The Role of Privacy Guarantees in Voluntary Donation of Private Health Data for Altruistic Goals

Ruizhe Wang
University of Waterloo
ruizhe.wang@uwaterloo.ca

Roberta De Viti
MPI-SWS
rdeviti@mpi-sws.org

Aarushi Dubey
University of Washington
aarushid@uw.edu

Elissa M. Redmiles
Georgetown University
elissa.redmiles@georgetown.edu

***Abstract*—While voluntary donation of private health information enables valuable research, privacy concerns often deter potential donors. Privacy Enhancing Technologies (PETs) aim to address these concerns, yet their effectiveness in encouraging data sharing remains unclear. This study conducts a vignette survey ($N = 494$) with participants recruited from Prolific to examine the willingness of US-based people to donate medical data for developing new treatments. It investigates four general guarantees offered across PETs: data expiration, anonymization, purpose restriction, and access control and two mechanisms for verifying these guarantees: self-auditing and expert auditing. This study also controls for the impact of confounds, including demographics and two types of data collectors: for-profit and non-profit institutions.**

Our findings reveal that respondents hold such high expectations of privacy from non-profit entities a priori that explicitly outlining privacy protections has little impact on their overall perceptions. In contrast, offering privacy guarantees elevates respondents' expectations of privacy for for-profit entities, bringing them nearly in line with those for non-profit organizations. Further, while the technical community has suggested audits as a mechanism to increase trust in PET guarantees, we observe limited effect from transparency about such audits. We emphasize the risks associated with these findings and underscore the critical need for future interdisciplinary research efforts to bridge the gap between the technical community's and end-users' perceptions regarding the effectiveness of auditing PETs.

I. Introduction

The altruistic use of personal health data holds immense potential for societal benefit. For example, the public health sector can leverage individuals' medical histories and health data to analyze epidemiological trends, enhance disease surveillance, improve risk prediction models, diagnose rare and emerging diseases, and accelerate the development of novel treatments [1], [2], [3], [4], [5], [6].

Beyond direct medical applications, health data sharing offers economic benefits. A unified data-sharing infrastructure could eliminate redundant clinical trials and research efforts, thereby reducing healthcare costs [7]. According to the European Commission, innovations that enhance health data accessibility could generate savings of billions of euros within the European healthcare sector alone [8]. These demands are further intensified by recent advances in artificial intelligence, particularly machine learning models, whose effectiveness heavily depends on the quality and quantity of data they are trained on.

Recognizing these converging medical, scientific, and economic imperatives, the National Academy of Medicine (NAM) has advocated for more open and collaborative approaches to health data sharing [9], emphasizing both the ethical obligation and practical necessity of broader data access frameworks.

However, despite these compelling benefits and institutional support for expanded health data sharing, privacy concerns remain a significant barrier to widespread data donation. Personal health data is routinely collected by healthcare providers [3], [5], mobile applications, and wearable devices [4], often flowing to third parties through data-sharing agreements [6] with limited transparency to the individuals who generated it. Studies consistently show that privacy concerns are a primary deterrent to individuals contributing their sensitive health information [10], [11], [12], [4], [5], particularly when hospitals permit data sharing with third parties without explicit patient consent [13].

In an effort to foster data collection and analysis while protecting the privacy of data contributors, there has been an increasing focus on deploying *privacy-enhancing* technologies (PETs) for data storage [14], [15], data processing [16], [17], [18], [19], [20], [21], [22], and machine learning (ML) training [23], [24]. Some systems explicitly consider data donation scenarios: Waldo [25] enables privacy-preserving queries over medical time-series data for remote patient monitoring; CoVault [24] supports expressive queries on sensitive data at scale under a strong threat model, and considers a national-scale epidemic analytics scenario; Anonify [22] provides decentralized dual-level anonymity for medical data donation; and Mycelium [26] supports differentially-private queries over large distributed graphs (e.g., disease-spread data) across millions of user devices.

However, despite these technological advances, implementing PETs alone does not guarantee increased data availability. Privacy protections must be expressed to potential donors in comprehensible terms that address their fundamental concerns, or these sophisticated mechanisms remain ineffective at encouraging data sharing.

In order to evaluate whether the guarantees offered by PETs ultimately impact people's willingness to donate their data,

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.230518
www.ndss-symposium.org

at least three steps are necessary. First, we must be able to effectively explain guarantees (either general guarantees or those of a particular PET) to the people they aim to protect: end-users. Prior work has focused on doing so for particular PETs, mainly differential privacy (DP) and end-to-end encryption (E2EE) [27], [28], [29], [30], [31], [32], [33], [34], [35], [36]. Second, using these explanations, we must evaluate whether the guarantees impact people's privacy concerns (or expectations), and in turn, their willingness to share data. Prior work has done such evaluations using self-report studies focused on individual PETs, again primarily DP [27], [37], [36] or E2EE [32], [38], [28]. Third, while rare, we must validate those results in the field to confirm that the effects observed replicate when studying actual versus intended behavior (e.g., [39], [40]).

In this work, we focus on addressing the first two steps outlined above: effectively explaining privacy guarantees and evaluating their impact on data-sharing intentions. We follow real-world medical consent frameworks [41], [42] and privacy regulations [43], [44], and individually investigate four fundamental guarantees across different technologies: anonymization, access control, data expiration, and purpose restriction (PG(1)–PG(4)), along with two verification processes: expert and self auditing (AG(1)–AG(2)). This approach helps identify which specific protections most influence privacy expectations and data-sharing decisions.

We examine the role of these PET guarantees and their audits in shaping privacy expectations and influencing data-sharing intentions within a specific context used in prior work (see e.g., [27], [36]): medical data donation. We ask:

RQ(1): How well do people understand and expect what is offered by the privacy-preserving guarantees PG(1)–PG(4) and auditing guarantees AG(1)–AG(2)?

RQ(2): How does the deployment of privacy guarantees and auditing influence people's willingness to donate their personal health data?

To address these questions, we conduct a vignette survey ($n = 494$) following the best practice methodology in prior work [27], [28], [35]. Each survey respondent is presented with a hypothetical opportunity to donate their health data to help develop a treatment for a specific chronic disease, along with how their data will be protected (PG(1)–PG(4) enforced (or not) by AG(1)–AG(2)). We control for confounding factors identified in prior work, including data-collection entity [6], [45] and socio-demographics [46], [47], [6].

We find that even when told nothing about PETs implemented by the entity, participants are 23% more likely, on average, to expect a non-profit to implement PG(1)–PG(4) and AG(1)–AG(2). As a result of these already high privacy expectations for non-profit organizations, we find that mentioning a specific privacy protection in the survey does not significantly enhance people's willingness to donate towards non-profit entities: even when no privacy protection is explicitly mentioned, 89% of the participants are willing to donate to a non-profit entity. In contrast, for-profit entities need to effectively *demonstrate* their privacy protections; indeed, explicitly mentioning privacy protections in the survey does increase privacy expectations of for-profit entities from 50% to the level of non-profit entities. Privacy expectations, in turn, influence the willingness to donate.

Furthermore, while the technical community has suggested *external audits* as a mechanism to increase trust in PET implementation, our initial inquiry suggests that more work is needed to explain the purpose and effectiveness of such audits to end-users. In fact, the effect of audit statements on people's willingness to donate is limited to a specific scenario involving for-profit entities and auditing to check that purpose restriction (PG(4)) is correctly implemented.

We argue that it is critical for non-profit entities to rigorously implement data privacy measures , as any future data leak could lead to a significant loss of trust. Prior research [40] underscores that users place greater value on *maintaining* their expected privacy than on *gaining* additional privacy they did not initially anticipate. In contrast, we highlight the risk of for-profit entities engaging in "privacy washing" [48], where statements about PETs are used to artificially raise privacy expectations to encourage data collection. At the same time, our findings reveal that respondents are perceptive of the limitations in general PET guarantees, particularly concerning protections against data breaches. This underscores the need for future research to explore how to effectively communicate the potential of stronger, emerging PET guarantees and address the skepticism of end-users.

The full appendix of the paper is available.

II. Related Work

In this section, we examine prior work on people's willingness to donate personal data, privacy concerns, the role of PETs, and educational or explanatory strategies for PETs.

Data donation. The analysis of personal health data is crucial for medical research, particularly in diagnosing emerging or rare diseases and developing new treatments. Indeed, the COVID-19 pandemic has further intensified the demand for personal health data [49]. Despite the importance of this data, the sensitive nature of health data poses significant obstacles to its collection [50]. Prior work indicates that individuals are generally willing to donate their data for altruistic purposes [51], [2], [52], though there is a notable reluctance when it comes to their health data specifically [53], [54]. This reluctance diminishes when the donor or their close family members are directly affected by the disease under study [55], [3], suggesting that non-privacy-related factors influence donation decisions.

Privacy concerns. Concerns about privacy and the misuse of donated data are prominent among potential donors [54]. These concerns include the risk of being identified, discriminated against, and having personal sensitive data misused or leaked [56], [2], [5], [57]. A common source of these concerns is distrust of the receiving entity [56], [58], [2], [59], often due to fears that the entity might share data with unauthorized third parties [60], [53]. This distrust is intensified when participants are unfamiliar with the recipient entity [56], [58], or if it is a governmental or for-profit organization [2], [59], [61],

[45], [62], [63]. Beyond this *intentional* data misuse, there are concerns about data leakage caused by hackers or unintentional mishandling [64]. People are reluctant to interact with entities having any history of data breaches, doubting their ability to safeguard sensitive information [64].

Mental models on PETs. PETs aim to address user privacy concerns by providing technical safeguards for sensitive data. While various PETs are available for health data donation contexts [65], their effectiveness depends largely on users' trust and understanding. Prior work indicates that participants with greater online privacy literacy tend to have more trusting attitudes in PETs [66], while those with limited knowledge often remain skeptical.

Studies have also found that non-experts systematically misunderstand privacy technologies [34], [67], [32]. For instance, many incorrectly conceptualize encryption as merely a form of access control [34] or confuse it with simple data encoding [33]. Additionally, non-expert end users may struggle to understand the consequences of inadequate privacy protections [68], [69], [28], [70], [71], [72]. Lerner et al. [73] identified an even more fundamental barrier: some participants express inherent skepticism about the existence of "true" privacy, illustrating how deep-seated misconceptions can fundamentally undermine the reassurance PETs are designed to provide.

As a result, even when privacy protections are present, people may have risk expectations that are misaligned with reality. We investigate the alignment between stated privacy guarantees and people's expectations for how their data will be protected as part of RQ1.

Explaining PETs. To address these misunderstandings and evaluate the impact of PETs on downstream factors such as privacy concerns (or expectations) and willingness to share data, technologists and researchers have worked to explain PETs to the public to address their unfamiliarity with privacy concepts [74]. Prior work focuses heavily on E2EE and DP [28], [30], [27], [31], [29], [32], [33], [34], [35], [36]. These efforts remain ongoing, as effectively and scalable setting privacy expectations remains a challenge. Methods found effective to explain PETs include visualizations [75], mental models [76], [77], nutrition labels [78], metaphors [79], short statements [80], and privacy games [81], [82]. However, there is a gap in the literature regarding techniques to explain auditing guarantees, despite the importance of auditing in verifying compliance with privacy promises.

Privacy guarantees (PGs) and auditing. Prior research on PETs and their accompanying PGs has primarily focused on evaluating specific technologies (e.g., DP [83], [27] or E2EE [32]) and their combined impact on privacy concerns. A smaller subset of studies has assessed how these technologies affect willingness to share data [83]. Additional research has examined public understanding of specific guarantees such as data retention [84], data anonymization [85], and secondary use permissions [86].

Our approach differs by systematically comparing four core PGs implemented across many PETs: data expiration, data anonymization, use restriction, and access control (see §I). We also separately investigate auditing mechanisms, which serve as verification procedures rather than direct guarantees. In particular, we examine whether offering a given PG influences people's willingness to donate health data to the recipient entity. Furthermore, we focus on the effect of two different auditing processes – expert auditing and self auditing – on enhancing privacy expectations and willingness to share data.

To our knowledge, this is the first study to systematically compare these four core PGs and evaluate how different auditing mechanisms affect user privacy expectations and willingness to donate health data.

III. Methodology

As mentioned in §I, we address research questions RQ1 and RQ2 through a user survey. In this section, we discuss the ethical considerations of our study (§III-A), outline the explored statements (§III-B) and survey design (§III-C), and describe the cognitive interview and pilot study process used to refine and improve participants' comprehension of the presented scenarios (§III-D). We also describe our participant selection strategy (§III-E), analysis procedure (§III-F), and the limitations of our methodology (§III-G).

