

“We Need a Standard”: Toward an Expert-Informed Privacy Label for Differential Privacy

Onyinye Dibia*

Onyinye.Dibia@uvm.edu

University of Vermont

Mengyi Lu

Mengyi.Lu@uvm.edu

University of Vermont

Prianka Bhattacharjee

Prianka.Bhattacharjee@uvm.edu

University of Vermont

Joseph P. Near

jnear@uvm.edu

University of Vermont

Yuanyuan Feng

Yuanyuan.Feng@uvm.edu

University of Vermont

Abstract

The increasing adoption of differential privacy (DP) leads to public-facing DP deployments by both government agencies and companies. However, real-world DP deployments often do not fully disclose their privacy guarantees, which vary greatly between deployments. Failure to disclose certain DP parameters can lead to misunderstandings about the strength of the privacy guarantee, undermining the trust in DP. In this work, we seek to inform future standards for communicating the privacy guarantees of DP deployments. Based on semi-structured interviews with 12 DP experts, we identify important DP parameters to communicate privacy guarantees for transparency and trust. We further elaborate why and how these parameters should be disclosed according to expert consensus.

1 Introduction

Differential privacy (DP) [21, 22], a mathematical framework for ensuring the privacy of individuals when analyzing sensitive data, has gained traction in real-world adoption. The U.S. Census Bureau [68], Israel’s Ministry of Health [33], Google [31], Apple [3], Microsoft [46], LinkedIn [42], and Facebook [25] have deployed DP in their datasets or systems to ensure data privacy.

Responsible deployment of DP requires careful disclosure of the specific parameters for the deployment’s privacy guarantee. Unfortunately, there is not yet a standard for communicating these parameters, which can vary greatly between deployments. Failure to disclose certain parameters or misleading descriptions of parameter values, can lead to misunderstandings about the strength of the privacy guarantee—even for DP experts. In one example, a confusing statement about the unit of privacy in a DP data release [6] led DP researchers to believe the privacy guarantee was weaker than claimed [36] and required clarification by the original authors [7].

While DP guarantees are mathematically rigorous, their complexity and context-dependence can make them difficult to interpret without sufficient detail. As a result, various stakeholders (e.g., auditors, data analysts, researchers, privacy professionals) need clear

documentation of how DP is implemented to evaluate privacy risks and utility trade-offs. For example, analysts must understand privacy parameters to budget queries appropriately, researchers and auditors need to verify whether implementations apply DP meaningfully. Even when simplified, these details can help build public trust by demonstrating the strength and accountability of the system. Our work is motivated by the need to support these audiences with transparent, structured, and accessible DP disclosures.

Prior work on DP communication [50] has studied *how to explain* one or two important parameters (e.g., the privacy parameter ϵ), without considering *what parameters* should be explained to accurately convey the privacy guarantee rigorously. Prior work also primarily examined non-expert end-users’ understanding, which calls for simplified communication of DP. While simplified formats support accessibility for the general public, they lack precision for technical users or fail to comprehensively convey DP guarantees. In contrast, our study targets *technical users* and focuses on *what parameters* are important: we seek expert opinions towards a standardized DP label, to accurately and completely communicate privacy guarantees.

Our study aims to answer two key research questions:

- **RQ1:** What parameters should be included in a DP label to communicate privacy guarantees of DP deployments?
- **RQ2:** For parameters identified in RQ1, what are the key considerations (e.g. normal ranges, target audiences, methods of presentation) to ensure effective communication?

To answer these questions, we conducted a qualitative semi-structured interview study with 12 DP experts to seek expert consensus. We asked experts about the most important parameters (numeric and otherwise) for communicating about DP guarantees; about the normal ranges or values for these parameters; about the suitability of each parameter for communicating with both experts and non-experts; and about how to present parameters to various audiences.

New insights. Our results indicate significant consensus about *what* to communicate: experts agreed that important parameters like ϵ , δ , and the unit of privacy are vital for transparency in DP deployments. Compared to the closest related work [20], our quantitative analysis enumerates new parameters considered important by experts but not covered in previous work (e.g. deployment model), and additionally describes why each parameter is important, insights about how to set parameter values, challenges associated with communicating about parameters with various audiences, and advice for presenting DP guarantees. Experts in our study raised

*Lead author.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies 2026(1), 43–64

© 2026 Copyright held by the owner/author(s).
<https://doi.org/10.56553/poops-2026-0004>



significant open questions about communicating these parameters to various stakeholders, and there was limited consensus around normal ranges for some parameters.

