



Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations

Thomas Baumer, *Nexis GmbH*; Tobias Reitinger, *Universität Regensburg*;
Sascha Kern, *Nexis GmbH*; Günther Pernul, *Universität Regensburg*

<https://www.usenix.org/conference/soups2024/presentation/baumer>

This paper is included in the Proceedings of the
Twentieth Symposium on Usable Privacy and Security.

August 12–13, 2024 • Philadelphia, PA, USA

978-1-939133-42-7

Open access to the Proceedings
of the Twentieth Symposium
on Usable Privacy and Security
is sponsored by USENIX.

Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations

Thomas Baumer
Nexis GmbH

Tobias Reittinger
University of Regensburg

Sascha Kern
Nexis GmbH

Günther Pernul
University of Regensburg

Abstract

Organizations tend to over-authorize their members, ensuring smooth operations. However, these excessive authorizations offer a substantial attack surface and are the reason regulatory authorities demand periodic checks of their authorizations. Thus, organizations conduct time-consuming and costly access reviews to verify these authorizations by human decision-makers. Still, these deciders only marginally revoke authorizations due to the poor usability of access reviews. In this work, we apply digital nudges to guide human deciders during access reviews to tackle this issue and improve security. In detail, we formalize the access review problem, interview experts ($n = 10$) to identify several nudges helpful for access reviews, and conduct a user study ($n = 102$) for the *Choice Defaults Nudge*. We show significant behavior changes in revoking authorizations. We also achieve time savings and less stress. However, we also found that improving the overall quality requires more advanced means. Finally, we discuss design implications for access reviews with digital nudges.

1 Introduction

The Open Web Application Security Project (OWASP) lists “broken access control” as the Top 1 vulnerability and discovers it in 94% of the tested web applications [36]. Excessive authorizations are one driver for this OWASP vulnerability, as these are granted without an actual need and thus open an unnecessary attack surface. More precisely and within this paper, we ask highly qualified Identity and Access Management (IAM) experts to estimate the ratio of excessive autho-

rizations in mid- and large-sized organizations. Our experts expect about a fifth to a quarter of the authorizations to be excessive and vulnerable ($M = 22.8\%$, $SD = 6.4\%$, $n = 10$).

To mitigate this vulnerability, regulative authorities demand organizations to evaluate their authorizations with periodic access reviews. Well-known regulations include SOX [52], Basel III [6], MARisk [12], or HIPAA [51]. In large organizations, this involves hundreds of access review deciders for six figures of authorizations [18, 39]. These deciders (e.g., department heads) evaluate these authorizations within their responsibility. Although accountable, deciders face a time-consuming and frustrating task, as their expertise and objectives might not primarily match with security. Responsible deciders must also avoid mistakes: While revoked authorizations can interrupt their organization shortly, falsely confirmed excessive authorizations drive security risks [25]. Research [18] shows in a real-world case study that deciders only revoke 1.2% of the reviewed authorizations instead of the expected one-fifth excessive ones. Besides this clear need for improvement, only a few papers [18, 22, 26, 39] study access reviews.

As shown by Jaferian et al. [26], crucial issues for access reviews are rooted in poor usability. Using digital nudges to guide decisions [53] is thus a promising approach to improve access reviews. However, we identify several research gaps: First, current research does not formalize access reviews. Second, it is unknown how digital nudges address access review challenges. Third, it is unclear whether digital nudges actually improve access reviews. We investigate these research gaps with the following research questions:

- Q1** *How to formalize the access review problem?*
- Q2** *How do access review challenges map to digital nudges?*
- Q3** *Does an applied digital nudge (the Choice Defaults Nudge) benefit the access review problem?*

This work follows a mixed methods approach in an exploratory sequential design. We use qualitative methods to define a formal and precise notation of the underlying problem (Q1) and to interview highly qualified experts ($n = 10$) about

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2024.
August 11–13, 2024, Philadelphia, PA, United States.

applying digital nudges for access reviews (Q2). Moreover, our quantitative methods use insights of Q1 and Q2 to conduct a user study ($n = 102$) with an application of the *Choice Defaults Nudge* for access reviews (Q3). Consequently, our methods lead to the following contributions:

- We are first to formalize the access review problem.
- We map the expected effects of digital nudges to access review challenges based on 10 expert interviews. We find that access reviews benefit from digital nudges.
- We show behavior changes leading to quality improvements and more revoked authorizations by applying the *Choice Defaults Nudge* within a user study ($n = 102$). Moreover, we achieve time savings and lower frustration.

The remainder of this work is outlined as follows: Section 2 covers the background of this work, including relevant terminology, access review challenges, digital nudges, and related work. Section 3 provides details about our mixed method approach. Subsequently, we present the three-fold results of our paper. In Section 4, we first formalize the access review problem. Second, we map digital nudges with the access review challenges through the expert interviews in Section 5. Third, we conduct a user study on the *Choice Defaults Nudge* for access reviews in Section 6. Following the results, we discuss the general findings of this work in Section 7. Finally, Section 8 concludes and gives an overview for future work.

2 Background

2.1 Terminology

Identity and Access Management (IAM) is a cornerstone of modern cybersecurity, as it manages users and their access to sensitive data and services of organizations. Therefore, IAM provides tools to administer, authorize, and authenticate identities. Regulative authorities acknowledge the relevance of IAM and demand proper security controls. Besides state-of-the-art authentication, one crucial control is to demonstrate that the users still require granted access. Access reviews are a typical tool to prove the actuality of the granted access. These access reviews are the main focus of this work.

Access reviews are a periodic and compliance-driven process to verify users' authorizations. A team of domain experts, managers, application owners, and security admins typically reviews the granted authorizations with their knowledge of current processes, people, and resources. Especially in large organizations, access reviews are labor-intensive. Because of the recurring workload of access reviews, an organization might not finish an access review before the next one starts. The primary goal of access reviews is *revoking excessive authorizations*. Secondary goals are the determination of responsibilities for authorizations, requesting missing authorizations, or organization-specific data quality requirements. [18,22,26]

Nudges help humans make choices in analogous and digital systems. While these individuals must make their choices freely, *choice architects* design *choice architectures* to support their decisions by *nudging* towards a desired option. A nudge is thus a characteristic, influencing a decision in the interest of the decider. An example of a nudge in a supermarket is making healthy food easily accessible while making the unhealthy one harder to reach. From an ethical perspective, a nudge does not prevent a human from making a specific choice and only influences the decision in the best interest of the human. A *digital nudge* applies the idea of nudges to information systems. With features of user interfaces for guidance, users can make their choices freely and supported by the best advice of the choice architecture. [27,45,50]

2.2 Access Review Challenges

Based on expert interviews, Jaferian et al. [26] summarize access review challenges (C1-C5). We utilize these challenges throughout the paper, and thus detail them in this section.

C1: Scale outlines the number of involved IAM entities for the access review. The scale of the users, roles, permission, accounts, or assignments quickly grows into large numbers [18,26,39], making careful considerations for organizing the access review's workload necessary. Furthermore, the heterogeneity of these entities within real-world organizations intensify this challenge [30].

C2: Lack of Knowledge refers to the understandability of roles and permissions [26,30,31]. IAM entities might not have telling names, comprehensive descriptions, or concepts like roles or permissions might not have been fully understood. Experts thus might take uninformed or *best guess* decisions, leading to a bias for keeping unnecessary granted authorizations, violating the Principle of Least Privilege (PoLP) [18]. Additionally, for large organizations, the knowledge about these entities is distributed (or even missing completely), making the advice of responsible domain experts necessary [30].

C3: Frequency describes a dilemma for the managers: access reviews are not their *actual* responsibility, but they are frequently asked for it [26]. The experts might not feel a need to participate, leading to failing access reviews. Ultimately, this may cause even more access reviews, since successfully executed access reviews are part of compliance and audits. Thus, while access reviews usually are only required yearly, some organizations execute them quarterly, hoping not to fail access reviews due to lack of participation [26].

C4: Human Errors are common due to the scale and manual execution of access reviews by human deciders. These experts ultimately decide about required or excessive access by applying the best of their knowledge. This process is, therefore, inherently error-prone, as decisions to the best of the experts' knowledge might be incorrect or uninformed [18,26].

C5: Exceptional Cases occur due to the scale and complexity of access reviews. Context knowledge is sometimes

required for an informed authorization decision. For example, some members of organizations might replace others while on leave, trainings or tests might require temporary access, etc. might cause disturbances [26].

2.3 Digital Nudges

Based on a literature survey, Jesse and Jannach [27] propose a taxonomy for digital nudges. The authors distinguish four primary categories with further sub-categories of digital nudges (N01-N13): decision information (N01-N04), decision structure (N05-N08), decision assistance (N09-N10), and social decision appeal (N11-N13). We refer to these nudges throughout the paper, and thus explain them in this section.

Decision Information tries to present information helpful for the decision-maker without altering the available choices. This category comprises information translation (N01), salience (N02), visibility (N03), and phrasing (N04).

- *N01: Information Translation* targets reducing the cognitive effort for a decision by simplifying information or decreasing vagueness and ambiguity [48].
- *N02: Information Salience* aims to raise or decrease the prominence of information, by visualizations or making information harder or easier to notice [11, 48].
- *N03: Information Visibility* fosters decision information. This category includes mechanisms to disclose [24, 28, 48], compare [11, 48] or warn with [24, 33, 48] (tailored [28, 33] or external [35, 49]) information.
- *N04: Information Phrasing* puts presented information in context to intervene with the decisions to make. This category comprises the utilization of heuristics or biases like anchoring [33, 34, 48], availability [44, 50], the endowment effect [11, 44], framing [11, 33, 48], loss aversion [34, 44, 48], priming [11, 48, 50], etc.

Decision Structure alters the decision arrangement, comprising the decisions' range & composition (N05), defaults (N06), consequences (N07), and required effort (N08).

- *N05: Range & Composition* groups and categorizes choices. Therefore, choice architects or the decision-makers themselves break large decision structures into smaller category partitions [28, 48, 53], to present these one after another [28, 33] or to make them more comparable to each other [28, 35]. Choice architects can also utilize ordering effects for the presented options [11, 48].
- *N06: Choice Defaults* is one of the most effective and well-studied nudges [24]. The nudge preselects choices without hindering the decision maker from actively making another choice. On the one hand, a decision-maker is more invested in an actively made decision [35, 48]. On the other hand, decision-makers

rather accept the preselected status quo than actively decide against it [11, 24, 35, 48].

- *N07: Option Consequences* add further yet rational insignificant effects to the choice without changing the overall economic incentives. These consequences include social outcomes or minor benefits & costs [24, 35].
- *N08: Option-related Effort* modifies the effort or ease to make decisions. This nudge includes capping [11, 48] or raising financial & physical effort [24, 35] of decision-makers choices to mitigate mindless actions. Furthermore, eased and more convenient choices speed up decisions, e.g., making desired choices more accessible [48].

Decision Assistance aids decision-makers to realize their intentions. This category includes the usage of reminders (N09) and commitment facilitation (N10).