A. Ethics

Our study was conducted under approval from our university's Ethics Review Board (ERB). The approval covered both the cognitive interviews and the survey. We implemented several measures to ensure ethical treatment of participants:

All survey data was collected anonymously through Qualtrics [87], with IP address collection disabled. Participants were presented with a consent form at the beginning of the survey that detailed the study purpose, data handling procedures, and compensation details. Participants could opt out of answering any demographic questions without affecting their compensation.

For cognitive interviews conducted via Zoom [88], participants were asked to avoid using their real names when joining the meetings to maintain anonymity. Only anonymized transcripts were retained for analysis after removing any personal identifiable information.

All research participants were recruited anonymously through Prolific [89], providing an additional layer of separation between researchers and participants' identities. The collected data was accessible only to the research team, and all data will be deleted after the completion of the research project in accordance with our ERB protocol.

Each participant received 1.20$ upon completing the survey. To determine this figure, we ran a test of 20 participants, who took 5m56s (median) to complete it. Thus, the actual compensation was 12$/hr, which aligns with Prolific recommendations. Interview participants were first recruited through a screening survey that collected demographic information and assessed willingness to participate in interviews. Initially, we compensated participants 1.25$ for completing this 5-minute screener (approximately 13.68$/hr), titled "Sign Up for Paid Interview on Data Donation." However, this above-average compensation led to only 8% interview attendance, as many users were motivated primarily by the screening fee.

To better align incentives and improve participation rates, we adjusted the screener compensation to 0.83$ (approximately 9.96$/hr), which better reflected platform norms for screening surveys. This adjustment increased our interview attendance rate to 21%. Participants who completed the 30-minute interviews were additionally compensated 15$.

B. Donation Scenarios

Many PETs involve complex mechanisms that cannot be adequately explained in brief consent forms without compromising informed consent principles. Additionally, consent forms must satisfy legal and ethical requirements while remaining accessible, as participants typically spend just minutes reviewing them [90], [91], [92].

To address this challenge, we align with real-world consent practices [42], [41] by measuring reactions to the high-level guarantees that PETs provide rather than explaining their technical implementations. This approach allows us to investigate fundamental privacy guarantees that appear consistently across different technologies—guarantees that are both meaningful to users and reflective of actual implementations:

PG(1): Anonymization [93], [94]: data is not linkable to its owner (as defined in [95]);

PG(2): Access control [4], [96], [97], [98]: data is accessible only by authorized people, specified in the data collection agreement;

PG(3): Data expiration [99], [100], [101]: data is discarded or inaccessible after a given expiration time, specified in the data collection agreement;

PG(4): Purpose restriction [24], [102]: data is only used for the stated collection purpose, specified in the data collection agreement;

In practice, implementing all guarantees simultaneously can be costly and challenging for real-world systems [43], [44]. This constraint is reflected in many consent frameworks, which often focus on providing a single guarantee—such as anonymization [42] or access control [41]. Following this practical approach, our study investigates the impact of each guarantee *individually*, rather than presenting them in combination. This targeted analysis enables us to identify *which* specific protections most significantly influence privacy expectations and, consequently, reported data-sharing intentions.

Beyond these privacy guarantees, the technical community has developed methods to verify that PETs function as intended. We also investigate how these auditing mechanisms affect user trust and willingness to donate data, specifically examining two approaches[1]:

AG(1): Expert auditing [103], [104]: engaging an aggregator-selected external advisory board to audit the system to verify that the PET functions as described in the data collection agreement;

AG(2): Self auditing [105], [106]: granting anyone, including donors or external advisors appointed by them, the ability to perform such audits.

We present all statements in plain text format, avoiding metaphors and visual aids due to their mixed effectiveness in privacy communications [107], [108] and to align with standard practice in medical and survey consent forms [41], [42]. Figure 1 and Figure 2 show the exact text presented (e.g., the factor levels).

Justification on selecting the included PETs is in §D.

C. Survey Structure

Our survey design is shown in Figure 3, and the full survey can be found in §A. First, we present to the respondents a *suvery introduction*, which we detail in Figure 4. We present a *donation scenario* using this introduction: we want to assess respondents' willingness to donate their health data to a *recipient entity* developing a new treatment for a specific *disease*. The type of entity (for-profit or non-profit) and the disease (cancer, diabetes, heart failure, high blood pressure, and stroke) are randomly selected; the disease options are taken from a list of common chronic diseases published by the Centers for Medicare and Medicaid Services (CMS) [109].

Then, we present to the respondents either a control (no statement) or an experimental statement. Experimental statements are composed of either one privacy statement, alone, or one privacy statement and one auditing statement. Privacy statements are uniformly selected at random from a pool of four privacy statements; auditing statements are selected at random from two auditing statements. Specifically, one statement from Figure 1 and one from Figure 2 are presented to each respondent.

Our survey questions assess respondents' self-reported:

- `scenario understanding`, used to filter out respondents that report not understanding the scenario (see §III-E): "How would you rate your understanding of the above scenario?" (4 point Likert scale: Fully understand - Not understand).
- `willingness to donate` their health data to the recipient entity: "In this scenario how likely would you be to donate your medical record?" (4 point Likert scale: Very likely - Very unlikely).
- `privacy expectations` regarding the specific privacy guarantees we investigate, measured by their agreement with the statements presented in Figure 5 using the same Likert scale as `willingness to donate`.[2] We assess participants' expectations about all guarantees, regardless of what statement they were presented.

[1] We considered but excluded government auditing as a separate category due to overlap with "expert auditing" and potential political bias (see §II).

[2] This question included an additional attention check with a sixth statement: "If I donate my data, I will meet Albert Einstein." Respondents who did not answer 'Very Unlikely' were removed from the experiment, as detailed in §III-E.

PG(1) Anonymization: The privacy-preserving technology removes any personal identifiable information at the time of data collection, so that the data stored by the recipient entity is not linkable to its data owner (i.e., the data is anonymous).

PG(2) Access control: The privacy-preserving technology restricts data access to the authorized scientists (within the recipient entity) which are working on the treatment for the given disease.

PG(3) Data expiration: The privacy-preserving technology discards the donated data, or makes it inaccessible after a given expiration time, which can be chosen by the data donor in the data collection agreement.

PG(4) Purpose restriction: The privacy-preserving technology ensures that the recipient entity can only use the donated data to develop a treatment for the given disease, and not in the context of any other research they may be working on (e.g., different disease).

Baseline: N/A: (No privacy guarantee is at all mentioned.)

Fig. 1: Privacy statements PG(1)–PG(4). One of the five statements (including the empty baseline statement) is randomly presented in the donation scenario immediately after the survey introduction.

AG(1): Expert auditing: An external advisory board of scientists and software engineers appointed by the recipient entity will regularly verify that the privacy-preserving technology is working as described. The results of this verification will be made public.

AG(2): Self-auditing: Anyone interested, including the respondent and the experts the respondent trusts, will be able to verify that the privacy-preserving technology is working as described. Anyone, including the respondent and the experts the respondent trusts, can make their verification results public.

Baseline: N/A: (No auditing statement is at all mentioned.)

Fig. 2: Auditing statements AG(1)–AG(2). One of the three statements (including the empty baseline statement) is randomly presented in the donation scenario after a non-control `privacy statement`.

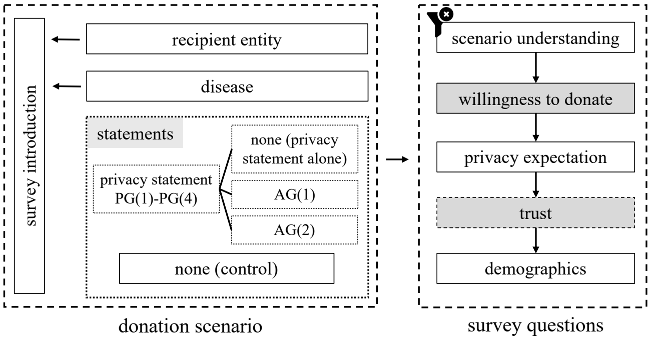

Fig. 3: Survey design. Each participant is presented with a donation scenario and a survey introduction with two variations: recipient entity and disease. `statements` (PG(1)–PG(4) and AG(1)–AG(2)) are also presented in the donation scenario. Participants are then asked five sets of questions (survey questions). When the control statements are presented, the `trust` (dashed) is omitted. The `willingness to donate` and `trust` questions (dark-shaded) are associated with an open-text question. The `scenario understanding` is used as a filter. Participants who did not understand the donation scenario were excluded from the analysis. A sample screen that participants would encounter is provided in Figure 6.

- `trust` that the recipient entity will implement the protection described by the `privacy statement`: "I trust the *entity* will handle my data as described." (4-point Likert scale Strongly agree - Strongly disagree).
- demographics and experiences, including age, gender, education level, donation history, technical background (i.e., education background and job field relationships with computer technologies), and egocentricity(i.e., whether respondents or their close relatives have the disease, as defined in [110], [111]).

We also ask respondents to explain their responses to the `willingness to donate` and `trust` questions. These explanations are collected through open-ended questions: "Why you are willing (or unwilling) to share your medical record with this *entity*?" and "Please explain why you do (or do not) trust that the *entity* will handle your data as described." We note that the `scenario understanding` and `trust` questions are only asked when a non-control privacy statement is presented.

A screenshot of the survey is presented in Figure 6 to illustrate the survey design.

D. Survey Refinement

We refine our survey thourgh a pilot study and multiple iterations of cognitive interviews, following the design of [90]. Inividuals are only allowed to participate in one of the studies.

We use the pilot study to test the survey design and the donation scenario. We found, for example, that no significant relationship between different diseases and donation willingness($p = 0.32$, $\chi^2 = 17.00$) in the pilot survey ($N = 81$ in the control condition). We thus present five common chronic diseases in the final survey (see §A) to ensure participants can relate to the donation scenario and catch positive egocentricity measurement. We list the major takeaways in §C.

We conducted 17 cognitive interviews with participants recruited through Prolific [89], selected from 49 users who completed our screening survey and indicated their willingness to participate in the interview. These interviews evaluated participants' understanding of the survey content and helped

Survey introduction

Imagine that an *entity* wants to develop a new treatment for *disease*. They need medical data from people with and without *disease* to develop the treatment. They ask you to donate your medical record to help develop the treatment. Your medical record contains your: (i) personal information, which may include information about your age, weight, gender, race; (ii) medical history, which may include information about allergies, illnesses, surgeries, immunizations, and results of physical exams and tests; and (iii) medical behavior, which may include information about medicines taken and health habits, such as smoking habits, diet and exercise.

Fig. 4: Survey Introduction. The *entity* type (for-profit or non-profit) and the *disease* (selected from a list of common chronic diseases) are randomly selected.

Anonymization

My full name or other personal identifiable information will be linked to the donated medical record.

Access control

Any employee at the recipient entity will be able to access the donated medical records.

Data expiration

The donated medical record will be deleted at a set point in time.

Purpose restriction

The donated medical records will be used for another purpose without my consent.

Expert auditing

A group of independent experts will verify whether the privacy-preserving technology works and publish a report on their findings.

Self-auditing

I will be able to hire someone to verify that my medical record is protected as described.

Fig. 5: To measure `privacy expectations` respondents reported their agreement with each statement listed above on a 4-point Likert Scale from "Very Likely" to "Very Unlikely".

facilitate consistent comprehension across respondents. Each interview involved presenting the survey to participants and actively seeking feedback on their interpretation of the data donation scenario, and the privacy and auditing statements. We additionally presented all statements to the interviewees and asked for their understanding and perception of each of them.

We continued refining the survey through cognitive interviews until no further constructive feedback was received. In the final interview, we observed generally good understanding of the statements and the donation scenario. For example, the participant described anonymization (PG(1)) as *"eliminating any sort of demographic or personal information affiliated with your data that could be associated with you if the data were to somehow be leaked"* and access control (PG(2)) as *"some sort of database that only certain individuals have access to by using a passcode and certain credentials to to log in."*

E. Survey Sampling

Following the recommendation of [112], we recruited participants through Prolific [89] and collected 560 responses. To ensure focused results, we restricted respondents to adults residing in the U.S. and requested a gender-balanced distribution. Our recruitment included 272 men, 275 women, 9 nonbinary individuals, and 4 participants who chose not to disclose.

Of these 560 respondents, we sequentially excluded the following from the analysis: 35 that submitted incomplete responses; 20 that failed our attention check (§III-C); and 11 that indicated they did not understand the survey scenario (using `scenario understand`). Thus, our final dataset comprises 494 respondents. To enhance statistical power, we limit the levels of age and education to binary groups based on the data distribution among our participants. We report their demographics in Table I and the number of respondents assigned to each condition in Table II.

F. Analysis

We analyzed the open-text questions about `willingness to donate` and `trust` (see §III-C) using inductive-thematic open coding. Two researchers independently coded each entry and generated a codebook from a random sample of at least 100 (20.2%) responses. Then, they composed a final codebook and double-coded all responses. Since all responses were double-coded and inconsistencies were resolved, we do not report inter-rater reliability (IRR) [113]. Ultimately, six responses did not fit this scheme and classified as 'Other'. The codebook of the open-text responses is presented in §B.

