

Comprehension from Chaos: Towards Informed Consent for Private Computation

Bailey Kacsmar
University of Alberta
Canada

Vasisht Duddu
University of Waterloo
Canada

Kyle Tilbury
University of Waterloo
Canada

Blase Ur
University of Chicago
USA

Florian Kerschbaum
University of Waterloo
Canada

ABSTRACT

Private computation, which includes techniques like multi-party computation and private query execution, holds great promise for enabling organizations to analyze data they and their partners hold while maintaining data subjects' privacy. Despite recent interest in communicating about differential privacy, end users' perspectives on private computation have not previously been studied. To fill this gap, we conducted 22 semi-structured interviews investigating users' understanding of, and expectations for, private computation over data about them. Interviews centered on four concrete data-analysis scenarios (e.g., ad conversion analysis), each with a variant that did not use private computation and another that did. While participants struggled with abstract definitions of private computation, they found the concrete scenarios enlightening and plausible even though we did not explain the complex cryptographic underpinnings. Private computation increased participants' acceptance of data sharing, but not unconditionally; the purpose of data sharing and analysis was the primary driver of their attitudes. Through collective activities, participants emphasized the importance of detailing the purpose of a computation and clarifying that inputs to private computation are not shared across organizations when describing private computation to end users.

CCS CONCEPTS

• Security and privacy → Usability in security and privacy.

KEYWORDS

Private Computation, Secure Multi-Party Computation, User Study

ACM Reference Format:

Bailey Kacsmar, Vasisht Duddu, Kyle Tilbury, Blase Ur, and Florian Kerschbaum. 2023. Comprehension from Chaos: Towards Informed Consent for Private Computation. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*, November 26–30, 2023, Copenhagen, Denmark. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3576915.3623152>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '23, November 26–30, 2023, Copenhagen, Denmark

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0050-7/23/11...\$15.00

<https://doi.org/10.1145/3576915.3623152>

1 INTRODUCTION

As data access and collection have grown, so have companies' attempts to leverage that data, with regulations trailing far behind. Collaborations between companies increasingly involve data sharing and disclosure. For example, Mastercard raised privacy concerns when it sold transaction data to Google to track whether Google ran digital ads that led to a sale at a physical store (ad conversion) [9].

Within such modern data-sharing practices, a **data subject** is an entity whose data is present in the data set, while a **data controller** is an entity holding a data set. Data controllers who are not themselves the data subject may have different privacy expectations or requirements compared to when data subjects themselves directly make data-sharing decisions. The data subject may not have understood their data could even be shared or sold [26, 47, 64, 82].

Private computation, encompassing complex cryptographic techniques like private set intersection (**PSI**) [13, 61] and multi-party computation (**MPC**) [33, 85], allows companies to analyze data while maintaining data subjects' privacy in many cases. Private computation is especially valuable for cases where the data is sensitive (e.g., health or financial data) [78], among mutually suspicious entities [12, 22], or when there are less open trust boundaries [78].

For example, PSI refers to a computation where two or more parties who each hold a private data set wish to collectively compute the intersection of their sets. The intersection can then be shared with one or more of the participating parties. For example, two companies could determine which users they have in common without disclosing the identities of the users not in common. PSI, as with many other private computations, can be implemented using homomorphic encryption or various other techniques. The privacy guarantees provided follow from the specific mechanisms used and are based on statistical assumptions or computational hardness.

While private computation is often substantially more computationally expensive and complex than its non-private analogue, there is an assumption that it is in some way *better*. For instance, it is presumed to be better for privacy that when PSI is used, data is only shared about clients the organizations have in common. To date, the degree to which users perceive private computation as better, or even feasible and plausible, has remained an open question.

Furthermore, despite a flurry of recent work investigating users' expectations of differential privacy [11, 45, 46, 84] and attempting to improve communication about differential privacy [17, 19, 29, 42, 53], users' attitudes about—and expectations for—the broader range of techniques subsumed under private computation has remained open. The only user-centered work on private computation [3, 74]

has investigated usability from an expert's, rather than an end user's, perspective. While explanations of differential privacy for end users often try to convey the intuition of adding noise or randomness to data or to a computation, the underlying mathematics of other private computation techniques lack an intuitive analogue, yet the guarantees and benefits are arguably more straightforward.

To recap, when an organization considers deploying private computation, two key attributes must be addressed. First, what privacy guarantees can actually be made to data subjects? Second, are those guarantees meaningful to the data subjects whose privacy they aim to protect? In this work, we investigate the second question through 22 semi-structured interviews.

Without knowing what data subjects understand and expect from private computation, one cannot develop tools that empower them to make informed choices. Thus, in this paper we ask and answer the following research questions (RQs):

- **RQ 1:** What do data subjects understand about private computation, and how can specific examples facilitate their understanding of the concept? *See Section 5.2.*
- **RQ 2:** How is a data subjects' willingness to share their data impacted when informed of private computation's protections and guarantees? *See Sections 5.3–5.4.*
- **RQ 3:** How do data subjects perceive private computation's risks (inference attacks and beyond)? *See Sections 5.5–5.6.*
- **RQ 4:** How are perceptions of companies influenced by their use of private computation? *See Section 5.7.*

In brief, we found the following implications for private computation in practice. First, data subjects are able to evaluate and understand the implications of private computation over their own data. Thus, neglecting to inform them of such practices is denying them autonomy over their own data. Second, while participants have an appreciation for the protections private computation can produce, they do not find these protections sufficient to overcome the need for both consent and transparency. That is, key details factor into participants' evaluation of acceptability (Section 5.4), and companies should communicate them. Third, participants are aware of unique high-risk threat models against which private computation cannot guarantee protection (Section 5.5). Thus, failing to communicate the implications of common private computation practices can create unintended risks for users and companies.

2 BACKGROUND

Private computation is the suite of techniques whose understanding by a broad range of users is this paper's focus. To provide context for user-centered communications, including highlighting the types of guarantees private computation provides, this section provides technical background on those techniques. Notions of private computation revolve around two key aspects: *what* is being protected, and from *whom*. The techniques guarantee particular protections as long as certain assumptions are met. The assumptions can be about potential adversaries, system complexities, or statistics. When these guarantees are not in place, private information may leak.

A private computation executes a function over an input to produce an output such that there are limits to what can and cannot be inferred by an adversary, even if the adversary possesses some form of additional data. The function enforces the limitations through

the use of mathematical protection mechanisms from cryptography (e.g., homomorphic encryption), statistical guarantees (e.g., differential privacy), or a combination of techniques. Such computations may be between two or more parties, and they may involve trusted third parties. What is being protected within private computation typically falls under one of the following two classes:

Class 1: Private Data Set, Public Results. Consider a scenario where one or more parties have a (joint) data set and want to release an analysis of the data set. For example, the Census Bureau may wish to release statistics about the population of a certain region. Abstractly, their analysis y is a function f of the data set D , i.e., $y = f(D)$. The party performing the analysis can employ a protection measure like **differential privacy (DP)** [25], which ensures that a single record in the data set D has bounded impact on the analysis y . That is, the output distribution of y shifts by at most a factor determined by a privacy parameter specified by the analyst. By bounding the impact of a single record, the individual records in the data set have a measure of protection against being revealed to anyone who accesses the results of the analysis. The analysis becomes a *private* version of the computation.

The data set D may be distributed among several parties (e.g., D_1, D_2). For example, a government may be interested in the wages of its student population and thus wish to intersect tax filings with universities' registration records. Here, the analysis y may be computed as a **secure multi-party computation (MPC)** [33, 85], which is a cryptographic protocol enabling the parties to compute the function $y = f(D_1, D_2, \dots)$ while ensuring that no party i learns anything except y and D_i using techniques like homomorphic encryption.

Class 2: Private Data Set, Public Subset. While the previous computations protected all individual data records while revealing the output of a computation, we now discuss approaches that instead aim to publicly (or selectively) reveal a subset of the data. Consider a case where parties want to learn additional information about their data or information about a relationship between data sets they each hold individually. For example, assume Google holds a set of digital ad views and Mastercard holds a set of credit card transactions [9]. Google may want to learn which ad views led to credit card transactions, while Mastercard may want to learn which transactions were preceded by an online ad. Abstractly, given a common identifier in the data, the two parties could learn the intersection of their sets. The process of learning this intersection while protecting the respective data sets is known as **private set intersection (PSI)** [31]. Using PSI, two or more parties can compute the intersection of their data without revealing data they possess outside of the intersection. Notably, PSI reveals no information about identifiers not in the other party's set, but fully reveals each identifier in common. Differential privacy can be used on the data sets for additional privacy [34], and extended forms of PSI can compute a function over the intersection [61].

Attacks on Private Computation. So far, we have defined what private computation protects. However, given that some information is revealed intentionally as part of a private computation, there are some risks. Recall that we reveal an analysis y as a function of a data set D : $y = f(D)$. Given y , it is possible for an adversary to compute the inverse of function f and obtain a set of possible

data sets D . This inverse can be computed when given only y , but the adversary may also have background knowledge in the form of a probability distribution over the possible data sets D , further restricting possible inputs and thus improving the attack.

Inference attacks, a subject of ongoing research, may pose significant privacy risks for subjects in the data set D . For statistical data sets, **de-anonymization attacks** or other information leakage can come via the execution of summation queries [49].

In the case of machine learning, attacks may use queries to the model and other attributes. We give a few examples from machine learning where the output y (given to the adversary) is a publicly released machine learning model (e.g., a neural network), the outputs produced by a distributed learning process like federated learning [52], or both. A **model inversion attack** [30, 37] computes the most likely input for one class of the model. For example, for a face recognition model this can be a picture of the recognized person. A **property inference attack** [32] computes a property of the records in the data set given a description of the property. For example, for a face recognition model this can be the ethnicity of the recognized person. A **membership inference attack** [72, 86] computes whether or not a given candidate was part of the data set D . For example, for a medical classification model, this can be whether or not a patient's record was included in the study.

Inference attacks are still feasible if the adversary cannot enumerate all possible data sets D because they only need to estimate the most likely inference. Differentially private protection mechanisms complicate inference attacks [86], but their theoretical analysis is complicated and error-prone [39].

3 RELATED WORK

Communicating Differential Privacy and MPC. As detailed in Section 2, private computation efforts use a technical mechanism to compute revealed outputs from protected inputs. The technical privacy mechanism of differential privacy, and its implications for end users, has received significant attention from the HCI community.

