nudge that matched the decision and the format of the choice was deployed based on previous work and pilot studies. In doing so, this study tested a social norm nudge (Choice of Cloud Service), positioning nudge (WiFi Choice), default nudge (Smartphone Encryption), and feedback nudge using color-coding (Password Creation). The nudges deployed in the study were all effective, but their impact differed in a variety of ways:

- **Type 1 vs. Type 2 nudges:** The feedback bar indicating password strength according to Hansen and Jespersen’s [38] differentiation would be classified a Type 2 nudge because the nudge makes consequences salient and attracts reflective attention, thereby primarily targeting System 2 via System 1. In contrast, the positioning nudge would be classified a Type 1 nudge as it only targets an automatic cognitive process, i.e. automatically picking the first option, that users might not necessarily be aware of.
- **Transparent vs. opaque nudges:** The social norm nudge was transparent to the users and was often mentioned as having influenced the decision process in the participants’ responses. The default nudge, however, might have been less transparent. Even though participants could see the pre-selected option, this was not explicitly mentioned and might therefore not have been noticed and perceived to be an intervention that could have influenced their decisions.
- **Active vs. passive nudges:** While the feedback bar encouraged users to actively test password changes to advance the feedback bar, the default nudge did not require active, cognitive involvement and deliberate action on the part of the participants.

While the differences did not seem to impact effectiveness, being aware of them is important, in terms of ethics and legislation. For example, the use of default nudges was recently subject to a discussion around cookie settings. New regulations require users to actively agree to cookie settings thereby forbidding the use of defaults in that context [32]. It would thus be interesting to follow up the discussion around active and passive decisions and to develop and test alternatives that fulfill the requirement of supporting active decision-making, such as using nudges that are supposed to target reflective reasoning such as the feedback bar deployed in this study.

6.3 Implications for Cybersecurity Researchers

We can now return to the four scenarios we tested:

6.3.1 Simple & Infrequent decisions. The nudges in the smartphone encryption decision made people more likely to choose the secure option, so for these kinds of decisions nudges seem to be indicated. To ensure that the

nudgee knows the implications of the choice, and to respect their need for autonomy, information should be provided when the choice is being made.

6.3.2 Simple & Frequent decisions. Also in terms of simple and frequent decisions, in this research, the WiFi choice, the combination of a simple nudge and information provision was most effective. Yet, the lack of durability of the effect suggests that it would be a good idea to always provide a nudge and salient information close to where the decision is to be made. This seems especially relevant as participants consulted security-irrelevant aspects, such as the WiFi name, in making their decisions when no other information was directly visible. Yet, security researchers should be careful when selecting the intervention (and evaluate long-term effects) to avoid habituation or resistance if users are repeatedly confronted with the intervention.

6.3.3 Complex & Infrequent decisions. When people make decisions where multiple aspects need to be considered, they will often focus on one particular aspect more than others [35]. This propensity can be exploited by making the nudge the aspect they pay most attention to. Our cloud service scenario used a social nudge but a wide variety of nudges could be used to help people to choose the most secure or privacy-respecting option. Information, on its own, helped support user decision-making but did not perform as well in encouraging secure decisions as the hybrid and simple nudges. One difficulty in designing nudges for complex choices lies in the assumption of relatively equivalent options. These might not be a given in real-life settings. One service or software might obviously outrank the secure option by having more features or due to the person knowing that their friends and colleagues use it. Thus, before applying a nudge, security researchers or designers should first analyse the features of the choice and other potential influences such as social aspects. Depending on the outcome, researchers could undertake measures to make the options more equivalent before applying a nudge or consider a more appropriate intervention.