Contributions. First, we synthesize 12 DP experts' consensus into an essential list of parameters that should be communicated for DP deployments, which lays the foundation for a comprehensive and theoretically sound DP communication standard. Second, we articulate the important properties and considerations for each parameter, which bridges the gap between theoretical guarantees and real-world DP deployments. Third, we discuss key challenges and open questions for communicating privacy guarantees, specifically regarding our proposed DP label format, which contributes to enhancing the transparency and trust in DP deployments.

2 Related Work

2.1 Communicating Differential Privacy

Differential privacy (DP) [21, 22] is a formal definition of individual privacy, typically achieved by adding Laplacian or Gaussian noise to computed statistics. Formally, two datasets D and D' are neighbors if they differ by one individual's data. A mechanism M satisfies (ϵ, δ) -DP if, for all neighboring datasets D and D' and possible outcomes S , $\Pr[M(D) \in S] \leq e^\epsilon \cdot \Pr[M(D') \in S] + \delta$. Here, ϵ (epsilon) is the privacy parameter (also called the privacy budget). Smaller ϵ values yield stronger privacy, and larger values yield weaker privacy. The δ (delta) parameter relaxes the definition for rare events, and is typically set very small.

Effectively communicating the nuances of DP is challenging. For example, Cummings et al. [17] found textual descriptions ineffective for communicating DP guarantees to end-users. Many studies with end-users have explored other ways to intuitively convey DP [5, 9, 29, 38, 44, 50, 69, 73–75]. Some studies focused on explaining the privacy risks associated with DP deployments. For example, Franzen et al. [29] found that, although quantitative risk communication formats can improve comprehension, users with low numeracy skills often struggle to grasp the explanations fully. Similarly, Nanayakkara et al. [50] tackled the challenge of explaining the epsilon (ϵ) parameter, which is often abstract and mathematically complex for non-expert users. They proposed using everyday metaphors and simplified risk explanations, to help users grasp how varying epsilon values affect their data privacy.

Other studies examined the effectiveness of visual tools in explaining DP to end-users [5, 9, 49, 69, 74, 75]. For example, Xiong et al. [75] found that explanatory illustrations and animations, such as data flow diagrams and noise heat maps, significantly improved user comprehension of privacy-utility trade-offs across central, local, and shuffer DP models. Karegar et al. [38] examined metaphors for explaining DP, such as blurring images. They found that metaphors improve DP understanding, but can cause misinterpretation.

In summary, prior work has largely studied *non-experts*, and focused on *how* to communicate DP guarantees. In contrast, our work studies *DP experts*, and focuses on *what to communicate*.

2.2 Privacy Labels

Inspired by standardized food nutrition labels, labels have been increasingly adopted to communicate privacy information [39, 58]. For instance, Kelley et al. [39, 40] developed one of the first privacy

nutrition labels to visually communicate how organizations collect, use, and share personal data. This simplified complex privacy policies and helped users locate relevant information. privacy labels have since been proposed for datasets [34], Internet of Things (IoT) devices [23, 24], and mobile applications [15]. However, studies have revealed that Apple's privacy labels for iOS apps can be difficult to locate and sometimes misleading [15, 43, 76], indicating designing effective privacy labels is challenging and requires rigorous user studies. Prior work has suggested using labels to explain DP [65, 72], but did not give concrete design proposals. This work also attempts to draw insights towards designing a comprehensive privacy label for DP, informed by expert opinions.

2.3 Transparency and Trust in DP

Trust in DP deployments depends crucially on transparency. Closest to our work, Dwork et al. [20] proposed the *Epsilon Registry*, a communal resource to support documentation of differential privacy deployments by sharing information about ϵ , design choices, and implementation context. They conducted qualitative interviews with 11 individuals from 7 organizations who were involved in DP deployments in 2016. Different from our study, they focused on specific properties of the participants' deployments, rather than DP transparency in general; they asked participants detailed questions specifically about epsilon (privacy loss) and privacy unit (granularity), but did not ask them to discuss other important parameters.