Researchers have aimed to explain differential privacy using a variety of techniques [17, 19, 29, 42, 53, 54] and to evaluate whether differential privacy improves users' willingness to share their data [11, 45, 46, 84]. Those efforts include attempts to convey risk using visuals, risk notifications, and metaphors. While past work has done an excellent job at investigating ways to communicate about differential privacy, these techniques are too narrow to apply to most other types of private computation. Differential privacy provides guarantees of the form "two neighboring data sets are indistinguishable within some probability," and understanding that guarantee requires first understanding the notion of neighboring data sets (i.e., those differing in one row). Private computation more generally does not focus on neighboring data sets. Furthermore, differential privacy's main privacy guarantees result from perturbing, or "adding noise" to, a data set. Whereas the aforementioned prior work on differential privacy aims to give non-technical users an intuition around "adding noise," the underlying mathematics of other types of private computation lack an intuitive analogue.

However, the guarantees and benefits of the other types of private computation we study are arguably more straightforward than differential privacy's guarantees related to neighboring data sets.

All forms of differential privacy provide a *statistical privacy guarantee*. As further described in Section 5.6, our participants raised concerns about such statistical protections; they felt that privacy guarantees should hold consistently. Other guarantees, such as the information theoretic or computational ones provided by other technical mechanisms, may be viewed more positively by the public and thus should be explained clearly to non-technical users.

Explanations and potential regulations must also take into account all relevant stakeholders. The limited prior work on private computation mechanisms other than differential privacy has focused on stakeholders other than the data subjects. For example, Agrawal et al. investigated the perspectives of specialists like industry professionals, researchers, designers, and policy makers [3]. They found that these specialist participants described private computation as a tool for enabling organizations to overcome 'legal gridlocks' related to data sharing. While these specialists acknowledged the importance of end users (data subjects), few prioritized end users' understanding of private computation, increasing the risk that private computation could be used for privacy theater [74]. Similarly, Qin et al. examined the usability of multi-party computation in terms of functionality [63], rather than through our lenses of users' perceptions and understanding.

Communicating Encryption. Whereas private computation uses advanced mathematics to compute a function while keeping the function inputs private, encryption uses advanced mathematics to encode data in a way that keeps it confidential. Researchers have studied users' mental models of encryption. For example, via 19 interviews, Wu and Zappala found that users often conceptualize encryption as "restrictive access control" [83]. Focusing on end-to-end encryption, Abu-Salma et al. found that surveyed users lacked confidence in their understanding of encryption and mistakenly believed that others could access information sent using end-to-end encryption [2]. Subsequent work has aimed to support users' mental models of encryption via improved descriptions [4, 5, 8, 23] or visualizations [77]. Further, due to gaps in their mental models, users often misunderstand the purpose of authentication ceremonies that help guarantee the security of end-to-end encryption [27, 36, 81].

Privacy Perceptions and Preferences. Previous work has frequently found users to be averse to their data being shared or sold [28, 41, 50, 51, 64, 71]. Private computation has the potential to counteract this aversion if its guarantees are communicated successfully. As a result, it is necessary to study users' awareness, understanding, and motivations of technical tools, including their implications for individual and societal privacy [7, 19, 57, 66, 76]. Information about individuals may be collected by employers, government entities, and friends. Which collector originally receives the information is one component of the 'context,' or social domain, in which information is shared. Recent work from Kacsmar et al. [40] found that different contexts, represented by the number and type of participating companies, have an observable influence on users' perceptions of data-sharing practices. Once information is moved to a different context, whether via use or disclosure, it can no longer be assumed to meet privacy expectations [35, 55]. Private computation involves two or more organizations contributing their data. That is, private computation inherently results in a change of context that can influence data subjects' perceptions and preferences.

Law and Policy. To the extent that law formalizes societal norms for enforcement, it is necessary to understand those norms. Legal notions of privacy are primarily framed in terms of individual protections from government and from corporations, with legal and financial penalties for non-compliance. The legal guarantees a company makes are typically communicated within complex privacy policies [18, 56, 67]. These guarantees are enforced, as much as they are, by local data privacy laws. For example, Canada has PIPEDA, the Personal Information Protection and Electronic Documents Act [59]. The United States has, among other laws, the Children’s Online Privacy Protection Rule (COPPA) [80], the Health Insurance Portability and Accountability Act (HIPAA) [79], and the California Consumer Privacy Act (CCPA) [75]. Member states of the European Union have the General Data Protection Regulation (GDPR) [82]. Regulations may impact individuals’ perceptions and thus necessitate recruiting participants from the same locale.

Designers of private computation protocols have suggested that these protocols can help “simplify the legal issues of information sharing” [62] and resolve privacy issues in various domains [20, 44, 60]. However, it takes time to change laws, whereas new technologies are in constant development. Thus, laws may not encompass current and future uses of private computation [48, 58]. Furthermore, our results demonstrate that private computation alone does not resolve privacy issues. Instead, it is critical for consent to be properly acquired, among other aspects of respecting users.

4 METHODS

As there has not been much prior work on users’ understanding of, and expectations for, the broad range of private computation methods we consider, we employ semi-structured interviews to allow us to follow up on participants’ responses and allow participants to ask for clarification. All participants received the same set of questions with the order shuffled as appropriate. Appendix B contains the interview guide. We refined our procedure through pilot studies with five participants. Questions that participants found confusing were either removed or clarified. We do not include responses from the pilot study in our results. The lead institution, located in a country without IRBs, has an institution-level Office of Research Ethics that approved our human-subjects study in an IRB-equivalent process. Ethics Board approval covered the design of the study, consent process, data analysis, and protection of the data collected. Only the researchers at the lead institution engaged in human-subjects research (specific study design, consent process, any interaction with human subjects). Furthermore, only the researchers at the lead institution had access to the data collected.

4.1 Procedure

We developed an interview protocol that addressed the research questions listed in Section 1. We designed our interview questions to gauge participants’ understanding and perceptions of key applications of private computation. We include a range of data leakage scenarios to understand how participants perceive risks.

Before starting, we reminded participants that participation was voluntary, that audio was being recorded, and that they were encouraged to ask questions throughout. The interview proceeded through the parts detailed in the rest of this section:

Expectations and Term Awareness. The interview began with baseline questions to establish participants’ existing perceptions. Participants were asked to “list some of the ways that you expect companies use data about you and others” and whether they had ever “come across” eight terms related to private computation that we presented in randomized order: “private computation,” “encryption,” “hashing,” “multi-party computation,” “differential privacy,” “federated learning,” “private machine learning,” and “secure computation.” Terms with which participants were familiar resulted in follow-up questions about where they had come across the term, what they thought its purpose was for companies and individuals, and a request to define the term in their own words.

Private Computation Definition. We then clarified “private computation” for participants by defining and comparing a non-private computation with a private computation. After participants had the opportunity to ask questions, they were asked to consider what they thought could be an example of “a computation where the result could be made public, but the inputs used to determine that result were sensitive and needed to stay private.”

Computation Scenario Perceptions. As one of the key parts of our investigation, we gathered participants’ perceptions of, and expectations for, private computation through discussing four scenarios in randomized order. Over the course of an interview, these scenarios create what is essentially the process of self-explanation for learning [6, 14, 15]. Self-explanation helps learners adjust their understanding of a topic through examples and explaining concepts back to others. Essentially, it is an inductive, generative process of learning private computation rather than a prescriptive process.

We presented participants with a selection of scenarios in which private computation could be suitably applied. Each scenario consisted of an overall description of the goal of the computation, as well as two ways this goal could be achieved. One way used a straightforward approach involving non-private computation as a baseline. We then presented an alternative approach that employed private computation, enabling participants to compare the two versions. For each scenario, we asked participants how acceptable they found each approach, as well as why. Their explanations and reasoning helped us identify what factors most influence perceptions of (non-)private computation. We also asked participants what differences they perceived between the straightforward computation and private computation in that scenario, how feasible they considered the private computation to be, and how the company performing data analysis might explain the private computation to users.

In terms of scenario selection, our goal was for each instance to reflect a known real-world deployment, encompass either a conventional “social good” goal or “profit-based” goal, be user-facing, and be something for which there existed clear non-private versions of the computation participants would likely have encountered previously. We chose our four scenarios—census data [1], wage equity [16], contact discovery [21], and ad conversion [9]—in consultation with our team’s cryptography experts based on their impression of the likelihood that private computation would actually be deployed in those scenarios in the real world based on cryptographic feasibility and privacy constraints. These four scenarios encompass three different private computation mechanisms. Both ad conversion and contact discovery are settings where PSI can be

deployed. Wage equity efforts can use MPC. Census data can use privacy preserving query procedures.

In more detail, the **wage equity scenario** described an organization collecting salary data with the goal of generating a report on inequities. The **ad conversion scenario** described a credit card company and an online company comparing their data with the goal of determining if digital ads lead to sales in physical stores. The **contact discovery scenario** described a social media company with the goal of determining whether a new user had contacts that already use the app. Finally, the **census scenario** described a government body collecting a range of data with the goal of informing policies and resource management, as well as making results public. See Appendix B for the full descriptions and interview guide.

Inference Attack Perceptions. We then presented participants with four descriptions corresponding to types of inference attacks. For each, we gave participants a series of examples of what specifically the company could learn, asking the participant to explain how acceptable they found that situation. For instance, in the case of a membership inference attack, we said, “One of the participating companies will additionally *be able to learn which specific records in the computed result correspond to you.*” The membership inference attack examples included the data set consisting of a set of members of a dating app, a set of frequent drug users, a set of low-income households, and a set of people with a specific health condition. For each example, participants were asked how acceptable it would be if the organizations involved could determine they were a member of the example data set, as well as to explain their reasoning. The other attacks corresponded to model inversion attacks, statistical inference attacks, and property inference attacks.

General Perceptions. At this point, participants had engaged with four private computation scenarios, as well as four types of inference attacks. To unite these ideas, we asked how the participants thought companies should be communicating to end users how they used data (with and without private computation), as well as what the companies’ responsibilities to their data subjects were.