6.3.4 Complex & Frequent decisions. This is a challenging area within which to nudge, as demonstrated e.g., by [71] studying different password nudges. In addition to considering the challenges associated with complex decisions, as detailed above, security researchers must also consider the challenges associated with the long-term effects of nudging. Based on our findings, a nudge is indeed indicated, although the poor durability of the intervention suggests that while people responded to the nudge they did not consider the long-term requirements of the nudge. Furthermore, any attempt to encourage stronger passwords must give the person a way to cope with the cognitive load and the need to retain that password securely. Combining the intervention with encouraging the use of a password manager might be a way of doing this. Applied to other complex and frequent decisions, this suggests that researchers should explore ways to (1) help users make a secure choice, and (2) reduce the effort associated with the decision and/or the frequency of the decision (and thus also the nudge intervention). After having considered a complex and frequent decision supported by a nudge intervention the first time, users should then have options or tools to facilitate subsequent choices. For example, in terms of privacy settings, users might have the option to make their previous choice a default setting.

7 THE ETHICS OF NUDGING

The results of this research suggest that increasing the transparency of nudge interventions does not diminish the power of the nudge on the decision. We achieved this transparency by complementing the nudge with educational measures that provide reasons for the intervention nudging the person towards a certain choice. Across all analyzed decision contexts, the number of people choosing the secure option increased when the transparent hybrid nudge was deployed.

In terms of the discussion surrounding the ethics of nudging, this research suggests that nudge interventions do not need to influence covertly solely by targeting processes and biases people might not necessarily be aware of. Such nudges might be ruled unacceptable and a form of manipulation. Combining simple nudges

with information enhances transparency, and thereby the political acceptability of nudge interventions. It seems possible to enhance the transparency of the simple nudges without compromising their effectiveness. As an example, both the positioning nudge (WiFi choice), that users might not have been aware of, and the popularity nudge (cloud service choice), often mentioned by participants in their free-text responses, influenced their choice of the more secure option. The passive default nudge used in the smartphone encryption condition was effective, but so was the nudge used in the password creation condition where users had to actively reason and adapt their password to adjust the feedback bar.

Previous work suggests that making nudge interventions transparent fosters nudgees' acceptance thereof, particularly if the choice they are being nudged towards aligns with their personal goals or is generally widely approved of. For example, a field experiment at a train station to encourage healthy food choices by positioning food in the store led to more healthy food purchases, regardless of whether the nudge intervention was opaque or disclosed to the customers [45]. This indicates, similar to our study, that nudges can indeed be effective if users are aware of their presence and influence and underlying rationale. Most customers expressed a positive attitude towards the intervention because they, too, wanted to make healthy food choices. In an interview study, Rapp *et al.* [68] found that many participants felt anxiety and worry when they perceived not being in control of a behavior change which can also be viewed as an argument for increasing the transparency of interventions. Furthermore, if the intervention and its goals are transparent to the nudgees, it might also be easier to detect a mismatch between nudge interventions and the nudgees' interests, and thus the opportunity to act if it becomes clear that the nudge intervention is not being deployed "for good" as judged by the nudgees themselves, the acid tests for nudges as argued by their creators [83].

In summary, our findings suggest that even if the simple and hybrid nudges are equally effective, the latter is preferable from an ethical perspective. Using a hybrid nudge may enhance the nudgee's ability to resist the nudge if the option they are being nudged towards does not align with their own interests.

It must be acknowledged that cybersecurity-related nudging is not always appropriate or desirable. Given the complexity and unknown nature of the actual user's context, nudges might well be harmful rather than helpful in some contexts. Sunstein [80] explains that a nudge could produce confusion or reactance in the target audience. This might happen if people consider the nudge to violate their autonomy [79]. For example, people may not want to be nudged towards a specific WiFi: they may prefer an automatic connection to save time and consider it to be their right to decide without prompting.

Finally, Sunstein explains that nudges, if not wisely designed and implemented, could lead people to engage in compensatory behaviors. For example, if a nudge is designed to lead people towards stronger passwords, it could seem to work. But, in the background, people are writing their passwords down, thereby weakening the password. It might be better to nudge people towards the use of a password manager so that they will be less likely to engage in compensatory behaviors.