- *N09: Reminders* actively put already available information into or out of the attention focus of the decision-makers. This nudge includes reminding of underlying goals, deadlines, and their relevance [11, 24, 33, 35, 48, 49] or stating social expectations for decisions [35].
- *N10: Commitment Facilitation* helps decision-makers to (timely) finish their asked for decisions. This nudge includes precommitment strategies (e.g., user-defined sub-goals) [24, 33, 35, 48] or public commitment (e.g., pressure by publicly communicating own goals) [11, 35].

Social Decision Appeal category focuses on the social implications of nudges, including the Messenger Reputation (N11), Social Reference Point (N12), and Empathy Instigation (N13).

- *N11: Messenger Reputation* considers the reputation of the messenger delivering the information for the nudge. On the one hand, the messenger effect nudges a decision-maker since the messenger provides a certain and influencing impression about itself. For example, an actually well-designed and important choice architecture might dilute its seriousness if it contains many spelling mistakes [44]. On the other hand, the reputation of a system can be improved when choice architects expect and forgive the errors of their decision-makers [28, 53].
- *N12: Social Reference Point* nudges a decision based on social opinions. E.g., the opinion of a majority (Argentum-Ad Populum), group (Group-Ad Populum) [16], or an opinion leader [35] can influence decision-makers. Additionally, deciders tend to follow a herd [34, 44, 48] and might desire a comparison with their peers influencing their own decisions [24, 33, 35].
- *N13: Empathy Instigation* uses feelings to influence deciders. For example, an avatar might smile or cry upon the choices of a decision-maker (moral suasion) [11, 48], or a choice architect can trigger reciprocity by doing something good for the decision-makers to nudge them into returning the favor with good choices [11].

2.4 Related Work

Access control ensures users can only act within their intended authorizations and is characterized by its necessary yet cumbersome maintenance. Related research on maintenance covers more efficient access control models, optimization, and general maintenance processes like access reviews. By evolving from access control matrices [41], the most dominant access control models are Role-Based Access Control (RBAC) [15, 37, 43] and Attribute-Based Access Control (ABAC) [23, 46] as these reduce maintenance costs. Modeling access control policies considers bottom-up, top-down, or hybrid approaches [14] but often overlook their actual optimization without recalculating them [31, 38]. Therefore, access control maintenance targets up-keeping authorizations in changing needs and environments based on IAM goals [25, 30]. This includes periodically reviewing and revoking excessive access [18, 22, 26], granting missing access [47, 54], and timely propagation [7] to maintain secure authorizations. This paper especially relates to work on maintenance by access reviews: Jaferian et al. [26] study its challenges and usability. Puchta et al. [39] show positive effects on using external data for access reviews. Groll et al. [18] assess decision quality. Hill [22] conducts a case study for HIPAA [51] compliant access reviews.

Digital nudges are a popular research topic, as shown by various surveys: While Bergram et al. [9] conduct a general literature review, Schaer and Stanoevska-Slabeva [44] analyze digital nudges in customer-journeys and Jesse and Jannach [27] in recommender systems. Additionally, a survey of Caraban et al. [11] covers a practical and ethical application. As an established means to shape human behavior, applications of (digital) nudges exist for many domains. Examples include e-commerce [2, 13], sustainable smart home [8], contract tracing [17], or cybersecurity. In detail, cybersecurity examples include digital nudges to prevent phishing [55] or increase password quality [29, 56, 57]. An application of digital nudges for access reviews remains open so far.

3 Methods

This research uses mixed methods in an exploratory sequential design. First, we use qualitative methods to formalize the Access Review Problem (ARP) (Q1) and relate access review challenges to digital nudges (Q2). Second, we use these qualitative insights in quantitative methods to study the effect of the *N06: Choice Defaults* for access reviews (Q3). Third, a discussion wraps up the findings. Figure 1 depicts our mixed methods. In the following, we detail each part.

3.1 Formalizing the Access Review Problem

While the access review challenges comprise a global view, we formalize the actual Access Review Problem (ARP) in

Section 4. Its goal is to understand the underlying problem better. This formalization targets a quantifiable and comparable foundation for the solution of the ARP. Thus, we argue access review as a transition between two authorization states, depicted as confusion matrices. This precise formalization of the ARP is the basis for the hypotheses of the user study.

3.2 Relation of Access Review Challenges to Beneficial Digital Nudges

Complementary challenges to the ARP are discussed in the literature, including scale, lack of knowledge, frequency, human errors, and exceptional cases [26]. Digital nudges are a promising approach to address the ARP and its challenges. But it is unknown, whether digital nudges can help and which effects can be expected from their application (Q2).

To better understand this relationship between access review challenges and digital nudges, we investigate and map access review challenges from Jaferian et al. [26] with the digital nudge taxonomy of Jesse and Jannach [27] by conducting semi-structured expert interviews based on the guidelines of Adams [1]. The interviewed industry experts provide practical experience in access control and reviews. Therefore, we target highly qualified professionals with at least five years of experience working with large IAM systems, periodical executed access reviews, and managing thousands of identities or consultants with practical experience for many enterprises. Of course, these highly qualified experts are not readily available, but we managed to acquire 10 of these experts through personal and professional contacts. The experts are located in Germany. We use their expertise for a well-grounded argumentation for the relationship between access review challenges and digital nudges. Section 5.1 details further on the method for the expert interviews.

3.3 User Study for the Choice Defaults Nudge

After laying out theoretical foundations for digital nudges and access reviews in Sections 4 and 5, we study the application of a selected digital nudge in-depth. The expert mapping of digital nudges and access review challenges suggests several digital nudges. To sharpen the scope of the use study, we select *N06: Choice Defaults* based on the following reasons:

- Literature considers *N06: Choice Defaults* among the most effective digital nudges [24].
- The expert interviews had strong positive and negative expectations, inviting a more detailed examination.
- We felt confident to apply the *N06: Choice Defaults* to an access review and study its effects precisely.

We design an access review, simulating a real case: experts often describe access reviews as repetitive, time-consuming, and tedious tasks, requiring a strenuous thought process to




Figure 1: Mixed methods approach for this study.

determine correct authorizations. We thus hand out each participant a one-pager about the case. Participants manage a fictitious marketing department containing three teams within the case: graphic design, social media, and event management. The instruction describes the functions and tasks of each team and explains the unwanted implications of excessive or missing authorizations. While it is theoretically possible to review each decision using the document correctly, it takes some thought to make a correct decision.

To study the *N06: Choice Defaults* in-depth, we use three distinct configurations for access reviews with the same data basis: default accept, default reject, and a neutral default. This directly compares the default accept and reject configuration with a neutral state. The default accept configuration preselects every decision with an accept, the default reject vice versa, and the neutral default does not preselect.

We acquire 102 participants from a university context in Germany and randomly assign them to one of the three *N06: Choice Defaults* configurations. We select our sample size based on similar papers (c.f. Caine [10], also for the expert interviews). The (under-)graduate students have mostly a background in business informatics and IT security, indicating that they know essential IT security concepts and enterprise information systems. Furthermore, the participants are unaware of the research objective on digital nudges. We raffle a €100 gift card to one lucky participant to motivate participation. The participants must log in with authenticated accounts to avoid repeated participation and enable remote participation. We pilot the study with fellow researchers. Section 6.1 provides further details for the method.

3.4 Ethical considerations

Our experts were informed and consented to an anonymous publication of parts of their interviews. We will not share the recordings and delete the data one year after the publication.

The Institutional Review Board (IRB) *German Association for Experimental Economic Research e.V* approved the user study to comply with ethical requirements for working with humans. The certificate is available online.¹

¹<https://gfew.de/en-ethik/HQwmKGTZ>

4 Q1: The Access Review Problem

To better understand access review and benchmark our user study design, we introduce a representation of granted authorizations and security policies within a confusion matrix depicting User Permission Assignments (UPAs). Figure 2 maps the actually granted authorizations with security policies. We assess authorizations as effective access grants (which may contain errors), while security policies define the conceptual access users should have (ground truth). We construct a classical confusion matrix by mapping these authorizations and security policies with a binary distinction. Thus, the effectively granted UPAs are Predicted Positive (PP) as $PP = TP + FP$, while $P = TP + FN$ should be granted. PN and N are vice versa not-granted UPAs. Therefore, the True Positives (TPs) describe UPAs, granted in reality and conceptually. The sensitivity $SEN = \frac{TP}{P}$ represents the rate of correctly granted UPAs. Vice versa, True Negatives (TNs) describe UPAs, not granted in reality and in concept. The specificity $SPC = \frac{TN}{N}$ represents the rate of correctly not-granted UPAs. Together, sensitivity and specificity express the balanced accuracy $BA = \frac{SEN+SPC}{2}$, equaling 100% in a perfect world without errors.

		Authorization	
		Positive PP	Negative PN
Security Policy	Positive P	TP	FN
	Negative N	FP	TN

Figure 2: Confusion matrix for UPAs.

However, type I (False Positives (FPs)) and type II (False Negatives (FNs)) errors are present in reality. On the one hand, FPs are granted authorizations not considered by security policies (excessive UPAs). These excessive authorizations drive security risks, as over-privileged users are a target for threat actors. The primary goal for access reviews is lowering FP, which is highlighted in Figure 2. On the other hand, FNs are mistakenly not granted authorizations (missing UPAs). An example of their impact is when users cannot do their legitimate tasks because they do not have access to the required systems. This causes dissatisfaction for the users and slows down processes. In a relative notation, the False Discovery Rate (FDR) describes the percentage of excessive UPAs FP

based on PP as $FDR = \frac{FP}{PP}$. Vice versa, the False Omission Rate (FOR) describes missing UPAs as $FOR = \frac{FN}{PN}$.

Thus, an access review can be understood as transitioning from one UPA set depicted as confusion matrix C^1 to another C^2 . The primary goal is to reduce the FDR while retaining or improving BA. We introduce definitions for Access Reviews (ARs) and the Access Review Problem (ARP) as:

Definition 4.1 (Access Review (AR)). Given a confusion matrix C^1 describing an UPA^1 set, an access review AR revokes a subset of the effectively granted authorizations $R \subset PP^1$. When executing AR a confusion matrix C^2 describes the resulting set as $UPA^2 = UPA^1 \setminus R$.

Definition 4.2 (Access Review Problem (ARP)). Design AR in such a way that a (human) deciders can review and revoke UPAs $R \subset PP^1$ according to their knowledge about security policies P^1 , that the FDR is reduced ($FDR^1 > FDR^2$), without lowering BA ($BA^1 \leq BA^2$). The ARP is solved on a $FDR^2 = 0\%$ without decreasing BA : $BA^1 \leq BA^2$.

The following hypotheses hence allow testing whether an access review design improves the ARP:

- H₀** An access review design does not improve the ARP as the FDR remains or rises $FDR^1 \leq FDR^2$ or BA remains or decreases $BA^1 \geq BA^2$.
- H₁** An access review design improves the ARP as the FDR decreases $FDR^1 > FDR^2$ and the BA raises $BA^1 < BA^2$.