Collective Activity. We concluded the interview with a collective (or connective) drawing exercise that built upon all topics participants engaged with throughout the study [73, 87]. We asked participants to pretend they were working at an organization that hoped to use private computation and then consider how they would choose to explain private computation to their customers or clients. Participants were able to write, draw, verbally respond, or use whatever other means of communication they preferred. After providing their own explanation, participants were shown all previous participants’ responses to the question and asked what they would add to that explanation and what (if anything) they would remove from it until they arrived at their final version of the explanation. We note that this collective approach integrates input from a range of participants without requiring synchronized timing or a shared location. However, a participant’s potential contribution differs based on when they participated, so participants’ responses and contributions should not be compared with each other.

Table 1: Participants’ demographics, including age range, gender, and highest education completed. Participants indicated whether they have an education or work experience in a tech-related field, as well as in cryptography in particular.

ID	Age	Gender	Education	Tech	Crypto
1	18-24	Woman	High School		
2	18-24	Woman	Bachelors		
3	35-44	Woman	High School		
4	45-54	Man	Bachelors		
5	25-34	Man	Grad School	✓	
6	55-64	Woman	Grad School		
7	18-24	Man	Some college	✓	
8	25-34	Woman	Bachelors		
9	25-34	Man	Bachelors		
10	25-34	Man	Grad School	✓	✓
11	45-54	Man	High School		
12	18-24	Man	Some college		
13	35-44	Woman	Bachelors		
14	25-34	Man	Some college	✓	
15	35-44	Man	Some college		
16	35-44	Man	Bachelors		
17	25-34	Man	Bachelors	✓	
18	35-44	Man	Grad School		
19	35-44	Woman	Some college		
20	55-64	Woman	Grad School		
21	25-34	Woman	Some college	✓	
22	25-34	Woman	Bachelors		

4.2 Participant Recruitment

We recruited participants based in the USA via the Prolific crowdsourcing service using a survey that included demographic information and when they could be available for a synchronous hour-long interview over a video call. We kept interviewing new participants until reaching saturation (no longer finding new themes). We seemed to have reached saturation with just under 20 interviews, but we performed a few extra to be sure. Participants received \$1.45 USD via Prolific for the initial scheduling survey (average time 4 minutes) and an additional \$30 USD for participating in the interview. While most interviews lasted between 50 and 60 minutes, the shortest was 40 minutes and the longest 90 minutes. These times include debugging technical issues (e.g., fixing a microphone).

4.3 Participant Distribution

As detailed in Table 1, we interviewed 22 participants falling in the following age ranges: 18-24 (4 participants), 25-34 (8), 35-44 (6), 45-54 (2), and 55-64 (2). Among participants, 10 identified as a woman and 12 as a man, with no other gender identities being used. Participants reported working in a variety of fields, including politics, libraries, environmental organizations, education, insurance, health, music engineering, technology, personal assistance, chiropractics, and marketing. Participants’ highest level of education completed included a graduate degree (5 participants), a bachelor’s or associate’s degree (8), some college without a degree (6), and high school (3). Further, six participants reported they “had an education in, or work in, the field of computer science, computer engineering, or IT.” One of those participants also reported that they “had an education in, or work in, the field of cryptography.” We note that the only restrictions on participation was age (18-65) and country of residence. The upper bound was due to requirements our Office of Research Ethics sets for including older participants. We chose not to exclude the participant who reported cryptography experience

as during the interview it became clear their familiarity was overstated. Their responses did not differ from those from participants without that reported background.

4.4 Data Analysis

We audio-recorded each interview. We automatically transcribed the audio via speech-to-text software. Afterwards, a member of the research team listened to each recording and corrected the automated transcriptions, as well as grouping responses by question.

We analyzed this qualitative data using an inductive approach, allowing themes to emerge. Two members of the research team extracted participant responses and then collaboratively clustered them according to similar sentiments and themes using the affinity mapping procedure [38, 43, 70]. Affinity mapping allowed us to employ a team-based, collaborative approach to iteratively identify all aspects participants articulated when discussing their understanding of private computation, as well as private computation's implications. As part of the iterative affinity mapping process, after the two researchers formed initial clusters of participant quotes, they reviewed each quote within a theme to see what they had in common and discuss whether the quotes contained any points not encapsulated by others within that theme. Through iteration, we ensured that unique insights were not overshadowed by more prevalent ones. This process enabled us to capture the full range of attributes participants considered, as well as those that most commonly influenced their opinions.

For example, among responses to the acceptability of the ad conversion case, we identified the following themes: consent, privacy, benefits to the company, and low (perceived) sensitivity. Responses to contact discovery brought out themes of consent as well as benefits, limitations, perceived risk, and data minimization preferences. We reviewed emergent themes with respect to commonalities and differences across scenarios and questions to better understand participants' priorities and concerns. These clusters correspond to the structure of the findings we report in Section 5.

4.5 Limitations

While we strived to ensure a diverse sample in many aspects, our participants represent a convenience sample and skew young (less than 20% were age 45+) and educated (69% had completed a bachelor's or graduate degree). Our participants are WEIRD (western, educated, industrialized, rich, and democratic), and we make no claims as to our results being representative of other population groups [69]. All of our scenarios are based upon typical cases in North America, where our participants live, and some examples may not be permitted by laws in other countries. Similarly, our scenarios may not cover data analysis tasks that might be both legal and common outside North America. Finally, as with other response-based studies, we acknowledge the potential for bias towards what participants perceive as socially desirable [65].

5 RESULTS

We present our results centered on our research questions. In terms of comprehension (RQ1), we present the development of participants' understanding of private computation from their first descriptions through the final explanation they constructed. In terms

of perceptions and influence on acceptability (RQ2 to RQ4), we evaluate any changes in perception between scenarios and the reasons participants reported for these changes. This approach enables us to compare the influence of phrasing versus the actual impact as the interview format allowed participants to frame their reasoning in their own words. Thus, we identify themes participants use in their decision-making process when considering data-sharing scenarios, describe how descriptions of private computation influence participants' perceptions of scenarios, and describe the impact private computation has on expectations for companies' responsibilities.

5.1 Initial Knowledge and Expectations

Participants' initial expectations for data usage could influence their perceptions of private computation. Thus, we present an overview of participants' incoming knowledge and expectations.

Expectations. Participants had expectations in terms of what data companies use (purchase history, demographics, search history, salary data, and user preferences), the purposes for which companies use the data (financial gain, improving services, forging social connections, and personalization), and companies' responsibilities with respect to the data (anonymization, preventing re-identification). P8 emphasizes that despite being aware of companies' practices, they do not necessarily approve of them:

“Even though I don't love that, I expect them to use it like for their marketing purposes...grow the bottom line of their business, to make money off of my data, and who I am as a person” (P8).

Participants have an expectation that companies are protecting the data entrusted to them, but P18 expressed concern that data usage practices may go beyond what they expect: “Of course, they may use it for other reasons which I'm not even aware of” (P18).

Relevant Preexisting Knowledge. As a proxy for identifying any preconceived notions participants may have about private computation, we showed participants a set of relevant terms (see Section 4.1). There was only one term for which all participants expressed familiarity: encryption. The only other term with any amount of recognition was hashing. However, hashing familiarity was limited to being a data-mapping strategy and not related to cryptographic hash functions. All other terms either had no participants reporting familiarity or participants being unable to place the origins of their familiarity beyond thinking they may have heard the phrase before. In these cases, the participants guessed they either came across the phrase in terms and conditions or in news articles. Thus, we limit detailing previous knowledge to the term encryption.

Source of Awareness. We surmise that the term encryption is thoroughly embedded in various facets of day-to-day life. Participants responded that they learned of encryption via leisure, education, employment, and when managing finances. However, encryption was not viewed as being particularly relevant to participants' lives:

“[It's] something that's used by techie people or politicians or people who are doing nefarious things. I don't think of encryption as guaranteeing things for individuals, like the lay public like myself” (P6).

Guarantees. On one side, participants expressed skepticism as to what tangible protections encryption can provide. Emphasis was made that there are “no guarantees” (P16) and that, while it may provide some protections, encryption does not make it impossible for malicious actors to access things. For participants that were more optimistic of the protections, encryption was viewed as a means of making it difficult for unauthorized people to access data.

Companies’ Purpose. Some participants responded that encryption is used to provide the “illusion of security” (P8), while others thought encryption is used to provide “customers safety with their data” (P21). Ultimately, whether they had confidence in the protections or not, participants reported that companies use encryption for protecting customer data, protecting proprietary information, gaining customers’ trust, or avoiding legal penalties.

Defining Encryption. In general, participants’ definitions of encryption were not fully comprehensive, but they did show an understanding of encryption at a conceptual level. Essentially, participants highlighted that encryption modifies the information to which it is applied. These changes were referred to as “scrambling” (P20) and “masking or disguising” (P15) the information. Further, participants reported that these changes have the goal of providing security to the information so that it cannot be read by unintended recipients. These responses regarding transformations echo what past work termed an “iterative” mental model of encryption [83].

5.2 Comprehension of Private Computation

We asked participants to define the term “private computation” at three points throughout the interview as a low-level assessment technique for evaluating learning and understanding of concepts [6, 14, 15]. We observe an increase in understanding via participants’ own explanations of private computation comparing their original response at the start of the interview to their final definition at the end. Over the course of the interview, participants responded to the four different example scenarios, impacting their understanding.

First Attempts. We first asked participants to define private computation in their own words at the beginning of the interview. Specifically, participants were shown an abstract definition and asked to think of an example that could fit the definition. This definition occurred before participants were shown any of the scenarios included in the study. Participants struggled to provide an initial definition of private computation. In fact, many participants were unable to come up with any definition. Of those that did provide a definition, they were generally brief and overlapped heavily with the initial definition we had provided.

In contrast, participants did come up with several examples in response to our prompt for “an example of a computation where the result can be made public, but the numbers used to determine the result are sensitive and need to stay private.” That said, not all participants came up with an example, some came up with more than one, and some participants changed their mind about their example. See Table 2 in Appendix A for the list of examples participants provided. The subject domains of the examples included salaries, research studies, and organizations’ financial data. The public outputs included aggregates, averages, company trends, and post-processed data. While not all of the examples were appropriate

settings for private computation, the participants identified a number of cases that already exist. In particular, participants identified examples that corresponded to two of the scenarios we used later in the study: census data and wage equity.