8 LIMITATIONS & FUTURE WORK

This research is subject to some limitations, which we acknowledge here. We furthermore present an outlook on future work.

First, the study was conducted using Amazon Mechanical Turk which is a somewhat artificial context. The context supports comparisons with related research, e.g., by [87]. It also allowed us to control for and purposefully influence certain variables, and to ensure that adequate information was provided to all participants in terms of the nudges used in the study to comply with ethical guidelines for psychological research. However, as the participants made artificial decisions with no direct real-life consequences, it is unclear whether the decisions will transfer to real-life decisions in a similar context. Controlled artificial settings and field studies have their

own advantages and disadvantages and can complement one another. For future research, it would be important to confirm our findings in more realistic settings and to test transferability and external validity.

Furthermore, the Amazon Mechanical Turk sample might differ from the general public. One example is the sample's reported security attitude and intention. The relatively high HAISQ values and the SEBIS scores exceeding the mean values reported for password creation in the original study [31, p.7] can be viewed as an indication for this.

Second, this study was undertaken as a first step in highlighting the consequences or long-term effects of nudge interventions on future decisions. However, within the limits of this study and the difficulties of having the same people responding several times, only one additional follow-up study was conducted after a two-week lapse. The results are highly relevant but warrant further investigation. It would be interesting to analyse more points of time for a longer duration to compare groups that have been provided with the nudge intervention repeatedly, as compared to only one exposure.

Third, the study comprised different decisions and nudge mechanisms to maximize comprehensiveness and relevance for the HCI community. Still, as different interventions were deployed, it was not possible to compare effects across decisions, but only within conditions. Future research examining the influence of context on the effects of nudges might benefit from studying the same intervention across different decision types or different interventions within one decision context.

Fourth, asking people for their perceived reasons for picking a certain option revealed some interesting insights. Yet, it would have added further value to directly ask people whether and how the interventions impacted their choices. This was not done in the main study to avoid priming effects in the follow-up study, nor did we do this in the follow-up study. Future studies should consider including a suitable compromise, such as only asking participants who do not intend to return for the follow-up study, or relying on the participants' memory (supported by screenshots) after they have completed the follow-up study.

Fifth, even though the order of the security decisions and conditions has been randomized to balance sequential effects, a general priming effect, in terms of security, cannot be ruled out. It is possible that the instructions or the use of words such as "insecure/secure" in certain decision contexts primed the participants to think about security and influenced their subsequent choices. Future studies, perhaps conducted in the wild, might consider placing the security decisions within a realistic context to decrease potential priming effects towards an explicit focus on security. For example, the decision to encrypt a smartphone might be placed within the larger context of setting up a new phone, as attempted in this study, or users might be asked to create a password for a relevant application that would actually be used to shift the focus towards the application.

Sixth, we plan to derive a framework for cybersecurity researchers and developers to convey principles to inform security interface design based on the insights we gained from this study.

9 CONCLUSION

This research explored the power of nudges on security decisions either with, or without, associated educational information, in different contexts. The four decision types were representative of complex vs. simple and frequent vs. infrequent decisions.

We commenced this study by asking four questions. We now briefly review our answers to these questions:

First, across several definitions, we established that a nudge comprises the following criteria: predictability of the influence and outcome, involvement of automatic cognitive processes, equality of choice costs, and retention of all pre-nudge choices.

Second, we found that the *hybrid nudge* (a nudge combined with information provision), was at least as, and sometimes more effective in encouraging secure choices as the simple nudge on its own, indicating that increasing the intervention's transparency did not diminish its efficacy.

Third, the detected effect was visible in all decision contexts including frequent vs. infrequent and simple vs. complex decisions. This indicates that these decision dimensions might not play a role for one-time deployment of nudges. However, the frequency of a decision should be considered with regards to future decisions (see below).