Our study emphasizes overall transparency in DP deployments and provides both deeper and broader results. First, our interview questions encourage participants to enumerate all important parameters without priming participants to focus on a predetermined set; our results include analysis of roughly twice as many parameters as Dwork et al. Second, our interviews investigated several in-depth aspects not covered by Dwork et al.: (a) *why* each parameter is important; (b) the typical ranges for parameter values; (c) the challenges of communicating each parameter, and the target audience for each; and (d) suggestions for how to communicate each parameter to various audiences. Third, Dwork et al.'s 2016 study drew from several early DP deployments then; participants in our study discuss important issues (e.g. deployment models, utility information, and empirical privacy metrics) that were not as well-understood nearly 10 years ago.

Dwork et al. [20] inspired several attempts to build actual *Epsilon Registries*, including the *Oblivious Privacy Deployments Registry* [59], which documents privacy parameters used in real-world DP deployments. Desfontaines' [19] informal list of deployments documents additional properties, including the unit of privacy. The Wikimedia Foundation [28] publishes DP page-view datasets, with documentation of the mechanisms used and pre- and post-processing procedures. Despite these efforts, gaps remain in how transparency is achieved for DP deployments, particularly in communicating critical parameters such as the unit of privacy, which plays a key role in defining privacy guarantees. Although some registries, such as the *Oblivious/OpenDP* registry [59], include relevant information under fields such as "scope", they do not always explicitly label or emphasize the unit of privacy.

Moreover, transparency can be counterproductive without clear communication. Cummings et al. [18] emphasized that incomplete

or unclear descriptions of DP can erode user trust. Building on these efforts, our work aims to identify *all* key parameters for responsible transparency in DP deployments. Our DP label informs future epsilon registries by avoiding transparency traps [18], such as focusing too heavily on certain parameters while overlooking broader privacy considerations.

2.4 Transparency with DP Stakeholders

Our study aims to elicit how to ensure transparency from expert opinions; in practice, DP deployments impact many different stakeholders with differing levels of expertise and different requirements for transparency. Here, we provide a high-level overview of prior work associated with different groups of stakeholders, how our results might apply to those groups, and the future research necessary to ensure transparency for each.

End-users. As described earlier, most studies investigating *how* to communicate about DP have focused on end-users—members of the general public with no expertise in DP—and the results suggest that communicating DP guarantees with this audience remains challenging. Our results on transparency are likely applicable to this group, but significant additional research is required to determine *how* to communicate with end-users.

Data users. Users of DP data are primarily concerned with the utility of DP data, rather than its privacy properties, influencing their expectations for transparency. Moreover, data users are typically not DP experts, which complicates DP communication efforts. To date, only a few studies have investigated communication of DP properties with data users. Williams et al. [70] investigated the utility requirements of economists via a survey; they found that economists have low awareness of DP, are not convinced it is needed, and are not sure how to decide whether a data release provides sufficient utility for downstream tasks. These results align with the response to the US Census Bureau's use of DP in the 2020 Decennial Census [1]. Users of Census data expressed concern over the necessity for DP and its impact on data quality [41, 48, 63]. Other work has focused on measuring utility for specific tasks, rather than communicating with data users about utility. For example, Rosenblatt et al. [62] study whether results from population-level studies would be reproducible if DP synthetic data were used; Hay et al. [32] compare the accuracy of DP algorithms on a variety of datasets for specific histogram tasks. Experts in our study included elements of utility in their recommendations for transparency, but were not primarily focused on the needs of data users. The limited existing work in this area suggests that additional research on how to communicate with DP data users to achieve transparency is needed.