In addition to presenting a descriptive analysis detailing the distribution of responses on the survey items that address our research questions, we construct logistic regression models to analyze factors related to two dependent variables: `privacy expectations` (RQ1) and `willingness to donate` (RQ2). For RQ1, the independent variables were the presence of a given privacy and/or auditing statement. The dependent variable was the `privacy expectation` corresponding to the `privacy`

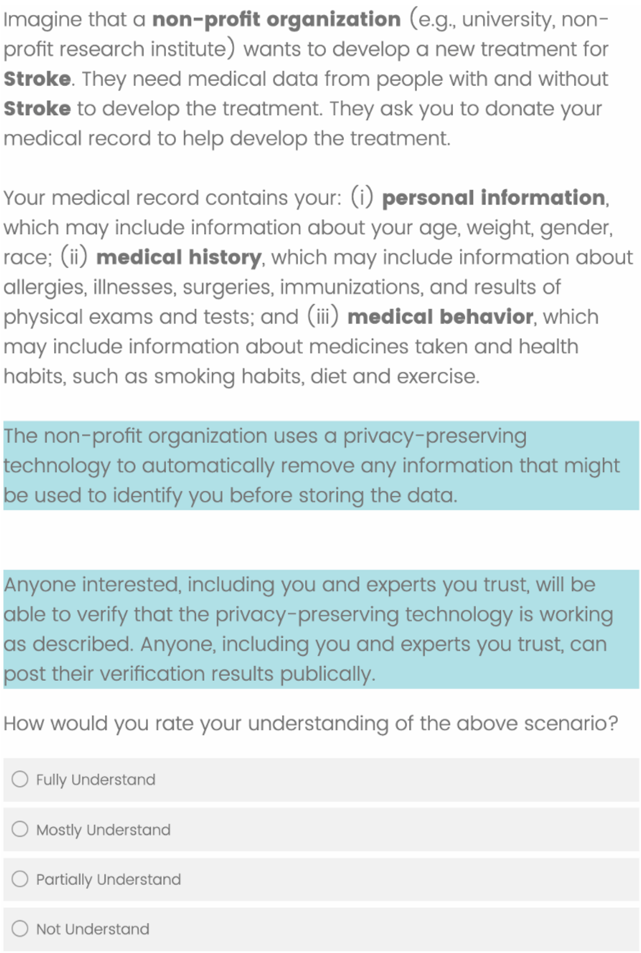

Imagine that a **non-profit organization** (e.g., university, non-profit research institute) wants to develop a new treatment for **Stroke**. They need medical data from people with and without **Stroke** to develop the treatment. They ask you to donate your medical record to help develop the treatment.

Your medical record contains your: (i) **personal information**, which may include information about your age, weight, gender, race; (ii) **medical history**, which may include information about allergies, illnesses, surgeries, immunizations, and results of physical exams and tests; and (iii) **medical behavior**, which may include information about medicines taken and health habits, such as smoking habits, diet and exercise.

The non-profit organization uses a privacy-preserving technology to automatically remove any information that might be used to identify you before storing the data.

Anyone interested, including you and experts you trust, will be able to verify that the privacy-preserving technology is working as described. Anyone, including you and experts you trust, can post their verification results publically.

How would you rate your understanding of the above scenario?

○ Fully Understand

○ Mostly Understand

○ Partially Understand

○ Not Understand

Fig. 6: Screenshot of the survey, which consists of three segments. The first segment presents the *survey introduction*, varying the *disease* and the *entity* across participants. The second segment presents the randomized *privacy statement* and *auditing statement*, highlighted with color coding for emphasis. The third segment poses the survey question that participants are required to answer.

TABLE I: Participant demographics. We note that the second category of each demographical attribute (e.g., no technical background) is considered the baseline scenario during analysis, except for "age", which is treated numerically.

Description	Category	n	%
Age	18 - 29	120	24.3%
	30 - 49	222	44.9%
	50 - 64	111	22.5%
	65+	41	8.3%
Gender	Woman	254	51.42%
	Man	231	46.76%
	Others	9	1.82%
Education	B.S. or above	358	72.47%
	Up to H.S.	136	27.53%
Technical Background	Yes	129	26.11%
	No	365	73.89%
Donation History	Yes	56	11.34%
	No	438	88.66%
Egocentricity	Yes	201	40.69%
	No	293	59.31%

TABLE II: Number of respondents assigned to each condition (i.e., who saw each `privacy statement` and `auditing statement`, or no `privacy statement` (bottom row) or no `auditing statement` (columns 4 and 7).

	For-Profit			*Non-Profit*		
	AG(1)	AG(2)	Ctrl.	AG(1)	AG(2)	Ctrl.
PG(1)	22	18	19	19	20	18
PG(2)	17	21	19	18	18	19
PG(3)	18	20	18	19	19	17
PG(4)	21	18	19	20	18	19
Ctrl.	-	-	19	-	-	21

`statement` presented to a given respondent. Using this model we compare the privacy expectations of respondents in the experimental conditions (those shown a `privacy statement`) with those of the control groups. Responses from respondents in an experimental conditions are only modeled in the analysis of the `privacy statement` they were shown.

For RQ2, the dependent variable was `willingness to donate`, and the independent variables were the presence of a given `privacy statement`, the `privacy expectation` for each guarantee, as well as demographics and experiences. We collapsed the independent variables into a binary measure, designating responses of "Likely" or "Very Likely" as True and all others as False. We did so to avoid the ambiguity introduced by intermediate scale points and to simplify the statistical model.

We categorized education into two groups: with and without a bachelor's degree. We took "no bachelor's degree" and "less than forty years old" as the reference categories, respectively. For gender, we took man as the reference category. Technical background and donation history are binary factors, and we took the negative response as the reference category. We built separate regression models for the two recipient entities: for-profit and non-profit.

G. *Limitations*

Although the four privacy-preserving guarantees examined in this study (§I) cover a broad range of PETs, they are not fully comprehensive and do not capture the full complexity of PETs or real-world threats, which may involve privacy issues beyond the scope of our investigation. Furthermore, our approach may not fully reflect the intricacies of real-world donation scenarios, where multiple guarantees may coexist. While such complexities lie outside the primary focus of our research, qualitative investigations have been conducted in prior work [2], [54], and quantitative exploration of these aspects remains an interesting direction for future research.

Despite our efforts to mitigate misunderstandings—through rigorous cognitive interviews, filtering respondents who reported not fully understanding the scenarios, and controlling for privacy expectations in our statistical analyses—the brief descriptions of each guarantee might still have led to partial or incorrect understanding among respondents. This limitation may have influenced our results.

Additionally, our study relies on self-reported data, which is subject to well-documented biases, known as the privacy paradox, where individuals' expressed privacy concerns often differ from their actual behaviors [114]. To address these limitations, we implemented several methodological safeguards. Firstly, we employed vignette-based scenarios, which research has shown to effectively mirror real-world behaviors and reduce hypothetical bias [115]. Secondly, our work is informed by prior research from 2019 [116] and replicated in 2022 [117] that demonstrated crowdsourced samples, specifically from Prolific [117], well approximate the security and privacy behavior, expectations, and knowledge of the general population for US adults between 18-50 who have at least some college education. While there may still be discrepancies between participants' expressed privacy concerns and their actual behaviors, research confirms that this overambitiousness is systematic [118], [92]. It establishes survey data as a reliable upper-bound approximation of real-world behavior.

Furthermore, statistically significant results do not inherently eliminate the risk of underlying biases that could skew the findings. Various unmeasured factors not accounted for in our analysis might have influenced the interpretations of the results. Lastly, our respondents were recruited via Prolific, which may limit the generalizability of our findings. The sample might not fully represent the diversity of the U.S. population, nor does it capture the perspectives of individuals from other countries, highlighting a key limitation of this work.

IV. Results

A. RQ1: Understanding of PGs

We first address RQ1 (see §I): "How well do people understand and expect privacy guarantees PG(1)–PG(4) and auditing guarantees AG(1)–AG(2)?" We analyze survey responses and their underlying rationales.

1) Analysis Setup: We compare `privacy expectations` on PG(1)—PG(4) and AG(1)–AG(2) among groups of respondents who received different privacy statements (listed in Figure 1) and auditing statements (listed in Figure 2). In Figure 7, we show the `privacy expectations` of respondents who were shown different `privacy statements`. To assess the statistical significance of the descriptive quantitative results presented above, we used logistic regression to analyze the relationship between the presence of a `privacy statement` in the scenario and respondents' `privacy expectations`. We summarize results, which we distinguish for non-profit and for-profit entities, in Figure 8.

2) Response Distribution and Statistical Analysis Results: We find that **privacy expectations differ based on the data-collecting entity.** In the control group, which was not shown any `privacy statement`, 26%-37% of respondents expected for-profit entities to provide anonymization, data expiration, access control, purpose restriction, expert auditing, and self-auditing. In contrast, a higher percentage of respondents (43%-76%) expected non-profit entities to employ these mechanisms, except for self-auditing (AG2). For instance, we note that respondent P430 (who received a for-profit scenario) expected privacy protections even though they were in the control group (and thus received no statement about privacy protection): *"I also assume that the data is looked at in the aggregate and likely no one at the company knows me"*. In the same situation, respondent P53 (who received a non-profit scenario instead) made a stronger assumption: *"I trust that the organization will uphold strict privacy and ethical standards."*

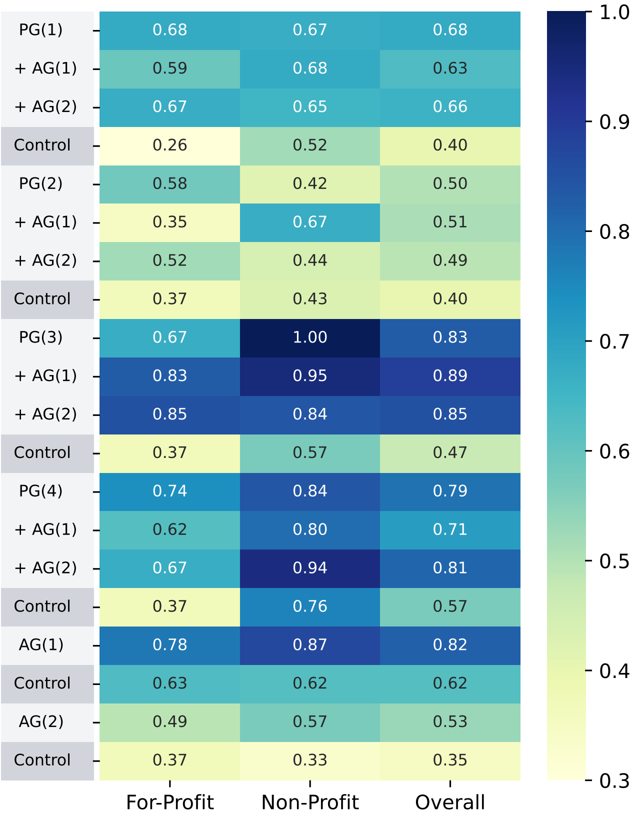

Fig. 7: Percentage of respondents who had a positive `privacy expectation` when shown a particular `privacy statement` in their donation scenario. The values in the table are the percentage of respondents in a given condition who had the `privacy expectation` (see Figure 5) that corresponded to the `privacy statement` they were presented. The 3rd column ("overall") reports results across both entities. For example, the right-most and top-most numerical cell indicates that 68% of participants in both entity scenarios who were shown PG(1) alone – with no `auditing statement` – expected their data to be anonymized. The agreements are binarized here to present the overall tendency.

In the for-profit scenario, all privacy statements, except for the access-control statement (PG2), significantly increase the corresponding privacy expectations (p-value < 0.05). In the non-profit scenario, where expectations are already high, only the privacy statement about data expiration (PG3) significantly increases expectations (from 57% to 100%). Comparing the for-profit and non-profit subsets of the descriptive results (1st and 2nd column of Figure 7), we observe that presenting a `privacy statement` was particularly beneficial in for-profit scenarios. Indeed, the presence of a `privacy statement` raised respondents' `privacy expectations` to nearly the same level as non-profit scenarios for anonymization (PG1, 68% and 67%) and access control (PG2, 58% and 42%). It also

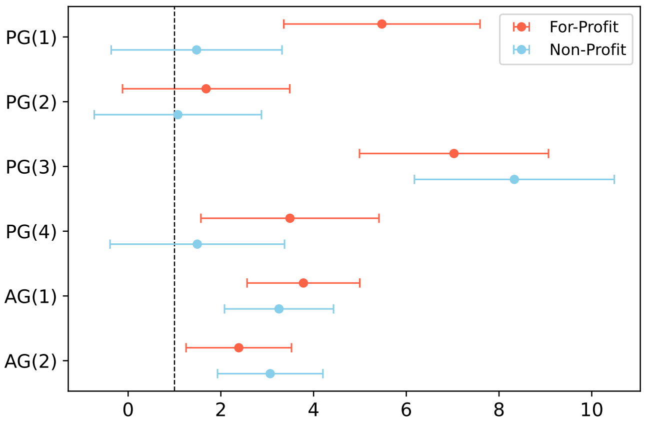

Fig. 8: Odd ratios and 95% confidence intervals of presenting different `privacy statements` on the respondents' `privacy expectations` toward the two kinds of recipient entities, for-profit and non-profit.

reduced the gap in expectations for data expiration (PG3, 67% and 100%)) and purpose restriction (PG4, 74% and 84%)).