Fourth, neither our most effective hybrid nudge's influence nor that of a single simple nudge or information provision continued to exert its influence in a follow-up study two weeks later. When the nudge intervention was absent, previous exposure did not continue to influence subsequent decisions.

Based on the outcome of this study, the following recommendations can be made:

- When contemplating using simple nudges, consideration should be given to combining it with information provision given that hybrid nudges are more transparent. They help nudges to understand *why* they are being nudged towards a particular option.
- Hybrid nudges, by being transparent to nudges, can foster active rather than passive decision making to maximize acceptability and to allow nudges quickly to detect alignment or mismatches between the choice architect's and their own interests.
- The nudge intervention's influence might not endure influencing future decisions where the nudge is not present in the choice architecture. Choice architects should be cognisant of the frequency of the decision being influenced i.e. frequent or infrequent. The frequency of the nudge's deployment should match the frequency of the decision to be made. The long-term effects of regularly-deployed nudges, or the design of interventions that *do* endure to influence future decisions, should be a topic of future research.

10 ACKNOWLEDGEMENT

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. Additionally, the second author thanks TU Darmstadt for allowing to visit on the Mercator programme funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 251805230/GRK 2050.

REFERENCES

- [1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–41.
- [2] Reed Albergotti. 2014. Facebook Rolls Out Privacy Checkups to All 1.3 Billion Users. Sep 4 <https://blogs.wsj.com/digits/2014/09/04/facebook-rolls-out-privacy-checkups-to-all-1-3-billion-users/> Accessed 13 May 2018.
- [3] R Alberto and V Salazar. 2012. Libertarian paternalism and the dangers of nudging consumers. *King's Law Journal* 23, 1 (2012), 51–67.
- [4] Hunt Allcott. 2011. Social norms and energy conservation. *Journal of Public Economics* 95, 9-10 (2011), 1082–1095.
- [5] Hunt Allcott and Sendhil Mullainathan. 2010. Behavior and energy policy. *Science* 327, 5970 (2010), 1204–1205.
- [6] Hazim Almuhiemedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! : A Field Study on Mobile App Privacy Nudging. ACM, New York, NY, USA, 787–796.
- [7] American Psychological Association. 2016. Ethical Principles of Psychologists and Code of Conduct. <http://www.apa.org/ethics/code/index.aspx>.
- [8] AndroidCentral. 2017. More Android phones are using encryption and lock screen security than ever before. <https://www.androidcentral.com/more-android-phones-are-using-encryption-and-lock-screen-security-ever> Accessed 04 December 2019.
- [9] Terrence August, Robert August, and Hyoduk Shin. 2014. Designing user incentives for cybersecurity. *Commun. ACM* 57, 11 (2014), 43–46.
- [10] Ian Ayres, Sophie Raseman, and Alice Shih. 2013. Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *The Journal of Law, Economics, and Organization* 29, 5 (2013), 992–1022.
- [11] Rebecca Balebako, Pedro G Leon, Hazim Almuhiemedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2011. Nudging users towards privacy on mobile devices. In *Proceedings of the CHI Workshop on Persuasion, Nudge, Influence and Coercion*. ACM, New York, NY, USA, 1–4.

