Credential Disclosure in (EU) Digital Identity Wallets: Compliance, Privacy Risks and Practical Mitigations

Sheila Zingg *ETH Zurich*
Daniele Lain *ETH Zurich*
Yoshimichi Nakatsuka *ETH Zurich*
Kari Kostiainen *ETH Zurich*
Stefan Bechtold *ETH Zurich*
Srdjan Čapkun *ETH Zurich*

Abstract

The European Union will introduce the EUDI Wallet by late 2026, which allows users to hold digital credentials (i.e., representations of physical official identity documents) on their devices. This will allow users to securely and privately disclose identity attributes to websites. Although such a system has many benefits, it also introduces risks caused by poor credential disclosure decisions. In this paper, we (i) conduct a large-scale survey on credential disclosure with users and experts and (ii) evaluate the effectiveness and feasibility of our Credential Assistant that displays expert recommendations and user opinions. Our results show that users are likely to both undershare (e.g., only ∼40% of users disclosed their marriage certificate to government websites) and overshare (e.g., ∼20% of users disclosed their official ID to news websites). This indicates that users struggle to understand compliance requirements and protect their privacy, which will impact the usability of the EUDI Wallet and lead to privacy violations, identity theft, and other abuses of leaked credentials. Finally, we show that our Credential Assistant significantly reduces users' credential disclosure mistakes from ∼15% to ∼7%. However, it does not fully eliminate poor credential disclosure decisions, indicating that stronger interventions may be necessary, especially for sensitive attributes.

1 Introduction

The need for online identity verification is increasing. Services are moving increasingly online, and many of those services have legal obligations to perform stringent identity checks (e.g., KYC for banks and payment providers, identity checks for lodging providers, identity checks for e-Visas). This move towards online services is also spreading to sectors that must verify data about users beyond their identity (e.g., a pharmacy verifying a prescription online). Furthermore, governments are increasingly requiring stricter enforcement of user age verification [8, 20, 40, 63].

EUDI Wallet To respond to the growing need for online identity verification, the European Union introduced the EU Digital Identity (EUDI) Wallet [21, 23] in 2024, with the aim of having a functional system operational in late 2026. The EUDI Wallet will allow users to hold digital credentials, which are representations of physical official documents, on their phone, in the EUDI Wallet app. Verifiers (e.g., websites) can request attributes contained in a credential from users through the EUDI Wallet, and users can accept those requests to disclose the attributes. This reduces the friction users experience from online identity verification, which currently often requires users to submit pictures or videos of documents. Additionally, this system can increase security, as photo and video identity verification can be tricked [87, 88]. Lastly, the system gives users more control over their data, allowing them to disclose individual attributes rather than send a full credential.

However, giving users full control over their data and reducing the friction of disclosing data also bring risks. Currently, if a malicious website wants to collect the official IDs of users, it would have to ask users to submit an image or video of their ID, which many users would not do because it requires effort. However, in the future, a malicious website could request this credential via the EUDI Wallet, and users could disclose the credential with just one click, which is more likely to succeed. This risk is not far-fetched, as prior research has shown that users are likely to accept any data request even if the data is not actually needed [2, 5, 15, 27, 30, 41, 64, 76] and it is common for websites to request more data than they need [56, 59, 65, 79]. However, data disclosure is not always optional (e.g., it is legally required for a bank to collect users' official IDs when they open a bank account). If users are too conservative and reject required data requests, their user experience may suffer. This, in turn, may make them more likely to accept all requests in the future, as they believe services will not work without the requested data, which would cause the EUDI Wallet to lose many of its benefits.

The risk of bad disclosure decisions due to the EUDI Wallet has not been ignored by governments. The EUDI framework

requires mandatory registration for websites that want to request EUDI credentials, during which they must enter all attributes they plan to request. The EUDI Wallet rejects any request that exceeds what the website entered during registration, and the website's access to the EUDI system may be suspended or canceled in such cases [26, Art. 9]. However, this safeguard does not protect against over-claiming: if a website registers more attributes than it actually needs, there is a risk that users will overshare personal data to avoid losing access to the website. This problem is similar to websites over-claiming what data they need in their privacy policies nowadays. While such over-claiming may run afoul of the GDPR's data minimization principle [75], the EUDI framework does not provide an explicit safeguard against such behavior. Furthermore, certifying the data requirements of all websites is challenging, as similar ideas (e.g., automatically certifying the permission requirements of apps) have failed.

This paper We conducted a large-scale survey with 1035 users and 27 experts from the fields of cybersecurity, law, policy, and ethics. In this survey, we measured which credentials users would disclose to websites and asked experts which credentials they deemed necessary for, which credentials they believe should never be disclosed to, and which credentials they would disclose to websites. The goal of this survey was to understand whether users make risky credential disclosure decisions and what credentials a website may need, according to experts. Furthermore, we evaluated the effectiveness and deployment feasibility of our Credential Assistant that shows expert recommendations and/or information about what other users would do based on the category a website belongs to. For this, we conducted a large-scale user study with 1002 participants, in which participants were shown simulated EUDI credential requests that they could accept or reject.

Our results indicate that users are likely to both undershare (e.g., only ∼40% of users disclosed their marriage certificate to government websites) and overshare (e.g., ∼20% of users disclosed their official ID to news websites) credentials. Furthermore, our results indicate that the oversharing risk will likely be higher in real-world deployments, as users overshared more when seeing a credential request, even if no dark patterns were used (e.g., for disclosing medical records with LinkedIn, ∼31% of users accepted the credential request, but ∼0% of surveyed users said they would disclose). Our results also show that the Credential Assistant was able to reduce the number of mistakes users make significantly, in most cases (e.g., it reduced the average number of mistakes users made from 15% to 7% when showing expert recommendations). However, the Credential Assistant did not have any nudging effect if it showed low confidence numbers (i.e., 51%-55% recommend to disclose). This may be an issue, as our results show low expert agreement in many cases where there is a good justification for a website to request a credential.

We conclude that our Credential Assistant is able to mitigate the risk of poor credential disclosure decisions; however, residual risk remains. Although 7% may seem low, on the scale of the EUDI, this still puts the sensitive identity information of tens of thousands of users at risk. Thus, more intrusive interventions (e.g., similar to browser warnings) may be needed to ensure that users do not ignore them.

Our Contribution To summarize, the main contributions of our paper are:

1. *Large-scale survey and user study* We conducted a survey with over 1000 users and over 20 experts to understand credential requirements for websites. We conducted a user study with over 1000 users to evaluate the effectiveness of our Credential Assistant for the EUDI using a wide range of website types and credentials.
2. *EUDI credential disclosure risks* We found evidence of both oversharing and undersharing, which may result not only in privacy violations, but also in identity theft and reduced user experience.
3. *Credential Assistant for the EUDI Wallet* We evaluated a scalable Credential Assistant that displays expert recommendations and user opinions, and show that it is effective in reducing the risk of poor disclosure decisions; however, high residual risk remains. Thus, considering the sensitivity and the planned scale of the EUDI, more intrusive interventions may need to be explored.

2 Background

The European Union will introduce the EU Digital Identity (EUDI) by the end of 2026 [21, 23]. This consists of an infrastructure for online identity verification and the EUDI Wallet app, in which EUDI credentials (digital representations of official identity documents, e.g., a digital ID) can be stored. A regulation [80], architecture reference framework (ARF) [22], and more than two dozen Implementing Acts [24] have been published, based on which EU member states must revise their laws and procedures to implement their EUDI Wallets.

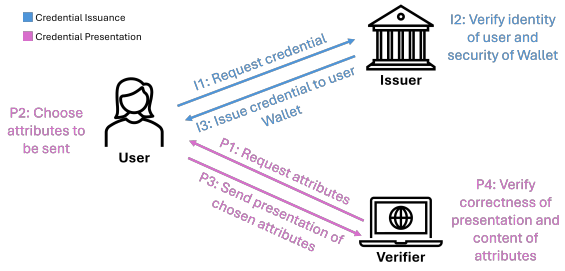

Figure 1: The parties and processes involved in the EUDI.

Fig. 1 shows an overview of the planned EUDI. There are three main parties: 1) the credential holder (which we will call

the user), 2) the issuer, and 3) the verifier or relying party. The issuer signs a credential containing attributes and issues it to a user. The user stores this credential on their own device in the EUDI Wallet app. Later, a verifier can request attributes from a user, and the user can accept or reject the request. If the user accepts the request, the EUDI Wallet generates a presentation that proves that the credential was issued to the user by a specific issuer and discloses the requested attributes contained in the credential. The verifier validates that the credential was issued by a trusted issuer and that the credential was issued to the user who presented the credential.

The following aspects of the EUDI are important for this paper. I) A user can present an EUDI credential to a verifier without communicating with any other party. II) Initially, the EUDI Wallet will only store digital ID cards, but it is planned to be able to hold various credentials. III) EUDI credentials are planned to consist of individually hashed attributes that can be disclosed individually. IV) The UI of the EUDI Wallet has not been decided, and each member state can choose how their version of the EUDI Wallet looks. V) All issuers and verifiers will be registered, allowing anyone to verify who issued a credential and who is requesting attributes. A verifier registration that contains all attributes that a verifier wants to request is planned, and the EUDI Wallet would reject any request that does not follow this registration [23, 25]. VI) Currently, there is no plan to allow websites to display the purpose for a request in the EUDI Wallet. However, adding this may be considered if helpful to users.

Threat Model We assume that all parties follow the EUDI protocol correctly; however, the verifier may request more data than needed. This assumption is reasonable, as websites often over-collect data [56, 59, 65, 79].

3 Main Results & Key Takeaways

Our research is divided into two areas: **A1 Disclosure Behavior**, where we measure users' credential disclosure decisions and identify risks. **A2 Decision Support**, where we evaluate the effectiveness of a scalable Credential Assistant for the EUDI. To explore these two areas, we conducted a survey and user study (details in Section 4). In this section, we discuss why these research areas are relevant and summarize the main results and key takeaways. Actionable recommendations based on our results are listed and discussed in Section 7.

3.1 A1 Disclosure behavior

Oversharing online has been shown in many contexts [2, 5, 15, 27, 30, 41, 64, 76]. However, to the best of our knowledge, no prior work has investigated user disclosure behavior for digital versions of official documents (i.e., digital credentials), and it is unclear whether the same conclusions can be drawn in this context. On the one hand, since users handle official physical documents in their day-to-day lives, one could assume that they can evaluate risks and benefits better compared to more abstract concepts like cookies. On the other hand, users have been shown to make more risky decisions online [74]. Furthermore, the usability loss from disclosing too little and the risk from disclosing too much are different for digital credentials than in other contexts. In this area, we investigate: What misconceptions do users have about credential disclosure? Do users make poor disclosure decisions for EUDI credentials?

Main results: First, our results show that users undershare and overshare EUDI credentials. We define undersharing as cases where there is a good reason for the website to request the credential, but many users would not disclose. Examples include: only ∼40% of users would disclose their marriage certificate to government websites (which will likely be required for, e.g., tax and immigration services), only ∼45% of users would disclose their visa to air travel, and only ∼54% of users would disclose their official ID to international ground travel websites (which will likely be required to simplify border control). We define oversharing as cases where there is no valid reason for the website to request the credential, but many users would still disclose. Examples include: ∼20% of users would disclose their official ID to news, ∼10% would disclose their visa to banking, and ∼6% would disclose their health insurance to payment service websites. Although some of these numbers may seem small, at the planned scale of the EUDI, this would affect thousands or even millions of users.

Furthermore, our results show that oversharing increases when users make false connections between the credential name and the attributes it contains, or the website name and the credentials it requires. For example, users often believed that the bank account validation could be used to pay even though this was not possible, leading them to overshare it (e.g., ∼34% for e-commerce websites). Furthermore, users connected encyclopedia websites such as Wikipedia with educational or professional credentials (e.g., ∼10% would disclose a diploma, ∼6% would disclose professional licenses) even though this connection does not actually exist. Additionally, we find that users were more conservative when deciding which credentials they would disclose to a website (what they think they should do) versus when responding to a request (how they actually act). This mirrors the privacy paradox, where users' stated privacy sensitivity does not match their actions [28, 51, 76]. For example, ∼31% of users accepted a request for medical records from LinkedIn compared to ∼0% of surveyed users who stated that they would provide medical records to LinkedIn. We observed this change in behavior even though our simulated requests did not use dark patterns, which are common on real-world websites [56, 59, 60, 65].