Auditing statements significantly increased expectations of corresponding audits in both for-profit and non-profit scenarios (p-value < 0.05). However, we note that **incorporating an audit statement does not significantly change `privacy expectations`**. The statistical significance of the auditing statements in Figure 8 indicates that respondents shown the audit statement correctly expected the corresponding type of audit, while respondents that were not shown the auditing statement did not expect this protection. As observed in Figure 7, presenting an `auditing statement` does not increase the number of respondents who expect the privacy guarantee implied by the `privacy statement` they were shown, despite the stronger assurances that audits imply.

3) Qualitative Analysis of Privacy Expectations via Open-Answer Responses: Across for-profit and non-profit scenarios, 69% and 85% of respondents, on average, expressed that their positive privacy expectations were formed based on the presence of the privacy statement, belief in privacy obligations of the entity, trust in the entity, or trust in the auditing process.

Satisfaction with the privacy guarantees. 16.3% and 20.8% of respondents in the for-profit and non-profit scenarios, respectively, were convinced by the received `privacy statement`. Some respondents, like P56, vaguely stated that the policy looks promising: *"I see their policy, and they have to follow their own policy."* Others felt explicitly safer and trusted the statement. For example, P494 noted: *"Trust is very important when it comes to medical data. I believe the organization has privacy policies that outline that they collect and will use my data. I also believe the organization will employ security measures to safeguard data."* Similarly, P555 highlighted the presence of the `privacy statement`: *"(...) purposely states there is software in place to conserve privacy."* and P516 felt their information was safe because of it: *"(...) with the privacy protection in place, they are isolating the data they need while basically 'throwing out' the rest by putting in under that protection. In essence, my information is safe, and they're only using what they said they'd use."* P208 also felt safer contributing to research knowing the privacy mechanisms were in place: *"I want to be able to contribute to research to better improve cancer treatment, and I feel safe if my data is protected through the mechanisms above."*

Belief in legal or reputational obligations. 16.3% and 10.6% of respondents in the for-profit and non-profit scenarios, respectively, believed that entities are forced to protect the donated data due to legislative requirements or reputational concerns. For example, P28 stated regarding for-profits: *"A for-profit organization wouldn't want to violate HIPAA, HITECH laws."* Additionally, P117 wrote *"I would trust them to do the right thing so they won't face lawsuits."* Regarding non-profits, P141 noted: *"I trust that the law will restrict any data leaks to third parties."*

Trust in the recipient entity. 13.9% and 17.5% of respondents in the for-profit and non-profit scenarios, respectively, expressed a general trust in the recipient entity without specifying reasons. For instance, P70 succinctly said: *"I feel they are reliable and trustworthy"* and P50 stated that *"I assume they take their research seriously, so they would handle the data carefully."* Other respondents, like P145, reported having no reason not to trust it: *"I have no reason to think they would do anything nefarious with my medical data."*

Trust in the auditing process. 7.8% and 7.9% of respondents in the for-profit and non-profit scenarios, respectively, found the auditing statement to add at least some reliability. For example, P202 showed some reservation but found confidence in the public nature of audits: *"I don't fully trust them, but I somewhat do, particularly if audits and verification of results are made public. That said, claiming that the advisory board is external is only partly reassuring, as it's appointed by the institute."* Additionally, P540 had trust in the entity's data handling because *"they let outsiders audit them"* and P51 (who received the expert-auditing statement) because of *"safelocks and checks in place"*. On the other hand, one interviewee doubted the expert auditing process and stated: *"somebody else says something doesn't mean that it's real"*. No survey respondent explicitly raised concern towards the auditing process.

Across for-profit and non-profit scenarios, 31% and 15% of respondents, on average, expressed negative privacy expectations because they were skeptical of the privacy statement or doubted whether the recipient entity would actually employ PETs as stated. Their qualitative responses offer insights on the underlying reasons:

Skepticism on the privacy statement and limits of privacy-preserving guarantees. 21.2% and 10.5% of respondents in the for-profit and non-profit scenarios, respectively, expressed general distrust in the feasibility of the privacy statement. In the non-profit scenario, P34 stated that *"no privacy technology is foolproof"*. Similarly, in the for-profit scenario, P45 wrote: *"I don't trust that the privacy-preserving tech would work."* More precisely, privacy statements were considered too ideal to be fully enforced. Some respondents believed that unauthorized employee access would be unavoidable. For instance, P71 said: *"Why would I trust someone other than my doctor with my medical records? These days especially, I don't trust anyone. There could be a breach or simply people I don't know from a hole in the wall will then have access to all my med records. Insane"!* Some others felt that data breaches and cyberattacks were inevitable. For instance, P83 stated that *"the primary reason for my distrust is due to past news of companies being hacked by people and their data getting leaked"*, while P124 noted that *"corporate data breaches are very common"*, and P80 echoed this sentiment, saying: *"I believe the [intentions] will be good, but data can be hacked"*. We observe that the auditing statement did not substantially instill trust on all accounts — with the exception of expert auditing for access control in non-profit scenario, and data expiration in for-profit scenario. For example, P474 mentioned: *"They can check their privacy technology all they want but when there is a breach it is done and info is stolen. After it fails then they say sorry and offer monitoring but the info is still stolen."*

Doubt on the recipient entity's motivation to employ PETs. 8.6% and 3.3% of participants in the for- and non-profit scenarios, respectively, claimed that these entities are inherently self-serving and lack the motivation to uphold privacy-preserving guarantees or implement such measures at all. Most criticism was directed at for-profit entities. For instance, P55 noted that for-profit entities *"will do what is profitable and not much more than that"*. Similarly, P119 stated *"(...) because it is a for-profit organization. I expect them to cut corners"* and P32 wrote that *"for-profit organization have low standard of morality"*. However, some respondents also expressed concerns about non-profit entities. For example, P190 in the non-profit scenario remarked: *"Medicine has become a business. My data is only useful to them if it helps them make more money. Money comes first before the actual well being of humans."*

B. RQ2: Willingness to donate health data

Next, we answer the second research question (RQ2, see §I): "How does the deployment of privacy guarantees and auditing influence people's willingness to donate their personal health data?" As in the previous subsection, we statistically analyze the responses and then report the qualitative rationales collected from the open-text questions.

1) Analysis Set Up: In Figure 9, we summarize the respondents' `willingness to donate` their health data to the presented recipient entity. The 1st column ("overall") reports results for both entity types, the 2nd column reports the results of the subset with for-profit entity, and the 3rd column reports the results of the subset with non-profit entity. The values represent the percentage of respondents willing to donate their heath data to the recipient entity.

We constructed a logistic regression model to understand the factors that influence `willingness to donate`. We first analyze and confirm the influence of the recipient entity and then separate the participants based on the recipient entities and analyze them in separate models. The results are summarized in Table III.

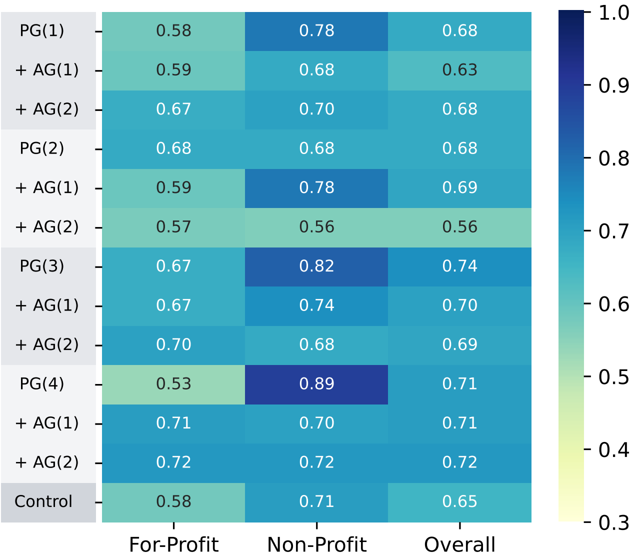

Fig. 9: Percentage of respondents willing to donate their personal health data to the recipient entity in each scenario (e.g., the right-most and top-most numerical cell indicates that 68% of participants in both entity scenarios that with PG(1) presented but no `auditing statement` were willing to donate.) The agreements are binarized here to present the overall tendency.

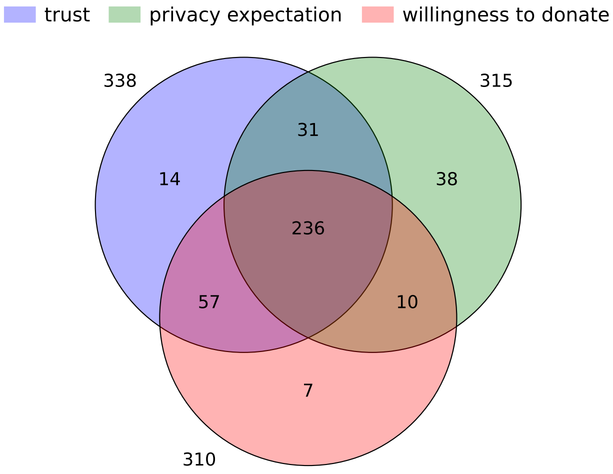

Fig. 10: Overlap in respondent's `willingness to donate`, `privacy expectation` for the statement they were shown, and `trust` that the entity would protect their data as described (respondents in the control – no `privacy statement` – are excluded). The numbers outside each circle summarize the total number of respondents who were e.g., willing to donate their data (310). $N = 454$.

In Figure 10, we visualize the overlap in experimental group respondent's `willingness to donate`, `privacy expectation` for the statement they were shown, and `trust` that the entity would protect their data as described.

TABLE III: Influences of `privacy expectations` and `demographics` on `willingness to donate` of the three participant groups. OR: odd ratio. CI: confidence interval.

Entity	Factor Levels	OR	95% CI	*p*-value
Both	Entity - FP	0.38	[0.24, 0.59]	**<0.001**
For-Profit	*Privacy Statement*			
	PG(1)	0.65	[0.25, 1.70]	0.381
	PG(2)	0.68	[0.27, 1.75]	0.429
	PG(3)	0.86	[0.32, 2.28]	0.756
	PG(4)	0.62	[0.22, 1.70]	0.352
	Privacy Expectation			
	PG(1)	1.09	[0.66, 1.80]	0.741
	PG(2)	1.60	[0.96, 2.69]	0.073
	PG(3)	1.70	[1.01, 2.88]	**0.047**
	PG(4)	3.25	[1.78, 5.87]	**<0.001**
	AG(1)	2.46	[1.40, 4.27]	**0.002**
	AG(2)	1.32	[0.76, 2.27]	0.317
	Demographics & Experiences			
	Education	0.73	[0.43, 1.26]	0.264
	Age	1.00	[0.98, 1.01]	0.765
	Gender	1.30	[0.80, 2.12]	0.304
	Tech Background	1.57	[0.88, 2.80]	0.127
	Egocentricity	0.69	[0.42, 1.14]	0.144
	Donation History	3.25	[1.31, 8.01]	**0.011**
Non-Profit	*Privacy Statement*			
	PG(1)	1.27	[0.52, 3.13]	0.603
	PG(2)	0.94	[0.37, 2.38]	0.901
	PG(3)	1.17	[0.45, 3.00]	0.746
	PG(4)	1.54	[0.61, 3.90]	0.369
	Privacy Expectation			
	PG(1)	1.72	[1.02, 2.89]	**0.043**
	PG(2)	1.46	[0.87, 2.46]	0.152
	PG(3)	1.27	[0.73, 2.25]	0.400
	PG(4)	5.21	[2.61, 10.38]	**<0.001**
	AG(1)	2.59	[1.35, 4.90]	**0.004**
	AG(2)	1.17	[0.68, 2.03]	0.563
	Demographics & Experiences			
	Education	0.46	[0.26, 0.80]	**0.007**
	Age	1.00	[0.98, 1.02]	0.897
	Gender	1.16	[0.69, 1.97]	0.560
	Tech Background	0.71	[0.40, 1.26]	0.237
	Egocentricity	1.36	[0.83, 2.23]	0.217
	Donation History	2.25	[1.06, 4.75]	**0.033**

2) Response Distribution and Statistical Analysis Results: There was no direct relationship between showing a `privacy statement` and respondents' `willingness to donate`. However, several `privacy expectations` positively correlate with `willingness to donate` for both non-profit and for-profit entities, except in specific cases like access control for both for- and non-profit and data expiration for the non-profit.Visualized in Figure 10, we observe high alignment among the three explored constructs, where 236 (52.0%) respondents gave all positive responses and 63 (13.9%) respondents gave all negative responses. Combined with the results of RQ(1), the findings suggest that **`willingness to donate` is influenced by `privacy expectations`, which are, in turn, impacted by `privacy statements`**.