- [12] Adrien Barton and Till Grüne-Yanoff. 2015. From libertarian paternalism to nudging - and beyond. *Review of Philosophy and Psychology* 6, 3 (2015), 341–359.
- [13] Yoav Benjamini and Yoel Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)* 57, 1 (1995), 289–300.
- [14] Avril Blamey, Nanette Mutrie, and Aitchison Tom. 1995. Health promotion by encouraged use of stairs. *British Medical Journal* 311, 7000 (1995), 289–290.
- [15] Elcomsoft Blog. 2017. Android Encryption Demystified. <https://blog.elcomsoft.com/2017/05/android-encryption-demystified/> Accessed 04 December 2019.
- [16] Jennifer Swindell Blumenthal-Barby and Hadley Burroughs. 2012. Seeking better health care outcomes: the ethics of using the “nudge”. *The American Journal of Bioethics* 12, 2 (2012), 1–10.
- [17] Jennifer s Blumenthal-Barby and Aanand D Naik. 2015. In defense of nudge–autonomy compatibility. *The American Journal of Bioethics* 15, 10 (2015), 45–47.
- [18] Jürgen Bortz and Christof Schuster. 2011. *Statistics for Human and Social Scientists: Limited Special Edition (Statistik für Human-und Sozialwissenschaftler: Limitierte Sonderausgabe)*. Springer, Berlin/ Heidelberg, Germany.
- [19] Thom Brooks. 2013. Should we nudge informed consent? *The American Journal of Bioethics* 13, 6 (2013), 22–23.
- [20] Patrick Brown. 2012. A nudge in the right direction? Towards a sociological engagement with libertarian paternalism. *Social Policy and Society* 11, 3 (2012), 305–317.
- [21] Ryan Calo. 2014. Code, Nudge or Notice? *Iowa Law Review* 99 (2014), 773.
- [22] Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. 2019. 23 Ways to Nudge: A Review of Technology-Mediated Nudging in Human-Computer Interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 503.
- [23] Pew Research Center. 2017. Many smartphone owners don’t take steps to secure their devices. <https://www.pewresearch.org/fact-tank/2017/03/15/many-smartphone-owners-dont-take-steps-to-secure-their-devices/> Accessed 04 December 2019.
- [24] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. 2013. Nudging people away from privacy-invasive mobile apps through visual framing. In *Proceedings of the IFIP Conference on Human-Computer Interaction*. Springer, Berlin/Heidelberg, Germany, 74–91.
- [25] Robert B Cialdini and Melanie R Trost. 1998. Social influence: Social norms, conformity and compliance. In *The Handbook of Social Psychology* (4 ed.), Daniel T. Gilbert, Susan T. Fiske, and Gardner Lindzey (Eds.). McGraw-Hill, New York, 151–192.
- [26] Jacob Cohen. 2013. *Statistical power analysis for the behavioral sciences*. Routledge, London, UK.
- [27] Russell DiSilvestro. 2012. What does not budge for any nudge? *The American Journal of Bioethics* 12, 2 (2012), 14–15.
- [28] Paul Dolan, Michael Hallsworth, David Halpern, Dominic King, Robert Metcalfe, and Ivo Vlaev. 2012. Influencing behaviour: The mindspace way. *Journal of Economic Psychology* 33, 1 (2012), 264–277.
- [29] Paul Dolan, Michael Hallsworth, David Halpern, Dominic King, and Ivo Vlaev. 2010. MINDSPACE: influencing behaviour for public policy. <https://www.instituteforgovernment.org.uk/publications/mindspace> Accessed 5 Dec 2019.
- [30] Marc Dupuis and Faisal Khan. 2018. Effects of peer feedback on password strength. In *Proceedings of the APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, New York, NY, USA, 1–9.
- [31] Serge Egelman and Eyal Peer. 2015. Scaling the security wall: Developing a security behavior intentions scale (sebis). In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 2873–2882.
- [32] EU GDPR Compliant. 2018. Cookies Consent under the GDPR. February 14 <https://eugdprcompliant.com/cookies-consent-gdpr/> Accessed 5 January 2020.
- [33] Adrienne Porter Felt, Robert W Reeder, Alex Ainslie, Helen Harris, Max Walker, Christopher Thompson, Mustafa Emre Acer, Elisabeth Morant, and Sunny Consolvo. 2016. Rethinking connection security indicators. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. Usenix, Berkeley, CA, USA, 1–14.
- [34] Thomas Franke, Christiane Attig, and Daniel Wessel. 2019. A personal resource for technology interaction: development and validation of the affinity for technology interaction (ATI) scale. *International Journal of Human-Computer Interaction* 35, 6 (2019), 456–467.
- [35] Gerd Ed Gigerenzer, Ralph Ed Hertwig, and Thorsten Ed Pachur. 2011. *Heuristics: The foundations of adaptive behavior*. Oxford University Press.
- [36] Paul A Grassi, Michael E Garcia, and James L Fenton. 2017. Digital identity guidelines. *NIST Special Publication* 800 (2017), 63–3.
- [37] Pelle Guldberg Hansen. 2016. The definition of nudge and libertarian paternalism: Does the hand fit the glove? *European Journal of Risk Regulation* 7, 1 (2016), 155–174.
- [38] Pelle Guldberg Hansen and Andreas Maaløe Jespersen. 2013. 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* 4, 1 (2013), 3–28.
- [39] Marian Harbach, Alexander De Luca, Nathan Malkin, and Serge Egelman. 2016. Keep on Lockin’ in the Free World: A Multi-National Comparison of Smartphone Locking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 4823–4827.