Key takeaways: The risk for users when making poor disclosure decisions in the EUDI is high. The risk of undershar-

ing may not be immediate, as it is likely that the website will inform the user that the credential is necessary, after which the user will choose to disclose. However, this back-and-forth reduces the user experience and may make users more likely to accept any credential request in the future, as they believe that websites will not work otherwise. The most apparent risk of oversharing is loss of privacy; however, other risks, such as identity theft, can result from the misuse of leaked credentials. The EUDI is not planned to be mandatory; therefore, there will be value for malicious parties to collect identity information that they can later use to impersonate their victims using traditional identity verification methods. The official ID is a credential that contains most information used in weak identification processes; however, it was also often overshared in our experiment. Thus, a malicious verifier could set up a fake website in which they request users' official ID credentials, which requires minimal effort. After that, they can use the collected identity information to impersonate their victims.

One example is *SIM swapping attacks*, where a malicious verifier receives a replacement SIM with the phone number of their victim. This, in turn, allows them to receive SMS 2FA codes instead of their victim, with which they can infiltrate additional accounts, e.g., bank accounts. Empirical research [33, 34, 53] has shown that mobile carriers often use weak identity verification when issuing replacement SIMs, especially for remote replacement. Furthermore, they could *extract financial and health data* through the call-in identity verification of health insurers and banks. Health insurers and banks often use weak identity verification (e.g., asking name, date of birth, address, residency number, last large payment received, etc.)[1] to verify the identity of a caller. Lastly, they could *open a bank account in the name of their victim*. While many banks use liveness checks when verifying the user's identity, research has shown that deepfakes can trick those checks [9,54]. Thus, a malicious party, who knows the content of a victim's official ID, can use this information to create a deepfake of their victim and impersonate them.

3.2 A2 Decision Support

Decision support systems have been studied and shown to be effective in multiple contexts [2–4, 14, 38, 47, 55, 86]. Most of those systems follow one of two approaches: 1) display information about what other users decided to do, and 2) provide information about privacy risks and what data must be disclosed to use a service. We believe that neither of these approaches can be applied to the EUDI without adjustments. 1) As discussed in A1, users made mistakes when deciding which credentials to disclose to websites. Thus, we believe that a system solely based on user opinions is insufficient, and expert opinions need to be leveraged. 2) Collecting expert evaluations for a large number of websites in a scalable way is not trivial. Additionally, the interdisciplinary nature of the EUDI means that multiple experts may need to be consulted, increasing this challenge. For this research area, we evaluate a scalable Credential Assistant that uses the category of a website as a proxy to display gentle, non-intrusive nudges, i.e., it displays expert recommendations and/or information about what other users would do. Using this Credential Assistant, we investigate: Is it sufficient to display information based on website categories? How effective is the Credential Assistant?

Main results: Our results, somewhat surprisingly, show that users did not significantly base their decision on the headquarters or traffic size of the website. The standard deviation between different websites was low for most website-credential pairs, and higher standard deviations were mostly seen in cases where disclosing the credential was controversial (around 50% disclosure rate). There were only two cases with high standard deviations that were not explained by other factors: the official ID for e-commerce (0.08) and the driver's license for international ground travel (0.14). This indicates that the website category is a good proxy for collecting and displaying information in this context, and thus, the Credential Assistant can be built in a scalable way, since experts do not have to evaluate each website individually.

We simulated a Credential Assistant that always provided the correct information to test if users could be nudged towards better decisions. With such a Credential Assistant, the number of mistakes users made significantly decreased regardless of whether they saw expert recommendations or user opinions (the average number of mistakes reduced to $\sim$7% from $\sim$15% in both cases). However, as discussed in A1, users often make mistakes, and thus, it is unlikely that a Credential Assistant based on user opinions would be accurate enough. Our results show that expert recommendations were more accurate, especially in cases where there is no justification for a credential request (71%-96% would recommend never to disclose). However, the expert consensus was lower (50%-71%) for many cases where the website may require the credential, which could be an issue, as in our study, the Credential Assistant did not have a nudging effect when it showed low confidence numbers (i.e., 51%-55%).

Key takeaways: Our Credential Assistant reduces the number of credential disclosure mistakes that users make, indicating that implementing such a support system will help users. However, a residual mistake rate of $\sim$7% remained, which is problematic considering the sensitivity of EUDI credentials and the scale of the planned EUDI. One reason may be that our Credential Assistant was fully non-intrusive, and thus, it was easy for users to ignore. More intrusive interventions (e.g., similar to browser warnings) may be needed to ensure that users are aware of the information that Credential Assistant provides. Another reason may be that Credential Assistant always showed the same type of display, regardless of the sensitivity of the credential and regardless of whether

[1] Tested with banks and health insurers in 3 European countries.

the user would have made a mistake or not, which may have made users less attentive. Thus, it may make sense to limit the display of Credential Assistant to sensitive credentials or to when the user is about to make a poor disclosure decision.

4 Experimental Setup

In this section, we first describe why we need a survey with users and experts and a user study to explore A1 and A2. After that, we describe the setup of the survey and user study.

4.1 Overview of our Methodology

We first discuss what experiments A1 and A2 require. Then, we combine them into a survey and a user study.

A1 Disclosure behavior The privacy paradox [28, 51, 76] has shown that users often act less privacy-consciously than their stated privacy preferences. Thus, disclosure behaviors should be investigated through what users *think they should disclose* and what users *disclose when asked*. To understand what users think they should disclose, we need a survey in which users pick the credentials that they would disclose to different websites. To understand what users would disclose when asked, we need a user study with simulated EUDI credential requests for which users receive no additional support. To accurately represent the planned EUDI, we need to assess a large range of websites and credentials. However, collecting user study data for such a large number of website-credential pairs would be unfeasible. Thus, we first conduct a survey to understand users' disclosure beliefs and then a user study to understand the difference between users' beliefs and actions.

A2 Decision Support We aim to understand 1) whether the website category is a good proxy to display information in the Credential Assistant, 2) whether different versions of Credential Assistant are effective at reducing the mistakes users make, and 3) whether the required data for Credential Assistant can be collected from experts. For 1, we need a survey in which users pick the credentials that they would disclose for a wide range of websites (including from different countries and with different traffic sizes) in each category. This will allow us to understand whether user opinions are uniform within a category (meaning that the website category is a good proxy). For 2, we need a user study with simulated EUDI credential requests in which users receive access to one of many Credential Assistant versions. This will allow us to measure how user behavior changes for each version. For 3, we need a survey with experts to understand for each website which credentials are required, risky, and make sense to disclose. This allows us to understand how many experts need to be consulted to get an accurate representation of expert recommendations and whether there is expert consensus.

Study Overview Putting this together we require: I) a survey with users where they pick which credentials they would disclose to a large range of different websites, II) a survey with law, cybersecurity, policy, and ethics where they provide expert evaluations regarding the credential risks and requirements for websites from different categories, and III) a user study where users with and without a Credential Assistant need to react to simulated EUDI credential requests. Since I and II are similar surveys, we will combine their description.

Ground Truth One challenge with measuring oversharing and undersharing is that there is no ground truth, since disclosure decisions are subjective. However, in some cases, there is no valid reason for a website to request a credential, and in other cases, there is a use case that makes disclosure sensible. Based on this, we categorized each website-credential pair into justified, unjustified, or uncertain and defined undersharing as not disclosing when justified, and oversharing as disclosing when unjustified. We validated our categorization with the expert survey results (details in Appendix A).

4.2 Survey on EUDI Credenials

Credential Choice We chose credentials based on the W3C verifiable credential use cases [83] and the use cases noted in the EUDI ARF [22]. We used 14 credentials, which covered a wide range of use cases. As the attributes of the credentials have not yet been decided, we chose a set of attributes for each credential. These attributes generally match the current physical documents, but we removed redundancies. For example, most documents currently contain the name of the owner; however, it is unclear if this would be the case in the EUDI, so we removed this redundancy.

Website Choice We chose website categories based on Cloudflare Radar's domain categories [18]. We used 17 website categories that covered a wide range of services. For users, where possible, the websites covered different traffic sizes, and were Europe-based unless all well-known websites of a category came from outside of Europe. For 3 website categories, we used websites from Europe, the USA, and China to investigate whether the country mattered to users.

Participants 27 experts completed our survey: 3 law, 15 cybersecurity, 5 policy, and 4 ethics. We recruited experts by reaching out to them directly via e-mail, and no payment was provided. To protect the identity of the experts, we did not ask demographic questions.

1107 users completed our survey. Participants were recruited through the Prolific platform and were paid ∼$6.7 for the 25 min survey. Participants were chosen with a balanced gender and age split, and were required to reside in Germany, Italy, or France. We removed submissions that were incomplete, in which the participant revoked consent, or that were of

low quality with an exceptionally fast completion time. After this, we were left with 1035 valid responses.

Survey Flow Steps in *italics* were not shown to experts.

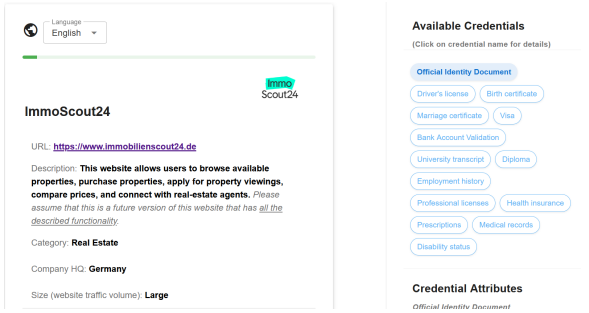

(a) Website description and credentials in the sidebar.

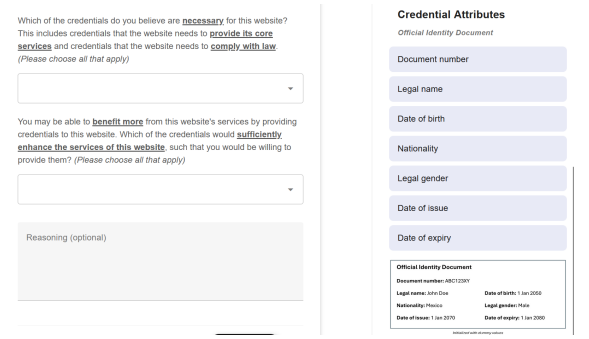

(b) Survey questions and the credential attributes in the sidebar.

Figure 2: Survey screens of the website evaluation task (more in Appendix C). The sidebar shows details on the credentials.

1. Consent form
2. *Pre-survey demographic questionnaire*
3. *Evaluation of credentials*

   In this task, users were asked to pick which of the credentials they own or owned as physical or electronic official documents. After that, for each document they picked, they were asked how often they use that document and how sensitive that document is to them. As the privacy paradox has shown that users often act differently than they think they should act [28,51,76], we decided to split the sensitivity question into two questions: 1) How comfortable are they with sharing this document? 2) When would they share the document even if uncomfortable?

4. Evaluation of websites

   Participants were first shown a short explanation about what the EUDI Wallet is and how EUDI credentials are planned to look and function. Then they were shown different websites (with some website details). For each website, participants were asked to choose: [law experts] the credentials that are necessary for the website, [cybersecurity, policy, and ethics expert] the credentials that should never be requested by the website, and [all] the credentials that they would disclose to the website. Participants could access details on all credentials in a sidebar. Fig. 2 shows the survey screens.

For experts, the survey was provided in English. For users, the survey was provided in English, German, French, and Italian. Translations from English to the other languages were performed using an LLM (Copilot [1]), and all translations were proofread by a native speaker to ensure that the meaning of the survey questions matched in all languages.

4.3 User Study on the Credential Assistant

Credential Assistant Overview We evaluate a Credential Assistant that shows gentle, non-intrusive nudges in the form of expert recommendations and/or user opinions. Since the EUDI Wallet is planned to support a large number of websites, the Credential Assistant must scale. However, collecting expert opinions for each individual website will not scale. Thus, we evaluate a Credential Assistant that uses the category of a website to display information and investigate if adding the website's traffic size and headquarters to refine the categories makes sense. The reason for this choice is that 1) the website's category largely determines what data it needs, 2) the website's reputation (for which traffic size is an indicator) [12,46] and headquarters [19,45,71] determine user trust, 3) there are a limited number of website categories, and thus, expert data collection is feasible, and 4) there are services (e.g., Cloudflare Radar [17]) that categorize websites.

Scenario Choice A scenario contained a website, a requested credential, and, potentially, a website-provided purpose. The scenarios were grouped into 6 sets, S1-S6, each of which was designed for a different test group. We ranked each credential's sensitivity based on the user survey responses and made a balanced choice of credential sensitivities for the scenarios in each set. S2/S4 contained a website-provided purpose. In S5/S6, the Credential Assistant displayed both expert recommendations and user opinions, while it only showed one of the two in the other scenarios. S2/S3/S5 used unjustified, S4/S6 used justified, and S1 used a mix of justified, unjustified, and uncertain website-credential pairs. Thus, we say that choosing yes in S2/S3/S5 and choosing no in S4/S6 is a mistake. S1 tested different confidence levels (i.e., percentages for the user opinion) for the Credential Assistant.