Privacy statements and expectations influence donation intentions, but they are not the only factors. While participants were generally less willing to donate to for-profit entities (OR=0.38, *p*-value < 0.001), other factors are also at play. Looking at control group participants (who received no privacy statements), we found that the difference in `willingness to donate` between for-profit and non-profit entities was only 8.8%, despite a much larger 20.3% difference in `privacy expectations` between these entities. This discrepancy suggests that **non-privacy-related factors, such as perceived donation benefits, also influence participants' `willingness to donate`**.

Furthermore, we observe that **respondents with bachelor's degrees are less willing to donate to a non-profit entity (OR = 0.46, *p*-value =** 0.007**)**, which may be related to their greater awareness of the complexities and potential vulnerabilities with `privacy statement`. They have heightened privacy concerns that need to be addressed. For example, P42 wanted to verify that the PET worked visually; P97 thought that only the stored data was encrypted, but the data transaction was not protected the same way; and P236 argued that: *"privacy-preserving technology that works today may not work in a few years"*.

Finally, we found that **prior donors are more willing to donate to both entities.** Respondents who have donated before were more willing to donate to both entities, with the OR value of 3.25 (*p*-value = 0.011) and 2.25 (*p*-value = 0.033) for the for- and non-profit entities, respectively.These respondents expressed strong willingness of supporting science and recognized the importance of medical data in developing new treatments. For example, P40 claimed: *"I feel that it is important information to share in hopes that they can find better ways to deal with diabetes"*. and P208 responded: *"Breakthroughs in science and moving forward in knowledge, our greatly benefited by such Data Collection."*

3) Qualitative Analysis of Donation Intention Via Open-Answer Responses: **Many respondents reasons for donating (or not) included a desire to support research, ego-centric connections, or desire for personal reward.** In the for-profit scenario, 50.2% of respondents' reasons for donating (or not) fell into these categories, in the non-profit scenario, 53.3% did.

One non-privacy-related reason for donating included is helping research. For example, P107 wrote that *"there are a number of things I could share to research such as money, time, data, and more. I would be willing to share my medical records to contribute to cancer research in hopes that people will be healthier and to enhance research future medicines"*. P97 stated: *"even if they are for-profit (which I don't like), any closer we get to developing better treatments and/or a potential cure to cancer is something I'd be willing to assist"*.

Ego-centric reasons, particularly familial connections, also motivated donations. P154 mentioned that the motivation was: *"my dad died when his heart failed on him, and I'd have rather that not had happened"*, and P233 noted that they *"would be willing to help to prevent someone else's mom from dying of a stroke like [theirs] did"*.

However, their existing (negative) experiences with for-profit healthcare entities also deter some of them from making donations. For example, P136 claimed they have heard multiple instances where entities promise not to sell information but circumvent this by using different terminology, effectively still selling the data. P244 also mentioned that their information has been leaked by a clinician in the past.

Some respondents desired personal reward for donating and thus were not interested in donating in our scenario: e.g., P92 explained that *"I would give them my data if I got paid for it. I would think twice if they want me to donate the data and make no money from it while the researchers will make money off the of research they compiled."* An alternative explanation could be the perceived imbalance in cost and profit between the donor and the recipient entity, where the donor provides data for free while the recipient entity profits from its use. For instance, participants P495, P463, and P459 all share this perception. However, we note that some individual participants alternatively were motivated to donate due to egocentricity, even to for-profit entities. For example, P45 *"just lost my father to heart failure thirteen days ago. (...) I would be more than willing to donate my medical records if it helps develop medicines to make hearts function better."*

Respondents expressed negative sentiments toward for-profit entities collecting data. The desire for personal reward related closely to lack of donation intent for for-profit organizations. In the for-profit scenario, 11.4% explicitly refused due to negative perceptions of such organizations. For instance, P191 and P202 respectively stated: *"I have a negative connotation with for profit organization"* and *"I feel like this type of organization already profits significantly off a large number of people (occasionally off of me as well) and as such, I do not want to give them direct permission to profit further off of me."* Furthermore, respondents felt that for-profits already benefit enough and were unwilling to contribute further. For instance, P36 wrote: *"For profit organization makes profit using the data. I'm never going to donate it. They can BUY it from me for a reasonable amount."*

Alternatively, a small group of 1.3% respondents explicitly mentioned that **they would donate to non-profits but remained concerned about potential data misuse**. For instance, P122 mentions: *'it's always going to be in the back of my head that there's a possibility that it's being sold or taken or used in some way I'm not okay with. Scares me a bit."* Then, P282 notes: *'it is difficult to police everyone in a non-profit organization. I have seen leaders in organizations act unethically at times, which makes me think that even a well-intended organization cannot fully control the actions of every employee or volunteer they have."*

Privacy-related donation considerations focused on data sensitivity and leakage. Concerns for not donating were mostly about data leakage, with respondents expressing that they 12.9% and 13.9% participants in the for-profit and non-profit scenario, respectively, were protective of their data and hence would not donate at all, as P545 noted that *"there are so many data leaks from so called safe places. No one can anticipate what hackers can do in the future."* Others cited previous data leak experiences as to why they would not donate, as P290 said: *"my medical data has been breached in the past by a clinician and I had to file a formal grievance against the health care system that employed her. For this reason, because my medical history was abused, I no longer have any trust and will never voluntarily consent to my medical history/information being shared with others."* Some respondents also expressed that they would only possibly donate to specific entities that have already proven trustworthiness to them, as P357 explained that *"regardless of privacy policy, the probability of a leak or misuse is high. unless it is an organization I have had personal interaction with, or am very familiar with, it is unlikely I would donate my medical records."*

Some participants expressed concern about the level of detail in medical data. P155 highlighted this concern: *"I would not feel comfortable with so much of my personal health information being shared with an organization that is not involved with my direct medical care, regardless of whether there is an advisory board or not"*; as did P229: *"It's a lot of detailed information that I'm worried if it somehow gets in the wrong hands, it could reveal a lot of private information about me."* Contrastingly, some respondents felt comfortable sharing data due to a lack of sensitive information (e.g., *"I don't have mayor illnesses or nothing to hide so Im ok with that"*) or because they already share information with other entities (*e.g., "Considering I already disclose this information to other organizations (for example, data collection on phone, data compiled through search history), other agencies likely have the information on hand already and so another organization having it is no different"*).

V. Concluding Discussion

Health data donation decisions involve multiple complex factors beyond privacy considerations. Non-privacy factors, such as recipient reputation and perceived societal benefit, often influence willingness to donate. Moreover, the effectiveness of PETs depends on people's belief on these protections. PETs' impact on donation willingness diminishes substantially when individuals misunderstand guarantees or doubt implementation integrity.

Prior work has highlighted the ways in which statements on explaining PETs can fail due to being vague or inaccurate. Our work extends this understanding by showing that even when people clearly understand the protection guarantees, the perceived effectiveness of these guarantees is often limited by non-privacy-related factors. Our participants reported strong pre-existing privacy expectations, *even in the absence of explicit privacy statements*. We observe that these expectations vary strongly based on the profit model of the entity. For for-profit entities, where baseline privacy expectations are lower, strong privacy statements can effectively elevate these expectations among users. In contrast, the impact on non-profit organizations is minimal because users already hold high privacy expectations for these entities. This disparity in privacy expectations between non-profit and for-profit entities reflects deeper psychological mechanisms that influence how people form judgments about organizational trustworthiness and privacy practices.

The Halo Effect. Participants' qualitative responses suggest that their strongly positive privacy expectations for non-profit entities arise from basic moral judgments: non-profits are perceived as more ethical and dedicated to serving the public good because they are not pursuing profit. This findings are consistent with a phenomenon termed "the halo effect", which describes how an overall positive impression of a particular entity (e.g., the good work that a non-profit does) can lead to ungrounded assumptions about how the organization operates [119], [120]. While prior work in the business literature focuses on how the halo effect influences expectations about business practices such as donation management, our work illustrates that the halo effect also influences people's expectations about the way non-profits handle data privacy.

However, these positive expectations do not always align with reality. Privacy safeguards are not consistently enforced or widely practiced in non-profit entities; for example, many health organizations rely on broad consent frameworks rather than adopting stricter purpose-restricted consent measures [121]. This gap between expectations and reality can result in privacy expectation violations, a major source of privacy concerns as explained by the theory of contextual integrity [122]. Indeed, a well-known non-profit mental health helpline shared it's data with a commercial entity, leading to significant public outcry [123]. Future work may seek to explore how to design privacy communications that set appropriate expectations in the presence of such effects, as a prior work has sought to do in other domains (e.g., [124]).

Privacy Washing. Conversely, our participants had far lower initial expectations of the privacy practices of for-profit entities due to their perceived singular focus on profit (a "horn" effect). As a result, information about data protection guarantees significantly raised their privacy expectations, and in turn, their willingness to donate their data.

This means for-profit entities can gain competitive advantage [125] by transparently implementing and communicating PETs to counter their trust deficit. Yet this dynamic creates vulnerability to "privacy washing" [126], [48], where organizations make vague or exaggerated privacy claims to deceptively raise expectations. A recent incident involving GoodRx exemplifies this problem: despite promising never to share personal health information, it shared sensitive data with third-party advertisers [127]. Such deception ultimately erodes trust and discourages future donations when discovered, deepening the very horn effect that motivated privacy washing.

Opt in Versus Opt out. Switching to an "opt-out" model for health data donation might seem appealing to overcome donation hesitancy, particularly for for-profit entities struggling against lower baseline trust. However, such approaches face legal barriers in the U.S. HIPAA [128] sets a high bar for health data sharing, requiring explicit "opt-in" authorization for most research uses.

The Role of Auditing. In theory, auditing could serve as a powerful mechanism to counterbalance both the halo effect and privacy washing by providing objective verification of privacy guarantees, regardless of organizational structure. Auditing should transform abstract guarantees into verifiable technical reality. Yet our findings reveal that auditing has surprisingly limited impact on elevating privacy expectations among participants, with the exception of scenarios involving purpose use restrictions and for-profit entities.

This limited effectiveness arises from participants' fundamental skepticism about technological infallibility., and shows a disconnect between expert and public perspectives on verification mechanisms, consistent with findings in prior work [32], [67], [34]. While security and privacy experts typically regard auditing as foundational to establishing trust and demonstrating compliance, our participants takes it as merely another layer of protection that could ultimately fail. As one respondent succinctly stated, "They can guarantee privacy all they want; things still get hacked." This skepticism aligns with emerging research suggesting that many users have developed a resigned attitude toward privacy, doubting that true privacy protection is achievable in practice [73].

Recommendations for Recipient Entities. Given participants' skepticism about privacy protections, recipient entities seeking health data donations should consider several approaches to address these concerns. First, they should recognize that users' prior experiences with data leakage and exposure to media coverage of breaches create a sense of inevitability about privacy risks. This fatalistic perspective requires more than just technical solutions.

For technical implementations, recipient entities should consider advanced technologies like multi-party computation (MPC) that can minimize breach impacts by distributing sensitive information across multiple entities, ensuring no single point of failure. However, these systems must be engineered carefully, as any failure could damage trust more severely than if simpler methods were compromised.

When communicating privacy guarantees, recipient entities should move beyond explaining basic techniques like E2EE. Our findings show that participants struggle to understand more complex concepts like auditing, despite their effectiveness. Recipient entities should develop communication approaches that better explain these sophisticated cryptographic concepts in accessible ways that demonstrate their benefits.

For-profit entities in particular should recognize their trust disadvantage and focus on transparent implementation and clear communication of robust privacy protections. At the same time, they should avoid "privacy washing" through vague or exaggerated claims, which can erode trust if discovered.

Recommendations for End-Users. Our research highlights several key points that individuals should consider before donating their health data. First, donors should evaluate recipient entities based on the data protections they explicitly state in policy documents rather than their organizational structure. The halo effect observed in our study demonstrates how implicit assumptions about certain types of entities' business practices can lead to unfounded assumptions about privacy practices, particularly for non-profit entities.