5 Q2: IAM Experts on Digital Nudges

5.1 Method Details

The interviews comprise three phases: an interviewee introduction, an explanation of access review challenges and digital nudges, and a workshop to generate the mapping of access review challenges and digital nudges. (i) The interviewee’s introduction collects data about their access review experience, their perspective on its challenges, and their estimation of excessive authorizations (FP). (ii) The explanation phase ensures essential knowledge about digital nudges, reminds the interviewee of access review challenges, and ensures a common vocabulary. We use the interviewees’ perspectives on access review challenges to explain to them the access review challenges of Jaferian et al. [26]. (iii) The procedure for querying the mapping for each considered digital nudge [27] follows this scheme: First, we explain the digital nudge in general and provide a suitable example for the interviewee. Afterward, we let the expert freely reflect on the effect of this digital nudge and its benefit to all access review challenges. Finally, we ask the expert to rate each access review challenge on a five-level Likert scale from very positive (+2) to very negative (-2). This rating scheme helps the expert to express

their arguments more comparable to each other. The complete interview script is available in Appendix A.1.

We interviewed 10 highly qualified experts with experience in conducting several Access Reviews (ARs) specialized for IAM by implementing IAM tools (engineers), responsible for managing thousands of users in IAM systems (inhouse), or advising clients (consultants). Table 1 protects their identities but depicts their high expertise for ARs. The interviews took an average of 60 minutes and were recorded, transcribed, coded, and evaluated. We translated relevant parts of the interviews into English during the coding process.

Table 1: Participants for expert interviews.

Interview	Experience				Sector
	Years	Clients	Users	ARs	
E01: IAM consultant	8	40		20	Multiple
E02: IAM consultant	5	15		10	Multiple
E03: IAM engineer	12	40		15	Multiple
E04: IAM inhouse	8	15	1k	40	Insurance
E05: IAM consultant	19	25		10	Multiple
E06: IAM consultant	13	40		25	Multiple
E07: IAM consultant	6	15		50	Multiple
E08: IAM inhouse	15	2	19k	4	Biotech
E09: IAM consultant	11	4		10	Banking
E10: IAM inhouse	7	1	13k	120	Insurance

We recorded the interviews with Microsoft Teams, transcribed them with Word, and summarized and coded them with Excel. For the coding, we use both deductive and inductive coding [3]. Since we already know the access review challenges [26], we first applied deductive coding based on these challenges for each digital nudge. This deductive coding already sorts large parts of the interviews in proven codes. However, we noticed that several augmentations exist within these codes. Thus, we also developed inductive codes for each nudge and challenge combination to capture the interviews comprehensively. For the rating of each nudge and challenge pair, we initially used the mean expert ratings. After coding and comparing the interviews, we slightly adapted the ratings, to balance well-reasoned arguments across the experts. The resulting codebook is available in Appendix A.2.

5.2 Results

This section presents the experts’ mapping. We build on the presented background of the access review challenges (C1-C5) in Section 2.2 and digital nudges (N01-N13) in Section 2.3. The resulting mapping is depicted in Table 2, whereas the challenges serve as columns and the digital nudges as rows. The cells summarized a rating for each challenge and nudge. In the following, we detail each digital nudge.

N01: The experts stress the benefits of comprehensible data. While C1 and C3 do not decrease, comprehensible data indirectly increases its learnability and comfort for the deciders, easing management eventually. For C2 and C4, the

Table 2: Nudges [27] and access review challenges [26].

Nudges	C1	C2	C3	C4	C5
N01: Information Translation	1	2	1	2	0
N02: Information Saliency	1	0	1	1	2
N03: Information Visibility	1	2	0	1	2
N04: Information Phrasing	0	-1	0	1	0
N05: Range & Composition	2	1	1	2	2
N06: Choice Defaults	2	-2	2	-2	0
N07: Option Consequences	0	-1	1	-1	-1
N08: Option-related Effort ↗	-1	1	-1	1	1
N08: Option-related Effort ↘	1	-1	1	-1	-1
N09: Reminders	0	1	2	-1	0
N10: Commitment Facilitation	1	0	1	1	0
N11: Messenger Reputation	1	2	1	2	2
N12: Social Reference Point	0	2	0	1	2
N13: Empathy Instigation	1	1	1	1	0

Note: Option-related effort is ↗ = increased, ↘ = decreased. The Likert scale spans from very positive +2 to very negative -2.

experts anticipate a strong positive effect, as comprehension is essential for C2: “If data is displayed more comprehensibly, it’s helpful for users with little knowledge [C2] about the decision.” (E06). Being comfortable with the data is relevant for C4: “If the user is comfortable with the displayed data, you can expect fewer human errors [C4].” (E07)

N02: The experts emphasize the focus: “In my opinion is the highlighting of C5 the only option to manage large data sets.” (E05) However, “it depends on the quality of the highlighting” (E03), since excessive or missing highlighting might draw away attention from relevant decisions. But upon sufficient and reliable quality, decision-makers can efficiently focus on the highlighted decisions or attributes and decide the remainder quicker (benefit for C1 and C3). Decision-makers “actually want to decide diligently but are hindered by its scale. These decision-makers could diligently and mindfully decide just the highlighted decisions in an efficiency tradeoff.” (E09)

N03: Showing additional data is crucial for C2 and C5 to make informed decisions while streamlining the focus to relevant attributes (C4). By only offering limited attributes for each decision in default, the management of C1 is eased. However, the user might not know the relevancy of specific hidden attributes as these move out of focus (C4).

N04: Our interview partners express reservations as decisions might not be based on rational knowledge but on biased phrasing (C2). However, for a well-executed implementation, its utilization can raise the access review acceptance (C4) as its relevancy could be communicated more effectively.

N05: The setup of meaningful partitions and sorting imposes overhead compared to just showing all decisions in one turn, thus worsening C1 and C3. However, the experts anticipate quite positive effects on all challenges. Similar sorted or clustered partitions leverage efficiencies as deci-

sions transfer to whole partitions. These efficiencies ease the management for C1 and C3 since the workload decreases, while more consistent and mindful decisions mitigate C2 and C4. Furthermore, clustering and appropriate communication of exceptional cases (C5) can positively influence.

N06: The experts discuss the strong effects of N06: *Choice Defaults*. Due to the reduced workload by the preselection, the experts rate a positive effect on C1 and C3. However, the experts worry that deciders adopt a preselected default without further thought, leading to uninformed (C2) and mindless (C4 and C5) decisions. While a mindful default prevents errors on uncertainty (like for C5) or on evident cases, just adopting the recommended default can become a fallacy, assuming the recommendation algorithm’s imperfections. This is especially an issue if the decision-makers trust the preselection so that they mindlessly adopt the default instead of a mindful decision. A falsely set default would then lead to a systematic bias, endangering the next audit relevant to compliance. In sum, the experts anticipate the potential of N06: *Choice Defaults*, but advise careful application.

N07: “In practice, negative consequences dominate. For example, we will tell your boss if you don’t finish your access review tasks within 14 days.” (E01) The experts acknowledge that creative and positive consequences could be feasible and reasonable, making frequent access reviews more comfortable (C3). However, they doubt there would be a game-changer in the long term because the effects would wear down over time (C3), and the decisions might be based on avoiding pressure or pursuing benefits (C4) instead of reason (C2 and C5). In this context, it is also worth noting that “disadvantaged individuals need special consideration” (E09) because finishing an access review in time might not be fair for these (C5).

N08: This nudge’s influence on the access review challenges is ambivalent, as it depends on whether the option-related effort is increased (↗) or decreased (↘). If the effort increases (vice versa for decrease), the users take more time to decide. For C1 and C3, this worsens the situation as the workload rises with its time consumption. Taking more time for a decision (e.g., requiring a reason for confirming a high-risk authorization) also benefits C2, C4, and C5, as the decider would need to consider a reason or reconsider the decision. But the experts also stressed the efficiency and acceptance of the access review, as some users easily become annoyed by increased effort: E.g., “We once required the users to set a note for the reviewed authorizations, but one user just put question marks for every note to bypass the input check.” (E04)

N09: “By a simple reminder [email], we observe more participation.” (E10) While reminders are especially relevant for C3 to communicate open tasks or instructions and goals for access reviews (C2), they can also pressure decision-makers to decide quickly but uninformed (C4). The audience and channel of reminders are also essential for C4. E.g., an inexperienced decider might require instructions or training. The experts also noted that reminders via an email channel

dominate in practice but are quickly perceived as spam. “Everybody wants something from all colleagues. Ironically, some colleagues even configure automated email filters which they won’t check afterward.” (E09) In this sense, a personal or multichannel address is most effective, but it is a considerable effort for the IAM team conducting the access review.

N10: The experts appreciate the autonomic commitment in combination with semantic partitioning (N05) of the decisions. An autonomic configuration of sub-goals and sub-deadlines suitable for the deciders benefits C1 and C3 as the deciders “perceive control over scale and frequency” (E06). This leads to more comfort, as sub-goals and sub-deadlines become meaningful for the deciders, mitigating C4.

N11: The experts stress the importance of this nudge: “Most important point; If the IAM team is not accepted, it is going be tough.” (E04) Furthermore, they note its failure in practice: “Access reviews are usually perceived negatively.” (E10) With a suitable messenger reputation, users will trust and endure the tedious tasks of the access reviews more, which is beneficial for C1 and C3. The experts also anticipate strong benefits for C2, C4, and C5 as the decision-makers will dare to ask or tell an approachable IAM team their relevant questions or mistakes: “If the IAM team is approachable, users communicate errors more eagerly or at all.” (E07)

N12: Similar to N11, if the social reference point sympathizes with the access review, decision-makers are likely to endure the tedious workload (C1 and C3). However, on low sympathy, the opposite effect might apply. The experts anticipate positive effects for C2, C4, and C5 because deciders discuss the access review: “For example, we introduced access review chat groups for business units. Decision-makers can talk about access reviews, like showing their own or seeing others’ progress, asking questions, etc.” (E07) In this sense, exceptional cases (C5) might become evident after a discussion and sharing knowledge about similar cases (C2), while noticing the colleagues’ progress might remind stragglers or expose them to peer pressure (C4).

N13: “On large scale [C1] and high frequency [C3], the decision-makers want to work with a pleasant tool.” (E06) Moral suasion and empathetic feedback (C2) can inform and convince the decision-maker about odd user behavior (e.g., mindlessly accepting all authorizations) without losing their motivation (C4). Reciprocity also fosters mitigation of C4 by “always addressing the positive side: the access review is meant to help you, the decision-maker, to compliantly and securely maintain your authorizations.” (E08)

Furthermore, the experts estimate a mean on excessive authorizations (FP) at 22.8% ($SD = 6.4\%$). Since we also asked our experts about common AR challenges, we confirm the AR challenges first published by Jaferian et al. [26].

In summary, our experts conclude positive and negative effects when using digital nudges. Table 2 summarizes these key takeaways. We hope to motivate future work with it as most digital nudges invite dedicated research on access reviews.