Test groups Participants were assigned to the control, the baseline, or the test group and saw 10 scenarios in total. Participants in the control group did not see the Credential Assistant. Participants in the baseline group saw a Credential Assistant that always had standard confidence, made no mistakes, and

did not contradict a website-provided purpose. Half of the baseline group saw specific expert types in the Credential Assistant, while the other half saw an unspecified expert. The test group was divided into 16 classes that each tested one condition in which the Credential Assistant showed unexpected behavior, i.e., low confidence, very high confidence, mistakes in the displayed information, or displayed information at odds with the website-provided purpose. Since it would not make sense for the Credential Assistant to always show unexpected information, participants in the test group saw 2 or 4 test scenarios, and otherwise baseline scenarios. For completeness sake, a detailed overview of the scenarios and test groups is shown in Appendix E.

Participants 1024 participants completed our user study. Participants were recruited through the Prolific platform and were paid ∼$6.7 for the 25 min study. Participants were chosen with a balanced gender and age split and were required to reside in the EU. We removed submissions that were incomplete or in which the participant revoked consent. After this, we were left with 1002 valid responses.

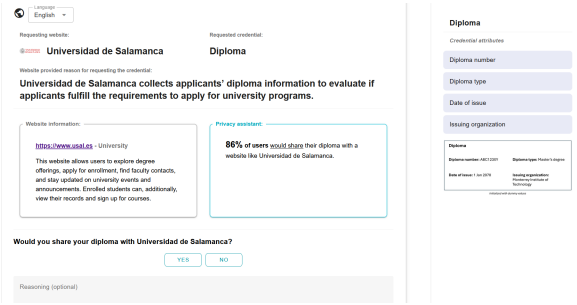

Figure 3: Study screen for the credential disclosure task with the Credential Assistant and a purpose (more in Appendix C). The sidebar shows the attributes of the requested credential.

User study Flow

1. Consent form
2. Pre-study demographic questionnaire
3. Credential disclosure task

   Participants were first shown a short explanation about what the EUDI Wallet is and how EUDI credentials are planned to look and function. Then, participants (except for the control group) were introduced to the Credential Assistant. After that, participants were shown multiple simulated scenarios of a website requesting a credential. Users did not have to navigate a real or simulated website. The task was to decide whether they would accept or reject the request. As the EUDI Wallet does not exist yet, we used a UI based on similar credential requests (e.g., OAuth). Fig. 3 shows the study screens.
4. Post study questionnaire

The study was provided in English, German, French, Italian, and Spanish. Translations from English to the other languages were performed using an LLM (Copilot [1]), and all translations were proofread by a native speaker to ensure that the meaning of the study questions matched in all languages. We used a commitment check at the beginning of the study, which has been shown to increase response quality [70].

5 Results

In this section, we discuss the results of our survey and user study based on the methodology in Section 4. Additional results are shown in Appendix B. As mentioned in Section 4.1, we categorized website-credential pairs into justified, unjustified, and uncertain depending on whether there is a valid reason for the website to request the credential. We validated our categorization with the results of the expert survey (details in Appendix A), and find that they match in most cases.

5.1 A1 Disclosure behavior

In this section, we identify and discuss examples of oversharing and undersharing in users' disclosure beliefs and actions.

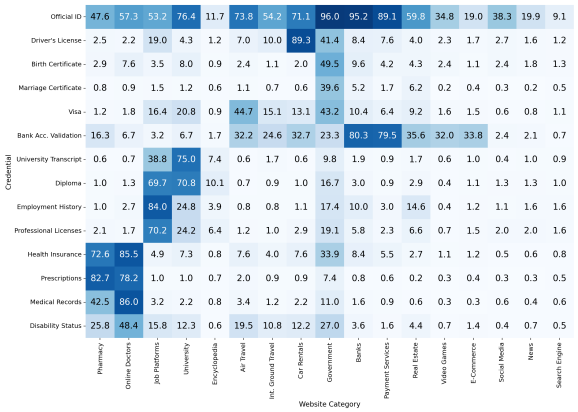

Figure 4: User survey results showing the percentage of users who would disclose a credential for each website category.

General undersharing and oversharing Fig. 4 presents the results of the user survey showing what percentage of participants would disclose a credential for each website category. While the numbers often follow expectations, there are multiple instances of undersharing and oversharing. Examples of undersharing include the following. Only ∼40% of users would disclose their marriage certificate, only ∼50% would disclose their birth certificate, and only ∼43% would disclose their visa to government websites, even though such documents are often required for government services, such as

taxes, immigration services, and marriage registrations. Furthermore, only ∼45% of users would disclose their visa to air travel websites, and only ∼54% would disclose their official ID to international ground travel websites, even though it is likely that this credential will be collected in the future to speed up compliance checks for border control requirements.

Examples of oversharing include the following cases, where there is no justification for disclosure. ∼20% of users would disclose their official ID to news websites, ∼9% to search engine websites, and ∼12% to encyclopedia websites. Furthermore, ∼6% of users would disclose their health insurance to payment service websites, and ∼7% would disclose their prescriptions to government websites. Although these percentages may seem low, the absolute number of credentials a verifier can collect can be high for high-traffic websites, enabling scalable collection of credentials without much effort. For instance, it is not unrealistic to assume that a search engine reaches a traffic volume of a million unique users. If ∼9% are willing to disclose their official ID, then a verifier can collect 90,000 official IDs. This is in line with prior research in other contexts that has shown that users were willing to disclose data even for small benefits [2, 5, 15, 27, 30, 41, 64, 76].

Naming confusion The results of the user survey presented in Fig. 4 also show oversharing when users connect the website or credential name to a use case that is not present. For example, encyclopedia websites (e.g., Wikipedia) would likely require no credentials at all, as they simply compile and allow the search for information. However, the names of such websites often contained words related to knowledge, which in turn can be connected to education and profession. This is likely why a higher number of participants said they would disclose university-related (∼10% for diploma, ∼7% for university transcript) or profession-related credentials (∼6% for professional licenses, ∼4% for employment history). Another example is the bank account validation credential, which did not contain the data needed to execute a payment. Still, many participants believed that they could pay with this credential ("... bank account verification is sometimes required when purchasing a premium subscription ..."), and thus, frequently said they would disclose it (∼32% for video game, ∼34% for e-commerce, ∼32% for air travel websites). These examples are especially concerning, as they indicate that a verifier could name their website in a specific way or leverage misconceptions about a credential to collect more data. To the best of our knowledge, this effect has not been investigated before.

Oversharing for credential requests For all user study scenarios where disclosure is unjustified, we compared the responses of the control group with the number of users who said they would disclose the same credential to the same website in our survey. This indicates whether users are likely to disclose more when responding to a request compared to what they think they should do, which is important, as this is closer

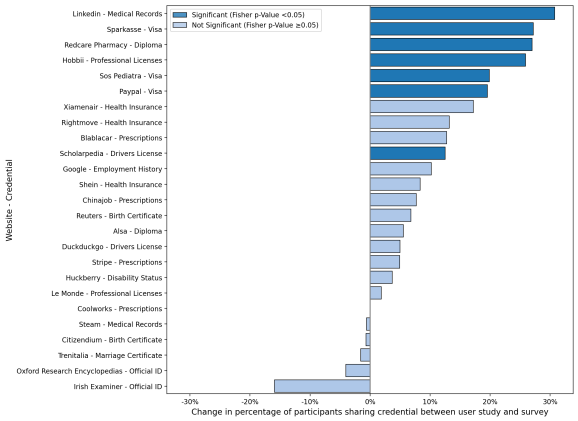

Figure 5: Comparison of how many participants said they would disclose a credential for a website in the user survey, as compared to how many control group participants in the user study disclosed the credential to the website. Positive percentages indicate increased disclosure in the user study. For all website-credential pairs, disclosure is unjustified.

to how the EUDI Wallet works. The results of this comparison are presented in Fig. 5, and show that in most cases, more users disclosed a credential when responding to a request, in many cases significantly so. Examples include: ∼31% of users disclosed their medical records to LinkedIn as compared to ∼0% of users saying they would do so, ∼29% of users disclosed their diploma with Redcare Pharmacy as compared to ∼2% of users saying they would do so, and ∼24% of users disclosed their visa with Sos Pediatra as compared to ∼4% of users saying they would do so. Furthermore, our results show that this increased oversharing occurred less often when there was a good justification for the credential request, indicating that it does not simply come from the study design (results in Appendix B.3 for completeness). We note that, in practice, users may be even less likely to reject credential requests, as websites may use dark patterns [56, 59, 60, 65] or refuse to provide services. Prior work has shown similar effects, such as that users act less privacy conscious than they state they are (the privacy paradox) [28, 51, 76], and that users usually fill out all fields, including optional fields, in online forms [52, 68]. However, to the best of our knowledge, increased oversharing when responding to requests has not yet been explored.

5.2 A2 Decision support

We evaluate a Credential Assistant that uses the category of a website to provide information to help users decide whether to disclose a credential. In this section, we first assess whether the website category is a good proxy to choose the information to display. After that, we evaluate whether it is feasible to collect the data needed for the Credential Assistant. Lastly, we measure whether the Credential Assistant is effective at help-

ing users make better EUDI credential disclosure decisions and discuss the impact of different design choices.

Impact of country and traffic size We evaluated how often user survey participants would disclose a credential to a Europe-based compared to a USA- or China-based website, and how often they would disclose a credential to a high-traffic (large) compared to a low-traffic (small) website (details in Appendix B.1). The results show that there is little difference between how often participants would disclose a credential for different countries or traffic sizes, with only a few entries showing a significant difference. Thus, even though some users mentioned considering the country (e.g., "Due to its connections to China, I would not entrust any data to this website", "I do not want to give any data to the USA, nor to its subsidiaries"), on average, users are unlikely to base their disclosure decisions on the country or traffic size. This is surprising, as surveys show that EU consumers feel more confident purchasing from their own country or the EU [19, 45, 71], and prior research has shown that reputation and perceived size matter for website trust [12, 46].

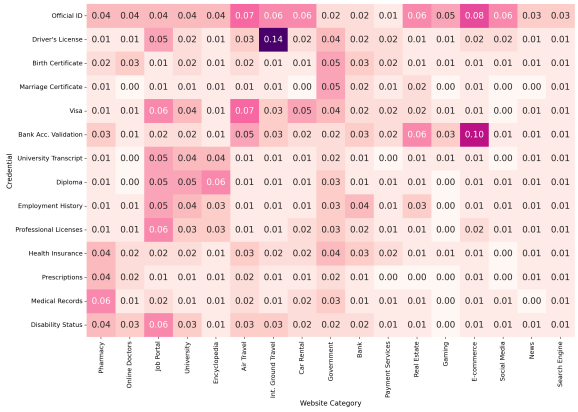

Figure 6: User survey results showing the standard deviation of the percentage of participants who would disclose a credential between different websites within a category.

In-category differences Fig. 6 presents the results of the user survey showing the standard deviation of the percentage of users who would disclose a credential between websites in a category. This shows whether users' disclosure decisions differ for each website, which would necessitate more sophisticated support. The results show that the standard deviation is low for most website category-credential pairs, indicating that most disclosure decisions are based on the website category. However, there are two notable exceptions. 1) The standard deviation for driver's license-international ground travel is high (0.14). This comes from BlaBlaCar, which has a bus and a ride-sharing service. Although we asked participants about the bus services (clearly stated in the survey), many participants may have thought about the ride-sharing service, for which a driver's license is needed to sign up as a driver. 2) The standard deviation for official ID-e-commerce is high (0.8). This is likely due to some e-commerce websites selling age-restricted items (for which an official ID may be needed for age verification) while others do not. Thus, a Credential Assistant that displays information based on the website category is likely sufficient; however, multi-service websites may need to be handled separately, and a tailored website categorization for the EUDI context may be needed.