Second, individuals should avoid assuming privacy protections exist beyond those explicitly stated by the entity requesting data. While entities are legally bound to follow both their explicit promises and applicable privacy laws, the baseline legal protections may be less comprehensive than what users expect for sensitive health data [129].

Finally, individuals should consider engaging privacy advocates before donating sensitive health data. Such consultation can help identify potential privacy risks that might not be immediately obvious and provide a more balanced perspective on the true privacy implications of donation.

Future Work. While our study provides valuable insights into privacy expectations and willingness to donate health data, several promising directions remain for future research. First, researchers could expand beyond our four privacy guarantees and two auditing mechanisms to incorporate a more comprehensive range of PETs, with a particular focus on understanding how advanced and complex guarantees can be effectively communicated to non-technical audiences and how they shape privacy expectations and willingness to donate. This could include exploring the interplay between different privacy guarantees and how they are perceived by individuals.

Second, future work should examine these privacy expectations across more diverse demographic groups, including individuals with a strong distrust of data-handling practices. Identifying and understanding these individuals is critical in fostering confidence in PETs and practices.

Third, our study focused on formal medical data donation scenarios, yet health data is increasingly collected through passive means via consumer devices and casual in-app permissions. As prior work emphasizes the importance of modality, timing, and context in privacy information design [125], future research could explore how such factors intersect with biases such as the halo effect to influence privacy expectations, trust, and willingness to share sensitive information.

Finally, the increasingly common partnerships between for-profit and non-profit entities in the medical domain warrant dedicated study. Building on prior work [52] that found deep skepticism toward hybrid entities in COVID-19 contact tracing, future work may seek to further interrogate how such partnerships, which are common in the medical domain, further complicate our understanding of people's privacy expectations and their sensitivity to statements of data protection.

Acknowledgments

All authors were supported in full or in part by the Max Planck Society during the course of this research. The fourth author additionally acknowledges support from a Google Research Scholar Award "Aligning Technical Data Privacy Protections with User Concerns" and NSF Award # 2429838.

References

[1] M. J. Bietz, C. S. Bloss, S. Calvert, J. G. Godino, J. Gregory, M. P. Claffey, J. Sheehan, and K. Patrick, "Opportunities and challenges in the use of personal health data for health research," *Journal of the American Medical Informatics Association*, vol. 23, no. e1, Sep. 2015.

[2] M. Aitken, J. de St. Jorre, C. Pagliari, R. Jepson, and S. Cunningham-Burley, "Public responses to the sharing and linkage of health data for research purposes: a systematic review and thematic synthesis of qualitative studies," *BMC Medical Ethics*, vol. 17, no. 1, Nov. 2016.

[3] G. Bartlett, B. Macgibbon, A. Rubinowicz, C. Nease, M. Dawes, and R. Tamblyn, "The importance of relevance: Willingness to share eHealth data for family medicine research," *Frontiers in Public Health*, 2018.

[4] E. Greene, P. Proctor, and D. Kotz, "Secure sharing of mhealth data streams through cryptographically-enforced access control," *Smart Health*, 2019.

[5] K. K. Kim, J. G. Joseph, and L. Ohno-Machado, "Comparison of consumers' views on electronic data sharing for healthcare and research," *Journal of the American Medical Informatics Association*, vol. 22, no. 4, pp. 821–830, Mar. 2015.

[6] M. G. Trinidad, J. Platt, and S. L. R. Kardia, "The public's comfort with sharing health data with third-party commercial companies," *Humanities and Social Sciences Communications*, vol. 7, no. 1, Nov. 2020.

[7] A. Szarfman, J. G. Levine, J. M. Tonning, F. Weichold, J. C. Bloom, J. M. Soreth, M. Geanacopoulos, L. Callahan, M. Spotnitz, Q. Ryan *et al.*, "Recommendations for achieving interoperable and shareable medical data in the usa," *Communications medicine*, p. 86, 2022.

[8] E. Commission, "European data governance act," https://digital-strategy.ec.europa.eu/en/policies/data-governance-act.

[9] D. Whicher, M. Ahmed, S. Siddiqui, I. Adams, C. Grossman, and K. Carman, "Health data sharing to support better outcomes," *Washington, DC: National Academy of Medicine*, 2020.

[10] E. M. Redmiles, "User Concerns & Tradeoffs in Technology-facilitated COVID-19 Response," *ACM Digital Government: Research and Practice*, 2020.

[11] T. H. Voigt, V. Holtz, E. Niemiec, H. C. Howard, A. Middleton, and B. Prainsack, "Willingness to donate genomic and other medical data: results from germany," *European Journal of Human Genetics*, 2020.

[12] H. Silber, F. Gerdon, R. Bach, C. Kern, F. Keusch, and F. Kreuter, "A preregistered vignette experiment on determinants of health data sharing behavior: Willingness to donate sensor data, medical records, and biomarkers," *Politics and the Life Sciences*, vol. 41, no. 2, 2022.

[13] J. Lewis, "Patient Data Sharing: The Public's Opinion," 2019, https://medium.com/swlh/patient-data-sharing-the-publics-opinion-6c385d6d7eda.

[14] R. A. Popa, C. M. Redfield, N. Zeldovich, and H. Balakrishnan, "Cryptdb: protecting confidentiality with encrypted query processing," in *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, 2011, pp. 85–100.

[15] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S. G. Choi, W. George, A. D. Keromytis, and S. M. Bellovin, "Blind Seer: A Scalable Private DBMS," in *2014 IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2014.

[16] A. Arasu, S. Blanas, K. Eguro, R. Kaushik, D. Kossmann, R. Ramamurthy, and R. Venkatesan, "Orthogonal Security With Cipherbase," in *6th Biennial Conference on Innovative Data Systems Research (CIDR'13)*, January 2013.

[17] S. Bajaj and R. Sion, "TrustedDB: A Trusted Hardware Based Database with Privacy and Data Confidentiality," in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '11. New York, NY, USA: Association for Computing Machinery, 2011, pp. 205–216.

[18] A. Baumann, M. Peinado, and G. Hunt, "Shielding Applications from an Untrusted Cloud with Haven," *ACM Trans. Comput. Syst.*, vol. 33, no. 3, Aug. 2015.

[19] P. Gupta, Y. Li, S. Mehrotra, N. Panwar, S. Sharma, and S. Almanee, "Obscure: Information-Theoretic Oblivious and Verifiable Aggregation Queries," *Proceedings of the VLDB Endowment*, vol. 12, no. 9, pp. 1030–1043, May 2019.

[20] A. R. Chowdhury, C. Wang, X. He, A. Machanavajjhala, and S. Jha, "Cryptϵ: Crypto-Assisted Differential Privacy on Untrusted Servers," *CoRR*, vol. abs/1902.07756, 2019.

[21] M. Al-Rubaie and J. M. Chang, "Privacy-preserving machine learning: Threats and solutions," *IEEE Security & Privacy*, 2019.

[22] S. Gaballah, L. Abdullah, M. Alishahi, T. Nguyen, E. Zimmer, M. Mühlhäuser, and K. Marky, “Anonify: Decentralized dual-level anonymity for medical data donation,” 03 2024.

[23] P. Wu, J. Ning, J. Shen, H. Wang, and E. Chang, “Hybrid trust multi-party computation with trusted execution environment.”

[24] R. De Viti, I. Sheff, N. Glaeser, B. Dinis, R. Rodrigues, J. Katz, B. Bhattacharjee, A. Hithnawi, D. Garg, and P. Druschel, “Covault: A secure analytics platform,” *arXiv preprint arXiv:2208.03784*, 2022.

[25] E. Dauterman, M. Rathee, R. A. Popa, and I. Stoica, “Waldo: A private time-series database from function secret sharing,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 2450–2468.

[26] E. Roth, K. Newatia, Y. Ma, K. Zhong, S. Angel, and A. Haeberlen, “Mycelium: Large-scale distributed graph queries with differential privacy,” in *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles*, 2021, pp. 327–343.

[27] R. Cummings, G. Kaptchuk, and E. M. Redmiles, ““I need a better description”: An investigation into user expectations for differential privacy,” in *ACM Conference on Computer and Communications Security (CCS)*, 2021.

[28] R. Abu-Salma, E. M. Redmiles, B. Ur, and M. Wei, “Exploring user mental models of End-to-End encrypted communication tools,” in *8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18)*, 2018.

[29] O. Akgul, R. Abu-Salma, W. Bai, E. M. Redmiles, M. L. Mazurek, and B. Ur, “From secure to military-grade: Exploring the effect of app descriptions on user perceptions of secure messaging,” in *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 2021, pp. 119–135.

[30] W. Bai, “User perceptions of and attitudes toward encrypted communication,” Ph.D. dissertation, University of Maryland, College Park, 2019.

[31] A. De Luca, S. Das, M. Ortlieb, I. Ion, and B. Laurie, “Expert and non-expert attitudes towards (secure) instant messaging,” in *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 2016, pp. 147–157.

[32] S. Dechand, A. Naiakshina, A. Danilova, and M. Smith, “In encryption we don’t trust: The effect of end-to-end encryption to the masses on user perception,” in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 401–415.

[33] R. Abu-Salma, M. A. Sasse, J. Bonneau, A. Danilova, A. Naiakshina, and M. Smith, “Obstacles to the adoption of secure communication tools,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 137–153.

[34] J. Wu and D. Zappala, “When is a tree really a truck? Exploring mental models of encryption,” in *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, 2018, pp. 395–409.

[35] A. Xiong, T. Wang, N. Li, and S. Jha, “Towards effective differential privacy communication for users’ data sharing decision and comprehension,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 392–410.

[36] P. Nanayakkara, M. A. Smart, R. Cummings, G. Kaptchuk, and E. M. Redmiles, “What are the chances? explaining the epsilon parameter in differential privacy,” in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 1613–1630.

[37] B. Bullek, S. Garboski, D. J. Mir, and E. M. Peck, “Towards understanding differential privacy: When do people trust randomized response technique?” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 3833–3837.

[38] W. Bai, M. Pearson, P. G. Kelley, and M. L. Mazurek, “Improving non-experts’ understanding of end-to-end encryption: An exploratory study,” in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 210–219.

[39] S. Dooley, D. Turjeman, J. P. Dickerson, and E. M. Redmiles, “Field evidence of the effects of pro-sociality and transparency on covid-19 app attractiveness,” 2022.

[40] A. Acquisti, L. K. John, and G. Loewenstein, “What is privacy worth?” *The Journal of Legal Studies*, vol. 42, no. 2, pp. 249–274, 2013.

[41] N. L. Health, “Example standardized language elements: Biomedical Research Consent Form,” https://med.nyu.edu/research/office-science-research/clinical-research/sites/default/files/research-study-development-templates/template-standard-language-biomedical.docx.

[42] M. C. of Wisconsin Institutional Review Board, “Information Letter Template,” https://www.mcw.edu/-/media/MCW/Departments/Human-Research-Protection-Program/Researchers/Consent-Form-Templates/Informational-Letter-Template-FINAL.docx.

[43] R. Neame *et al.*, “Effective sharing of health records, maintaining privacy: a practical schema,” *Online J Public Health Inform*, vol. 5, no. 2, p. 217, 2013.

[44] A. S. Kazley, A. N. Simpson, K. N. Simpson, and R. Teufel, “Association of electronic health records with cost savings in a national sample,” *Am J Manag Care*, vol. 20, no. 6, pp. e183–e190, 2014.

[45] R. Hendricks-Sturrup and C. Y. Lu, “An assessment of perspectives and concerns among research participants of childbearing age regarding the health-relatedness of data, online data privacy, and donating data to researchers: Survey study,” *Journal of Medical Internet Research*, vol. 25, p. e41937, Mar. 2023.

[46] E. Hargittai, E. M. Redmiles, J. Vitak, and M. Zimmer, “Americans’ willingness to adopt a covid-19 tracking app,” *First Monday*, vol. 25, no. 11, p. online, 2020.

[47] S. Dooley, D. Turjeman, J. P. Dickerson, and E. M. Redmiles, “Field evidence of the effects of privacy, data transparency, and pro-social appeals on covid-19 app attractiveness,” in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–21.

[48] A. M. Cirucci, “Oversharing the super safe stuff:“privacy-washing” in apple iphone and google pixel commercials,” *First Monday*, 2024.

[49] B. K. Scott, G. T. Miller, S. J. Fonda, R. E. Yeaw, J. C. Gaudaen, H. H. Pavliscsak, M. T. Quinn, and J. C. Pamplin, “Advanced digital health technologies for covid-19 and future emergencies,” *Telemedicine and e-Health*, vol. 26, no. 10, pp. 1226–1233, 2020.

[50] R. A. Fahey and A. Hino, “Covid-19, digital privacy, and the social limits on data-focused public health responses,” *International Journal of Information Management*, vol. 55, p. 102181, 2020.