Second Attempts. Later in the study, we again asked participants to define private computation. At this point, they had seen all four private computation scenarios and the cases corresponding to inference attacks. For the second explanation, we informed participants that they could use any medium, including drawing a picture, verbal explanations, and writing. Participants’ second attempt was overwhelmingly more successful than their first. Every participant provided a definition, though with their chosen medium varying; see Figure 1 for a selection of responses. Each definition was reasonably accurate, even if not completely comprehensive. Participants included in their descriptions what is being learned and what is being protected as important. Other aspects they suggested to include were how it will benefit the client and what the computation actually is. In addition to their explanation, participants noted attributes they considered critical to quality explanations. Participants particularly emphasized transparency and honesty. Participants also recommended including examples (especially as figures), summaries, and placing visual emphasis on critical points.

Final Explanation. Participants’ final definition is the one they derived after seeing the collective answer from previous participants. In other words, each participant was shown the explanation derived by consensus by the previous participants. They were then asked what they would add or remove to the current explanation with consideration to their own initial response to the prompt. Earlier in the study, participants made more dramatic changes, and they often incorporated large portions of their own explanation with smaller components of the current collective explanation. As the study progressed, participants made fewer and smaller changes, adding finesse as they identified attributes they considered valuable for an explanation being directed at the public.

When they made changes to the derived explanation, participants expressed the importance of clarity, accuracy, and conciseness. Participants emphasized the value of being concise, but that it needs to be balanced with accuracy. For example, P17 noted that the original example would actually not protect the inputs:

“The only thing I noticed is like, in this example, it’s obvious the data is too small, that you can tell like the ages of specific men and women just because there’s only two men and two women” (P17).

Ultimately, participants made changes to improve clarity across steps in the illustration (consent, input, storage, output) and to emphasize the purpose of private computation. For example, P6 found the term “privacy” failed to encapsulate what is being done and instead suggested using the term “secure computation.” They expressed concern that there is a dichotomy between privacy and using customer data such that private computation could never really represent what is being done: “If you’re using my data, then there’s no privacy... if there’s privacy, then you’re not using my data” (P6). This phrasing choice, which took private computation to secure computation was never reverted by later participants. Further,

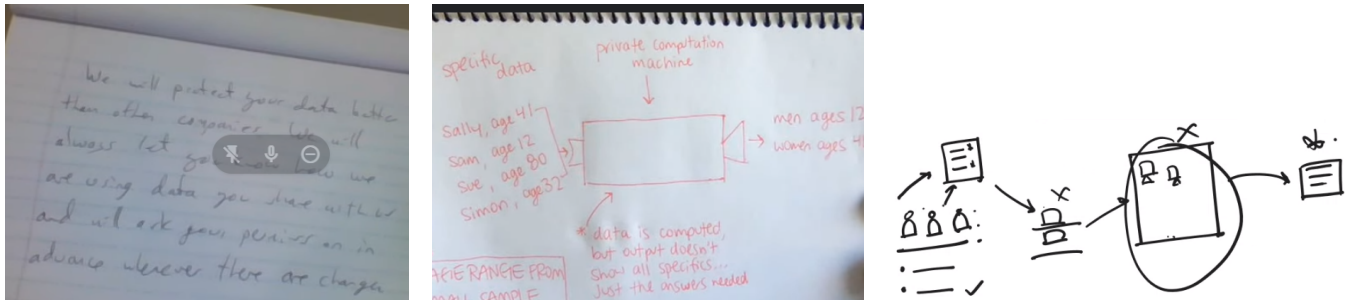



Figure 1: Participants used a range of mediums to convey private computation. Responses included written or typed text, drawn images (digital and paper), and verbal definitions. The above illustrations are from P6, P8, and P10, respectively.

Secure computation is a way that a company analyzes your data. The final analysis will be made public at [access location]. However, your specific data is protected and cannot be traced back to you, nor can your specific data points be traced back to you. The analysis will be specifically [example], and this is being done because [purpose].



This is the information we're getting from you, but, rest assured, only Part Three will be shown. You can trust us to keep your information private. <If true>This information will only be used for this project and nothing else in the future.

Figure 2: Final explanation of private computation derived from all participants via collective (connective) drawing.

other participants who noticed the explanation started with a different term expressed support for the change and that “secure sounds better” (P8). To improve clarity, P8 introduced a visual example to the explanation. This illustration remained a core component of the final explanation, with other participants making small adjustments. Ultimately, though, later participants expressed an appreciation for the visual (P9, P10, P16-P18, P20).

The final explanation after all 22 participants, shown in Figure 2, encompassed attributes participants emphasized throughout the interview process. Within the final answer, there is an explanation providing an overview of the concept, an example that walks the reader through the process (including permission to use the data being requested), the purpose of the computation, and a description of the guarantees being claimed.

As they constructed their explanations, participants did not focus on wanting to know the details of the mathematical mechanism used to achieve the guarantees. Participants trusted that the functionality was feasible without the details, leaving no need for

complicated metaphors to prove it (see Section 5.3). This decision focuses communication of private computation on aspects that are relevant, actionable, and understandable to the populace [68].

Based on the process of collectively creating the final explanation, participants wanted to know the inputs, the outputs, the guarantees, and most of all the purpose of a computation. Notably, the final derived explanation specified what was being done and why, provided an illustrated example, and gave a brief explanation of the implications the computation could have for users. These components are aligned with the themes that emerged when participants explained the acceptability of the four private computation scenarios, detailed later in this section. This consistency suggests these attributes are critical for obtaining informed consent to private computation. The remainder of this section revisits these components and provides insight into why participants considered them relevant.

5.3 General Impact of Private Computation

For our second research question, we found the following key points. Private computation may influence data subjects’ willingness to share their data. However, this influence is not without limits. Participants expressed confidence in the ability for private computation to provide the guarantees described in the scenarios. In many of the presented scenarios, private computation made participants look more favorably upon data sharing. However, as will be discussed in Section 5.4, private computation is not able to completely overcome factors previous work has found to matter to participants (e.g., purpose and consent).

Feasibility of Private Computation. Participants overwhelmingly considered the private computations described in each scenario to be feasible. Not only did participants think the scenarios were possible, but they thought such computations may already be happening (e.g., P12 and P13 commenting on census data). Participants did express concern, however, that companies may not be truthful about what they do with the information they collect (e.g., P22 commenting on contact discovery) and therefore thought it required some sort of enforcement. As one participant emphasised, feasibility was not the critical factor: “...it’s you know, whether there are guards in place, it’s do we have cops to make sure that they’re going to do what they’re supposed to do” (P16).

Participants acknowledged that private computations could be more expensive than non-private computations, which was stated in the scenario descriptions where appropriate. When they considered

the costs, participants included both the company's perspective and their personal views. While participants noted that companies may incur costs from using such computations (P4 and P11), this was not considered a valid reason not to protect users' privacy. Participants even advocated that companies should spend more money on such projects to ensure that users are safe and secure (P2, P20, and P22).

Initial Perceptions of Scenarios. Within our sample, participants generally perceived some scenario goals more positively than others. Specifically, the scenarios for wage equity and census data were generally perceived positively, with responses clustering on the acceptable end of the scale and with few respondents considering these goals unacceptable. The scenarios for ad conversion and contact discovery, however, were viewed less positively. For both, responses clustered on the unacceptable end of the scale. For instance, after they considered the contact discovery description, P14 responded that: *"I want some privacy. I don't need 100%, but I'd like a little bit at least if that's not asking too much"* (P14).

Potential to Impact Acceptability. For each scenario, participants viewed one description corresponding to a non-private computation and subsequently another description corresponding to a private computation. The private computation for both the ad conversion and the contact discovery scenarios saw a positive change in acceptability. Furthermore, wage equity had the most significant improvement with no participants reporting the private computation scenario to be unacceptable.

In the private computation scenarios, the stipulations restricting the amount of data revealed and ensuring that companies cannot use the data for any other purposes were cited as improvements over the non-private analogues: *"Even less of the data...data that is not relevant at all, they modify it to not make it available and I think that's, that's very thoughtful"* (P9). When considering the above attributes participants responded that *"it feels a little bit more protected that way"* (P12), *"aligns a smidge more with my values"* (P8), and *"sounds like another layer of security"* (P19).

Overall, the descriptions corresponding to a private computation tended to improve participants' perceptions of acceptability: *"They're not, you know, over exploiting what they're getting"* (P22). The exception with respect to acceptability was the scenario for census data: *"The second one [describing private computation] is kind of saying the same thing...they're trying to make it sound a little bit better"* (P19). However, even for the other scenarios, the improvement was not unconditional. Participants expressed concern for aspects that private computation does not, or cannot, address. Ultimately learning something is the goal of any private computation, and that is not something that can be changed: *"At the end of the day, they're still like learning specific things about me"* (P7).

Impact on Acceptability Due to Misconceptions. While some participants expressed exceptional insight into the risks and implications of private computation, others felt reassurance from its attributes. Unfortunately, not all of the attributes that gave participants reassurance provide the actual protections participants expected. We identified two main concepts that participants found reassuring but are known not to provide the guarantees attributed to them. The first concept that provides false assurances is aggregation. For example, P6 described the protection from aggregation as: *"When*

it's aggregated, it's lost. It cannot be disassembled. And private does not communicate that in any way shape, or form to me" (P6). This confidence in averages and aggregation is unfortunately misplaced as there are a number of ways a malicious party could carefully select queries such that they can learn about an individual [49]. The incorrect idea that one can "blend into the crowd" via averages and aggregates without risk was also evident in participants' responses to the assorted inference attacks they were shown. That is, participants tended to find property inference attacks more acceptable than attacks that targeted an individual.

Other concepts that provided false assurances were law, policy, and standardization. The assumption that the practices are "legal" or "industry standard" influenced acceptability. For example, P4 specifically stated that if the practice is not an industry standard, then the acceptability would decrease. For example, P16 concluded that if companies disclose such practices in their terms and conditions, it must be legal: *"I don't know in the real world if this is legal to do. I would assume it's legal if it's in their terms"* (P16). However, while participants expressed confidence that the law protects against improper data-sharing practices, this belief was not universal. Some participants stated that such practices do *"not sound ethical even if it's legal"* (P11).

5.4 Bounded Impact of Private Computation

For each scenario, we asked participants how acceptable the scenario was and how companies should explain private computation if they use it. Across scenarios, participants expressed a range of conditions that influence the acceptability. These conditions demonstrate limitations for private computation in terms of influencing data subjects' willingness to share their data.