Legal and regulatory stakeholders. As described by Garrido et al. [30], legal and regulatory compliance officers represent a key group of stakeholders in transparency for DP guarantees. A significant amount of prior work has examined the relationships between DP and legal requirements for privacy Nissim and Wood [56] and Sokolovska and Kocarev [66] developed frameworks relating legal privacy requirements with technical requirements, and analyze the extent to which DP's technical properties satisfy these requirements. Holzel [35] and Cohen and Nissim [14] perform similar analyses, but specifically for GDPR. Chin and Klienfelter [12] and Cohen et al. [13] analyze specific use cases (Facebook ads and the

US Decennial Census, respectively). The body of prior work suggests that DP can play an important role in fulfilling legal and regulatory privacy requirements, but that (1) the specifics details of the DP guarantee is important for fulfilling legal requirements, and (2) additional context beyond the DP guarantee itself is also important to ensure compliance. Our transparency results align with the technical features of DP considered important for legal compliance in prior work, suggesting that transparency for DP experts may overlap considerably with transparency for legal experts. Additional research with legal and regulatory experts is needed to identify gaps between these requirements and to investigate how DP guarantees should be communicated with this group.

3 Method

We conducted a qualitative interview study with 12 DP experts. Our semi-structured interview design was motivated by the open-ended nature of our research questions. One of our goals was to identify the important parameters for transparency in DP deployments, and key considerations associated with them; an open-ended, qualitative study is best suited for answering these questions [57]. This section details the participant recruitment, the interview procedures, the data analysis process, and ethical considerations of our study.

3.1 Sample & Recruitment

Recruitment: We used the purposive sampling method to recruit DP experts who are at least 18 years old and have a minimum of five years of research experience or practical involvement with DP. Given the niche expertise required for this study, we first approached potential qualified participants via targeted recruitment emails sent to our professional network and the OpenDP mailing list (where DP experts gather). We also used snowball sampling to seek peer recommendations of additional qualified candidates from enrolled participants. We successfully recruited ten experts via targeted email and two additional experts through peer recommendation. Our study invitation email includes an eligibility survey (Appendix A) focused on their professional background in DP. Our recruitment method and sample size align with prior qualitative studies in DP area [20, 30, 64]. Similarly, we were limited in sample size by the challenge of recruiting experts in DP. See Section 6 for additional discussion of this limitation.

Our sample: 12 DP experts from 10 different organizations participated in our study between May and August 2024. On average, these experts had 7.5 years of experience in DP. We assigned participant numbers P01–P12 to these experts. Table 1 details the expert demographics in terms of work experience, job roles, and gender, as well as the industries they are engaged with. See Table 6 in Appendix C for additional expert demographics.

Our participants include two professors and one postdoc at research-focused universities whose primary publication area is DP; five individuals who have played key roles in the design or development of US government DP deployments; and six individuals who have played key roles in the design or development of industry DP deployments. All are currently recognized experts in the area of DP. Where available, we list the H-index of each participant in Table 1 as an (imperfect) proxy for expertise.

ID	Experi- ence (Years)	Role	Gen- der	PhD (Y/N)	H- index
P01	10+	Academic Researcher	F	Yes	25-30
P02	5 - 10	Academic Researcher	M	No	30+
P03	10+	Data Scientist	F	Yes	5-10
P04	5	Data Privacy Engineer	M	No	N/A
P05	5	Academic Researcher	F	Yes	5-10
P06	5	Academic Researcher	F	Yes	25-30
P07	5 - 10	Data Privacy Engineer	M	Yes	15-20
P08	5 - 10	Data Scientist	F	Yes	10-15
P09	5	Data Privacy Engineer	M	No	5-10
P10	5 - 10	Software Engineer Building DP Tools	M	No	N/A
P11	5 - 10	Data Scientist	M	No	5-10
P12	5 - 10	Data Privacy Engineer	M	No	5-10

Table 1: Expert Demographics, n = 12

3.2 Interview Procedure

The research team informed participants about the study's data collection, processing, and storage procedures, obtained their consent to participate, and assessed their eligibility for the study via an eligibility survey (details in Appendix A). The research team conducted the interviews via Microsoft Teams and made audio recordings after obtaining verbal consent from the experts. The semi-structured interviews followed a three-part format:

- (1) The first part of the interview focused on identifying key differential privacy (DP) parameters that experts considered essential for inclusion in DP systems and deployments. Experts were asked to elaborate on the significance of each parameter and whether specific parameters are often overlooked or underestimated in DP discussions.
- (2) The second part of the interview explored why each parameter should be included in our proposed DP label, examining its typical or "normal" range within the context of DP, its relevance to the general and technical audiences. Experts were also asked to assess the applicability and effectiveness of these ranges, providing examples where necessary.
- (3) The third part of the interview explored the design, structure, and presentation of the DP label. We seek experts' feedback on our proposal for a layered design, as well as their insights on how each parameter should be displayed, where the labels should be placed for maximum visibility, and other recommendations for improving the label's layout and accessibility.