6 Q3: Choice Defaults in Access Reviews

6.1 Method Details

In the data set of the user study (Appendix B.1), we let participants review (accept or remove) granted UPAs $PP = TP + FP$ (legit TP or excessive FP), leading to UPA revoke operations only. Not granted UPAs $PN = FN + TN$ (missing FN or legit TN) are not considered. After piloting, we determined 160 UPAs serving as decisions to align an estimated study duration of 20-30 minutes and not to deter participation. Therefore, the crafted data set comprises 160 UPAs (PP) split into 80 legitimate ones (TP) and 80 excessive ones (FP), clearly distinguished by a case study document (see Appendix B.2). Figure 3 summarizes the initial UPAs as a confusion matrix.

		Authorization	
		$PP = 160$	$PN = 232$
Security Policy	$P = 80$	$TP = 80$	$FN = 0$
	$N = 312$	$FP = 80$	$TN = 232$

Figure 3: Confusion matrix for the case of the user study.

We configure and execute the access reviews with the commercial tool NEXIS4². The tool can import our data set, configure *N06: Choice Defaults*, execute large-scale access reviews, and collect relevant data points. Figure 4 displays a simplified screenshot of the review process. Further screenshots for all groups are available in Appendix B.3. For data collection, we make three observations for each access review participant: their decisions for the 160 UPAs, their time consumption, and their self-assessment for the NASA Task Load Index (TLX) [20]. (i) The tool stores each binary decision out-of-the-box, leading to a total of 16,320 manual decisions for 102 participants and 160 UPAs. (ii) We measure the time consumption for each participant by comparing the events for starting the access review and confirming the final completion prompt. (iii) After completion, we ask the participants to fill out a questionnaire for the NASA TLX [20] to capture their perceived workload. These questions are based on a Likert scale (-3 to +3) and include:³

- Mental Demand: How mentally demanding was the task?
- Temporal Demand: How hurried or rushed was the pace of the task?
- Performance: How successful were you in accomplishing what you were asked to do?
- Frustration Level: How insecure, discouraged, irritated, stressed, or annoyed were you?

²<https://nexis-secure.com/en/>

³We omitted the questions for physical demand and effort, as these are not applicable or relevant for our study.

Employee		Permission
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Moore, Evelyn F:\Documents\Social_Media_Strategy\
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Moore, Evelyn Approval vacation requests
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Miller, Sophia Book tradefair / exhibition stands

Figure 4: Simplified screenshot of the access review.

During the post-processing of the study, we used Microsoft Excel and R⁴ for data cleansing or data analysis. Data cleansing primarily comprises capping the time consumption for the AR to 60 minutes, as some participants took a break. We calculate the means, standard deviations, and non-parametric ANOVA of the AR confusion matrix, time-consumption, and NASA TLX indices. For our exploratory analysis of correlations (Spearman) and local regressions, we utilize a pair plot generated in R (see Appendix Figure 9). Supporting the open data idea, we publish all data to replicate our results on GitHub: <https://github.com/AccessReview/Availability>.

6.2 Results

This Section summarizes our observations of the user study (see Table 3). A post-hoc power analysis based on ANOVA for our three groups ($n = 34$) and an $\alpha = .05$ results in effect powers of .13 for a small effect ($f = .1$), .6 for a medium effect ($f = .25$), and .95 for a large effect ($f = .4$).

For all 102 participants, the mean review time t for the 160 decisions is $t = 22$ minutes with $SD = 13$ minutes. Deciders of all groups used to over-accept authorizations, amounting to a total accept rate of $1 - \frac{R}{PP} = 56.1\%$ (rather than a $SEN = 50\%$). H_0 is rejected for 99 of 102 reviews. The remaining 3 participants failed to achieve an ARP improvement. All participants' mean BA increased from 87.2% to 91.2% ($SD = 7.9\%$). The false discovery rate FDR , which represents the amount of excessive authorizations, was reduced from 50.0% to 21.6% ($SD = 14.7\%$). This improvement came at the cost of some erroneous revokes, leading to a mean FOR of 2.9% ($SD = 3.5\%$). In sum, most participants improved the ARP. The result data shows that two deciders behaved as “spammers” by either blindly accepting all authorizations (one decider in the accept group) or blindly rejecting them all (one decider in the reject group). These participants are among the three who failed to improve the ARP. While the data set is too small to make this finding statistically significant, it seems evident that the spammers just adopted the default.

The neutral configuration group is a control group for the default accept and reject nudge. Users from this group took a mean time of $t = 26$ minutes ($SD = 15$) and accepted 57.8% of the authorizations. The neutral group estimated the temporal demand as slightly low, with a mean score of -.8. On average,

neutral users stated the mental demand to be slightly high (.9) and their frustration to be neutral to slightly high (.5). They estimated their performance to be slightly above average (.9). The achieved BA is 91.9% ($SD = 5.8\%$), with the error rates FDR of 21.0% ($SD = 10.7\%$) and FOR of 2.6% ($SD = 2.5\%$).

The accept group only took $t = 19$ minutes ($SD=10$). With a time save of 24.3% to the neutral group. While the perceived TD was unchanged at -.8, both FL and MD were reduced by almost one point to a score of -.2 ($\Delta = -.7$) and .2 ($\Delta = -.7$). The accept rate was slightly higher than in the neutral group with 58.7% (+.9%). The default accept group achieved a BA of 92.3% ($SD = 5.3\%$), scoring .4% higher than the neutral one. The error rates were also marginally better than in the neutral group with $FDR = 20.8\%$ ($\Delta = -.2\%$, $SD = 9.3\%$) and $FOR = 2.2\%$ ($\Delta = -.4\%$, $SD = 2.6\%$).

Like the accept group, deciders of the reject group finished quicker than the neutral group with $t = 21$ minutes ($\Delta = -16\%$, $SD = 13$). Again, the estimated TD of -.4 did not reflect this ($\Delta = +.4$), but the stated FL and MD were reduced to -.6 ($\Delta = -1.1$) and -.2 ($\Delta = -1.1$). Unlike the accept group, however, the reject group showed a considerably reduced accept rate of 51.8% (-6.0%), which is very close to the initial $SEN = 50\%$. Unfortunately, the increased willingness to revoke did not improve the results: The deciders revoked fewer excessive authorizations than the neutral group ($FDR = 22.9\%$, $\Delta = +1.9\%$, $SD = 21.4\%$) and more correct ones ($FOR = 3.9\%$, $\Delta = +1.3\%$, $SD = 4.7\%$). With $BA = 89.4\%$ ($SD = 11.2\%$), BA was still improved regarding the initial state ($\Delta = +2.2\%$), but worse than the neutral configuration ($\Delta = -2.5\%$).

We ran a non-parametric Kruskal-Wallis test ($\alpha = .05$) to check for the significance of our observations between the three groups. We detect differences for the number of revokes R ($p = 0.039$), indicating that $N06: Choice Defaults$ did affect users' willingness to accept or reject authorizations. We also confirm differences for MD ($p = .049$) and FL ($p = .038$), indicating that lower stress perceptions result from the applied $N06: Choice Defaults$. We used Dunn's test for pairwise comparisons, showing that the neutral and reject groups differ for MD ($p.adj = .045$) and FL ($p.adj = .038$). However, the quality metrics BA , FDR , and FOR did not differ significantly between the study groups, which is unsurprising since the data set balances TP and FP at 80.

A test for Spearman correlation showed no significant correlation between review duration t and any of the quality metrics (BA , FDR , FOR), indicating that quality did not depend on the time spent. The data shows a significant positive correlation between the deciders' frustration level FL and t for the total population (.286) and the neutral group (.403), as well as between the stated mental demand MD and t (total: .237, neutral: .423). FL and MD are strongly correlated for all groups (total: .646, neutral: .589, accept: .672, reject: .664). We follow that deciders did not strictly distinguish between MD and FL and that longer reviews are perceived as more frustrating and/or mentally demanding. Interestingly, the per-

⁴<https://www.r-project.org/>

Table 3: General summary of the user study, including arithmetic means and standard deviations.

Group	n	Fails	t		R		BA		FDR		FOR		MD		TD		PF		FL	
			M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD
Initial	-	-	-	-	-	-	.872	-	.500	-	.000	-	-	-	-	-	-	-	-	-
Total	102	3	22	13	70.3	19.2	.912	.079	.216	.147	.029	.035	.3	1.6	-.6	1.5	1.1	1.6	-.1	1.7
Neutral	34	0	26	15	67.5	12.4	.919	.058	.210	.107	.026	.025	.9	1.4	-.8	1.5	.9	1.5	.5	1.7
Accept	34	1	19	10	66.1	19.5	.923	.053	.208	.093	.022	.026	.2	1.5	-.8	1.5	1.2	1.6	-.2	1.8
Reject	34	2	21	13	77.2	22.8	.894	.112	.229	.214	.039	.047	-.2	1.8	-.4	1.6	1.1	1.7	-.6	1.6

Note: *M* for arithmetic mean and *SD* for standard deviation. *n* for the participant count and *Fails* for executions in rejecting H_0 . *BA*, *FDR* and *FOR* for measuring the Access Review Problem (ARP). *t* for the time-consumption of the AR and *R* for the amount of rejected UPAs. *MD* (Mental Demand), *TD* (Temporal Demand), *PF* (Performance) and *FL* (Frustration Level) for the NASA TLX.

ceived temporal demand *TD* did not correlate with *t*, possibly due to a missing baseline of a “normal” review duration. The result data showed a strong positive correlation between the perceived performance *PF* and actual performance *BA* (total: .607, neutral: .635, accept: .605, reject: .639), and a negative one between *PF* and the error rates *FDR* (total: -.336, neutral: -.516, reject: -.422; accept: not significant) and *FOR* (total: -.541, neutral: -.568, accept: -.480, reject: -.603). Therefore, the deciders had a realistic estimation of their performance. The result data also showed significant negative correlations between *FL* / *MD* and *BA* as well as positive ones between *FL* / *MD* and the error rates *FDR* / *FOR*, each for some groups. However, the causality remains unclear if deciders who find the task more difficult experience more stress, more stressed deciders deliver poor results, or both. Figure 9 (Appendix) shows the Spearman correlation and the local regressions.

► **Key takeaways of our user study:** (i) Almost all deciders improved the ARP. (ii) The required time differed substantially but was unrelated to quality (*BA*). (iii) *N06: Choice Defaults* led to reduced time effort and stress perception. (iv) A default reject led to more rejects. (v) A simple *N06: Choice Defaults* did not affect quality (*BA*) significantly but influenced the number of rejects. In detail, however, more increase in false rejects is tolerable as false accepts legitimate excessive authorizations leading to a false sense of security. (vi) Deciders’ self-assessed performance correlates significantly with *BA*, indicating the deciders’ realistic self-assessment.