Motives Matter. When explaining how acceptable they found a scenario, participants said they considered the goals and intentions of the company (P22) and whether they considered the reasons to be just and fair (P11). Goals that benefited society tended to shift their responses toward acceptability. Goals that corresponded to corporate gain tended to shift their responses toward unacceptability. The scenarios for census data and wage equity were viewed as benefiting society. Participants called census data *"crucial information gathering"* (P8). When they viewed the census description, participants were influenced by their trust in the government, the importance of the census for society, and how such data is used: *"If the government is going to spend money, it may as well be based on some data rather than shooting from the hip"* (P6). Similarly, the wage equity description was considered to provide an important societal benefit that prioritized fairness and countered discrimination:

"Wage equity should be a goal of a civilized society and companies aren't going to do that on their own, so third party organizations come in to try to ameliorate some of the inequity" (P13).

Compared to when the organization's goal was viewed as benefiting society, scenarios where the computation benefited the company were received less positively: *"This is based on making more money, they're not considering the actual person involved"* (P11). In particular, the ad conversion scenario was seen as exploitative: *"Want to determine whether...ads are effective? Well, you're still in business, right?...That's enough"* (P16). Some participants expressed

that they understood why the company would want to perform such computations to determine if money spent on advertising was used effectively. Participants that expressed such understanding were still divided; while some thought it was fair, others thought companies should determine effectiveness without using additional personal data: *“Companies should have their own analytics... to figure out their own conversions”* (P21).

Regulate the Restrictions. In the census case, the use of private computation actually increased the number of participants that considered the scenario to be unacceptable or completely unacceptable. Participants expressed concern both about the aspect of “any query” being permitted as well as about how query restrictions would be determined. Participants worried that companies would exploit such restrictions such that *“it’s more like withholding information”* (P18). As a result, they wanted to know *“who is making the decisions regarding the information that’s permitted”* (P8).

Participants’ views were dependent on who makes the restrictions as well as what is restricted. P16 spoke about the importance of allowing the public to replicate results themselves whenever possible. They supported protecting individuals, but emphasized the importance of balancing protections and transparency: *“If we’re talking strictly numbers I lean towards all information available. There shouldn’t be any math problem that is hidden”* (P16). This view was shared by other participants who also emphasized that the acceptability of such restrictions is highly conditional:

“Depending on what information is permitted, it might be good for somebody to know something that they’re not permitting through the system, or it might be bad to let people know something” (P13).

Finally, some participants considered both descriptions to provide insufficient protections and desired additional restrictions (P5 and P10). These participants suggested a hybrid version of the descriptions to produce what they considered to be a more privacy-preserving version. Specifically, to address their concerns, they suggested a query variant that only allows aggregate-based (or average-based) queries while also preventing inferences.

Divulge the Details. Identifying what information individuals prioritized in their decision making is key to ensuring that the necessary information is communicated in the future. Participants mentioned a number of details they indicated as influencing acceptability. In particular, participants who responded that a scenario was neutral or unacceptable emphasized that further information was required before the scenario could be considered acceptable. First and foremost, participants wanted to know when their data was being used: *“That [the data] is being used. What’s being done with it. The other company that is involved, that is having access to it. If it’s going to be ongoing”* (P17). Participants also wanted to know specifically how the data is being used. They wanted to know who is performing the computations and why they are being done. They wanted to know for how long the data is kept, how the data is protected (including the limits of those protections), and the implications for their privacy if their data is used in these ways.

For some participants, a failure to provide details or implement any of the protections the organization claims are reasons to decline to participate in private computation. In other words, even when

private computation is employed, participants care about appropriate flows of information [55]. Participants want to be allowed to judge if a flow is appropriate for themselves. To do so, they require details with respect to the information flows.

Consent Above All. Participants’ desire to be informed about information flows would also give them autonomy over their data:

“Every time your data is used in some kind of computation, you should be specifically alerted by the company. They shouldn’t be able to do private computations... without you being aware of it” (P13).

A theme that emerged across all scenarios was consent, as well as the importance of choice and communication as part of meaningful consent: *“If they don’t prompt you, then completely unacceptable. If they do prompt you, then completely acceptable”* (P17). Further, P1 and P16 both emphasized that consent is not a one-time thing. Companies need to be informing individuals periodically, or *“every step of the way”* (P16), about how their data is being used. As part of this process, the company needs to ensure that the data subject continues to consent: *“When they sign up for the credit card and periodically, they should be reminded that all of their data is, you know, being sold to other companies”* (P1).

In cases where participants may want to withdraw consent, the means to do so should be clear and accessible. Companies need to be *“giving simple directions of, you know, where to go to opt out on the application”* (P4). Such directions support individuals who change their mind about data use, as well as those who did not initially understand or intend to agree: *“If a person finds out they signed something they really didn’t understand, they can have a way to retract their permissions”* (P13).

The final attribute participants emphasized as critical for consent is the use of clear and transparent communication. That is, companies need to be proactive and not just rely on legal contracts to avoid liability. For instance, communication about data use should not be buried in terms and conditions nor obfuscated by legalese:

“Be more upfront about how they’re using our data instead of varying it in like really wordy terms and conditions in language that the average person like myself... can’t understand very well” (P1).

5.5 Risks for Unique Threat Models

In addition to the risks discussed toward the end of the study (e.g., inference attacks), participants highlighted other risks they perceived as possibilities. Participants questioned the implications of private computation and identified a number of risks associated with certain deployment contexts. Both P13 and P19 mentioned risks associated with the goals of the scenarios, regardless of the use of private computation. Individuals can be in situations where computing connections could put someone’s safety at risk. For instance, in the contact discovery scenario, P19 expressed concern that such connections could reveal someone’s internet presence to an abusive ex or someone for whom they have a restraining order:

“[Via] common contacts now he all of a sudden has a friend who has her information and now he has her information. If through the tangled web you could be able to find people... that’s a growing problem” (P19).

Such risks are not necessarily resolved with a technical solution, such as private set intersection, but instead highlight the importance of informing users and respecting their own risk assessments.

5.6 Inference Attacks and Acceptability

Before we presented any of the inference attacks, one participant independently brought up the concern that organizations might make inferences: *“If you’re only giving limited information, you might wonder if they’re gonna acquire other personal information about you from that”* (P22). Participants also expressed concern that they *“can’t really figure out... the implication”* (P6) of computations or *“how it could be exploited”* (P15). The concern is that companies may request limited information, but learn more via other means.

When presented with specific examples of information leakage, two risks associated with inference attacks were most concerning to participants. First, participants worried about any instance where an individual is identified (e.g., membership inference attacks). Second, they worried about any instance where a group of people could be discriminated against (e.g., in certain property inference attacks). Across all inference attacks, the perceived sensitivity of the data affected acceptability. Location data, health data, sexual orientation, and religion are cases where the type of data was deemed to be especially sensitive. Of particular concern was health data. Participants, who were all located in the United States, expressed concern that their insurance company would get this information:

“If that information then got shared with like my insurance company [they] would then decide to raise my rates because maybe I am at an increased risk for heart disease” (P1).

Among participants, there was concern that the inferences made through the attacks could be used in malicious ways and to propagate bias and discrimination: *“What this data is going to be used for, the state of it, should be used to propel humanity forward. Not hold, not keep people back”* (P16).

With respect to the inference attacks, some participants viewed all such attacks as unacceptable because the companies were *“not supposed to have that information, period”* (P6). However, we did observe that inferences that target groups rather than individuals were viewed less negatively. Properties of groups were generally perceived to be somewhat more acceptable. However, this trend was conditional upon the specific property and that property’s potential implications for individuals and society. For instance, if the property could be used to *“manipulate the populace”* (P13) or was *“rude”* or *“discriminating”* (P22), participants found it less acceptable.

For conditional attacks, information leaks only occur probabilistically. However, this was not necessarily viewed positively by participants: *“It’s based on what that record is in relation to even if it needs to be protected”* (P16). Many found it unacceptable regardless of the percentages and stated that the percentage was irrelevant. The three participants who reported a tipping point placed it at a 50%, 25%, or 1-2% chance the exact record would be learned.

5.7 Expectations for Responsibilities

While private computation positively impacted participants’ perception of the scenarios, these perceptions were impacted to a greater degree by other factors. The absence of attributes like transparency

and communication would lead to a more negative reception even when employing private computation:

“It takes more effort, though, and time on the companies to do that. But if they’re willing to, I think it might add a lot to their, you know, trust in the credentials of that company” (P22).

Participants identified responsibilities for companies, governments, and even themselves as individuals. They felt companies have the greatest responsibility with respect to the law, protecting user data, and treating data with respect. Governments’ responsibility was to protect individuals by creating and enforcing policy.

Proactive and Transparent Communication. An individual’s ability to protect themselves is almost inconsequential without support. For example, after expounding on how a company’s priority is financial gain, one participant expressed concern for how data subjects are supposed to learn what they need to have data autonomy: *“How do I protect myself and who teaches me how to protect myself? Who’s responsible for teaching me how to protect myself?”* (P6).

When using customer data, companies need to be upfront about their actions, yet also provide greater granularity of control: *“It’s my responsibility at this point, quite honestly, which is really hard because it’s very confusing”* (P6). For example, rather than giving data subjects a vague description, companies can be more specific: *“It doesn’t really give much more information on what type of data is being used”* (P12). That is, participants suggested having companies detail what is being protected and what risks persist even when employing a privacy-preserving technique.

Respectful Treatment. Participants expected companies to protect the data entrusted to them using the “best” security measures available as that data is not just some abstract input to compute over. In other words, they wanted companies to re-humanize the data entrusted to them as all of the data they hold corresponds to an *“actual individual person with a name, a face”* (P9). Participants expected companies to treat data with respect. Reckless treatment of data can have real consequences for people:

“I think the ultimate responsibility is to use it with caution. To protect people’s privacy. It’s up to the company to make sure they only share to the extent the person allowed them to.” (P9)

Respecting the people who are represented by the data requires companies to exercise clear communication. Without transparency into data-sharing practices, data subjects lack autonomy.

Participants also felt companies need to acquire explicit and ongoing permission for the collection and use of data regardless of the use of private computation. One participant even hypothesized that data-sharing practices would be more positively received if there were less obfuscation and manipulation: *“A big social outcry... that could really be prevented if they were open from the very beginning. If people just knew, they wouldn’t be so spooked by it”* (P9).