To support a shared understanding of the study's framing, we began each interview by presenting visual examples of nutrition-style labels (e.g., food and iOS privacy labels) and explained our goal of developing a similar label for DP deployments. We clarified that the term "parameter" was defined broadly as: any feature, metric, or concept that experts believed was important for understanding or communicating a DP guarantee. This included, but was not limited to, numeric values like ϵ or δ . Participants were encouraged to consider contextual and structural elements (e.g., the unit of privacy, privacy loss budgeting, or utility constraints), as well as

any other information they deemed critical for interpreting a DP deployment.

To ensure clarity of the questions and estimate the duration of the interview, we piloted the interview questions piloted with two graduate students in our research group who were familiar with DP not involved in the study design. After piloting, we reduced the number of questions and adjusted the order of some questions. The finalized interview guide appears in (Appendix B).

3.3 Data Analysis

To analyze the interview transcripts, we used a hybrid thematic analysis process [8, 27, 61] to systematically identify recurring themes related to our research questions. This approach combined deductive coding, guided by predefined DP concepts, with inductive coding to allow new themes to emerge from the data, ensuring that both established DP principles and novel insights were captured.

The interviews were transcribed using Microsoft Teams' automated transcription feature and carefully reviewed and corrected to ensure accuracy and completeness. Once familiarized with the data through repeated readings, two researchers independently open coded the 12 transcripts. Although the predefined DP concepts informed the researcher's understanding, the coding process prioritized emergent themes from the transcript contents.

Following the independent coding phase, the two researchers collaborated to merge their individual codebooks into a combined version. Where disagreements arose, they discussed the differences and, if necessary, sought input from the broader research team to reach a consensus. Once a unified codebook was finalized for each transcript, the researchers grouped the codes by parameter and synthesized them into broader themes. This process allowed for systematic organization of the data while preserving its complexity and nuance. Finally, the researchers revisited the data to validate the themes, ensuring that they accurately reflected the content of the interview. This collaborative and iterative approach ensured the analysis was rigorous, coherent, and faithfully represented the insights provided by the experts.

3.4 Ethical Considerations

This study received approval from our university's Institutional Review Board (IRB). All participants received information about the study's objectives, data collection, processing, and storage procedures and provided their consent through the eligibility survey (Appendix A). They also verbally reaffirmed their consent at the start of each interview. To protect participants' privacy, only general demographic information and professional background were collected, and identifiable data were anonymized during transcription. Audio recordings from the interviews were transcribed, coded with unique identifiers, and securely stored. All participants received equal compensation of \$40 in the form of an electronic gift card upon completing the interview. Four experts chose to donate their compensation to charity.

4 Results

To answer RQ1, we present the key DP parameters that experts deemed essential for communicating privacy guarantees, and synthesize why they are important.(Section 4.1). Then, we answer RQ2

Parameter	§
Privacy parameters	4.1.1
Unit of privacy	4.1.2
Utility information	4.1.3
Mechanism used	4.1.4
Algorithm hyperparameters	4.1.5
Deployment model	4.1.6
Empirical privacy metrics	4.1.7
Privacy interpretation-semantic	4.1.8
Other parameters	4.1.9

Table 2: Parameters Mentioned by Experts

by presenting our qualitative analysis results of the communication challenges (Section 4.2), typical ranges for parameter values (Section 4.3), audience relevance (Section 4.4), and recommended presentation formats for these parameters (Section 4.5).

4.1 Important DP Parameters

We identified nine key categories of parameters that experts considered essential for defining and communicating DP guarantees: privacy parameters (e.g., epsilon, delta), unit of privacy, utility information, mechanism used, algorithm hyperparameters, deployment model, empirical privacy metrics, privacy interpretations or semantics, and other parameters.