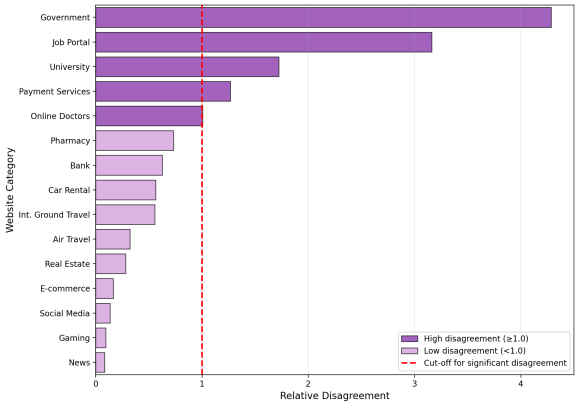

Figure 7: Expert survey results showing the disagreement among experts for each website category. Disagreement was computed by normalizing the variance of the number of credentials experts would disclose to websites in each category.

Viability of data collection and display As our previous results have shown, the website category is likely a good proxy to display information in the Credential Assistant, and thus, we discuss whether expert opinions can be collected and displayed for each website category. Websites can be categorized into a limited number of categories, and services that categorize websites already exist (e.g., Cloudflare Radar [17]). As discussed, the EUDI may require a slightly different website categorization than is currently available; however, changing the categorization topology is feasible. Thus, it is possible to determine for each website which category it belongs to and to display the correct expert recommendation once collected.

Using the observed expert-to-expert inter-class correlation of $\sim$0.057, the number of experts needed on average across all website categories and credentials for $\pm$5 percentage-point precision at 95% confidence is 24, which aligns with expert interview studies that often use 10-30 experts [6, 13]. It is likely feasible to ask $\sim$24 experts to evaluate each website category, especially since categories can be prioritized based on their importance. However, for some categories, expert disagreement is high, as shown in Fig. 7. For websites with above-average disagreement, it is unlikely that experts will converge on a majority opinion. For such categories, a larger panel of experts (100+, depending on the desired accuracy)

may be needed to get an accurate representation of the percentage of experts who would recommend disclosing a credential.

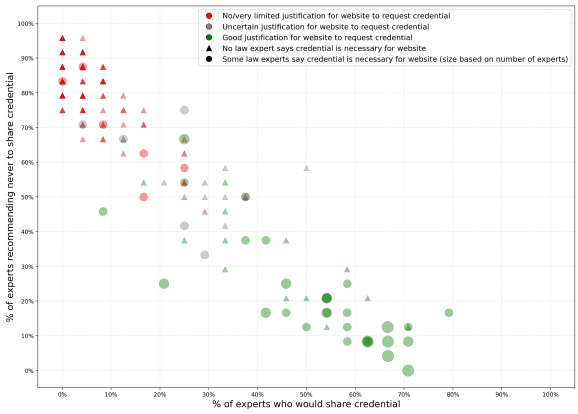

Figure 8: Expert survey results showing how many experts would disclose a credential to a website vs. how many experts recommend never to disclose that credential to that website.

To better understand when expert opinion is aligned and when it is not, we further analyze the expert responses. Fig. 8 presents the results from the expert survey showing the percentage of experts who would disclose compared to the percentage who would recommend never to disclose a credential to a website. Additionally, it shows the number of law experts who state it is necessary to disclose the credential to the website. The credential-website pairs are categorized as justified, unjustified, and uncertain as described in Section 4.1. The figure shows that for the unjustified pairs, expert opinion is largely aligned, with a high percentage of experts recommending never to disclose, and a low percentage of experts willing to disclose a credential. However, for the justified pairs, responses are more mixed, with in many cases, ∼40-60% of experts stating that they would disclose the credential, and mixed results on whether the credential is necessary or not. This indicates that expert recommendations will be effective at warning users against dangerous disclosure decisions, but less effective at helping users decide when disclosure may be beneficial. This is unfortunate, as users may be uncertain in those cases, and prior research has shown that users lean on expert advice when uncertain [11, 32, 42, 61, 69]. However, we note that more research is needed to understand if the low percentages for justified pairs come from the lack of context in our survey or from a general hesitancy of experts to recommend disclosure, for instance, because risks cannot be fully assessed or because use cases are still emerging. More detailed figures are shown in Appendix B.1.

Effectiveness of Credential Assistant Fig. 9a presents the results of the user study that compares the average percentage of users who disclose a credential with a website in the control group (i.e., no Credential Assistant) with the users who saw the correct information in the Credential Assistant. The results show a significant decrease in the users' average mistake rate for the users that received support from the Credential Assistant for 3 out of 4 designs. In those cases, the average mistake rate dropped by ∼8% (e.g., from ∼15% to ∼7% when seeing expert recommendations). We note that this reduction in mistakes occurred even though the Credential Assistant was fully non-intrusive and showed a similar display for all credentials regardless of sensitivity. Thus, even though a residual risk of ∼7% may be unacceptably high, it is likely that more pinpointed and intrusive interventions will lead to a further significant reduction in the number of mistakes. These results are in line with prior research in other contexts, which has shown that users change their behavior based on nudges [2–4, 14, 38, 55, 86]. However, we add to this a specific focus on helping users avoid clear mistakes.

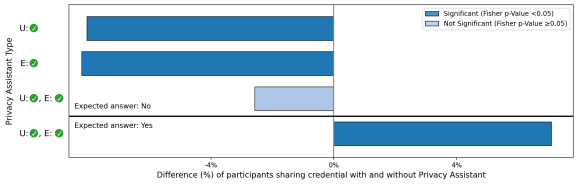

(a) Credential Assistant only shows correct information.

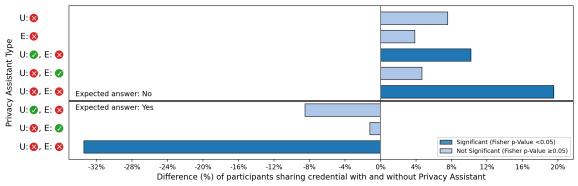

(b) Credential Assistant shows mistakes in the information.

Figure 9: User study results comparing the percentage of users disclosing a credential in the control group with the users with a Credential Assistant. A positive number indicates that more users with a Credential Assistant disclosed the credential. The Credential Assistant displays: U = user opinion, E = expert recommendation, ✓ = correct data, ✗ = mistake.

Potential of misleading by mistakes We expect that it is highly unlikely that mistakes get displayed to users if a tool like the Credential Assistant gets implemented in the EUDI Wallet, since governments and user protection agencies would have strong interests in quality outputs. However, it cannot be fully ruled out that mistakes could occur, and thus, we tested whether mistakes in the Credential Assistant display could lead users to make more mistakes (i.e., herding effects). Fig. 9b presents the results of the user study that compares the average percentage of users who disclose a credential with a website in the control group (i.e., no Credential Assistant) with the users who saw the Credential Assistant displaying mistakes. While the users with the Credential Assistant showing mistakes always made more disclosure mistakes than the

control group, this effect was only significant for 3 out of the 8 Credential Assistant versions, 2 of which showed mistakes in both the user and expert information, which is highly unlikely. Without those cases, the average increase in mistakes is around ∼5% (e.g., from ∼9% to ∼14% when seeing wrong user information and the correct expert recommendation).

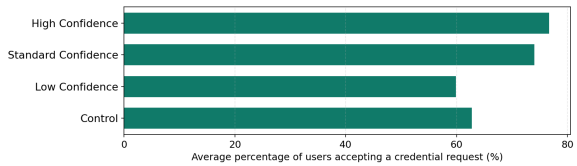

Figure 10: User study result for the average percentage of users who disclosed a credential for different Credential Assistant confidence levels. The control group had no Credential Assistant. Low confidence = 51-55% of users would disclose, standard confidence 81-85% of users would disclose, and high confidence = 91-95% of users would disclose.

Confidence levels Fig. 10 presents the user study results showing how many users, on average, disclosed a credential for different Credential Assistant confidence levels. The results show that the average percentage of users disclosing a credential is significantly higher (∼74% vs. ∼63%) for the users who saw a Credential Assistant with standard confidence (i.e., showing 81-85% of users would disclose) than for the users without a Credential Assistant. Furthermore, the average percentage of users disclosing a credential is similar (∼63% vs. ∼60%) for the users without a Credential Assistant and the users who saw a Credential Assistant with low confidence (i.e., showing 51-55% of users would disclose). Lastly, the average percentage of users disclosing a credential is similar (∼74% vs. ∼77%) for the users who saw a Credential Assistant with standard confidence and the users who saw a Credential Assistant with high confidence (i.e., showing 91-95% of users would disclose). This indicates that the Credential Assistant only has a nudging effect if the confidence is sufficiently high; however, the nudging effect does not increase beyond some confidence level. This is in line with prior research in other contexts that has shown that higher confidence nudges have a greater impact [58, 66]

Demographics Fig. 11 presents the user study results showing how the disclosure decisions of users in the 18-29, 30-49, and 50+ age groups changed based on the Credential Assistant. Our results indicate that users in different age groups prefer different information. For instance, the Credential Assistant showing the disclosure decisions of other users had a stronger nudging effect on the 18-29 age group than the Credential Assistant showing expert recommendations. However, for the 30-49 and 50+ age groups, the effect was opposite. This aligns with prior research showing that dependence on peers

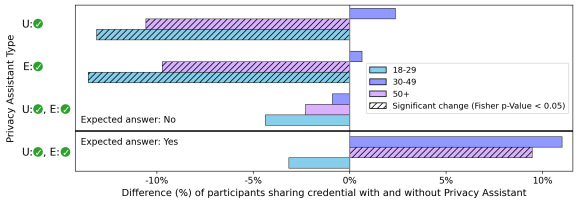

(a) Credential Assistant only shows correct information

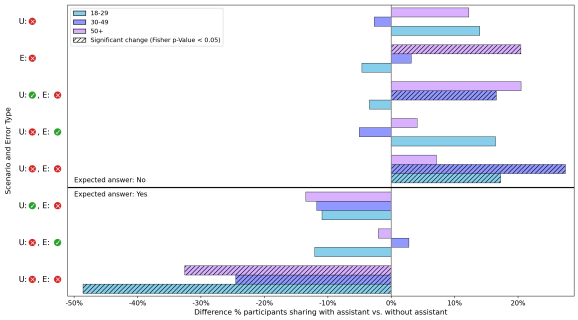

(b) Credential Assistant shows mistakes in the information

Figure 11: User study results comparing the percentage of users disclosing a credential in the control group with the users with a Credential Assistant for the age groups 18-29, 30-49, and 50+. A positive number indicates more users who saw a Credential Assistant disclosed the credential. The Credential Assistant displays: U = user opinion, E = expert recommendation, ✓ = correct data, ✗ = mistake.

decreases and trust in experts increases with age [50, 73, 84]. Thus, tailored nudges based on demographics may be more effective, as has been proposed in some prior research [16, 62].

Purpose Fig. 12 presents the results of the user study showing how often users disclosed a credential to a website for different website-provided purposes, and different Credential Assistant types. Our results show that a vague purpose almost always leads to less disclosure (e.g., ∼71% with purpose vs. ∼74% without purpose for a justified request with a correct Credential Assistant). On the other hand, showing an extensive purpose almost always increases the amount of disclosure, regardless of whether the credential request was justified or not (e.g., ∼22% with purpose vs. ∼16% without purpose for an unjustified request without a Credential Assistant). However, if the Credential Assistant incorrectly recommends not to disclose the credential, even though the request was justified, then the number of users willing to disclose decreases significantly, even with an extensive website-provided purpose (∼69% without Credential Assistant vs. ∼77% with correct Credential Assistant vs. ∼53% with mistaken Credential Assistant). This indicates that even an extensive website provided purpose cannot counteract mistakes in the Credential Assistant, and that allowing websites to add purpose statements in the EUDI Wallet could be risky if those statements

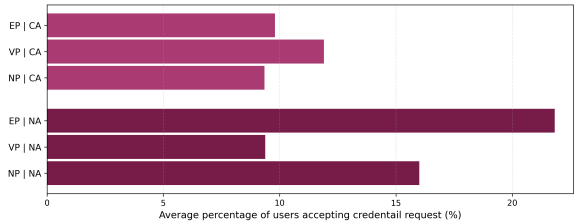

(a) Credential request not supported by the website provided purpose.

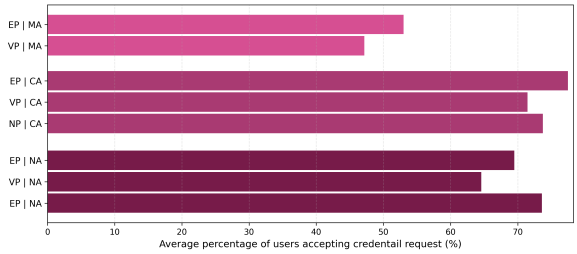

(b) Credential request justified by the website provided purpose.