7 Discussion

7.1 Acceptance Bias

Participants of the user study tend to accept existing authorizations. Existing research already documents and analyzes over-granting in real-world scenarios [18, 47, 54]. However, such scenarios involve strongly imbalanced data (see expert interviews: $1 - 22.8\% = 77.2\%$ of authorizations are estimated to be correct), social implications (a revoke acts against the interests of a real person), and unequal visibility of the two error types. Erroneous revokes are detected quickly, and the decider alone is responsible, while erroneous accepts are not

immediately visible and all previous approvers share the responsibility for also not resolving the error. With an initial *SEN* of 50% and no personal repercussions, the study had none of these biases and made no implication that acceptance is favorable to revocation. Still, deciders accept authorizations too often, with an average accept rate of 57.8% in the neutral group (see Section 6.1). While the study data does not explain this behavior, a possible explanation might be that the status-quo bias discourages deciders from revoking [42]: Following a real-world scenario, the study description states that participants need to review *existing* authorizations, which would be revoked upon rejection. The existence of a general status quo bias could also explain the relatively weak effect of the default accept bias on the accept rate: Study participants with default accept or reject nudge configuration needed to change an existing preselection to make an active decision and are thus also confronted with a status quo bias. If a status quo bias is already the reason for over-accepting in the neutral group, the effect of the default accept nudge would only repeat an already present bias. In contrast, the default reject nudge creates a new status quo that nudges the deciders in the opposite direction. The explanation seems plausible based on the study results, as the accept rate of the default accept group is closer to the neutral group (58.7%), and the accept rate of the default reject group is closer to the actual 50% (51.8%).

7.2 Implications for Access Review Challenges

► **Decider motivation affects quality (C4):** As described in Section 6.2, the user study participants had a reasonable estimation of their own performance. The user study design is fair, with a planned execution time of 20-30 minutes and no hurdles for *N01: Information Translation* or *N03: Information Visibility*. Still, some deciders submitted results with relatively low quality. The correlations between perceived stress (*FL*, *MD*) and quality (*BA*, *FDR*, *FOR*) may also indicate that decider motivation was an important factor. It must be assumed that poor decider motivation contributes stronger in real-world scenarios with larger scale and poorer information basis, indicating that nudges targeting decider motivation (*N09-N13*) may be a valuable contribution to AR quality.

► **Longer reviews are more demanding (C1, C4):** The user study results showed significant correlations between the review duration t and the perceived stress (FL , MD), underlining the importance of a reasonable scale. While the user study already confirms that *N06: Choice Defaults* considerably reduces review time, *N05: Range & Composition* also seems promising. Choice architects should take care not to overwhelm deciders with too many decisions. Distributing review responsibilities to many instead of a few decision-makers might be helpful. Considering *N10: Commitment Facilitation* or splitting reviews into multiple suitable sub-reviews carried out at different times or limiting them to unreviewed or changed authorizations could also improve quality.

► **N06: Choice Defaults effectivity does not seem to depend on decision difficulty (C2, C4, C5):** We tried to assess whether the impact of *N06: Choice Defaults* depends on the difficulty of a decision. For this purpose, we grouped the user study decisions by the 160 UPAs and their respective study group (neutral, default accept, default reject), resulting in $3 * 160$ groups of 34 review decisions. We then calculated the error rate and standard deviation for the decisions in the neutral group as indicators of the decision difficulty or uncertainty of UPA. To measure the effect of the default accept nudge for any UPA, we subtract the number of accepts in the neutral group from the amount of accepts in the default accept group. The resulting difference is the amount of *additional* accepts achieved by the nudge. The default reject effectivity was calculated as equivalent to the difference of rejects in the neutral and default reject groups. A Spearman correlation test with a $\alpha = .05$ significance level showed no significant correlation between the indicators for a decision’s difficulty and the amount of additional accepts or rejects. The lack of correlation indicates that the effectivity of *N06: Choice Defaults* does not directly depend on the difficulty of a decision.

► **Spammers are an error source (C4):** Unlike the user study but in reality, a ground truth of detecting low-quality AR results is not available. Hence, it is helpful to identify “spammers” (deciders actually not trying to achieve an ARP improvement). The user study results suggest two possible ways to determine low-quality AR results: (i) While the review duration t did not correlate significantly with the quality metrics, we found that for the $n = 6$ deciders only taking $t = 6$ minutes or less, the mean BA ($M = 77.1\%$, $SD = 22.1\%$) drops a considerably $\Delta = -14.1\%$ comparing to BA of all participants ($M = 91.2\%$, $SD = 7.9\%$). (ii) Two spammers acted obviously ignorant by blindly accepting or rejecting all authorizations. In real-world scenarios, it might be helpful to use thresholds that, when undercut, classify the review as spam. We do not propose to dismiss such results categorically: it could be correct to accept all authorizations, or a decider could be quick. However, such deciders could be explicitly addressed to improve their result quality, e.g., by applying a custom nudge (like *N13: Empathy Instigation*) or requesting another person to check their decisions.

Table 4: Virtual best and worst advice.

Group	n	R	BA	FDR	FOR
Initial	-	-	.872	.500	.000
Total	102	70.3	.912	.216	.029
Virtual Best Advice	34	71.2	.931	.178	.023
Virtual Worst Advice	34	72.1	.885	.238	.043

Note: n for the participant count and R for the mean of rejected UPAs. BA , FDR and FOR are means for measuring the Access Review Problem (ARP).

► **Deciders have the last say (C4):** We re-grouped the user study decisions to simulate reviews with only correct and only incorrect *N06: Choice Defaults* (compare smart defaults [4, 5]). In reality, every decider had to make 160 decisions, of which 80 were TP (should be accepted) and 80 were FP (should be removed). This means that the default accept group had a correct preselection for exactly 80 authorizations, whereas the default reject group had a correct preselection for the other 80 ones. By virtually re-grouping these decisions, we create two sets of $34 * 160$ decisions each, for which one contains only correct default preselections and the other contains only incorrect ones. We then calculated the quality metrics BA , FDR , and FOR for both groups. Unsurprisingly, the virtual best advice group scored a higher overall quality than each of the three real study groups with $BA = 93.1\%$, and the lowest error rates with $FDR = 17.8\%$ and $FOR = 2.3\%$. The virtual worst advice group scored worse than all real groups with $BA = 88.5\%$, $FDR = 23.8\%$, and $FOR = 4.3\%$. However, the virtual best advice group’s results are closer to those of all real groups than a perfect result, for which BA would be 100% and both error rates would be 0%. Similarly, the virtual worst advice group did not perform terribly but, in fact, still achieved a mean improvement in the ARP. Results for both groups show that users are affected by the *N06: Choice Defaults* and that the quality of the applied nudge affects the quality of the AR result. However, deciders have the last say and may choose not to follow a default, attenuating the worst assumptions of some interviewed experts. Table 4 summarizes the figures for both virtual groups.

7.3 Two Undesired Responsibility Shifts

Real-world access reviews (without nudge support) assume reflective decision-makers in transparent environments, leading to two assumptions: reflection and transparency [11]. However, the expert interviews and the user study discard both assumptions. For the reflection assumption, experts report several instances of human errors (C4), and the user study shows that deciders are affected by *N06: Choice Defaults*. Additionally, the deciders make errors despite having all the necessary data (even for the best advice in Table 4). For the transparency assumption, experts report the troublesome endeavor to present the information needed (*N01-N03*, *N09*) as too many or too few details lead to an unclear big picture.

Hansen and Jespersen [19] evaluate ethical considerations for nudge applications by the nudge’s transparency and the decider’s reflective or automatic mode of thinking. As mentioned earlier, access reviews should strive for transparency and reflective decisions. Access reviews in the real world and those with nudges can fail one of these: the real-world access reviews can lack transparency, and the nudged ones can lack reflective choices. On the one hand, real-world access reviews force reflective decisions as overwhelmed deciders actively need to choose, leading to a lack of transparency and constructing an unpleasant ethical situation. While reflective choices make the deciders fully responsible for their actions, the sheer scale (C1) and frequency (C3) put so many decisions on the table that the actual big picture for the access review becomes non-transparent. Therefore, the deciders have to bear the responsibility for a volume of decisions above their capabilities as human decision-makers, raising ethical concerns. On the other hand, the access reviews with the *N06: Choice Defaults* stay more transparent but allow for less reflective decisions, leading to a responsibility split. As soon as scale (C1) and frequency (C3) make the deciders give up on reflective choices, the choice architect shares responsibility for the decision-makers adopting its defaults.

In summary, neither burdening the deciders with the responsibility of choices they do not comprehend nor splitting the responsibility between the choice architect and the deciders are desired responsibility shifts for access reviews.

7.4 Design Implications for Usability

Following Hansen and Jespersen [19], design implications for future access reviews (with digital nudges) involve facilitating meaningful decisions based on transparency and reflective choices. When applied properly, digital nudges empower deciders to make confident and meaningful decisions with transparent and honest guidance [19]. Most importantly, this implies perceiving access review deciders and their decisions not as hyper-rational but as human, including their strengths and flaws [21, 40]. In the following, we derive three implications for usability based on our results.

► **Partition meaningfully:** Several experts find *N05: Range & Composition* relevant as it allows for meaningful partitions of access review decisions. Partitions effectively mitigate the deciders’ scale perception and give a context for grouped decisions. Additionally, this allows abstract decisions for the whole partition. For example, deciding to revoke all authorizations of a person can be one meaningful decision instead of rejecting each of its authorizations one by one. Our experts name meaningful ways to partition decisions within access reviews, e.g., people leaving an organization, specific applications, critical authorizations, known past changes, organization-specific attributes, or processes. Ways to determine these partitions can range from choice architects’ or deciders’ experience to AI-based clustering.

► **Apply partition-specific digital nudges:** Digital nudges can be applied individually and combined for each partition. Based on the expert interviews, various digital nudges are suitable. For example, *N06: Choice Defaults* can preselect accepting security-uncritical authorizations (e.g., utility software) or rejecting security-critical ones (e.g., server access). Additionally, security-critical authorizations can be highlighted with a warning by *N02: Information Salience*. Thus, digital nudges can improve each partition’s usability to guide access review deciders, also considering individual organizational contexts.

► **Query performance perception:** In the user study results, we find a strong correlation in all groups for the objective quality metric *BA* and the deciders’ performance self-assessment *PF*. It shows that our user study participants had a reasonable perception of their performance. In contrast, a real-world access review cannot determine *BA* easily, as the underlying ground truth is unknown. This implies querying the deciders’ performance self-assessment (*PF*) can be a valid and easy-to-implement estimator for the access review’s quality (*BA*).

In summary, transparent digital nudges can guide human decision-makers to make meaningful, confident, and reflective choices. While the positive and negative effects of nudging require careful consideration, their anticipated effects are useful and promising tools for access review designs.

8 Conclusion

In this paper, we investigated digital nudges for access reviews. We formalized the access review problem. Subsequently, we interviewed highly qualified IAM experts to map the expected effects of digital nudges on access review challenges. Furthermore, we conducted a user study with *N06: Choice Defaults*. We found its influence on deciders’ behavior in revoking authorizations. Additionally, we achieve time savings (up to 24.3%) and lower frustration. A simple *N06: Choice Defaults* did not significantly influence the overall quality, but it can shift the decisions to more revokes. While these revokes cause some false rejects, false accepts would be worse as they create a false sense of security by legitimating excessive authorizations. For future work, we invite researchers to study the ARP, to investigate other digital nudges of Table 2 or their combinations, or to replicate this study with a larger sample size or smart defaults [4, 5]. In sum, digital nudges are a promising tool to improve access reviews but need careful application.