Table 2 presents these categories, which serve as an organizing structure for the wide range of parameters discussed during the interviews. A comprehensive list of individual parameters and how they map to these categories is included in Appendix D. In the following subsections, we elaborate on each parameter category, focusing on its role in enhancing transparency in DP deployments.

4.1.1 Privacy Parameters

Privacy parameters are core metrics in the DP literature for defining the strength of privacy guarantees, balancing privacy protection and data utility [21, 22]. DP experts in our study agreed on three key privacy parameters: Epsilon (ϵ), Delta (δ), and other privacy loss measures (such as rho (ρ) for zCDP [10]).

The privacy parameter assigned to a data release is often referred to as the privacy budget [71], which quantifies the allowable privacy loss in a system. In practice, parameters like epsilon and delta are central to real-world DP deployments and disclosed in the US Census Bureau's 2020 data release [2] and other industry deployments [4, 26]. Experts in our study unanimously emphasized the need to include privacy parameters in a DP label for transparency. As one expert remarked, *"These are the things that's gonna give experts a good understanding of the 'real' privacy guarantees of the system."*

4.1.1.1 Epsilon (ϵ)

Epsilon (ϵ) is the cornerstone parameter in DP [21], quantifying the indistinguishability between datasets and measuring privacy strength. It is the tuning knob for balancing privacy and accuracy as it determines how much noise to be introduced [71]. Smaller values provide stronger privacy but reduced utility [53] making epsilon

essential for evaluating DP systems, and defining privacy-utility tradeoffs.

Direct Impact on Privacy Guarantees: Eight experts highlighted epsilon as a key parameter directly impacting privacy guarantees. P01 stated, *"Epsilon is the main parameter that controls the privacy guarantee. So I think it's like really important, like at a high-level value".* Another described epsilon as a *"knob that you can tune, and it's sort of the center of a lot of decision-making that happens"* (P05), highlighting its central role in DP systems.

Enabling Interpretation and Comparability: Five experts highlighted epsilon's role in interpreting and comparing DP systems. One expert remarked, *"...it gives you some way of comparing things"* (P07). Another expert described epsilon as *"a Rosetta Stone across a whole bunch of different kinds of differential privacy"* (P09), underscoring its importance in bridging various DP frameworks.

4.1.1.2 Delta (δ)

Delta (δ) complements epsilon (ϵ) by bounding the probability of privacy failure in worst-case scenarios. While ϵ quantifies indistinguishability, δ accounts for residual risks, offering a fuller picture of a DP mechanism's privacy guarantees [52].

Ensuring Robust Privacy Guarantees: Three experts emphasized delta's role in capturing worst-case privacy risks. P01 explained, *"...in the worst case, delta captures how much your data could be exposed in the clear."* Another noted, *"Epsilon and delta together represent the degree of privacy being offered"* (P03).

Prevents Privacy Theatre: Two experts stated that delta safeguards against weak or misleading DP deployments. P01 explained, *"you should include delta to prevent against malicious uses, like people hiding things they are doing".* P05 warned, *"if you have a delta that's super big, then it's just a poor deployment...not offering the protections that people think they're getting."* Setting an explicit bound on privacy failure ensures transparency and upholds DP's credibility.

4.1.1.3 Other Privacy Loss Measures

Beyond epsilon (ϵ) and delta (δ), alternative frameworks like ρ -zero concentrated DP [10] and Rényi DP [47] can offer improved privacy loss quantification.

Contextualizing Privacy Guarantees: P01 and P07 emphasized that incorporating alternative privacy loss measures improves the accuracy and robustness of privacy guarantees. P01 explained, *"If you stick to epsilon as the primary parameter, composition under frameworks like rho or Rényi DP often results in a looser bound... Including these parameters ensures the guarantees are fully contextualized."*

Improving Privacy Loss Accounting: Several experts noted that alternative privacy loss measures help in quantifying privacy loss across multiple data releases. P08 stated, *"You can quantify the privacy loss... and compose everything to still understand the combined privacy loss."* By providing additional granularity, these parameters enhance the interpretability of privacy guarantees for technical stakeholders.