Figure 12: User study results showing the average percentage of users disclosing credentials to websites for different website-provided purposes and Credential Assistants. NP = no purpose, VP = vague purpose, EP = extensive purpose. NA = no Credential Assistant, CA = Credential Assistant with correct information, MA = Credential Assistant with mistakes.

are not verified. Our research adds a further indication that an extensive purpose can lead to more disclosure, which has some prior work showing more [55, 77] and other prior work showing less [49, 72] disclosure. We add to this, initial results on whether a purpose can counteract bad nudges, which has, to the best of our knowledge, not yet been explored.

6 Limitations

Simulated environment Since the EUDI Wallet and the surrounding ecosystem do not yet exist, it was not possible to test how often users disclosed credentials in real life. Instead, we opted for a simulated environment with fixed scenarios. We see two main ways in which the simulated environment may have affected the participants. First, users may have disclosed more freely in the simulated environment as they were not providing any real data, and thus, there was no risk of their data being misused. Second, participants may have been able to judge the necessity of disclosure more thoroughly in the simulated environment, as they did not see dark patterns and would not be locked out of services. A measurement study after the EUDI Wallet is deployed would be required to verify the extent to which these factors play a role.

Lack of context We did not provide a list of services that the website provides or information about why the website is requesting a credential. This choice was made because we wanted to assess what credentials participants would disclose across all website services, and we did not want to bias participants with our descriptions of the website's services. However, participants may not have thought about some credential use cases, which can explain some of the lower numbers in the results. For example, only a few participants were willing to disclose their official ID to video game websites; however, if we had given users the context that they are buying an age-restricted game, this number would likely be higher.

Selective disclosure & disclosure time The EUDI will allow selective disclosure of attributes (i.e., one can only disclose the name attribute of the identity card credential), and disclosure would generally happen when accessing a service requiring an attribute, instead of when accessing the website (i.e., the above-18 attribute will only be requested when purchasing alcohol and not when visiting the e-commerce website). Including this into our study would have made it too complex, as we already measured a large number of website categories, credentials, and Credential Assistant options. However, this means that more research is needed to understand what attributes users would be willing to disclose to websites, and if those attributes match what the website needs to operate. Additionally, more research is needed to understand for what services users are willing to disclose credentials and when such requests should be presented.

7 Recommendations for the EUDI

Add a Credential Assistant to the EUDI Wallet Our results show that users are likely to make disclosure mistakes with the EUDI Wallet, which could be a high risk to users, and may reduce the trust in the system. Furthermore, our results show that the Credential Assistant can reduce the number of mistakes users make. Thus, we recommend adding a Credential Assistant to the EUDI Wallet. This shows users the information right as they make their disclosure decision, which is when nudging is most effective.

Deploy nudges selectively, make them prominent Prior research has shown that the effect of warnings decreases over time as users become used to those warnings [7, 81, 82]. This is likely to happen for the Credential Assistant if users see nudges for every credential request, which may make them miss the important nudges. Thus, we recommend only showing nudges in clear cases (i.e., when the credential is required, or when disclosure is risky). However, as these are the cases where a user may face negative consequences from a poor disclosure decision, the nudge should be intrusive (e.g., similar to browser warnings) to ensure that users are aware of the possible consequences. Furthermore, it may make sense to time the warnings such that users see them when they are

about to make a poor disclosure decision and to change the warning design based on the sensitivity of the disclosed data.

Adapt Credential Assistant for different user groups Our research indicates that different demographic groups have different preferences for the Credential Assistant (e.g., older users prefer expert recommendations while younger users prefer user opinions). Based on this, we recommend adapting nudges based on user group preferences.

Check verifier registry using Credential Assistant As described in Section 2, the EUDI system will include a mandatory registry for verifiers, in which they must register the EUDI attributes they plan to request. Although the EUDI Wallet and the registrars ensure that websites can only request the attributes that they have registered, this safeguard does not protect against over-claiming: the EUDI framework lacks any systematic, scalable expert review of whether a website genuinely needs the attributes it registers and of the risks such requests pose to users. We believe that a tool like the Credential Assistant could help to partially automate a review of the verifier registry, by comparing the registered attributes with the Credential Assistant information, and flagging registrations that ask for credentials for which the Credential Assistant would recommend not to disclose. Experts could then focus on reviewing the flagged registrations.

Use surveys instead of usage data to assess user opinions Our results show that users are more likely to overshare when they see a credential request, and usage data would only contain cases where a user received a credential request from a website. Furthermore, it is possible for websites to use dark patterns and paid users to increase the number of users who disclose a credential, skewing the data in their favor. Thus, it is better to survey a balanced set of users to identify which credentials users would disclose to websites.

Prioritize transparency for increased public trust Prior research has shown that trust is important for technology adoption [10]. However, open letters [35, 36] and comments in our survey mention concerns about the EUDI (e.g., "I don't like the idea of a digital id system", "I worry about government overreach in this area", "There is no way i will get a digital id so you can be tracked and monitored by the EU"). Thus, it is important to help users understand the EUDI so they can trust it. The same is true for the Credential Assistant. Some participants noted that they did not trust the expert recommendations ("It's impossible to trust those so-called experts because I don't know who they are." "I would not trust the opinions of other users or whatever panel of experts were selected, unless it was some kind of reputable institute."). Thus, it is important that the source of information shown in the Credential Assistant is clearly communicated to users.

8 Related Work

Oversharing Studies have found that users are willing to use options providing less privacy even for low compensation for app [30], website [5, 41, 44], and location [27] data. Additionally, studies have shown that users often click accept on dialogs without understanding them [37, 57]. Furthermore, studies have shown that users often make worse privacy decisions due to misconceptions. For example, users often underestimate negative consequences, do not think about privacy unless a violation occurred, believe their privacy is protected by the presence of privacy policies, or believe that more control over data equals more privacy [2, 3, 15]. Lastly, the privacy paradox shows that users often believe that they are privacy conscious, but act counter to that belief [28, 51, 76].

Nudging Nudging systems give users additional information without forcing them to make a specific decision. Studies have shown that nudging is effective and helps combat privacy misconceptions and biases [2, 43]. Nudging systems have been evaluated and shown to be effective in contexts, such as cookies [39], app permissions [55], file sharing [29], social media posts [85], configurations (e.g., privacy settings) [48, 67], and location sharing [78]. However, studies have also shown nudging risks and challenges. For example, herding effects can lead to worse decisions as users decide solely based on the majority opinion [38]. Lastly, studies have shown that the timing of nudges matters significantly for their effectiveness [31].

9 Conclusions

In this paper, we investigated the disclosure decisions of users regarding digital identity data using the EUDI Wallet. Our results show that users are likely to overshare (e.g., ∼20% of users would disclose their official ID to news websites) and undershare (e.g., only ∼40% of users would disclose their marriage certificate to government websites). We find that oversharing is especially dangerous, as EUDI credentials contain highly sensitive information that, when leaked, can be misused for identity theft. We evaluate a scalable Credential Assistant that displays expert recommendations and/or user opinions and show that it significantly reduces the number of mistakes users make from ∼15% to ∼7%. This result is encouraging, especially as our Credential Assistant was fully non-intrusive. However, at the scale of the EUDI, a residual risk of ∼7% still puts thousands or even millions of users at risk, which is unacceptable. More intrusive interventions (e.g., similar to browser warnings) for sensitive disclosure decisions may be needed. Furthermore, interventions may need to be timed so that users see them when they are about to make a mistake. Future research should investigate whether such stronger interventions can reduce the disclosure risk to acceptable levels.

Ethical Considerations

We received IRB approval for our survey and user study. When writing that application at the start of our research and throughout the project, we thoroughly considered the ethical implications of our work and believe that this work did not cause harm or encourage negative outcomes. We identify the following stakeholders of our research: 1) The eventual users of the EUDI, 2) The lawmakers developing the EUDI, 3) The Prolific participants in our survey and user study, 4) The expert participants in our survey, and 5) The research team.

1) The EUDI is already under development, and it will be deployed in 2026. Users of the EUDI will need to make complex decisions about when to disclose their credentials and when not to. Prior research has shown that making constant decisions about privacy is challenging for users, leading them to make decisions that they regret. The goal of our work is to develop a system that can protect users, which we believe to be inherently positive. A second goal of our system is to inform users when disclosure may be necessary, which can increase the benefit that users gain from the EUDI. Furthermore, the system we evaluate only nudges users and does not force them to make a specific decision, which respects the user's right to make their own decisions. Of course, a nudging system may be used to influence users maliciously by manipulating data. However, we address this concern by analyzing to what extent users can be nudged in a wrong direction and by discussing the importance of trust for the EUDI.

2) A risk from a paper like this would be that it reduces trust in the EUDI, which could harm lawmakers developing the system. We mitigate this risk by providing a fair assessment of a specific aspect of the EUDI and also discussing its benefits. Furthermore, we evaluate a potential solution to mitigate the identified risks. Lastly, we give a set of recommendations for the EUDI that can mitigate the identified risks. As the EUDI has not yet been deployed, lawmakers can address these risks, which could strengthen trust in the system.

3) Prolific participants were informed about the goals and content of the survey/user study through a consent form. No deception was deployed, and thus, all participants gave informed consent. Participants were compensated for their work at a rate of $\sim$\$6.7 for a 25 min survey, which exceeds the minimum wage in any country from which we recruited. Compensation was not based on participants' performance, and participants who chose to abandon the study and withdraw consent were still compensated. All responses we received from Prolific participants were pseudonomized, and we were not able to link responses to specific individuals. Additionally, we will not publish the demographics and free-text responses to ensure that the dataset does not incidentally leak the information of any participant. As such, we believe the identity of all study participants is fully protected, and thus, there is no risk to them from the responses they provided. Lastly, we made sure to tell participants that the EUDI does not exist yet, and thus, all questions are hypothetical. Similarly, we made sure to tell participants that all data requests and all Credential Assistant recommendations were made up. Thus, we believe the risk that participants misunderstood our study as actual recommendations instead of fictional examples is low.

4) Expert participants were informed about the goals and content of the survey through a consent form. No deception was deployed, and thus, all experts gave informed consent. We did not compensate experts; however, we offered to keep them updated on the findings of our research. As we recruited experts from fields that are likely to be interested in the findings, we believe this to be sufficient. The responses received from experts were pseudonomized; however, as we reached out to a limited number of experts directly, the research team may be able to make links to individual experts. We did not record any links we may have made and did not report them in any way. Lastly, we deleted the list of experts that we reached out to and never disclosed this information to anyone.

5) Our research did not deal with any topics that could harm the research team. As all user studies were based on solid ethical foundations with IRB approval, we believe there is no risk to the research team from this work.

Open Science

We published all our artifacts in the repository at https://anonymous.4open.science/r/credential_disclosure_in_digital_wallets-11BA. The repository contains a README.md file describing the content and how to set up the database and run the code. The artifacts are:

1) All pseudonomized responses collected from the survey and user study. The free-text and demographic responses are removed to protect the anonymity of the participants as promised in our IRB application and consent forms. Instead of the individual demographic responses of participants, we publish aggregated demographic data (i.e., the number of participants in each demographic category and the aggregated responses per demographic category).

2) The Python notebooks used to read out data from the databases and to generate the tables and charts in this paper.

3) All material used to generate the survey and user study. This includes the full list of websites with the website information, all credentials with their attributes, all user study scenarios, and all questions asked to the participants.

References

[1] Github copilot. URL: https://docs.github.com/en/copilot/get-started/what-is-github-copilot/.

[2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3):1–41, 2017.

[3] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[4] Alessandro Acquisti, Leslie K John, and George Loewenstein. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2):160–174, 2012.

[5] Alessandro Acquisti, Leslie K John, and George Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274, 2013.

[6] Ayman Alarabiat and Isabel Ramos. The delphi method in information systems research (2004-2017). *Electronic Journal of Business Research Methods*, 17(2):pp86–99, 2019.

[7] Bonnie Brinton Anderson, Jeffrey L Jenkins, Anthony Vance, C Brock Kirwan, and David Eargle. Your memory is working against you: How eye tracking and memory explain habituation to security warnings. *Decision Support Systems*, 92:3–13, 2016.

[8] Arcom. Technical guidelines on age verification for the protection of persons under 18 from online pornography. URL: https://www.arcom.fr/en/find-out-more/legal-area/legal-resources/technical-guidelines-age-verification-protection-persons-under-18-online-pornography.

[9] Cybercrime Atlas. Unmasking cybercrime: Strengthening digital identity verification against deepfakes. URL: https://reports.weforum.org/docs/WEF_Unmasking_Cybercrime_Strengthening_Digital_Identity_Verification_against_Deepfakes_2026.pdf.