Availability

For transparency and future research, we make the case study, all collected data, and the analysis of the user study open-source (<https://github.com/AccessReview/Availability>). In detail, we publish the instructions and data set of the case study, participants’ results ($n = 102$), their choices ($n = 16,320$), and the *R* code to replicate our statistical evaluations.

Acknowledgments

The German Federal Ministry of Education and Research supported the research leading to these results as part of the DEVISE project (<https://devise.ur.de>).

This work would not have been possible without the help of our 10 interviewed experts and 102 user study participants. The experts invested a total of 10 hours and the participants a total of 42 hours to support our endeavor. Thank you!

References

- [1] William C. Adams. *Conducting Semi-Structured Interviews*, chapter 19, pages 492–505. John Wiley & Sons, Ltd, 2015.
- [2] Marvin Auf der Landwehr, Maik Trott, and Christoph von Viebahn. Consumers choice? fostering sustainability in grocery deliveries through digital nudging. In *Twenty-Ninth European Conference on Information Systems (ECIS 2021)*, ECIS 2021, page 1–16. Association for Information Systems, 2021.
- [3] Theophilus Azungah. Qualitative research: deductive and inductive approaches to data analysis. *Qualitative Research Journal*, 18(4):383–400, Jan 2018.
- [4] Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. A data-driven approach to developing iot privacy-setting interfaces. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces, IUI '18*, page 165–176, New York, NY, USA, 2018. Association for Computing Machinery.
- [5] Paritosh Bahirat, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. Overlooking context: How do defaults and framing reduce deliberation in smart home privacy decision-making? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*, New York, NY, USA, 2021. Association for Computing Machinery.
- [6] Basel Committee on Banking Supervision. Basel III: A global regulatory framework for more resilient banks and banking systems, June 2011.
- [7] Thomas Baumer, Mathis Müller, and Günther Pernul. System for cross-domain identity management (scim): Survey and enhancement with rbac. *IEEE Access*, 11:86872–86894, 2023.
- [8] Michelle Berger, Elias Greinacher, and Linda Wolf. Digital nudging to promote energy conservation behavior - framing and default rule in a smart home app. In *Thirtieth European Conference on Information Systems (ECIS 2022)*, ECIS 2022, page 1–16. Association for Information Systems, 2022.
- [9] Kristoffer Bergram, Marija Djokovic, Valéry Bezençon, and Adrian Holzer. The digital landscape of nudging: A systematic literature review of empirical research on digital nudges. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, CHI '22*, New York, NY, USA, 2022. Association for Computing Machinery.
- [10] Kelly Caine. Local standards for sample size at chi. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, page 981–992, New York, NY, USA, 2016. Association for Computing Machinery.
- [11] Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, page 1–15, New York, NY, USA, 2019. Association for Computing Machinery.
- [12] Federal Financial Supervisory Authority (BaFin). Rundschreiben 05/2023 (BA) - Mindestanforderungen an das Risikomanagement - MaRisk, June 2023.
- [13] Sandro Franzoi and Jan vom Brocke. Sustainability by default? nudging carbon offsetting behavior in e-commerce. In *Thirtieth European Conference on Information Systems (ECIS 2022)*, ECIS 2022, page 1–15. Association for Information Systems, 2022.
- [14] Ludwig Fuchs and Günther Pernul. HyDRo – hybrid development of roles. In *Information Systems Security*, pages 287–302. Springer Berlin Heidelberg, 2008.
- [15] Ludwig Fuchs, Günther Pernul, and Ravi Sandhu. Roles in information security – a survey and classification of the research area. *Computers & Security*, 30(8):748–769, 2011.
- [16] Cristina Gena, Pierluigi Grillo, Antonio Lieto, Claudio Mattutino, and Fabiana Vernerio. When personalization is not an option: An in-the-wild study on persuasive news recommendation. *Information*, 10(10), 2019.
- [17] Abdul Muqet Ghaffar and Thomas Widjaja. Framing as an app-design measure to nudge users toward infection disclosure in contact-tracing applications. In *Thirty-first European Conference on Information Systems (ECIS 2023)*, ECIS 2023, page 1–16. Association for Information Systems, 2023.
- [18] Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Monitoring access reviews by crowd labelling. In Simone Fischer-Hübner, Costas Lambri-noudakis, Gabriele Kotsis, A. Min Tjoa, and Ismail

Khalil, editors, *Trust, Privacy and Security in Digital Business*, pages 3–17, Cham, 2021. Springer International Publishing.

- [19] Pelle Guldborg Hansen and Andreas Maaløe Jespersen. Nudge and the manipulation of choice: A framework for the responsible use of the nudge approach to behaviour change in public policy. *European Journal of Risk Regulation*, 4(1):3–28, 2013.
- [20] Sandra G. Hart. Nasa-task load index (nasa-tlx); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9):904–908, 2006.
- [21] Jonas Hielscher, Uta Menges, Simon Parkin, Annette Kluge, and M. Angela Sasse. “Employees who Don’t accept the time security takes are not aware Enough”: The CISO view of Human-Centred security. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 2311–2328, Anaheim, CA, August 2023. USENIX Association.
- [22] Linda Hill. How automated access verification can help organizations demonstrate HIPAA compliance: A case study. *J Healthc Inf Manag*, 20(2):116–122, 2006.
- [23] Vincent C. Hu, David Ferraiolo, Rick Kuhn, Arthur R. Friedman, Alan J. Lang, Margaret M. Cogdell, Adam Schnitzer, Kenneth Sandlin, Robert Miller, Karen Scarfone, et al. Guide to attribute based access control (abac) definition and considerations (draft). Technical report, National Institute of Standards and Technology, 2014.
- [24] Dennis Hummel and Alexander Maedche. How effective is nudging? a quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, 80:47–58, 2019.
- [25] Matthias Hummer, Sebastian Groll, Michael Kunz, Ludwig Fuchs, and Günther Pernul. Measuring identity and access management performance - an expert survey on possible performance indicators. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy*, pages 233–240. SCITEPRESS - Science and Technology Publications, 2018.
- [26] Pooya Jaferian, Hootan Rashtian, and Konstantin Beznosov. To authorize or not authorize: Helping users review access policies in organizations. In *Proceedings of the Tenth USENIX Conference on Usable Privacy and Security*, SOUPS ’14, page 301–320, USA, 2014. USENIX Association.
- [27] Mathias Jesse and Dietmar Jannach. Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3:100052, 2021.
- [28] Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. Beyond nudges: Tools of a choice architecture. *Marketing Letters*, 23(2):487–504, Jun 2012.
- [29] Shelia M. Kennison, Ian T. Jones, Victoria H. Spooner, and D. Eric Chan-Tin. Who creates strong passwords when nudging fails. *Computers in Human Behavior Reports*, 4:100132, 2021.
- [30] Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. Maintain high-quality access control policies: An academic and practice-driven approach. In Vijayalakshmi Atluri and Anna Lisa Ferrara, editors, *Data and Applications Security and Privacy XXXVII*, pages 223–242, Cham, 2023. Springer Nature Switzerland.
- [31] Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Optimization of access control policies. *Journal of Information Security and Applications*, 70:103301, 2022.
- [32] Stefan Meier, Ludwig Fuchs, and Günther Pernul. Managing the access grid - a process view to minimize insider misuse risks. In *11th International Conference on Wirtschaftsinformatik (WI2013)*, pages 1051–1065, 2013.
- [33] Christian Meske and Tobias Potthoff. The dinu-model - a process model for the design of nudges. In *European Conference on Information Systems*, pages 2587–2597, 06 2017.
- [34] Tobias Mirsch, Christiane Lehrer, and Reinhard Jung. Making digital nudging applicable: The digital nudge design method. In *International Conference on Information Systems*. AIS, 2018.
- [35] Robert Münscher, Max Vetter, and Thomas Scheuerle. A review and taxonomy of choice architecture techniques. *Journal of Behavioral Decision Making*, 29(5):511–524, August 2015.
- [36] OWASP Top 10 team. Owasp top10, 2021. Accessed: 11/15/23.
- [37] Simon Parkinson and Saad Khan. A survey on empirical security analysis of access-control systems: A real-world perspective. *ACM Comput. Surv.*, 55(6), dec 2022.
- [38] Alexander Puchta, Fabian Böhm, and Günther Pernul. Contributing to current challenges in identity and access management with visual analytics. In Simon N. Foley, editor, *Data and Applications Security and Privacy*

- XXXIII, pages 221–239, Cham, 2019. Springer International Publishing.
- [39] Alexander Puchta, Sebastian Groll, and Günther Pernul. Leveraging dynamic information for identity and access management: An extension of current enterprise iam architecture. In *Proceedings of the 7th International Conference on Information Systems Security and Privacy - ICISPP*, pages 611–618, Online Streaming, 2021. INSTICC, SciTePress.
- [40] Ita Ryan, Utz Roedig, and Klaas-Jan Stol. Unhelpful assumptions in software security research. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS '23*, page 3460–3474, New York, NY, USA, 2023. Association for Computing Machinery.
- [41] Pierangela Samarati and Sabrina Capitani de Vimercati. Access control: Policies, models, and mechanisms. In Riccardo Focardi and Roberto Gorrieri, editors, *Foundations of Security Analysis and Design*, pages 137–196, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [42] William Samuelson and Richard Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1:7–59, 1988.
- [43] Ravi S. Sandhu. Role-based access control. portions of this chapter have been published earlier in sandhu et al. (1996), sandhu (1996), sandhu and bhamidipati (1997), sandhu et al. (1997) and sandhu and feinstein (1994). In Marvin V. Zelkowitz, editor, *Advances in Computers*, volume 46 of *Advances in Computers*, pages 237–286. Elsevier, online, 1998.
- [44] Armando Schär and Katarina Stanoevska-Slabeva. Application of digital nudging in customer journeys - A systematic literature review. In *25th Americas Conference on Information Systems, AMCIS 2019, Cancún, Mexico, August 15-17, 2019*. Association for Information Systems, 2019.
- [45] Christoph Schneider, Markus Weinmann, and Jan vom Brocke. Digital nudging: Guiding online user choices through interface design. *Commun. ACM*, 61(7):67–73, jun 2018.
- [46] Daniel Servos and Sylvia L. Osborn. Current research and open problems in attribute-based access control. *ACM Comput. Surv.*, 49(4), jan 2017.
- [47] Bingyu Shen, Tianyi Shan, and Yuanyuan Zhou. Improving logging to reduce permission Over-Granting mistakes. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 409–426, Anaheim, CA, August 2023. USENIX Association.
- [48] Cass R. Sunstein. The council of psychological advisers. *Annual Review of Psychology*, 67(1):713–737, 2016. PMID: 26393867.
- [49] Barnabas Szaszi, Anna Palinkas, Bence Palfi, Aba Szollosi, and Balazs Aczel. A systematic scoping review of the choice architecture movement: Toward understanding when and why nudges work. *Journal of Behavioral Decision Making*, 31(3):355–366, 2018.
- [50] Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Nudge: Improving decisions about health, wealth, and happiness. Yale University Press, New Haven, CT, US, 2008.
- [51] United States Congress. Health Insurance Portability and Accountability Act of 1996, 1996.
- [52] United States Congress. Sarbanes-Oxley Act of 2002. Corporate responsibility, 2002.
- [53] Markus Weinmann, Christoph Schneider, and Jan vom Brocke. Digital nudging. *Business & Information Systems Engineering*, 58(6):433–436, Dec 2016.
- [54] Tianyin Xu, Han Min Naing, Le Lu, and Yuanyuan Zhou. How do system administrators resolve access-denied issues in the real world? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 348–361, 2017.
- [55] Sarah Y. Zheng and Ingolf Becker. Checking, nudging or scoring? evaluating e-mail user security tools. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 57–76, Anaheim, CA, August 2023. USENIX Association.
- [56] Samira Zibaei, Dinah Rinoa Malapaya, Benjamin Mercier, Amirali Salehi-Abari, and Julie Thorpe. Do password managers nudge secure (random) passwords? In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 581–597, Boston, MA, August 2022. USENIX Association.
- [57] Samira Zibaei, Amirali Salehi-Abari, and Julie Thorpe. Dissecting nudges in password managers: Simple defaults are powerful. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 211–225, Anaheim, CA, August 2023. USENIX Association.