Government Regulation and Enforcement. Participants also expressed a level of reassurance toward a scenario in which companies comply with government regulations. However, the nature of these regulations was not always clear to them. While some participants called for clearer regulations, some directly called for the practice of

companies selling data to be made illegal: “*They need to stop selling our information in general... Passing that information to a company, I just think it should be illegal*” (P19). Most participants felt private computation does not impact companies’ legal responsibilities:

“Health is a sensitive topic and there are already legal protections for health information and so on... I don’t see how why this addition of technology should change those protections” (P16).

Participants made suggestions as to how the law can be enforced via independent third parties. For instance, P21 suggested a third party could perform compliance checks and P1 suggested an independent entity could review points critical to consent. The third party could also determine the best way to communicate to users about how their data is being used. They could also determine what information users need to make informed choices about their data.

6 DISCUSSION

Across participants, each individual demonstrated increased understanding (via explanatory evaluation) and communicated to the researchers factors related to private computation that influenced their perceptions of these practices. The reasoning expressed by our participants included both traditional aspects for data sharing (purpose and transparency) as well as technical guarantees (statistical-inference protection, property-inference protection, and membership-inference protection). In this section, we discuss how to better communicate to data subjects about private computation.

For Researchers. The use of private computation often improved participants’ perceptions of the acceptability of data sharing. In other words, participants recognized the value of applying private computation. However, these improvements were neither universal nor unconditional. Private computation did not resolve participants’ concerns in all scenarios. We recommend that future work build on the description our participants collectively created (Figure 2) while aiming to improve communication about private computation.

The description participants created focused on the purpose and implications of data sharing, rather than the complex mathematical underpinnings. Notably, we found that participants did not feel they needed to understand how private computation worked mathematically to find it plausible and feasible. Earlier work on encryption similarly found that users trust the mechanism works without necessarily understanding the mathematics, albeit while holding some misconceptions [23, 24, 77, 83]. As long as users trust the entities using private computation, our study suggests that private computation can make more types of data sharing acceptable.

However, private computation’s protections both have limits and create trade-offs. For example, for private set intersection, a malicious participating entity could fraudulently add non-members to its own list to determine whether those individuals are in another entity’s database. In this sense, the privacy guarantees depend on the honest participation of each entity. Given that our participants closely scrutinized the purpose of data sharing in evaluating acceptability even with private computation, future communications might further highlight the need to trust participating entities.

As mentioned earlier, differential privacy provides probabilistic privacy guarantees, whereas other types of private computation often provide more straightforward guarantees. Future work ought

to compare users’ perceptions of these types of guarantees more directly, such as whether participants would prefer their data be protected by differential privacy or other types of private computation. Researchers should also evaluate whether and why probabilistic guarantees are appropriate and sufficient for their systems.

For Lawmakers and Policymakers. Regulations covering private computation should account for how descriptions of such practices influence data subjects’ willingness to share their data, potentially more so than the actual guarantees. For example, confidence in the protections of aggregated computations and averages may be misplaced [49]. To ensure that dishonorable organizations do not use this confidence to propagate dark patterns [10], regulations must require that companies communicate data sharing’s implications. It is impossible to express all possible implications that could result from a computation. Nonetheless, laws should require companies to make explicit what types of protections are impossible or unlikely.

In prior work, some experts felt private computation could help organizations overcome ‘legal gridlocks’ related to sharing data [3]. In contrast, one of our key results is that private computation was not a panacea for participants’ concerns. While participants generally preferred the private computation variant we showed over its non-private analogue, their attitudes still depended most heavily on the purpose of data sharing and consent processes. As a result, laws and regulations ought to consider private computation as a best practice for data sharing despite its potentially heavy computational costs, rather than a silver bullet enabling previously unacceptable flows of personal information.

For Companies. Private computation techniques are a powerful tool that can increase trust from their users when used as a data-minimization technique. That is, a company should employ appropriate private computation tools for data analyses that are already part of their workflow. When the company adds new types of data collection or flows, private computation alone is insufficient. Communication should be transparent, accessible, and clear. The onus is on the companies to ensure they obtain informed consent. While meaningful consent is a challenge to achieve, it goes a long way toward fostering user trust. Even when using private computation, companies must communicate with the same level of transparency, including details related to how the computation is used and what the company might learn from the computation.

7 CONCLUSION

While technical solutions are a powerful tool for protecting data, such protections do not directly correspond to personal privacy protections. The data being protected in these scenarios is not just an abstract concept, but instead is a placeholder for individuals with real lives and all the complexities that entails for their threat models. Researchers, data collectors, and policy makers need to remember that the protections provided by protocols and constructions do not—and cannot—encompass the full range of risks experienced by individuals in society. Technical privacy solutions must be conscious of the space in which they may be deployed. As we found in our interview study, technical solutions do add value, but that value must not be overstated. The data on which we compute so abstractly is very concrete for the people whose lives generated it.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of NSERC via grants RGPIN-05849 and IRC-537591, as well as the Royal Bank of Canada.

REFERENCES

- [1] John M. Abowd. 2018. The US Census Bureau Adopts Differential Privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [2] Ruba Abu-Salma, Elissa M. Redmiles, Blase Ur, and Miranda Wei. 2018. Exploring User Mental Models of End-to-End Encrypted Communication Tools. In *Proceedings of the 8th USENIX Workshop on Free and Open Communications on the Internet*.
- [3] Nitin Agrawal, Reuben Binns, Max Van Kleek, Kim Laine, and Nigel Shadbolt. 2021. Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [4] Omer Akgul, Ruba Abu-Salma, Wei Bai, Elissa M. Redmiles, Michelle L. Mazurek, and Blase Ur. 2021. 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*.
- [5] Omer Akgul, Wei Bai, Shruti Das, and Michelle L. Mazurek. 2021. Evaluating In-Workflow Messages for Improving Mental Models of End-to-End Encryption. In *Proceedings of the 30th USENIX Security Symposium*.
- [6] Thomas A. Angelo and K. Patricia Cross. 2012. *Classroom Assessment Techniques*. Jossey Bass Wiley.
- [7] Erinn Atwater, Cecylia Bocovich, Urs Hengartner, Ed Lank, and Ian Goldberg. 2015. Leading Johnny to Water: Designing for Usability and Trust. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security*.
- [8] Wei Bai, Michael Pearson, Patrick Gage Kelley, and Michelle L. Mazurek. 2020. Improving Non-Experts' Understanding of End-to-End Encryption: An Exploratory Study. In *Proceedings of the 2020 IEEE European Symposium on Security and Privacy Workshops*.
- [9] Mark Bergen and Jennifer Surane. 2018. Google and Mastercard Cut a Secret Ad Deal to Track Retail Sales. *Bloomberg*. <https://www.bloomberg.com/news/articles/2018-08-30/google-and-mastercard-cut-a-secret-ad-deal-to-track-retail-sales>.
- [10] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proceedings on Privacy Enhancing Technologies* 2016, 4 (2016), 237–254.
- [11] Brooke Bullek, Stephanie Garboski, Darakhshan J. Mir, and Evan M. Peck. 2017. Towards Understanding Differential Privacy: When Do People Trust Randomized Response Technique?. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*.
- [12] Jan Camenisch and Gregory M. Zaverucha. 2009. Private Intersection of Certified Sets. In *Proceedings of the International Conference on Financial Cryptography and Data Security*.
- [13] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. 2018. Labeled PSI from Fully Homomorphic Encryption with Malicious Security. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*.
- [14] Michelene T.H. Chi, Miriam Bassok, Matthew W. Lewis, Peter Reimann, and Robert Glaser. 1989. Self-Explanations: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science* 13, 2 (1989), 145–182.
- [15] Jennifer L. Chiu and Michelene T.H. Chi. 2014. Supporting Self-Explanation in the Classroom. In *Applying Science of Learning in Education: Infusing Psychological Science Into the Curriculum*. Society for the Teaching of Psychology, 91–103.
- [16] City of Boston. 2013. Boston: Closing the Wage Gap. https://www.cityofboston.gov/images_documents/Boston_Closing%20the%20Wage%20Gap_Interventions%20Report_tcm3-41353.pdf.
- [17] Amy Corman, Rachel Canaway, Chris Culnane, and Vanessa Teague. 2022. Public Comprehension of Privacy Protections Applied to Health Data Shared for Research: An Australian Cross-Sectional Study. *International Journal of Medical Informatics* 167 (2022).
- [18] Lorrie Faith Cranor. 2012. Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.
- [19] Rachel Cummings, Gabriel Kaptchuk, and Elissa M. Redmiles. 2021. "I Need a Better Description": An Investigation Into User Expectations For Differential Privacy. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*.
- [20] Emiliano De Cristofaro, Jihye Kim, and Gene Tsudik. 2010. Linear-Complexity Private Set Intersection Protocols Secure in Malicious Model. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security*.
- [21] Emiliano De Cristofaro, Mark Manulis, and Bertram Poettering. 2013. Private Discovery of Common Social Contacts. *International Journal of Information Security* 12, 1 (2013), 49–65.
- [22] Emiliano De Cristofaro and Gene Tsudik. 2010. Practical Private Set Intersection Protocols with Linear Complexity. In *Proceedings of the International Conference on Financial Cryptography and Data Security*.
- [23] Verena Distler, Carine Lallemand, and Vincent Koenig. 2020. Making Encryption Feel Secure: Investigating How Descriptions of Encryption Impact Perceived Security. In *Proceedings of the 2020 IEEE European Symposium on Security and Privacy Workshops*.
- [24] Verena Distler, Marie-Laure Zollinger, Carine Lallemand, Peter Roenne, Peter Ryan, and Vincent Koenig. 2019. Security-Visible, Yet Unseen? How Displaying Security Mechanisms Impacts User Experience and Perceived Security. In *Proceedings of ACM CHI Conference on Human Factors in Computing Systems*.
- [25] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third Theory of Cryptography Conference*.
- [26] Benjamin Fabian, Tatiana Ermakova, and Tino Lentz. 2017. Large-Scale Readability Analysis of Privacy Policies. In *Proceedings of the International Conference on Web Intelligence*.
- [27] Matthias Fassl, Lea Theresa Gröber, and Katharina Krombholz. 2021. Exploring User-Centered Security Design for Usable Authentication Ceremonies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [28] Casey Fiesler and Blake Hallinan. 2018. "We Are the Product" Public Reactions to Online Data Sharing and Privacy Controversies in the Media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [29] Daniel Franzen, Saskia Nuñez von Voigt, Peter Sörries, Florian Tschorsch, and Claudia Müller-Birn. 2022. Am I Private and If So, How Many? Communicating Privacy Guarantees of Differential Privacy with Risk Communication Formats. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*.
- [30] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*.
- [31] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. 2004. Efficient Private Matching and Set Intersection. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*.
- [32] Karan Ganju, Qi Wang, Wei Yang, Carl A. Gunter, and Nikita Borisov. 2018. Property Inference Attacks on Fully Connected Neural Networks using Permutation Invariant Representations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*.
- [33] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*.
- [34] Adam Groce, Peter Rindal, and Mike Rosulek. 2019. Cheaper Private Set Intersection via Differentially Private Leakage. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 6–25.
- [35] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking Data Out of Context to Hyper-Personalize Ads: Crowdworkers' Privacy Perceptions and Decisions to Disclose Private Information. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [36] Amir Herzberg, Hemi Leibowitz, Kent Seamons, Elham Vaziripour, Justin Wu, and Daniel Zappala. 2020. Secure Messaging Authentication Ceremonies Are Broken. *IEEE Security & Privacy Magazine* 19, 2 (2020), 29–37.
- [37] Briland Hitaj, Giuseppe Ateniese, and Fernando Pérez-Cruz. 2017. Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*.
- [38] Karen Holtzblatt and Hugh Beyer. 1997. *Contextual Design: Defining Customer-Centered Systems*. Elsevier.
- [39] Thomas Humphries, Simon Oya, Lindsey Tulloch, Matthew Rafuse, Ian Goldberg, Urs Hengartner, and Florian Kerschbaum. 2023. Investigating Membership Inference Attacks Under Data Dependencies. In *Proceedings of the 36th IEEE Computer Security Foundations Symposium*.
- [40] Bailey Kacsmar, Kyle Tilbury, Mitti Mazmudar, and Florian Kerschbaum. 2022. Caring About Sharing: User Perceptions of Multiparty Data Sharing. In *Proceedings of the 31st USENIX Security Symposium*.
- [41] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. "My Data Just Goes Everywhere": User Mental Models of the Internet and Implications for Privacy and Security. In *Proceedings of the Eleventh Symposium on Usable Privacy and Security*.
- [42] Farzaneh Karegar, Ala Sarah Alaqra, and Simone Fischer-Hübner. 2022. Exploring User-Suitable Metaphors for Differentially Private Data Analyses. In *Proceedings of the Eighteenth Symposium on Usable Privacy and Security*.
- [43] Jiro Kawakita. 1991. *The Original KJ Method*. Technical Report. Tokyo: Kawakita Research Institute.
- [44] Lea Kissner and Dawn Song. 2005. Privacy-Preserving Set Operations. In *Proceedings of the Annual International Cryptology Conference*.
- [45] Patrick Kühnreiter, Viktoriya Pak, and Delphine Reinhardt. 2022. Replication: The Effect of Differential Privacy Communication on German Users' Comprehension