- [40] Daniel M Hausman and Brynn Welch. 2010. Debate: To nudge or not to nudge. *Journal of Political Philosophy* 18, 1 (2010), 123–136.
- [41] Crawford Hollingworth and Liz Barker. 2017. BE360: Protecting Consumers from ‘SLUDGE’. 28 November <https://www.research-live.com/article/features/be360-protecting-consumers-from-sludge/id/5031182> Accessed January 2020.
- [42] Julian House, Elizabeth Lyons, and D Soman. 2013. *Towards a taxonomy of nudging strategies*. Rotman School of Management, University of Toronto.
- [43] Eric J Johnson, Suzanne B Shu, Benedict GC Dellaert, Craig Fox, Daniel G Goldstein, Gerald Häubl, Richard P Larrick, John W Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504.
- [44] Daniel Kahneman and Patrick Egan. 2011. *Thinking, fast and slow*. Vol. 1. Farrar, Straus and Giroux, New York, NY, US.
- [45] Floor M Kroese, David R Marchiori, and Denise TD de Ridder. 2015. Nudging healthy food choices: a field experiment at the train station. *Journal of Public Health* 38, 2 (2015), e133–e137.
- [46] DJ Leiner. 2014. SoSci survey (Version 2.5. 00-i).
- [47] Y Lin, M Osman, and R Ashcroft. 2017. Nudge: Concept, Effectiveness, and Ethics. *Basic and Applied Social Psychology* 39, 6 (2017), 293–306.
- [48] P Lindhout and Genserik Reniers. 2017. What about nudges in the process industry? Exploring a new safety management tool. *Journal of Loss Prevention in the Process Industries* 50 (2017), 243–256.
- [49] David R Marchiori, Marieke A Adriaanse, and Denise TD De Ridder. 2017. Unresolved questions in nudging research: Putting the psychology back in nudging. *Social and Personality Psychology Compass* 11, 1 (2017), e12297.
- [50] Philipp Mayring. 2014. *Qualitative content analysis: theoretical foundation, basic procedures and software solution*. SSOAR Open Access Repository, Klagenfurth, Austria. <https://nbn-resolving.org/urn:nbn:de:0168-ssoar-395173>
- [51] Donald McMillan, Alistair Morrison, and Matthew Chalmers. 2013. Categorised ethical guidelines for large scale mobile HCI. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 1853–1862.
- [52] Adam W Meade and S Bartholomew Craig. 2012. Identifying careless responses in survey data. *Psychological Methods* 17, 3 (2012), 437.
- [53] Gabriela Michalek, Georg Meran, Reimund Schwarze, and Özgür Yildiz. 2016. *Nudging as a new "soft" policy tool: An assessment of the definitional scope of nudges, practical implementation possibilities and their effectiveness*. Technical Report. Economics Discussion Papers.
- [54] Gregory Mitchell. 2004. Libertarian paternalism is an oxymoron. *Nw. UL Rev.* 99 (2004), 1245–1277.
- [55] Philippe Mongin and Mikael Cozic. 2018. Rethinking nudge: not one but three concepts. *Behavioural Public Policy* 2, 1 (2018), 107–124.