Appendix

A Expert Interviews

A.1 Interview Script

I. Intro Section

Interview partner

- *What is your job position at organization XY?*
- *What is your IAM experience (years, clients, access review projects, managed identities, etc.)?*

Access review and its problems

- *Estimate the ratio of excessive granted access.*
- *Name 2-3 major challenges for access reviews.*

II. Explanation Section

- *Explain to the participant the access review challenges of Jaferian et al. [26]. Connect them to the major challenges of access review the participant named before.*
- *Explain to the participant digital nudges in general.*

III. Workshop Section

Mapping digital nudges and access review challenges

For each nudge in Table 5

1. *Explain the nudge and give an example fitting for the interview participant's environment.*
2. *The participant then freely reflects on the digital nudge and their relationship on access review challenges.*
3. *Finally, the participant rates each access review challenge, anticipating a very positive (+2), positive (+1), neutral (0), negative (-1), or very negative (-2) effect.*

Table 5: Digital Nudges [27] presented to the experts for mapping them to access review challenges [26].

<i>Nudges</i>	<i>C1</i>	<i>C2</i>	<i>C3</i>	<i>C4</i>	<i>C5</i>
<i>Decision Information</i>					
<i>N01: Information Translation</i>					
<i>N02: Information Salienc</i>					
<i>N03: Information Visibility</i>					
<i>N04: Information Phrasing</i>					
<i>Decision Structure</i>					
<i>N05: Range & Composition</i>					
<i>N06: Choice Defaults</i>					
<i>N07: Option Consequences</i>					
<i>N08: Option-related Effort</i>					
<i>Decision Assistance</i>					
<i>N09: Reminders</i>					
<i>N10: Commitment Facilitation</i>					
<i>Social Decision Appeal</i>					
<i>N11: Messenger Reputation</i>					
<i>N12: Social Reference Point</i>					
<i>N13: Empathy Instigation</i>					

Wrap-up

- *Name your TOP 3 digital nudges benefiting access review challenges.*

A.2 Codebook

We apply deductive and inductive coding to the expert interviews. The feasibility of digital nudges (based on the collection of Jesse and Jannach [27]) for access reviews suffice as interview questions. The access review challenges of Jaferian et al. [26] suffice as deductive codes, which we applied a priori to the interviews. Therefore, we trained and asked the interview partners about these challenges and asked for a Likert scale-based rating (2 (best), 1, 0, -1, -2 (worst)). The experts answered with different arguments, for which we extracted inductive codes. The rating for digital nudge, challenge and the inductive codes are detailed in the codebook (Table 6).

B User Study

B.1 Data Set

For the user study, we used a crafted data set (160 UPAs). We can pinpoint which UPAs are correctly (TP) and incorrectly (FP) assigned. Figure 5 (using a grid representation based on [32]) depict the data set. A processable format is available at GitHub.

B.2 Ground Truth Document

Access Review Case Study

You work as a busy head of the marketing department in a large industry company with many concurrent projects to maximize the income for your company. Your time is limited, and you have marketing goals to fulfill.

The security teams reminded you via email that your company is legally required (compliance) to review the permission assignments for the employees in your department. You must follow the **principle of least privilege**: Employees must have permissions required for their job, but not more. If you decide to revoke an excessive permission for one of your employees, the employee will no longer be able to access the associated resources by tomorrow.

While the security team points out that any excessive permission poses a security threat, you are aware that missing ones might prevent your employees from working until they re-obtain it via a time-consuming help-desk or self-service request.

The marketing department consists of three teams:

I. Graphic design team

- Create and edit images for the company's media and advertisement presence. This includes banners, logos, websites, or campaign designs that are used in advertisements or social media posts.
- Require a Photoshop license to work.

II. Social media team

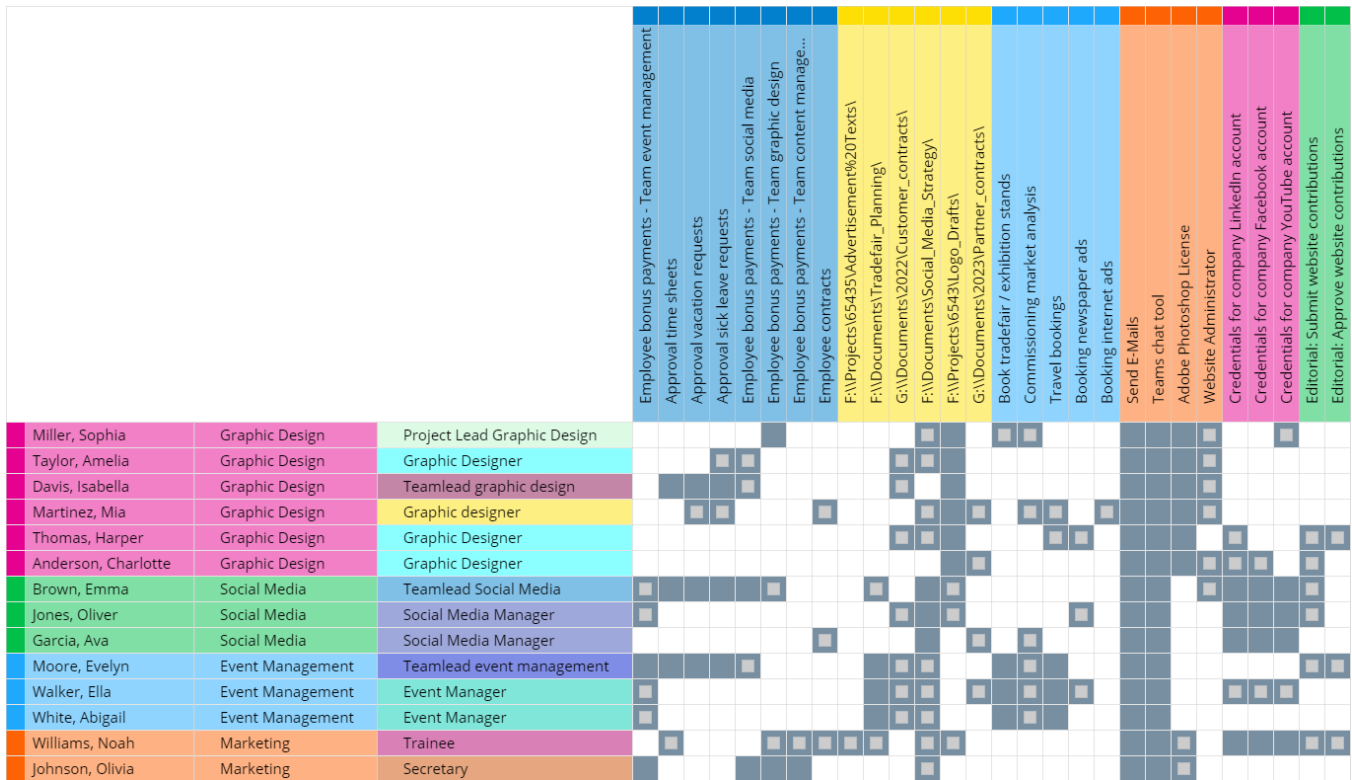


Figure 5: Grid visualization [32] of the user study data set. Blue cells resemble TP, gray ones FP, white ones TN, and FN were not present in the data set.

- Manage the company’s social media accounts.
- Need to communicate with potential customers, candidates for recruiting, and partners online.

III. Event management team

- Organize trade fairs and partner events across West and Central Europe.
- Book trade fair stands.
- High self-organization; often need to attend remote events without long preparation.

IV. Department hierarchies

- Every team is led by a team lead who overlooks the employee’s attendance and work results.
- Team leads have an annual budget for bonus payments, which they can distribute among their team members based on last year’s performance. The secretary reads the specified bonus payments defined by the team leads from the HR system and arranges for the salary to be posted.

- The department’s trainee used to intern in the graphic design team. Now, he is working in the social media team.

V. Misc

- Everybody communicates with MS Teams and Outlook.
- You can sort the columns.

B.3 Screenshots Access Review

We used three configurations of the access reviews with the same data basis. Figure 6, the neutral default, has two white buttons without a preselection. Figure 7 displays the default accept with a preselected *Approve*. Figure 8 shows the default reject with a preselected *Remove*.

B.4 Statistical Analysis

Figure 9 depicts a pair plot for each metric separated for their group. Green shows the default accept group, red the default reject, and blue the neutral one. The upper right part depicts Spearman correlations. The stars indicate the significance levels as "****": $p < .001$; "***": $p < .01$; "**": $p < .05$, and "." $p < .1$. The lower left depicts local regressions. Finally, the diagonal, the first row and column show metric distributions.

1. Please review the authorizations of your employees (0 %)

Search

	Employee ▲	Department	Job title	Permission	Application System
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Send E-Mails	Microsoft Active Directory
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Credentials for company LinkedIn account	Social Media Management System
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Credentials for company Facebook account	Social Media Management System
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	F:\Projects\6543\Logo_Drafts\	Fileshares
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	G:\Documents\2023\Partner_contracts\	Fileshares
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Website Administrator	Microsoft Active Directory
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Adobe Photoshop License	Microsoft Active Directory
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Editorial: Submit website contributions	Website Content Management System
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design	Graphic Designer	Teams chat tool	Microsoft Active Directory
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media	Teamlead Social Media	Approval sick leave requests	HR System
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media	Teamlead Social Media	Teams chat tool	Microsoft Active Directory
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media	Teamlead Social Media	Employee bonus payments - Team social ...	HR System

1/14

Figure 6: Screenshot of the neutral group for the user study.

1. Please review

	Employee ▲	Department
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media

Figure 7: Screenshot for the accept group.

1. Please review

	Employee ▲	Department
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Anderson, Charlotte	Graphic Design
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media
<input type="button" value="Approve"/> <input type="button" value="Remove"/>	Brown, Emma	Social Media

Figure 8: Screenshot for the reject group.

Table 6: Codebook for expert interviews.