and Data Sharing Attitudes. In *Proceedings of the Eighteenth Symposium on Usable Privacy and Security*.

[46] Patrick Kühntreiber and Delphine Reinhardt. 2021. Usable Differential Privacy for the Internet-of-Things. In *Proceedings of the 2021 IEEE International Conference on Pervasive Computing and Communications Workshops and Other Affiliated Events*.

[47] Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. 2020. The Privacy Policy Landscape After the GDPR. *Proceedings on Privacy Enhancing Technologies* 2020, 1 (2020), 47–64.

[48] Maureen Mahoney. 2020. *California Consumer Privacy Act: Are Consumers' Digital Rights Protected*. Technical Report. Consumer Reports. https://advocacy.consumerreports.org/press_release/consumer-reports-study-finds-significant-obstacles-to-exercising-california-privacy-rights/.

[49] Francesco M. Malvestuto, Mauro Mezzini, and Marina Moscarini. 2006. Auditing Sum-Queries to Make a Statistical Database Secure. *ACM Trans. Inf. Syst. Secur.* 9, 1 (2006), 31–60.

[50] Peter Mayer, Yixin Zou, Florian Schaub, and Adam J. Aviv. 2021. “Now I’m a Bit Angry”: Individuals’ Awareness, Perception, and Responses to Data Breaches that Affected Them. In *Proceedings of the Thirtieth USENIX Security Symposium*.

[51] Aleecia M. McDonald and Lorrie Faith Cranor. 2010. Americans’ Attitudes About Internet Behavioral Advertising Practices. In *Proceedings of the Ninth Annual ACM Workshop on Privacy in the Electronic Society*.

[52] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*.

[53] Priyanka Nanayakkara, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers. 2022. Visualizing Privacy-Utility Trade-Offs in Differentially Private Data Releases. *Proceedings on Privacy Enhancing Technologies* 2 (2022), 601–618.

[54] Priyanka Nanayakkara, Mary Anne Smart, Rachel Cummings, Gabriel Kaptechuk, and Elissa M. Redmiles. 2023. What Are the Chances? Explaining the Epsilon Parameter in Differential Privacy. In *Proceedings of the 32nd USENIX Security Symposium*.

[55] Helen Nissenbaum. 2019. Contextual Integrity Up and Down the Data Food Chain. *Theoretical Inquiries in Law* 20, 1 (2019), 221–256.

[56] Thomas B. Norton. 2016. The Non-Contractual Nature of Privacy Policies and a New Critique of the Notice and Choice Privacy Protection Model. *Fordham Intell. Prop. Media & Ent. LJ* 27 (2016), 181.

[57] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. 2018. Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (2018), 5–32.

[58] Sean O’Connor, Ryan Nurwono, Aden Siebel, and Eleanor Birrell. 2021. (Un)clear and (In)conspicuous: The Right to Opt-Out of Sale under CCPA. In *Proceedings of the 20th Workshop on Privacy in the Electronic Society*.

[59] Office of the Privacy Commissioner of Canada. 2019. PIPEDA in Brief. https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/pipeda_brief/.

[60] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. 2015. Phasing: Private Set Intersection Using Permutation-Based Hashing. In *Proceedings of the 24th USENIX Security Symposium*.

[61] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. 2018. Efficient Circuit-Based PSI via Cuckoo Hashing. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*.

[62] Benny Pinkas, Thomas Schneider, and Michael Zohner. 2018. Scalable Private Set Intersection Based on OT Extension. *ACM Trans. Priv. Secur.* 21, 2, Article 7 (2018).

[63] Lucy Qin, Andrei Lapets, Frederick Jansen, Peter Flockhart, Kinan Dak Albab, Ira Globus-Harris, Shannon Roberts, and Mayank Varia. 2019. From Usability to Secure Computing and Back Again. In *Proceedings of the Fifteenth Symposium on Usable Privacy and Security*.

[64] Emilee Rader. 2014. Awareness of Behavioral Tracking and Information Privacy Concern in Facebook and Google. In *Proceedings of the Tenth Symposium On Usable Privacy and Security*.

[65] Elissa M. Redmiles, Ziyun Zhu, Sean Kross, Dhruv Kuchhal, Tudor Dumitras, and Michelle L. Mazurek. 2018. Asking for a Friend: Evaluating Response Biases in Security User Studies. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*.

[66] Karen Renaud, Melanie Volkamer, and Arne Renkema-Padmos. 2014. Why Doesn’t Jane Protect Her Privacy?. In *Proceedings of the International Symposium on Privacy Enhancing Technologies*.

[67] John A. Rothchild. 2017. Against Notice and Choice: The Manifest Failure of the Proceduralist Paradigm to Protect Privacy Online (Or Anywhere Else). *Clev. St. L. Rev.* 66 (2017), 559.

[68] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* 21, 3 (2017), 70–77.

[69] Jonathan Schulz, Duman Bahrami-Rad, Jonathan Beauchamp, and Joseph Henrich. 2018. The Origins of WEIRD Psychology. Available at SSRN 3201031 (2018).

[70] Raymond Scupin. 1997. The KJ Method: A Technique for Analyzing Data Derived from Japanese Ethnology. *Human Organization* 56, 2 (1997), 233–237.

[71] Irina Shklovski, Scott D. Mainwaring, Halla Hrunn Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*.

[72] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership Inference Attacks Against Machine Learning Models. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy*.

[73] Jesper Simonsen and Toni Robertson. 2013. *Routledge International Handbook of Participatory Design*. Vol. 711. Routledge.

[74] Mary Anne Smart, Dhruv Sood, and Kristen Vaccaro. 2022. Understanding Risks of Privacy Theater with Differential Privacy. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 342, 24 pages.

[75] State of California Department of Justice. 2018. California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa>.

[76] Peter Story, Daniel Smullen, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2021. Awareness, Adoption, and Misperceptions of Web Privacy Tools. *Proceedings on Privacy Enhancing Technologies* 3 (2021), 308–333.

[77] Christian Stransky, Dominik Wermke, Johanna Schrader, Nicolas Huaman, Yasemin Acar, Anna Lena Fehlhaber, Miranda Wei, Blase Ur, and Sascha Fahl. 2021. On the Limited Impact of Visualizing Encryption: Perceptions of E2E Messaging Security. In *Proceedings of the Seventeenth Symposium on Usable Privacy and Security*.

[78] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. 2019. A Hybrid Approach to Privacy-Preserving Federated Learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*.

[79] United States Congress. 1996. Health Insurance Portability and Accountability Act of 1996. *Public Law* 104 (1996), 191.

[80] U.S. Federal Trade Commission. 2022. Children’s Online Privacy Protection Rule (“COPPA”). <https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa>.

[81] Elham Vaziripour, Justin Wu, Mark O’Neill, Daniel Metro, Josh Cockrell, Timothy Moffett, Jordan Whitehead, Nick Bonner, Kent Seamons, and Daniel Zappala. 2018. Action Needed! Helping Users Find and Complete the Authentication Ceremony in Signal. In *Proceedings of the Fourteenth Symposium on Usable Privacy and Security*.

[82] Paul Voigt and Axel Von dem Bussche. 2017. *The EU General Data Protection Regulation (GDPR): A Practical Guide*. Springer.