[51] K. Pilgrim and S. Bohnet-Joschko, “Effectiveness of digital forced-choice nudges for voluntary data donation by health self-trackers in germany: Web-based experiment,” *J Med Internet Res*, vol. 24, no. 2, p. e31363, Feb 2022.

[52] J. S. Seberger and S. Patil, “Post-covid public health surveillance and privacy expectations in the united states: scenario-based interview study,” *JMIR mHealth and uHealth*, vol. 9, no. 10, p. e30871, 2021.

[53] A. Middleton, R. Milne, M. Almarri, S. Anwer, J. Atutornu, E. Baranova, P. Bevan, M. Cerezo, Y. Cong, C. Critchley, J. Fernow, P. Goodhand, Q. Hasan, A. Hibino, G. Houeland, H. Howard, S. Hussain, C. Ingvoldstad, V. Izhevskaya, and K. Morley, “Global public perceptions of genomic data sharing: What shapes the willingness to donate dna and health data?” *American journal of human genetics*, vol. 107, 09 2020.

[54] I. Mori, “Public attitudes to commercial access to health data,” https://www.ipsos.com/sites/default/files/publication/5200-03/sri-wellcome-trust-commercial-access-to-health-data.pdf, 2016.

[55] D. Goodman, C. O. Johnson, D. Bowen, M. Smith, L. Wenzel, and K. Edwards, “De-identified genomic data sharing: the research participant perspective,” *Journal of Community Genetics*, vol. 8, no. 3, pp. 173–181, Apr. 2017.

[56] E. R. Weitzman, S. Kelemen, L. Kaci, and K. D. Mandl, “Willingness to share personal health record data for care improvement and public health: a survey of experienced personal health record users,” *BMC Medical Informatics and Decision Making*, vol. 12, no. 1, May 2012.

[57] J. B. McCormick and M. A. Hopkins, “Exploring public concerns for sharing and governance of personal health information: a focus group study,” *JAMIA Open*, vol. 4, no. 4, Oct. 2021.

[58] N. Howe, E. Giles, D. Newbury-Birch, and E. McColl, “Systematic review of participants’ attitudes towards data sharing: a thematic synthesis,” *Journal of Health Services Research & Policy*, vol. 23, no. 2, pp. 123–133, Apr. 2018.

[59] J. Stockdale, J. Cassell, and E. Ford, ““giving something back”: A systematic review and ethical enquiry into public views on the use of patient data for research in the united kingdom and the republic of ireland,” *Wellcome Open Research*, vol. 3, p. 6, Jan. 2019.

[60] B. Kacsmar, K. Tilbury, M. Mazmudar, and F. Kerschbaum, “Caring about sharing: User perceptions of multiparty data sharing,” in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 899–916.

[61] N. Garrison, N. Sathe, A. Antommaria, I. Holm, S. Sanderson, M. Smith, M. McPheeters, and E. Clayton, “A systematic literature review of individuals’ perspectives on broad consent and data sharing in the united states,” *Genetics in Medicine*, vol. 18, no. 7, Jul. 2016.

[62] P. Khanra, “Bridging health data donation,” 2023.

[63] American Medical Association, “Patient perspectives around data privacy,” 2022. [Online]. Available: https://www.ama-assn.org/system/files/ama-patient-data-privacy-survey-results.pdf

[64] A. H. Seh, M. Zarour, M. Alenezi, A. K. Sarkar, A. Agrawal, R. Kumar, and R. A. Khan, “Healthcare data breaches: Insights and implications,” *Healthcare (Basel)*, vol. 8, no. 2, p. 133, May 2020.

[65] S. Jordan, C. Fontaine, and R. Hendricks-Sturrup, “Selecting privacy-enhancing technologies for managing health data use,” *Frontiers in Public Health*, vol. 10, p. 814163, 2022.

[66] D. Harborth and S. Pape, “How privacy concerns, trust and risk beliefs, and privacy literacy influence users’ intentions to use privacy-enhancing technologies: The case of tor,” *ACM SIGMIS Database: The DATABASE for Advances in Information Systems*, vol. 51, no. 1, pp. 51–69, 2020.

[67] M. Oates, Y. Ahmadullah, A. Marsh, C. Swoopes, S. Zhang, R. Balebako, and L. F. Cranor, “Turtles, locks, and bathrooms: Understanding mental models of privacy through illustration,” *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 4, pp. 5–32, Aug. 2018.

[68] J. Tang, H. Shoemaker, A. Lerner, and E. Birrell, “Defining privacy: How users interpret technical terms in privacy policies,” *Proceedings on Privacy Enhancing Technologies*, 2021.

[69] L. Velykoivanenko, K. S. Niksirat, N. Zufferey, M. Humbert, K. Huguenin, and M. Cherubini, “Are those steps worth your privacy? fitness-tracker users’ perceptions of privacy and utility,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 4, dec 2022.

[70] N. Gerber, V. Zimmermann, and M. Volkamer, “Why johnny fails to protect his privacy,” in *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2019, pp. 109–118.

[71] A. Pham and C. Castro, “The moral limits of the market: the case of consumer scoring data,” *Ethics and Information Technology*, vol. 21, pp. 117–126, 2019.

[72] K. Vaniea, E. Rader, and R. Wash, “Mental models of software updates,” *International Communication Association*, pp. 1–39, 2014.

[73] B. Lerner, D. Passey, N. Sperber, and S. Knight, “OP043: The evolving attitude towards privacy and security of personal genomic data,” *Genetics in Medicine*, vol. 24, no. 3, p. S369, Mar. 2022.

[74] L. F. Cranor, P. Guduru, and M. Arjula, “User interfaces for privacy agents,” *ACM Trans. Comput.-Hum. Interact.*, vol. 13, no. 2, June 2006.

[75] C. Stransky, D. Wermke, J. Schrader, N. Huaman, Y. Acar, A. L. Fehlhaber, M. Wei, B. Ur, and S. Fahl, “On the limited impact of visualizing encryption: Perceptions of e2e messaging security,” in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021, pp. 437–454.

[76] G. Stewart and D. Lacey, “Death by a thousand facts: Criticising the technocratic approach to information security awareness,” *Information Management & Computer Security*, vol. 20, no. 1, pp. 29–38, 2012.

[77] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, “Expectation and purpose: understanding users’ mental models of mobile app privacy through crowdsourcing,” in *Proceedings of the 2012 ACM conference on ubiquitous computing*, 2012, pp. 501–510.

[78] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor, “Standardizing privacy notices: an online study of the nutrition label approach,” in *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, 2010, pp. 1573–1582.

[79] F. Karegar, A. S. Alaqra, and S. Fischer-Hübner, “Exploring user-suitable metaphors for differentially private data analyses,” in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022.

[80] Z. Papacharissi and J. Fernback, “Online privacy and consumer protection: An analysis of portal privacy statements,” *Journal of Broadcasting & Electronic Media*, vol. 49, no. 3, pp. 259–281, 2005.

[81] H. Tupsamudre, R. Wasnik, S. Biswas, S. Pandit, S. Vaddepalli, A. Shinde, C. J. Gokul, V. Banahatti, and S. Lodha, “Gap: A game for improving awareness about passwords,” in *Serious Games*, S. Göbel, A. Garcia-Agundez, T. Tregel, M. Ma, J. Baalsrud Hauge, M. Oliveira, T. Marsh, and P. Caserman, Eds. Cham: Springer International Publishing, 2018, pp. 66–78.

[82] K. Karl and Y. Tao, “Correcting overconfidence in online privacy: experimenting with an educational game,” *Information, Communication & Society*, vol. 26, no. 5, pp. 990–1007, 2023.

[83] A. C. Valdez and M. Ziefle, “The users’ perspective on the privacy-utility trade-offs in health recommender systems,” *International Journal of Human-Computer Studies*, vol. 121, pp. 108–121, 2019.

[84] A. Murillo, A. Kramm, S. Schnorf, and A. D. Luca, “"if i press delete, it’s gone" - user understanding of online data deletion and expiration,” in *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 329–339.

[85] G. Haddow, A. Bruce, S. Sathanandam, and J. C. Wyatt, “‘nothing is really safe’: a focus group study on the processes of anonymizing and sharing of health data for research purposes,” *Journal of evaluation in clinical practice*, vol. 17, no. 6, pp. 1140–1146, 2011.

[86] J. Juga, J. Juntunen, and T. Koivumäki, “Willingness to share personal health information: impact of attitudes, trust and control,” *Records Management Journal*, vol. 31, no. 1, pp. 48–59, 2021.

[87] Qualtrics, “Qualtrics xm: The leading experience management software,” 2020. [Online]. Available: https://www.qualtrics.com

[88] zoom, “Zoom: One platform to connect,” 2024. [Online]. Available: https://www.zoom.com/

[89] Prolific, “Prolific | quickly find research participants you can trust,” https://www.prolific.com, 2024.

[90] E. De Sutter, S. Verreydt, K. Yskout, D. Geerts, P. Borry, A. Outtier, M. Ferrante, C. Vandermeulen, N. Vanmechelen, B. Van der Schueren *et al.*, “Using provocative design to foster electronic informed consent innovation,” *BMC Medical Informatics and Decision Making*, 2022.

[91] J. M. Nathe and E. F. Krakow, “The challenges of informed consent in high-stakes, randomized oncology trials: a systematic review,” *MDM Policy & Practice*, vol. 4, no. 1, p. 2381468319840322, 2019.

[92] F. Kreuter, G.-C. Haas, F. Keusch, S. Bähr, and M. Trappmann, “Collecting survey and smartphone sensor data with an app: Opportunities and challenges around privacy and informed consent,” *Social Science Computer Review*, vol. 38, no. 5, pp. 533–549, 2020.

[93] S. Murthy, A. Abu Bakar, F. Abdul Rahim, and R. Ramli, “A comparative study of data anonymization techniques,” in *2019 IEEE 5th Intl Conference on Big Data Security on Cloud (BigDataSecurity)*, 2019.

[94] A. Shahid, M. H. Bazargani, P. Banahan, B. Mac Namee, T. Kechadi, C. Treacy, G. Regan, and P. MacMahon, “A two-stage de-identification process for privacy-preserving medical image analysis,” in *Healthcare*, vol. 10, no. 5. MDPI, 2022, p. 755.

[95] R. Chevrier, V. Foufi, C. Gaudet-Blavignac, A. Robert, and C. Lovis, “Use and understanding of anonymization and de-identification in the biomedical literature: scoping review,” *Journal of medical Internet research*, vol. 21, no. 5, p. e13484, 2019.

[96] V. C. Hu, D. Ferraiolo, D. R. Kuhn *et al.*, *Assessment of access control systems*. US Department of Commerce, National Institute of Standards and Technology, 2006.

[97] Q. Ni, E. Bertino, J. Lobo, C. Brodie, C.-M. Karat, J. Karat, and A. Trombeta, “Privacy-aware role-based access control,” *ACM Transactions on Information and System Security (TISSEC)*, 2010.

[98] R. Tourani, S. Misra, T. Mick, and G. Panwar, “Security, privacy, and access control in information-centric networking: A survey,” *IEEE communications surveys & tutorials*, vol. 20, no. 1, pp. 566–600, 2017.

[99] E. Politou, E. Alepis, and C. Patsakis, “Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions,” *Journal of Cybersecurity*, vol. 4, no. 1, 03 2018.

[100] S. Sarkar, J.-P. Banatre, L. Rilling, and C. Morin, “Towards enforcement of the eu gdpr: Enabling data erasure,” in *2018 IEEE International Conference on Internet of Things (iThings)*, 2018, pp. 222–229.

[101] R. Perlman, “File system design with assured delete,” in *Third IEEE International Security in Storage Workshop (SISW’05)*, 2005.

[102] J.-W. Byun, E. Bertino, and N. Li, “Purpose based access control of complex data for privacy protection,” in *Proceedings of the tenth ACM symposium on Access control models and technologies*, 2005.

[103] J. Li, J. Li, D. Xie, and Z. Cai, “Secure auditing and deduplicating data in cloud,” *IEEE Transactions on Computers*, 2015.

[104] C. Wang, S. S. Chow, Q. Wang, K. Ren, and W. Lou, “Privacy-preserving public auditing for secure cloud storage,” *IEEE transactions on computers*, vol. 62, no. 2, pp. 362–375, 2011.

[105] J. E. Holt and K. E. Seamons, “Logcrypt: forward security and public verification for secure audit logs,” *Cryptology ePrint Archive*, 2005.

[106] L. Chen, Q. Fu, Y. Mu, L. Zeng, F. Rezaeibagha, and M.-S. Hwang, “Blockchain-based random auditor committee for integrity verification,” *Future Generation Computer Systems*, vol. 131, pp. 183–193, 2022.