N	C	Likert	Inductive Codes
N01	C1	1	Understandability (E02, E03, E04, E09); No effect (E03, E06, E08, E10); Feel-Good (E02, E03); Uniqueness (E04, E09); Structure (E09)
N01	C2	2	Understandability (E02, E05, E06, E07, E08, E09, E10); Mental Load (E03, E05, E07, E09); Acceptance (E05, E07, E10); Wording (E05, E06, E07)
N01	C3	1	Recognition (E01, E04, E05, E06, E09); Learning (E04, E05, E09); Feel-Good (E05, E06)
N01	C4	2	Understandability (E02, E05, E07, E08, E09, E10)
N01	C5	0	Understandability (E05, E06); Recognition (E04)
N02	C1	1	Focus (E01, E02, E04, E05, E09); No effect (E06, E08, E10)
N02	C2	0	Focus (E06, E09, E10); No effect (E03, E07, E09)
N02	C3	1	Economic Efficiency (E01, E02, E04, E07, E09); Focus (E01, E02); Acceptance (E09)
N02	C4	1	Focus (E03, E05, E06, E07, E08, E09, E10); Algorithm-Quality (E03, E06, E09); Backlash (E03, E06, E09)
N02	C5	2	Focus (E03, E04, E05, E07, E10); Algorithm-Quality (E09)
N03	C1	1	More relevancy (E03, E04, E05, E07, E09); Less confusion (E03, E04, E05, E09); No reduction of decisions (E06, E08)
N03	C2	2	Showing more data (E01, E05, E08, E09, E10); Relevancy (E03, E04, E06)
N03	C3	0	Run-time (E05); Recognition (E07)
N03	C4	1	Mistake mitigation (E05, E07, E09, E10); Focus (E06, E07)
N03	C5	2	Showing more data (E01, E07, E09, E10); Need to know (E07, E10)
N04	C1	0	Insecurities of decision-maker (E09); Sense of responsibility (E07)
N04	C2	-1	Context-Awareness (E05, E07, E08, E09, E10); Bias (E04, E05, E09), Base direction (E05, E09)
N04	C3	0	Acceptance (E07)
N04	C4	1	Acceptance (E06; E08; E09; E10); Focus (E06, E09, E10); Pressure (E02)
N04	C5	0	Focus (E06, E07, E10)
N05	C1	2	Similarities (E01, E03, E04, E05, E07, E08, E09); Overhead (E08, E10)
N05	C2	1	Focus (E04, E08, E09, E10); Audience (E08, E09)
N05	C3	1	Economic Efficiency (E03, E05, E06); More Tasks (E09, E10)
N05	C4	2	Focus (E01, E03, E05, E06, E07, E09, E10); Similarities (E01, E05, E06, E07); Smaller Batches (E09, E10)
N05	C5	2	Exceptional Case Detection and View (E02, E03, E04, E07, E09)
N06	C1	2	Less work (E01, E02, E04, E05, E06, E09); No reduction of decisions (E07, E10)
N06	C2	-2	Recommendation Fallacy (E02, E04, E05, E06, E07, E09, E10); Recommendation Support (E06, E09)
N06	C3	2	Less work (E01, E02, E04, E05, E06, E08, E09)
N06	C4	-2	Less diligence/Focus (E01, E02, E04, E05, E06, E07, E09, E10); Recommendation Fallacy (E02, E04, E05, E06, E07, E09, E10)
N06	C5	0	Not in Focus (E05, E07, E10); Default handling (E06, E09); Special treatment (E04)
N06	Misc		Not Compliant (E01, E03, E07, E09); Needs good recommendation (E01, E03, E08, E09); Is it really a decision? (E03, E07, E09)
N07	C1	0	Speed (E01, E09)
N07	C2	-1	Recommendation Fallacy (E09)
N07	C3	1	Speed (E01, E04, E09); Gamification (E04, E05, E09); Feel-Good (E04, E07); Acclimatation (E09)
N07	C4	-1	Pressure (E01, E03, E07, E09, E10); Recommendation Fallacy (E06, E07, E09, E10); Less diligence (E01, E03, E07)
N07	C5	-1	Recommendation Fallacy (E07, E09); Fairness for disadvantaged individuals (E09)
N08	C1	-1 / 1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C2	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10)
N08	C3	-1 / 1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C4	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C5	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10)
N09	C1	0	No effect (E03, E07); More participation (E04, E10);
N09	C2	1	Instructions and Goals (E02, E03, E04, E05, E07, E08, E09, E10); Spam (E01, E02, E03, E07, E09, E10);
N09	C3	2	Spam (E01, E02, E03, E07, E09, E10); Attention (E04, E06, E07, E09, E10)
N09	C4	-1	Revisit (E03, E07, E08); Pressure (E05); Multi-Channel (E07, E09, E10); Audience (E03, E09)
N09	C5	0	Open Task (E01); No effect (E07)
N10	C1	1	Combination with N05 - Commitment for partitions (E02, E04, E06, E07, E08, E09, E10)
N10	C2	0	Autonomic planning and understanding (E04, E05, E08, E09)
N10	C3	1	Combination with N05 - Sub-Deadlines for partitions (E05, E07, E08, E09, E10); Comfort (E02, E06, E07, E08, E09, E10)
N10	C4	1	Focus (E04, E06, E07, E08, E10); Comfort (E02, E06, E07, E08, E09, E10)
N10	C5	0	Focus (E07, E10)
N11	C1	1	Endurance (E01, E02, E03, E04, E05, E07, E09, E10); Trust (E02, E03, E07, E08, E09)
N11	C2	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10)
N11	C3	1	Endurance (E01, E02, E03, E04, E05, E07, E09, E10); Trust (E02, E03, E07, E08, E09)
N11	C4	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10); Acceptance (E01, E02, E03, E04, E08, E10)
N11	C5	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10)
N12	C1	0	Endurance (E02, E03, E06, E09); Backlash (E09)
N12	C2	2	Approachable Peer-Group (E02, E03, E04, E07, E08, E09, E10)
N12	C3	0	Endurance (E02, E03, E06, E09); Backlash (E09)
N12	C4	1	Acceptance (E02, E03, E06, E07, E09, E10); Peer-Pressure (E03, E10)
N12	C5	2	Approachable Peer-Group (E02, E03, E04, E07, E08, E09, E10)
N12	Misc		Similarity to N11 messenger reputation (E01, E05)
N13	C1	1	Feel-Good (E02, E04, E06, E07, E09)
N13	C2	1	Feedback on odd behavior (E02, E03, E04, E07, E09)
N13	C3	1	Feel-Good (E02, E04, E06, E07, E09)
N13	C4	1	Feel-Good (E02, E05, E07, E08); Focus (E02, E04, E07, E09)
N13	C5	0	Feel-Good (E05)

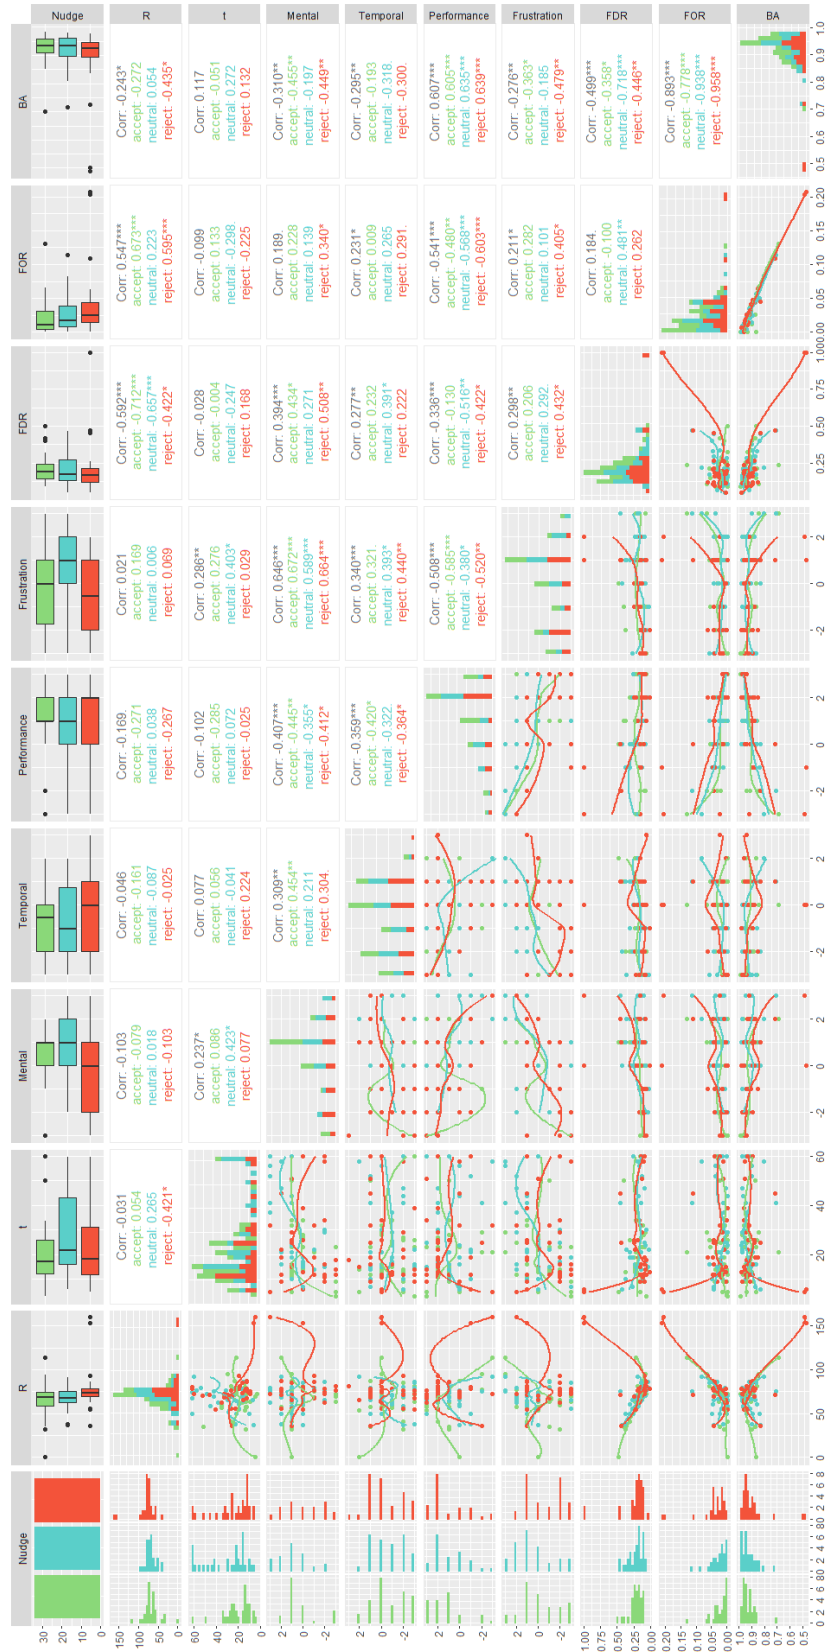


Figure 9: Pair plot of correlations (Spearman) and local regressions for the user study.