[10] Tammy Bahmanziari, J Michael Pearson, and Leon Crosby. Is trust important in technology adoption? a policy capturing approach. *Journal of Computer Information Systems*, 43(4):46–54, 2003.

[11] Phoebe E Bailey, Tarren Leon, Natalie C Ebner, Ahmed A Moustafa, and Gabrielle Weidemann. A meta-analysis of the weight of advice in decision-making. *Current Psychology*, 42(28):24516–24541, 2023.

[12] Yakov Bart, Venkatesh Shankar, Fareena Sultan, and Glen L Urban. Are the drivers and role of online trust the same for all web sites and consumers? a large-scale exploratory empirical study. *Journal of marketing*, 69(4):133–152, 2005.

[13] Ian Belton, George Wright, Aileen Sissons, Fergus Bolger, Megan M Crawford, Iain Hamlin, Courtney Taylor Browne Lūka, and Alexandrina Vasilichi. Delphi with feedback of rationales: How large can a delphi group be such that participants are not overloaded, demotivated, or disengaged? *Technological Forecasting and Social Change*, 170:120897, 2021.

[14] Andrew Besmer, Jason Watson, and Heather Richter Lipford. The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, pages 1–10, 2010.

[15] Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. Misplaced confidences: Privacy and the control paradox. *Social psychological and personality science*, 4(3):340–347, 2013.

[16] Stefano Calboli and Bart Engelen. Ai-enhanced nudging in public policy: why to worry and how to respond. *Mind & Society*, pages 1–19, 2025.

[17] Cloudflare. Cloudflare radar. URL: https://radar.cloudflare.com/.

[18] Cloudflare. Domain categories. URL: https://developers.cloudflare.com/cloudflare-one/traffic-policies/domain-categories/.

[19] European Comission. 2025 consumer conditions scoreboard. URL: https://commission.europa.eu/document/2816337b-4fd1-4db2-a71c-d14a206a5a93_en.

[20] European Commission. Commission releases enhanced second version of the age-verification blueprint. URL: https://digital-strategy.ec.europa.eu/en/news/commission-releases-enhanced-second-version-age-verification-blueprint.

[21] European Commission. A digital id and personal digital wallet for eu citizens, residents and businesses. URL: https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/694487738/EU+Digital+Identity+Wallet+Home.

[22] European Commission. Eudi architecture and reference framework. URL: https://eudi.dev/1.1.0/arf/.

[23] European Commission. European digital identity wallet. URL: https://eudi.dev/latest/.

[24] European Commission. The european digitalidentity regulation. URL: https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/915931811/The+European+Digital+Identity+Regulation.

[25] European Commission. Commission implementing regulation (eu) 2025/848 of 6 may 2025 laying down rules for the application of regulation (eu) no 910/2014 of the european parliament and of the council as regards the registration of wallet-relying parties. *Official Journal of the European Union*, L 2025/848:1–15, 2025. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202500848.

[26] European Commission. Commission implementing regulation (eu) 2025/848 of 6 may 2025 laying down rules of the application of regulation (eu) no 910/2014 of the european parliament and of the council as regards the registration of wallet-relying parties. *Official Journal of the European Union*, L 2025/848:1–15, 2025. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202500848.

[27] Dan Cvrcek, Marek Kumpost, Vashek Matyas, and George Danezis. A study on the value of location privacy. In *Proceedings of the 5th ACM workshop on Privacy in electronic society*, pages 109–118, 2006.

[28] Tobias Dienlin, Philipp K Masur, and Sabine Trepte. A longitudinal analysis of the privacy paradox. *New Media & Society*, 25(5):1043–1064, 2023.

[29] Paul DiGioia and Paul Dourish. Social navigation as a model for usable security. In *Proceedings of the 2005 symposium on Usable privacy and security*, pages 101–108, 2005.

[30] Serge Egelman, Adrienne Porter Felt, and David Wagner. Choice architecture and smartphone privacy: There's a price for that. In *The economics of information security and privacy*, pages 211–236. Springer, 2013.

[31] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything? the effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–328, 2009.

[32] Pardis Emami Naeini, Martin Degeling, Lujo Bauer, Richard Chow, Lorrie Faith Cranor, Mohammad Reza Haghighat, and Heather Patterson. The influence of friends and experts on privacy decision making in iot scenarios. *Proceedings of the ACM on human-computer interaction*, 2(CSCW):1–26, 2018.

[33] ENISA. Beware of the sim swapping fraud! URL: https://www.enisa.europa.eu/news/enisa-news/beware-of-the-sim-swapping-fraud.

[34] ENISA. Countering sim-swapping. URL: https://www.enisa.europa.eu/sites/default/files/publications/ENISA_REPORT-Countering_SIM_Swapping.pdf.

[35] Epicenter.works. Open letter to president and vice presidents of eu member states, 2023. Signed by 39 organizations. URL: https://epicenter.works/fileadmin/import/open_letter_eidas_2023-01_0.pdf.

[36] Epicenter.works. Open letter to swedish presidency and permanent representations of eu member states, 2023. Signed by 24 organizations. URL: https://epicenter.works/fileadmin/import/cso-eidas-open_letter_2023.pdf.

[37] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–14, 2012.

[38] Jeremy Goecks, W Keith Edwards, and Elizabeth D Mynatt. Challenges in supporting end-user privacy and security management with social navigation. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, 2009.

[39] Jeremy Goecks and Elizabeth D Mynatt. Supporting privacy management via community experience and expertise. In *Communities and Technologies 2005: Proceedings of the Second Communities and Technologies Conference, Milano 2005*, pages 397–417. Springer, 2005.

[40] Gov.UK. Online safety act: explainer. URL: https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer.

[41] Il-Horn Hann, Kai-Lung Hui, Tom Lee, and Ivan Png. Online information privacy: Measuring the cost-benefit trade-off. *ICIS 2002 proceedings*, page 1, 2002.

[42] Nigel Harvey and Ilan Fischer. Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational behavior and human decision processes*, 70(2):117–133, 1997.

[43] Almut Herzog and Nahid Shahmehri. User help techniques for usable security. In *Proceedings of the 2007*

symposium on Computer human interaction for the management of information technology, pages 11–es, 2007.

[44] Joshua B Hurwitz. User choice, privacy sensitivity, and acceptance of personal information collection. In *European data protection: Coming of age*, pages 295–312. Springer, 2012.

[45] Consumers International. World consumer rights day 2018 briefing: e-commerce backgrounder. URL: https://www.consumersinternational.org/media/154916/e-commerce-overview-report.pdf.

[46] Sirkka L Jarvenpaa, Noam Tractinsky, and Michael Vitale. Consumer trust in an internet store. *Information technology and management*, 1(1):45–71, 2000.

[47] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A" nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, 2009.

[48] Bart P Knijnenburg and Alfred Kobsa. Helping users with information disclosure decisions: potential for adaptation. In *Proceedings of the 2013 international conference on Intelligent user interfaces*, pages 407–416, 2013.

[49] Bart P Knijnenburg and Alfred Kobsa. Making decisions about privacy: information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 3(3):1–23, 2013.

[50] Lisa J Knoll, Jovita T Leung, Lucy Foulkes, and Sarah-Jayne Blakemore. Age-related differences in social influence on risk perception depend on the direction of influence. *Journal of adolescence*, 60:53–63, 2017.

[51] Spyros Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security*, 64:122–134, 2017.

[52] Kat Krol and Sören Preibusch. Control versus effort in privacy warnings for webforms. In *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*, pages 13–23, 2016.

[53] Kevin Lee, Benjamin Kaiser, Jonathan Mayer, and Arvind Narayanan. An empirical study of wireless carrier authentication for {SIM} swaps. In *Sixteenth symposium on usable privacy and security (soups 2020)*, pages 61–79, 2020.

[54] Changjiang Li, Li Wang, Shouling Ji, Xuhong Zhang, Zhaohan Xi, Shanqing Guo, and Ting Wang. Seeing is living? rethinking the security of facial liveness verification in the deepfake era. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 2673–2690, 2022.

[55] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Twelfth symposium on usable privacy and security (SOUPS 2016)*, pages 27–41, 2016.

[56] Jamie Luguri and Lior Jacob Strahilevitz. Shining a light on dark patterns. *Journal of Legal Analysis*, 13(1):43–109, 2021.

[57] Dominique Machuletz and Rainer Böhme. Multiple purposes, multiple problems: A user study of consent dialogs after gdpr. *arXiv preprint arXiv:1908.10048*, 2019.

[58] Robin Martin, Antonis Gardikiotis, and Miles Hewstone. Levels of consensus and majority and minority influence. *European Journal of Social Psychology*, 32(5):645–665, 2002.

[59] Arunesh Mathur, Gunes Acar, Michael J Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on human-computer interaction*, 3(CSCW):1–32, 2019.

[60] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. What makes a dark pattern... dark? design attributes, normative considerations, and measurement methods. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–18, 2021.

[61] Ethan A Meyers, Martin H Turpin, Michał Białek, Jonathan A Fugelsang, and Derek J Koehler. Inducing feelings of ignorance makes people more receptive to expert (economist) opinion. *Judgment and Decision Making*, 15(6):909–925, 2020.

[62] Stuart Mills. Personalized nudging. *Behavioural Public Policy*, 6(1):150–159, 2022.

[63] AP News. Supreme court upholds texas law aimed at blocking kids from seeing pornography online. URL: https://apnews.com/article/supreme-court-porn-age-verification-texas-12a73197796fe8c4bef0d888259543cf.

[64] Patricia A Norberg, Daniel R Horne, and David A Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1):100–126, 2007.

[65] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark patterns after the gdpr: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–13, 2020.

[66] Graham Overton, Ioannis Evangelidis, and Joachim Vosgerau. People believe if 90% prefer a over b, a must be much better than b. are they wrong? *Journal of Consumer Research*, 52(1):135–156, 2025.

[67] Sameer Patil, Xinru Page, and Alfred Kobsa. With a little help from my friends: can social navigation inform interpersonal privacy preferences? In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 391–394, 2011.

[68] Sören Preibusch, Kat Krol, and Alastair R Beresford. The privacy economics of voluntary over-disclosure in web forms. In *The Economics of Information Security and Privacy*, pages 183–209. Springer, 2013.

[69] Song Qi, Owen Footer, Colin F Camerer, and Dean Mobbs. A collaborator's reputation can bias decisions and anxiety under uncertainty. *Journal of Neuroscience*, 38(9):2262–2269, 2018.

[70] Qualtrics. Improve data quality by using a commitment request instead of attention checks. URL: https://www.qualtrics.com/articles/strategy-research/attention-checks-and-data-quality/.

[71] Christine Riefa. The challenge of protecting eu consumers in global online markets. 2017.

[72] Fuming Shih, Ilaria Liccardi, and Daniel Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 807–816, 2015.

[73] Laurence Steinberg and Kathryn C Monahan. Age differences in resistance to peer influence. *Developmental psychology*, 43(6):1531, 2007.

[74] John Suler. The online disinhibition effect. *Cyberpsychology & behavior*, 7(3):321–326, 2004.

[75] European Data Protection Supervisor. Techdispatch: Digital identity wallets, 2025. URL: https://www.edps.europa.eu/system/files/2025-12/25-12-16_techdispatch-digital-identity-wallet_en.pdf.

[76] Monika Taddicken. The 'privacy paradox'in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of computer-mediated communication*, 19(2):248–273, 2014.

[77] Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 91–100, 2014.

[78] Eran Toch. Crowdsourcing privacy preferences in context-aware applications. *Personal and ubiquitous computing*, 18(1):129–141, 2014.

[79] Van Hong Tran, Aarushi Mehrotra, Ranya Sharma, Marshini Chetty, Nick Feamster, Jens Frankenreiter, and Lior Strahilevitz. Dark patterns in the opt-out process and compliance with the california consumer privacy act (ccpa). In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–25, 2025.

[80] European Union. Regulation (eu) 2024/1183 of the european parliament and of the council of 11 april 2024 amending regulation (eu) no 910/2014 as regards establishing the european digital identity framework. *Official Journal of the European Union*, L 2024/1183:1–56, 2024. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401183.

[81] Anthony Vance, David Eargle, Jeffrey L Jenkins, C Brock Kirwan, and Bonnie Brinton Anderson. The fog of warnings: how non-essential notifications blur with security warnings. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 407–420, 2019.

[82] Anthony Vance, Jeffrey L Jenkins, Bonnie Brinton Anderson, Daniel K Bjornn, and C Brock Kirwan. Tuning out security warnings. *MIS Quarterly*, 42(2):355–380, 2018.