- [56] NSW Government. 2016. NSW Behavioural insights team. <https://www.dpc.nsw.gov.au/programs-and-services/behavioural-insights/> Accessed 5 Dec 2019.
- [57] Thomas RV Nys and Bart Engelen. 2017. Judging Nudging: Answering the Manipulation Objection. *Political Studies* 65, 1 (2017), 199–214.
- [58] European Federation of Psychologists’ Association. 2005. Meta-Code of Ethics. https://www.bdp-verband.de/binaries/content/assets/beruf/efpa_metacode_en.pdf.
- [59] Takahiro Ohshima and Akira Kanaoka. 2015. Password Strength Meters using Social Influence. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. Usenix, Berkeley, CA, US, 1–2. Poster.
- [60] Folke Ölander and John Thøgersen. 2014. Informing versus nudging in environmental policy. *Journal of Consumer Policy* 37, 3 (2014), 341–356.
- [61] Magda Osman. 2004. An evaluation of dual-process theories of reasoning. *Psychonomic Bulletin & Review* 11, 6 (2004), 988–1010.
- [62] Magda Osman. 2016. Nudge: How Far Have We Come? (*Economia. History, Methodology, Philosophy* 6 (2016), 557–570. Issue 4.
- [63] Kathryn Parsons, Dragana Calic, Malcolm Pattinson, Marcus Butavicius, Agata McCormac, and Tara Zwaans. 2017. The human aspects of information security questionnaire (HAIS-Q): two further validation studies. *Computers & Security* 66 (2017), 40–51.
- [64] Charlie Pinder, Jo Vermeulen, Benjamin R Cowan, and Russell Beale. 2018. Digital behaviour change interventions to break and form habits. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 3 (2018), 15.
- [65] Thomas Ploug and Søren Holm. 2015. Doctors, patients, and nudging in the clinical context-four views on nudging and informed consent. *The American Journal of Bioethics* 15, 10 (2015), 28–38.
- [66] Android Open Source Project. n.d.. Encryption. <https://source.android.com/security/encryption/full-disk> Accessed 04 December 2019.
- [67] Fahimeh Raja, Kirstie Hawkey, Steven Hsu, Kai-Le Clement Wang, and Konstantin Beznosov. 2011. A brick wall, a locked door, and a bandit: a physical security metaphor for firewall warnings. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*. ACM, New York, NY, USA, 1.
- [68] Amon Rapp, Maurizio Tirassa, and Lia Tirabeni. 2019. Rethinking Technologies for Behavior Change: A View from the Inside of Human Change. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 4 (2019), 1–30.
- [69] Imran Rasul and David Hollywood. 2012. Behavior change and energy use: is a ‘nudge’ enough? *Carbon Management* 3, 4 (2012), 349–351.
- [70] Karen Renaud and Marc Dupuis. 2019. Cyber security fear appeals: unexpectedly complicated. In *Proceedings of the New Security Paradigms Workshop*. Costa Rica, 42–56.