[107] L. Schaewitz, D. Lakotta, M. A. Sasse, and N. Rummel, “Peeking into the black box: Towards understanding user understanding of e2ee,” in *Proceedings of the 2021 European Symposium on Usable Security*, ser. EuroUSEC ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 129–140.

[108] A. Demjaha, J. Spring, I. Becker, S. Parkin, and A. Sasse, “Metaphors considered harmful? an exploratory study of the effectiveness of functional metaphors for end-to-end encryption,” in *Proceedings 2018 Workshop on Usable Security*. Internet Society, 2018.

[109] C. for Medicare and M. Services, “Medicare Chronic Conditions,” 2024, https://data.cms.gov/medicare-chronic-conditions.

[110] N. Grgić-Hlača, A. Weller, and E. M. Redmiles, “Dimensions of diversity in human perceptions of algorithmic fairness,” *arXiv preprint arXiv:2005.00808*, 2020.

[111] A. C. Plane, E. M. Redmiles, M. L. Mazurek, and M. C. Tschantz, “Exploring user perceptions of discrimination in online targeted advertising,” in *USENIX Security*, 2017.

[112] J. Tang, E. Birrell, and A. Lerner, “Replication: How well do my results generalize now? the external validity of online privacy and security surveys,” in *Eighteenth symposium on usable privacy and security (SOUPS 2022)*, 2022, pp. 367–385.

[113] N. McDonald, S. Schoenebeck, and A. Forte, “Reliability and inter-rater reliability in qualitative research: Norms and guidelines for cscw and hci practice,” *Proceedings of the ACM on human-computer interaction*, vol. 3, no. CSCW, pp. 1–23, 2019.

[114] S. Barth and M. D. de Jong, “The privacy paradox – investigating discrepancies between expressed privacy concerns and actual online behavior – a systematic literature review,” *Telematics and Informatics*, vol. 34, no. 7, pp. 1038–1058, 2017.

[115] E. M. Redmiles, Y. G. Acar, S. Fahl, and M. L. Mazurek, “A summary of survey methodology best practices for security and privacy researchers,” 2017.

[116] E. M. Redmiles, S. Kross, and M. L. Mazurek, “How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples,” in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019, pp. 1326–1343.

[117] I. Adjerid, E. Peer, and A. Acquisti, “Beyond the privacy paradox: Objective versus relative risk in privacy decision making,” *MIS Quarterly*, vol. 42, no. 2, p. 465–488, Feb. 2018.

[118] E. M. Redmiles, Z. Zhu, S. Kross, D. Kuchhal, T. Dumitras, and M. L. Mazurek, “Asking for a friend: Evaluating response biases in security user studies,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 1238–1255.

[119] R. E. Nisbett and T. D. Wilson, “The halo effect: Evidence for unconscious alteration of judgments.” *Journal of personality and social psychology*, vol. 35, no. 4, p. 250, 1977.

[120] I. de Bruin Cardoso, A. R. Russell, M. Kaptein, and L. Meijs, “How moral goodness drives unethical behavior: empirical evidence for the ngo halo effect,” *Nonprofit and Voluntary Sector Quarterly*, 2024.

[121] J. Kaye, E. A. Whitley, D. Lund, M. Morrison, H. Teare, and K. Melham, “Dynamic consent: a patient interface for twenty-first century research networks,” *European journal of human genetics*, 2015.

[122] H. Nissenbaum, “Privacy in context: Technology, policy, and the integrity of social life,” in *Privacy in Context*. Stanford University Press, 2009.

[123] A. S. Levine, “Suicide hotline shares data with for-profit spinoff, raising ethical questions,” *Politico, January*, vol. 28, 2022.

[124] S. Burton, L. A. Cook, E. Howlett, and C. L. Newman, “Broken halos and shattered horns: Overcoming the biasing effects of prior expectations through objective information disclosure,” *Journal of the Academy of Marketing Science*, vol. 43, pp. 240–256, 2015.

[125] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor, “A design space for effective privacy notices,” in *Eleventh symposium on usable privacy and security (SOUPS 2015)*, 2015, pp. 1–17.

[126] MainWP, “Unveiling the facade: Understanding the phenomenon of privacy washing,” 2023, accessed: 2024-06-06.

[127] F. T. Commission, “FTC Enforcement Action to Bar GoodRx from Sharing Consumers Sensitive Health Info for Advertising,” 2023, https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-enforcement-action-bar-goodrx-sharing-consumers-sensitive-health-info-advertising.

[128] “Summary of the hipaa privacy rule,” https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html.

[129] S. Pearman, E. Young, and L. F. Cranor, “User-friendly yet rarely read: A case study on the redesign of an online hipaa authorization,” *Proceedings on Privacy Enhancing Technologies*, vol. 2022, no. 3, 2022.

[130] R. Wang, R. De Viti, A. Dubey, and E. M. Redmiles, “The role of privacy guarantees in voluntary donation of private data for altruistic goals,” *arXiv e-prints*, pp. arXiv–2407, 2024.

[131] P. Ohm, “Broken promises of privacy: Responding to the surprising failure of anonymization,” *UCLA l. Rev.*, vol. 57, p. 1701, 2009.

[132] S. M. Narayan, N. Kohli, and M. M. Martin, “Addressing contemporary threats in anonymised healthcare data using privacy engineering,” *npj Digital Medicine*, vol. 8, no. 1, p. 145, 2025.

[133] I. R. S. T. S. Ann and K. V. S. E. Witchel, “Airavat: Security and privacy for mapreduce,” in *Usenix Org*, 2011, pp. 297–312.

[134] A. Narayanan and V. Shmatikov, “Robust de-anonymization of large sparse datasets,” in *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 2008, pp. 111–125.

[135] R. A. Popa, E. Stark, S. Valdez, J. Helfer, N. Zeldovich, and H. Balakrishnan, “Building web applications on top of encrypted data using mylar,” in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, 2014, pp. 157–172.

[136] K. Hill, “Another privacy problem for google: Engineer allegedly snooped in teens’ accounts.” [Online]. Available: https://www.forbes.com/sites/kashmirhill/2010/09/14/another-privacy-problem-for-google-engineer-allegedly-snooped-in-teens-accounts/

[137] R. Geambasu, T. Kohno, A. A. Levy, and H. M. Levy, “Vanish: Increasing data privacy with self-destructing data.” in *USENIX security symposium*, vol. 316, 2009, pp. 10–5555.

[138] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, “Hippocratic databases,” in *VLDB’02: Proceedings of the 28th International Conference on Very Large Databases*. Elsevier, 2002, pp. 143–154.

[139] R. Agrawal, R. Bayardo, C. Faloutsos, J. Kiernan, R. Rantzau, and R. Srikant, “Auditing compliance with a hippocratic database,” in *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, 2004, pp. 516–527.

[140] J. King and L. Williams, “Secure logging and auditing in electronic health records systems: What can we learn from the payment card industry position paper,” *Usenix HealthSec’12*, 2012.

[141] M. A. Shah, M. Baker, J. C. Mogul, R. Swaminathan *et al.*, “Auditing to keep online storage services honest.” in *HotOS*, 2007.

[142] R. A. Popa, J. R. Lorch, D. Molnar, H. J. Wang, and L. Zhuang, “Enabling security in cloud storage {SLAs} with {CloudProof},” in *2011 USENIX Annual Technical Conference (USENIX ATC 11)*, 2011.

Appendix

A. Survey

Omitted due to space; see the extended version [130].

B. Code Book

Omitted due to space; see the extended version [130].

C. Major Changes Following the Pilot Study

We conducted an initial pilot study and supplemented this with a series of cognitive interviews to identify areas for improvement in our survey. The key modifications included:

Simplify Statements and Understanding. The initial version of the PG and AG statements can be found in our full appendix. However, we receive feedback from the pilot study that the original statements are too abstract and technical. For example, one participant described the anonymization statement as: “at first glance it sounds good but I would say it’s a little bit big” and “I’m kind of stumped about that one but um is it so so big” for the data expiration statement.

We observed participants struggled to understand the statements. We thus include the understanding question in the survey to filter out participants who do not understand the statements.

Separate Auditing Statements. In the original study design, we included one auditing statement focused on external auditing. However, our pilot study revealed that participants had concerns about collusion between external auditors and recipient entities. To address this, we introduced a distinct self-auditing statement reflecting current best practices in data governance frameworks and provides higher level of assurance.

Presentation Enhancement. The original survey presented all statements in the same visual format, which our pilot study revealed could cause participants to overlook the privacy statements. To address this issue, we applied a colored background to key statements, creating clear visual distinction from the surrounding text (see Figure 6).

Demographics Collection. Based on pilot study feedback, we refined our demographics collection by removing income-related questions that were present in the original questionnaire. Additionally, we introduced questions about participants' educational background and professional field to better assess their familiarity with computer science concepts and provide more context for interpreting their responses.

D. Discussion on Included PETs

Anonymization is a foundational privacy guarantee aimed at stripping or obfuscating personal identifiers in data so individuals cannot be readily identified. In privacy-preserving systems, anonymization enables beneficial use of sensitive data without exposing identities, though it involves a trade-off between privacy and utility [131]. Simple de-identification has proven inadequate: numerous studies show ostensibly anonymized datasets can be re-linked with external information to re-identify individuals [132], [133]. To address these risks, stronger techniques such as differential privacy have been developed, providing formal guarantees by injecting statistical noise into query results [134]. Anonymization complements other privacy safeguards by reducing data identifiability, ensuring that even if data is accessed or leaked, it carries less risk of exposing personal information.

Access Control restricts data access to authorized users, embodying the principle of least privilege in privacy-preserving systems. In healthcare, this ensures only appropriate medical personnel can view sensitive records. In cloud storage, access control becomes more complex as users must trust external servers; without safeguards, administrators could potentially access confidential data [135]. To mitigate this risk, researchers have developed cryptographic approaches like end-to-end encryption that prevent even cloud providers from reading stored data without permission [136]. Access control complements other privacy measures by acting as the first line of defense, limiting who can access data and reducing the attack surface for privacy violations.

Data Expiration (also called data retirement or assured deletion) is a privacy guarantee that limits how long sensitive information remains accessible in a system. The idea is that personal data should automatically become irretrievable after it is no longer needed, preventing indefinite retention that could later lead to misuse or breaches. This is particularly important in healthcare, where regulations and ethics mandate retaining patient information only as long as necessary for treatment or research; enforcing expiration helps uphold those limits even on persistent cloud backups. Technically, implementing assured deletion is challenging, but systems like Vanish [137] demonstrated a feasible approach: encrypting data with a key that automatically disappears from a global distributed hash table after a set time, thereby making the ciphertext permanently unreadable once the timer expires. The ability to make data "self-destruct" in this way reflects the notion that the right and ability to destroy data are essential to protect fundamental societal goals like privacy and liberty [137]. By shrinking the window of exposure, data expiration complements other privacy-preserving measures: even if sensitive data is stolen or improperly accessed, it will not persist indefinitely, greatly limiting long-term privacy risks.

Purpose Restriction limits personal data use to only the specific purposes for which it was collected. In healthcare, this prevents patient data collected for treatment from being repurposed for marketing without consent. Implementation typically involves binding data to metadata about permitted uses and enforcing appropriate checks. Researchers have proposed making purpose a first-class parameter in database systems - as seen in Hippocratic databases that enforce "limited use" principles [138]. Purpose restrictions are often supported by accountability mechanisms that verify data was only accessed in ways consistent with its intended purpose [139]. This guarantee complements anonymization and access control by ensuring that even authorized data access remains confined to legitimate contexts, thereby maintaining patient trust and meeting legal privacy obligations [24].

Auditing is widely recognized as a cornerstone of privacy-preserving systems, providing accountability and fostering user trust in sensitive data environments. In domains like healthcare, where patient records are highly sensitive, robust audit trails help maintain compliance with privacy regulations (e.g., HIPAA) and deter misuse of data. For example, secure logging in electronic health record systems can record an "irrefutable trace" of each user's activity, discouraging unauthorized access and aiding breach investigations [140]. Likewise, in cloud storage, auditing mechanisms (often involving independent or cryptographically verifiable checks) ensure that service providers uphold data confidentiality and integrity commitments. Third-party or automated audits allow customers to "assess and expose risk," ultimately giving providers incentives to improve their services and reducing the risk of data lapses over time [141]. Major academic venues have highlighted these needs, with numerous systems incorporating secure audit logs and accountability frameworks as core features. For instance, researchers have proposed using cryptographic attestations as tamper-evident audit records so that clients (or regulators) can verify a cloud provider's behavior [142], and patient-centric accountability schemes that track how medical data is shared to expose any inappropriate access. By enabling such transparent oversight, auditing complements privacy-preserving techniques – ensuring that even as data remains protected (via encryption or access control), any access or policy violation leaves a verifiable trail. This capability is essential for legal compliance and for maintaining public confidence in both healthcare information systems and cloud data services.