[83] W3C. Verifiable credentials use cases. URL: https://www.w3.org/TR/vc-use-cases/.

[84] Jinping Wang, Maria D Molina, and S Shyam Sundar. When expert recommendation contradicts peer opinion: Relative social influence of valence, group identity and artificial intelligence. *Computers in Human Behavior*, 107:106278, 2020.

[85] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2367–2376, 2014.

[86] Rick Wash and Molly M Cooper. Who provides phishing training? facts, stories, and people like me. In *Proceedings of the 2018 chi conference on human factors in computing systems*, pages 1–12, 2018.

[87] Lilei Zheng, Ying Zhang, and Vrizlynn LL Thing. A survey on image tampering and its detection in real-world photos. *Journal of Visual Communication and Image Representation*, 58:380–399, 2019.

[88] Zheng Zheng, Qian Wang, and Cong Wang. Spoofing attacks and anti-spoofing methods for face authentication over smartphones. *IEEE Communications Magazine*, 61(12):213–219, 2023.

A Ground Truth Details

To test whether users make credential disclosure mistakes, we had to define what a mistake is. However, since disclosure decisions are inherently subjective, there was no absolute ground truth that we could use as a basis. Thus, using our own assessment, we categorized each website-credential pair into justified, unjustified, and uncertain based on whether the website has a valid reason (i.e., a valid use case where having access to the credential is necessary for the website or improves its services) to request the credential. We then compared all justified and unjustified website-credential pairs that we used in the user study with the results of the expert survey to validate our categorization. Table 1 shows the justified pairs, and Table 2 shows the unjustified pairs with the results of the expert survey.

Except for three pairs, our categorization matches the expert survey responses. The pairs where the categorization does not match can be explained as follows: 1) International Ground Travel-Official ID: Currently, the official ID is not collected; however, it is likely that this will change in the future to comply with border control requirements, especially for international travel across the Schengen borders. 2) Social Media-Official ID: Governments are tightening requirements for social media sites to verify users' age and put in place age restrictions, which would make providing an official ID mandatory. 3) Bank-Visa: It is likely that the visa was confused with a residency permit. Banks usually collect a residency permit; however, a visa is not necessary.

B Additional Results

In this section, we provide complementary figures for the results discussed in Section 5 and discuss additional results from our survey and user study.

B.1 Additional Figures

Fig. 13 presents the user survey results showing how often participants stated that they would disclose a credential to a Europe-based compared to a USA- or China-based website. Fig. 14 presents the user study results showing how often participants would disclose a credential to high-traffic (large) websites compared to low-traffic (small) websites. The figures show some significant differences for Europe vs. the USA and only one for Europe vs. China. The significant differences are also mixed between cases where there is more disclosure for Europe-based websites and some that show less disclosure. Thus, it seems like users do not heavily base their decisions on the country in which the website is based.

Fig. 15a presents the expert survey results showing the percentage of cybersecurity, policy, and ethics experts who would disclose a credential to a website. Fig. 15b presents the expert survey results showing the percentage of cybersecurity, policy, and ethics experts who recommend never to disclose a credential to a website. Fig. 15c presents the expert survey results showing the percentage of law experts who state that a credential is necessary for a website. In all figures, the justified scenarios from the user study are shaded in green, and the unjustified scenarios from the user study are shaded in red. Fig. 16 visualizes the percentage of experts who recommend never to disclose a credential and the percentage of experts who would disclose a credential. The x-axis has fixed percentage bins (each bin contains 10 percentage points). All website-credential pairs were categorized as unjustified (red), justified (green), and uncertain (gray), as described in Section 4.1. For each category of the total number of website-credential pairs in the category, the percentage of experts who would disclose a credential in each percentage bin was plotted as a solid line. For each category, of the total number of website-credential pairs in the category, the percentage of experts who would recommend never to disclose a credential in each percentage bin was plotted as a solid line. For example, if the justified category has 44 website-credential pairs, and for 22 of them, the percentage of experts willing to disclose the credential is between 40% and 50%, then there would be a dot on the 40-50 x-axis at the 50% y-axis for the solid green line.

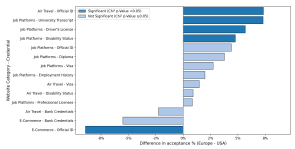

(a) Europe-based websites vs. USA-based websites

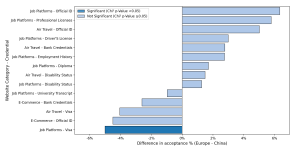

(b) Europe-based websites vs. China-based websites

Figure 13: Results of the user survey showing how many participants would disclose a credential for Europe-based websites vs. China-/USA-based websites. Limited to the website categories where different countries were considered and to the website-credential pairs where at least 10% of participants would disclose the credential for either country. Positive percentages indicate increased disclosure for Europa-based websites.

B.2 Credential Sensitivity & Frequency of Use

Table 3 presents the results of the user study showing for each credential, how many participants own or owned a physical or electronic document that corresponds to this credential. Furthermore, it shows, on average, for the participants who own the document, how frequently they use that document and how comfortable they feel using that document. Lastly, it shows how many participants did not choose the highest comfort response for a document, and for those participants, the average score of when they would still share the document. All scores were instantiated with a text to explain what the

		Expert		
Website Category	**Credential**	**Would Disclose (%)**	**Recommend not Disclose (%)**	**Is Necessary (%)**
Pharmacy	Health Insurance	54%	17%	33%
Pharmacy	Prescriptions	63%	8%	67%
Online Doctors	Prescriptions	67%	13%	100%
Online Doctors	Medical Records	71%	8%	67%
Government	Marriage Certificate	50%	13%	33%
Job Portal	Employment History	71%	13%	33%
University	University Transcript	54%	21%	67%
University	Diploma	54%	17%	33%
Air Travel	Visa	58%	25%	33%
Int. Ground Travel	Official ID	38%	38%	33%
Car Rental	Driver's License	67%	8%	100%
Social Media	Official ID	17%	54%	0%

Table 1: Overview of website-credential pairs and expert evaluations for scenarios categorized as having a good justification for the credential request.

score meant.

B.3 Direct Request for Justified Scenarios

For all user study scenarios in which disclosure is justified, we compared the responses of the control group with the number of users who said that they would disclose the same credential to the same website in our survey. The results are presented in Fig. 17. Compared to the unjustified scenarios, we see a lower percentage of significant increases in oversharing ($\sim$13% for justified vs. $\sim$28% for unjustified). Furthermore, in the justified scenarios, we see a scenario where users disclosed significantly less when seeing a request, which never happened in the unjustified scenarios. This indicates that while increased disclosure when seeing a request occurs both when the request is justified and unjustified, the effect is stronger in the unjustified cases, which lead to oversharing. Thus, requests lead more to oversharing when disclosure is unjustified than to the correct disclosure behavior when disclosing is justified.

B.4 Expert Types

Fig. 18 presents the user study results showing how often participants disclosed their credentials in the group without a Credential Assistant, in the group with a Credential Assistant, showing the recommendation of a generic expert, and in the groups with a Credential Assistant, showing the recommendation of a specific expert type. The results are split into the scenarios: unjustified request with the Credential Assistant only showing the correct expert recommendation, unjustified request with the Credential Assistant showing the correct expert recommendation and user opinion, and justified request with the Credential Assistant showing the correct expert recommendation and user opinion. None of the differences between the responses for different expert types is significant, and there is no clear pattern for how the participants reacted to the expert types in the different scenarios. This indicates that the type of expert in the expert recommendation does not matter to users for EUDI credential disclosure decisions.

B.5 Demographic Disclosure Differences

Figure 19 presents the user survey results showing how often participants in the 18-29 age group stated that they would disclose a credential for a website compared to the 50+ age group. The results indicate that users in the 50+ age group are significantly less likely to disclose credentials.

C Additional Experiment Screenshots

In this section, we show additional screenshots from our survey and user study as described in Section 4.

Fig. 20 shows the credential selector and the credential task questions in the user survey. Fig. 21 shows the instruction page for the website task in the survey.

Fig. 22 shows the instruction pages of the credential task in the user study. The first page describes the EUDI and EUDI credentials, and the second page (not shown to the control group) introduces the Credential Assistant. Fig. 23 shows the screens of the scenario task for different options for Credential Assistant (i.e., no Credential Assistant, single statement, double statement), and with and without a website-provided purpose.

		Expert		
Website	**Credential**	**Would Disclose (%)**	**Recommend not Disclose (%)**	**Is Necessary (%)**
Pharmacy	Diploma	0%	96%	0%
Online Doctors	Visa	0%	92%	0%
Job Portal	Prescriptions	4%	92%	0%
Job Portal	Medical Records	4%	92%	0%
Air Travel	Health Insurance	8%	92%	0%
Int. Ground Travel	Diploma	0%	88%	0%
Int. Ground Travel	Prescriptions	4%	88%	0%
Bank	Visa	25%	58%	0%
Real Estate	Health Insurance	0%	79%	0%
Gaming	Medical Records	0%	83%	0%
E-commerce	Professional Licenses	0%	83%	0%
E-commerce	Health Insurance	0%	83%	0%
E-commerce	Disability Status	4%	83%	0%
News	Official ID	17%	63%	33%
News	Birth Certificate	4%	75%	0%
News	Professional Licenses	4%	71%	0%
Payment Services	Visa	17%	71%	0%
Payment Services	Prescriptions	4%	88%	0%

Table 2: Overview of website-credential pairs and expert evaluations for scenarios categorized as having no or a very limited justification for the credential request.

D Participant Demographics

In this section, we show the demographic breakdown of the participants of our user survey and user study.

Fig. 24 shows the demographic breakdown of the participants of the user survey. Experts are not shown, as they were not asked to provide demographic data to protect their anonymity. Residency options were Germany, France, and Italy, as participants were required to reside in one of those countries.

Fig. 25 shows the demographic breakdown of the participants of the user study.

E Details on Experiment Groups

An overview of the control, baseline, and test groups is shown in Section 4.3. More details are shown in Table 4. The control group participants were not divided into classes and could see any of the control scenarios. The baseline participants were divided into a class that saw expert types and a class that did not. Both classes could see any of the baseline scenarios, but the class with expert types would see the "(E types)" version when an expert recommendation was shown. The test participants were divided into 16 classes, T1.1-T6.3, as shown in the table. Participants in normal frequency classes saw 2 test scenarios, and participants in high frequency classes saw 4 test scenarios. Otherwise, the participants saw baseline scenarios that did not have expert types (i.e., "(E types)" scenarios excluded) and that did not come from the same set. This is because the baseline and test scenarios in a set were the same website-credential pairs with different Credential Assistant versions or different website-provided purposes, and we did not want participants to get confused by seeing different data for the same pair. For the same reason, a participant who saw any version for a credential-website pair would not be shown another version for the same pair. For example, a participant in group T3.3 would see two test scenarios from S3 and 8 baseline scenarios from any of S1, S2, S4, S5, S6. Another example, a participant in the baseline group, who saw the B3.1 version of the first scenario of S3, will not see the B3.2 version of the first scenario of S3 (i.e., they are blocked from seeing the first scenario of S3 again).

Credential	Own Credential # Participants	Avg Frequency (1=low, 5=high)	Avg Comfort (1=high, 6=low)	Not Fully Comfortable # Participants	Avg Still Share (1=lr, 5=mr)
Official ID	1012	2.46	3.58	923	2.92
Driver's License	866	3.42	3.38	730	3.04
Birth Certificate	876	4.42	3.92	789	3.15
Marriage Certificate	262	4.33	3.82	232	3.05
Visa	331	3.44	3.61	291	3.09
Bank Acc. Validation	752	3.44	4.61	736	3.54
University Transcript	628	3.86	2.98	453	2.4
Diploma	720	4.05	2.77	472	2.38
Employment History	494	3.83	3.09	374	2.18
Professional Licenses	214	3.7	3.03	152	2.47
Health Insurance	921	2.43	3.3	771	2.38
Prescriptions	733	2.63	3.45	650	2.42
Medical Records	566	3.82	4.1	533	2.94
Disability Status	51	3.2	3.59	46	2.5

Table 3: User survey results showing how many participants own a credential (physically or electronically), how frequently they use it, how comfortable they are with sharing, and how often they would share if uncomfortable. Under "Not Fully Comfortable" is the number of participants who did not choose 5 for their comfort level. Each number was instantiated with a text explaining what frequency/comfort/still share when uncomfortable at that level meant. lr = least restrictive, mr = most restrictive.