[83] Justin Wu and Daniel Zappala. 2018. When Is a Tree Really a Truck? Exploring Mental Models of Encryption. In *Proceedings of the Fourteenth Symposium on Usable Privacy and Security*.

[84] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. 2020. Towards Effective Differential Privacy Communication for Users’ Data Sharing Decision and Comprehension. In *Proceedings of the 2020 IEEE Symposium on Security and Privacy*.

[85] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets (Extended Abstract). In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science*.

[86] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. 2018. Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting. In *Proceedings of the 31st Computer Security Foundations Symposium*.

[87] Theodore Zamenopoulos and Katerina Alexiou. 2018. *Co-Design As Collaborative Research*. Bristol University/AHRC Connected Communities Programme.

A ADDITIONAL TABLE

Table 2: Responses to “an example of a computation where the result can be made public, but the numbers used to determine the result are sensitive and need to stay private.”

Example data	Private Data	Public Output
(P1) Individual income, education completed	Individuals’ incomes	Mean income by education
(P2) Voting	Individuals’ votes	Result counts
(P3) Research study	Participants	Study data
(P5) Voting	Individuals’ votes	Eligible voters
(P6) Income, location	Households’ income	Mean income in a region
(P7) Salaries	Individuals’ salary	Average salary
(P9) Financial organizations’ data	Customer data	Financial trends
(P10) Telescope data	Raw data	Post-processed data
(P12) Personal data	i.e. age, demographics	Averages
(P13) Netflix views	Viewer distributions	Report on top service
(P17) Salaries	Individuals’ salary	Average salary
(P18) Political surveys	Individual responses	Aggregated conclusions
(P19) Profits	Beneficiaries	Donations
(P21) Elections	Individuals’ responses	Poll numbers
(P21) Infection disease studies	Collected data	Results

B INTERVIEW GUIDE

The order of the terms (a-h), the four scenarios (wage equity, census data, ad conversion, contact discovery), the four cases (one to four), and the examples within each case (a to d) were randomized.

Welcome. Today we are going to be talking about a topic that may be new to you. We're currently studying public sentiments and understanding of novel data science techniques. We're interested in learning about what people expect and what questions they want addressed if their data is being used for data science by a company. The interview process helps us to understand these expectations and based on them, to make design recommendations for other researchers and policy makers. Please let us know at any point if you have questions. Before we start, I just want to make sure you have a something to write with/on, pen and paper. Throughout the interview, we're going to go through four types of questions, some general, some about terminology, some about types of data sharing, and some about explaining how data is used. On average I expect this interview to take 60 minutes. Do you have any questions or concerns before we start?

To get us started, I'm going to ask you a general question on the topic. For the question, just state as many answers as come to mind and let me know when you're done. Please list some of the ways that you expect companies use data about you and others.

Next, we are going to talk about approaches to data sharing that focus on 'how' the data is shared. We are going to go through a series of terms and I'll ask you if you are familiar with them, and some follow up questions: (a) *Private Computation*; (b) *Encryption*; (c) *Hashing*; (d) *Multi-party Computation*; (e) *Differential Privacy*; (f) *Federated Learning*; (g) *Private Machine Learning*; (h) *Secure Computation*

Have you come across the term [(a) through (h)] before?

- (1) (if yes) Where have you come across the term before?
- (2) (if yes) What kind of guarantees do you think it provides to individuals? Some examples?
- (3) (if yes) What do you think the purpose or goal is for a company using this?
- (4) Please try to define the term in your own words

We're now going to introduce the term private computation. A **computation** is just a calculation (generally in math). For instance, determining the largest number from a list, determining the average, determining a sum. A **private computation**, is a computation that tries to limit the information revealed by the result. It attempts to perform a computation (such as an average, sum, max), and share the result without anyone learning the values used to find the result.

- (1) What do you think is an example of a computation where the result can be made public, but the numbers used to determine the result are sensitive and need to stay private? Follow up: what is sensitive and what is not in the example.
- (2) How would you describe private computation in your own words?

We are now going to talk about some different ways companies can work with client data.

Scenario 1 (wage equity): An organization aims to identify salary inequities across demographics. They reach out to individuals and employment organizations about their salary data. The organization conducts an analysis over the salary data and produces a report on salary inequities. The organization acquires the data for the analysis such that... How acceptable is the organization's goal? Scale: (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)

- (1) ...salary data is shared directly. They receive the salary information of individuals from the individuals or employers via a web-based tool.
- (2) ...salary data is submitted in a modified form privately (with technical and legal protections) via a web-based multi-party computation (MPC) tool. The technical protections prevent the identification of individuals' salary input from the final report. It also protects those who contributed their salary information from being connected to the salary information they provided (though does not prevent it from being known that they were a contributor). Using this technique can be more expensive for the analysis and they cannot use the data for any other purpose.

Scenario 2 (census data): Census data is acquired from citizens of the country by the governing body. It includes information with respect to their age, gender, occupation, income, place of residence. The governing body analyses the data it acquires to inform policies and resource management. It can also make the results of the census available to researchers or the public by... How acceptable is the organization's goal? Scale: (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)

- (1) allowing aggregate/statistical queries (e.g. averages, sums, etc.) over the original data.
- (2) allowing any query, but restricting individuals making queries from performing queries that allow them to make inferences/learn more information than is permitted. This means that some questions cannot be answered by querying the data.

Scenario 1 (ad conversion): An online ad company wants to determine whether ads shown to its users lead to sales in physical stores. They reach out to a credit card company, which has transaction data for physical stores to compute whether there are purchases connected to their ads. The two companies perform the computation such that... How acceptable is the organization's goal? Scale: (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)

- (1) ...they each share their data sets. The credit card company shares the purchase data in physical stores. The online company computes the correlation to online identities locations & online ad views.
- (2) ...the credit card company shares a modified version of their records. The credit card company shares the modified data such that the online company can only identify the financial records that correspond to its users. That is, the information on the other credit card clients (that do not use the online service) is not available to the online company. Using this technique can be more expensive for the company and they cannot use the data for any other purpose.

Scenario 4 (contact discovery): A social media app wants to connect users that are already contacts with one another. The social media app has a list of contact information (its users) and the new user has a list of contact information (their friends etc). The app wants to determine the common contacts between the new user and the existing app users (the intersection). Note that not all of the new users contacts may use the social media app and not all users of the app are contacts with the new user. The social media app can connect the new user to existing users by performing a computation such that... How acceptable is the organization's goal? Scale: (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)

- (1) ...the new user shares all their personal contact information with the social media app.
- (2) ...the new user shares a modified version of their personal contact information. The new user shares the modified data such that the social media company can only identify the new users' contacts that already use the social media app. That is, the other contacts (who do not use the social media app) are not available to the social media app. Using this technique can be more expensive for the company and they cannot use the data for any other purpose.

For each of [A], [B], [C], and [D], we asked the following questions:

- (1) How acceptable is it if the company uses (a)? Explain. (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)
- (2) How acceptable is it if the company uses (b)? Explain. (completely unacceptable, unacceptable, neutral, acceptable, completely acceptable)
- (3) What differences do you expect there should be (if any) if a company chooses to use (b) instead of (a)...
 - (a) ...in general? ...in terms of how companies inform their clients that their data is being used? ...in terms of what companies inform their clients about when their data is being used?
- (4) How feasible/possible do you think it is for a company to use (b) instead of (a)
- (5) How should a company be explaining the technique (b) to their clients if they use it?

Case 1: One of the participating companies will additionally *be able to learn which specific records in the computed result correspond to you. How acceptable is it if the records that correspond to you are...*

- a) ...your salary information? Explain.
- b) ...your credit history (e.g., credit score, mortgage status)? Explain.
- c) ...your location history (e.g., coordinates corresponding to your home, place of employment, etc.) Explain.
- d) ...your genetic markers (e.g., for heart disease, cancer, etc.)? Explain.

Case 2: One of the participating companies will additionally *be able to learn if records of you were used to perform the computation. How acceptable is it if the records they learn correspond to you in a data set of...*

- a) ...low-income households (and thus learn that you are in a low income household)? Explain.
- b) ...dating app members (and thus learn that you use that dating app)? Explain.
- c) ...people with a specific health condition e.g., diabetic, high-blood pressure, autoimmune diseases (and thus learn that you have that specific health condition)? Explain.
- d) ...frequent drug users e.g., alcohol, marijuana, others (and thus learn that you are a frequent user of that drug)? Explain.

Case 3: One of the participating companies will learn *properties for groups*. A group could be people with glasses or any other attribute corresponding to a group of people such as demographics. How acceptable is it if a company can learn, for example...

- a) ...glasses owners prefer shopping online? Explain.
- b) ...women prefer shopping online? Explain.
- c) ...glasses owners have poorer spending habits than non-glasses owners? Explain.
- d) ...women have poorer spending habits than non-women? Explain.

Case 4: When two companies perform the private computation, if one of the participating companies possesses other additional information (e.g. statistics) *they can infer the exact value of a record used in the computation*. How acceptable is it if a company can always learn whether an exact record was contributed by the other organization? Explain.

- a) How acceptable is it if a company can always learn whether an exact record was contributed by the other organization? Explain.
- b) Is it more or less acceptable if a company can accurately learn the record contributed by a different company only 75% of the time? Explain. ...50% of the time? Explain. ...25% of the time? Explain.
- c) To you, at what point (percentage) does this become unacceptable/acceptable? Explain.
 - How does it impact the acceptability if additional information has to be known to learn the values?
 - How does the information that needs to be known influence the acceptability?
 - How does the likelihood the additional information is known influence the acceptability?
- (1) In general, how do you think companies should be communicating to their customers/clients about how they use customer/client data in general?
- (2) In general, how do you think companies should be communicating to their customers/clients about how they use customer/client data if they use private computation for the process?
- (3) In general, what do you think are companies' responsibilities when using your data in these computations? Follow up depending on response: in terms of data protection responsibilities?

The last thing we are going to do is an exercise called co-design. Even though you may have just learned about these techniques, we want you to think about how you would communicate these techniques to someone. There are no right or wrong answers. Imagine you work for a company that wants to use private computation. How would you communicate these practices to your clients? You can draw, write, verbally explain, etc. [Show participant the previous suggestion.] What would you add to or remove from yours based on it? What would you add to or remove from the previous one?