- [71] Karen Renaud, Joseph Maguire, Verena Zimmermann, and Steve Draper. 2017. Lessons Learned from Evaluating Eight Password Nudges in the Wild. In *Proceedings of the LASER Workshop*. USENIX, Berkeley, CA, USA, 25–37.
- [72] Karen Renaud and Verena Zimmermann. 2018. Ethical guidelines for nudging in information security & privacy. *International Journal of Human-Computer Studies* 120 (2018), 22–35.
- [73] Karen Renaud and Verena Zimmermann. 2019. Nudging folks towards stronger password choices: providing certainty is the key. *Behavioural Public Policy* 3, 2 (2019), 228–258.
- [74] Patrick Rössler. 2017. *Content Analysis (Inhaltsanalyse)*. Vol. 2671. UTB.
- [75] Evan Selinger and Kyle Powys Whyte. 2012. What counts as a nudge? *The American Journal of Bioethics* 12, 2 (2012), 11–12.
- [76] Diana K Smetters and Rebecca E Grinter. 2002. Moving from the design of usable security technologies to the design of useful secure applications. In *Proceedings of the Workshop on New Security Paradigms (NSPW)*. ACM, New York, NY, USA, 82–89.
- [77] Keith E Stanovich and Richard F West. 2000. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences* 23, 5 (2000), 645–665.
- [78] Cass R Sunstein. 2015. Nudges Do Not Undermine Human Agency. *Journal of Consumer Policy* 38, 3 (2015), 207–210.
- [79] Cass R Sunstein. 2017. Forcing People to Choose Is Paternalistic. *Mo. L. Rev.* 82 (2017), 643–667.
- [80] Cass R Sunstein. 2017. Nudges that fail. *Behavioural Public Policy* 1, 1 (2017), 4–25.
- [81] Cass R Sunstein and Richard H Thaler. 2003. Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review* 70 (2003), 1159–1202. Issue 4.
- [82] Behavioural Insights Team. 2011. Behavioural insights team annual update 2010–11. Cabinet Office: London, UK. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/60537/Behaviour-Change-Insight-Team-Annual-Update_acc.pdf Accessed 14 October 2020.
- [83] Richard H Thaler and Cass R. Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, CT, US.
- [84] The British Psychological Society. 2014. Code of Human Research Ethics. <http://www.bps.org.uk/publications/policy-and-guidelines/research-guidelines-policy-documents/research-guidelines-poli>.
- [85] James Turland, Lynne Coventry, Debora Jeske, Pam Briggs, and Aad van Moorsel. 2015. Nudging towards security: Developing an application for wireless network selection for Android phones. In *Proceedings of the British HCI conference*. ACM, New York, NY, USA, 193–201.
- [86] James Kevin Turland. 2016. *Aiding information security decisions with human factors using quantitative and qualitative techniques*. Ph.D. Dissertation. School of Computing Science, Newcastle University.
- [87] Blase Ur, Felicia Alfieri, Maung Aung, Lujo Bauer, Nicolas Christin, Jessica Colnago, Lorrie Faith Cranor, Henry Dixon, Pardis Emami Naeini, Hana Habib, Noah Johnson, and William Melicher. 2017. Design and evaluation of a data-driven password meter. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 3775–3786.
- [88] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2012. How does your password measure up? the effect of strength meters on password creation. In *Proceedings of the USENIX Security Symposium (USENIX Security)*. USENIX, Berkeley, CA, USA, 65–80.
- [89] Anthony Vance, David Eargle, Kirk Ouimet, and Detmar Straub. 2013. Enhancing password security through interactive fear appeals: A web-based field experiment. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS)*. IEEE, New York, NY, USA, 2988–2997.
- [90] Emanuel von Zezschwitz, Malin Eiband, Daniel Buschek, Sascha Oberhuber, Alexander De Luca, Florian Alt, and Heinrich Hussmann. 2016. On Quantifying the Effective Password Space of Grid-based Unlock Gestures. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia (MUM)*. ACM, New York, NY, USA, 201–212.
- [91] 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 ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, New York, NY, USA, 2367–2376.
- [92] Brian Wansink. 2004. Environmental factors that increase the food intake and consumption volume of unknowing consumers. *Annual Review of Nutrition* 24 (2004), 455–479.
- [93] Markus Weinmann, Christoph Schneider, and Jan vom Brocke. 2016. Digital nudging. *Business & Information Systems Engineering* 58, 6 (2016), 433–436.
- [94] Mark White. 2013. *The manipulation of choice: Ethics and libertarian paternalism*. Palgrave Macmillan, New York, NY, US.
- [95] Karen Yeung. 2016. The forms and limits of choice architecture as a tool of government. *Law & Policy* 38, 3 (2016), 186–210.