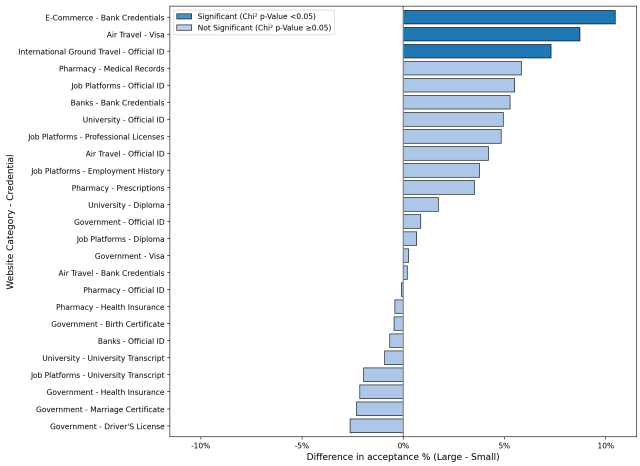

Figure 14: Results of the user survey showing how many participants would disclose a credential for high-traffic (large) websites vs. low-traffic (small) websites. Limited to the website categories where different website sizes were considered and to the website-credential pairs where at least 30% of participants would disclose the credential for either country. Positive percentages indicate increased disclosure for large websites.

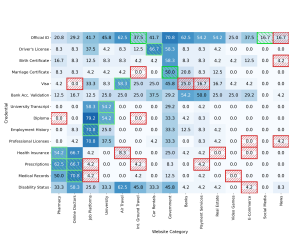

(a) Expert would disclose

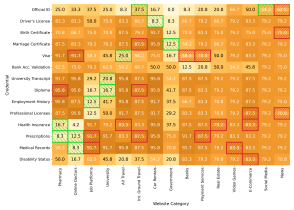

(b) Experts recommend never to disclose

(c) Law experts deem necessary

Figure 15: Results of the expert survey showing how many cybersecurity, policy, and ethics experts recommend never to disclose a credential to a website and how many would disclose the credential to the website. Also shown is how many law experts state that the credential is necessary for the websites. Fields shaded in green are the phase 2 scenarios with a justification and fields shaded in red are the phase 2 scenarios with no justification for the credential request.

Set	# scen.	Correct answer	With purpose	Opinion	Control	Baseline	Test
S1	10	N/A	No	U	C1.1: No CA	B1.1: CA 81-85%	T1.1: CA 51-55% T1.2 (high freq.): CA 51-55% T1.3: CA 91-95%
S2	5	No	Yes	U	C2.1: No purpose, no CA C2.2: Vague purpose, no CA C2.3: Ext. purpose, no CA	B2.1: No purpose, U = no	T2.1: Vague purpose, U = no T2.2: Ext. purpose, U = no
S3	10	No	No	U or E	C3.1: No CA	B3.1: U = no B3.2: E = no B3.3 (E types): E = no	T3.1: U = yes T3.2 (high freq.): U = yes T3.3: E = yes
S4	5	Yes	Yes	U	C4.1: No purpose, no CA C4.2: Vague purpose, no CA C4.3: Ext. purpose, no CA	B4.1: No purpose, U = yes B4.2: Vague purpose, U = yes B4.3: Ext. purpose, U = yes	T4.1: Vague purpose, U = no T4.2: Ext. purpose, U = no
S5	10	No	No	U and E	C5.1: No CA	B5.1: U and E = no B5.2 (E types): U and E = no	T5.1: U and E = yes T5.2: U = yes, E = no T5.3: E = yes, U = no
S6	10	Yes	No	U and E	C6.1: No CA	B6.1: U and E = yes B6.2 (E types): U and E = yes	T6.1: U and E = no T6.2: U = no, E = yes T6.3: E = no, U = yes

Table 4: All user study scenarios in the control, baseline, and test groups. In each set, all scenarios are the same with only the purpose or Credential Assistant changing. CA = Credential Assistant, U = User, E = Expert

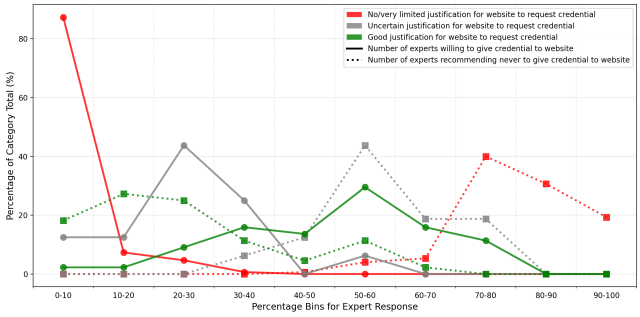

Figure 16: Results of the expert survey showing how many cybersecurity, policy, and ethics experts recommend never to disclose a credential to a website and how many would disclose the credential to the website. Red = unjustified request, green = justified request, gray = uncertain requests, solid line = experts would disclose, dotted line = expert recommend not to disclose. x-axis = percentage bins, y-axis = percentage of total number of website-credential pairs in a category for which the expert result falls in that category.

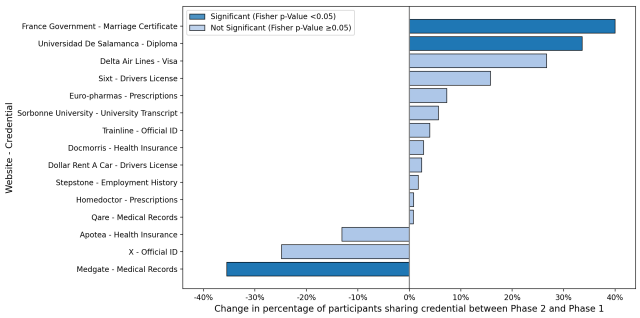

Figure 17: Comparison of how many participants said they would disclose a credential for a website in the user survey, as compared to how many control group participants in the user study disclosed the credential to the website. Positive percentages indicate increased disclosure in the user study. For all website-credential pairs, disclosure is justified.

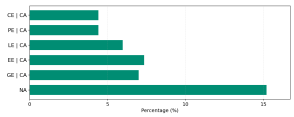

(a) Credential Assistant only shows expert opinion with 5-10% yes. Unjustified credential request.

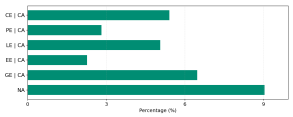

(b) Credential Assistant shows expert and user opinion with 5-10% yes. Unjustified credential request.

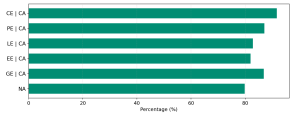

(c) Credential Assistant shows expert and user opinion with 85-90% yes. Justified credential request.

Figure 18: User study result showing the percentage of users who disclosed their credentials in the group with no Credential Assistant and the group with a Credential Assistant displaying the correct information with different expert types.

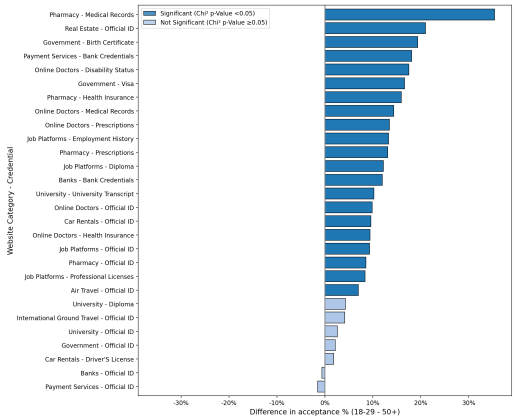

Figure 19: Results of the user survey comparing the percentage of users who stated they would disclose a credential to a website in the 18-29 and the 50+ age groups. Limited to the website-credential pairs where at least 50% of participants in either age group would disclose the credential. Positive percentages indicate increased disclosure in the 18-29 age group.

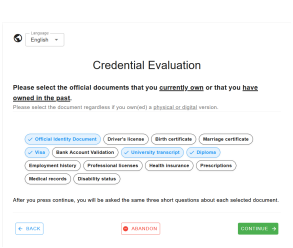

(a) Credential selector.

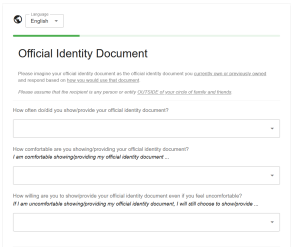

(b) Credential task questions.

Figure 20: User survey screens for the credential task. Participants first choose the credentials they own and then answer the questions for each credential.

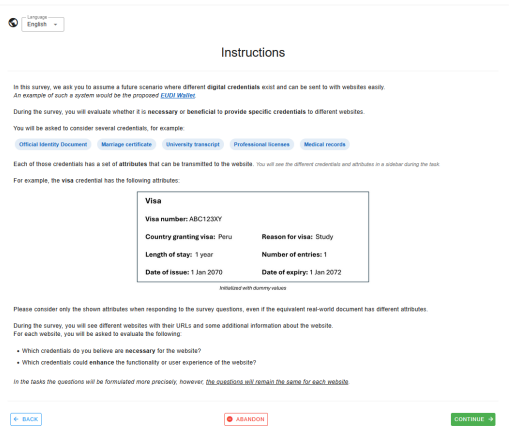

Figure 21: Survey screens showing the instruction page for the website task. The version shown is for users (experts saw different evaluation criteria at the end of the instructions).

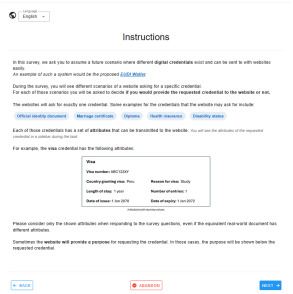

(a) Instruction page describing EUDI.

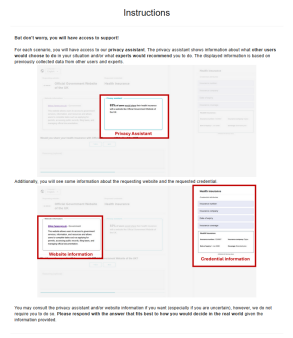

(b) Instruction page introducing Credential Assistant.

Figure 22: User study screens showing the instruction pages for the scenario task.

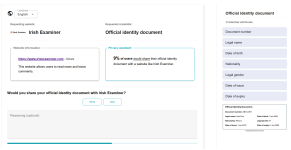

(a) Scenario with Credential Assistant with a single statement.

(b) Scenario with Credential Assistant with a double statement.

(c) Scenario without Credential Assistant.

(d) Scenario with website-provided purpose.

Figure 23: Survey screens for the scenario task with different versions of Credential Assistant and with and without a website-provided purpose. The sidebar always shows all attributes of the requested credential.

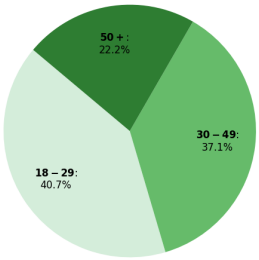

(a) User survey participants age group breakdown.

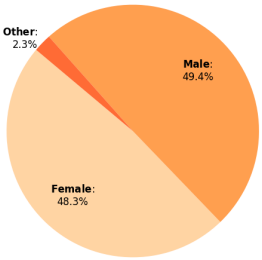

(b) User survey participants gender breakdown.

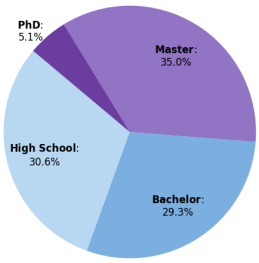

(c) User survey participants highest level of education breakdown.

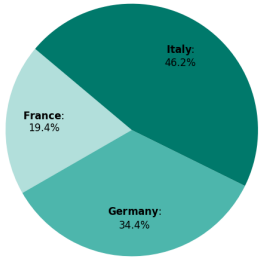

(d) User survey participants place of residency breakdown.

Figure 24: Demographic breakdown of the user survey participants.

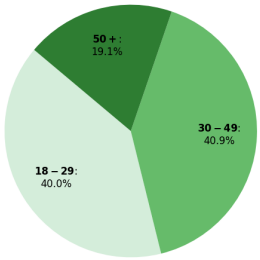

(a) User study participants age group breakdown.

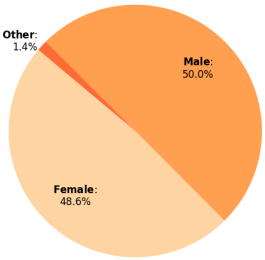

(b) User study participants gender breakdown.

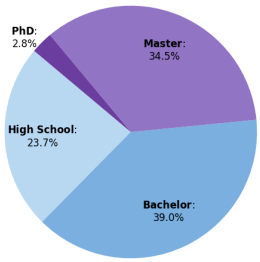

(c) User study participants highest level of education breakdown.

Figure 25: Demographic breakdown of the user study participants.