Enhancing Transparency and Trust: P04 and P11 emphasized that including these parameters enhances transparency and strengthens trust in DP systems. P04 remarked, *"Estimates of privacy loss parameters should be included... If you don't include those, it's either*

defeating the purpose or being dishonest." This underscores their role in ensuring the integrity of privacy guarantees and fostering confidence in DP deployments.

4.1.2 Unit of Privacy

The unit of privacy defines what and who is being protected in a DP system, shaping the scope and granularity of privacy guarantees [53]. It varies across applications—ranging from record-level (or event level) to user-level privacy—and influences both protection levels and utility trade-offs. Experts discussed its role in ensuring transparency, optimizing utility, and the challenges of interpretation.

Defines Scope of Protection and Ensures Transparency: Eight experts emphasized that specifying the unit of privacy clarifies the scope of protection. P03 noted, "*Unit of privacy ... tells you what the guarantee applies to.*" P04 reinforced this, stating, "*DP is giving a privacy guarantee. You should know who the guarantee applies to. The privacy guarantee is meaningless if you don't have that.*" Explicitly defining it also pushes organizations to clarify their privacy commitments, as P07 pointed out: "*It forces companies to actually think about it... It would push companies to say things like, 'We do user-level DP,' but what we actually mean is device-level DP.*" This transparency helps prevent "privacy theater" and ensures that the stated guarantees align with actual protections.

Utility Implications: The unit of privacy helps optimize DP mechanisms for practical use. P10 described it as "*a very useful tool as an algorithm designer to tune utility.*" By specifying the granularity of privacy, that is, privacy at different levels (e.g., user, session), organizations can tailor mechanisms to provide robust protections while maintaining usability and utility. This granularity ensures that privacy guarantees are not only meaningful but also adaptable to the specific needs of different applications.

4.1.3 Utility Information

Utility information emerged as essential for DP deployments, encompassing metrics related to data accuracy after privacy mechanisms are applied. Experts highlighted its role in balancing privacy guarantees with meaningful data use.

Ensures Transparency: Utility metrics help prevent overly noisy or misleading outputs. P03 noted, "*It keeps algorithm developers honest... forcing them to analyze their results so they're not producing wildly noisy, problematic data.*"

Influences Incentive to Share: Some experts emphasized utility's role in shaping individuals' willingness to share data. P05 stated, "*there's actually a pretty strong case for including it even to data subjects... it influences how they feel about the decision to share their data or how they feel about the privacy that they're being offered*"

Facilitates Evaluation of Data Quality: Utility metrics, such as error rates, are crucial for evaluating data quality. P09 explained, "*These things are key for an analyst who wants to understand the actual data quality and how close they can assume it will be to the underlying dataset.*" P08 emphasized that such metrics are "*super nice for the data user*".

4.1.4 Mechanism Used

The mechanism in DP dictates how privacy guarantees are implemented [53], shaping both precision and interpretability. It determines the way in which noise is applied, data is synthesized, or queries are answered, affecting the reliability of privacy assurances.

Precision of Guarantee Definition: The choice of mechanism defines privacy guarantees based on statistical properties and use cases. P01 noted, "*Epsilon, delta may not bind as tightly for every mechanism. The nature of the guarantees in a more nuanced way depends on the mechanism that is used.*" P12 connected mechanisms' role to their ability to provide statistically rigorous definitions of privacy stating, "*This is directly required to actually come up with a bound on the statistical power of a test.*"

Enables Interpretation and/or Comparison: Mechanisms facilitate understanding and comparison of privacy guarantees across different deployments. P10 explained, "*You sort of need to know what the privacy definition is to interpret the parameters. It's helpful to characterize mechanisms in ways that don't require providing the full mechanism details.*" This supports technical users in evaluating trade-offs.

4.1.5 Algorithm Hyperparameters

Algorithm hyperparameters are crucial in differential privacy (DP) implementations because these parameters, such as clipping bounds, noise multipliers, and batch sizes, should be included in transparency efforts like a DP label to inform stakeholders about privacy and data quality.

Transparency: Hyperparameters impact data usefulness and assess result quality. As P04 noted, "*It tells you how confident you should be in the results... Putting those in context of how good the release is, is also important.*" P08 added, "*Including clipping or clamping parameters ensures that the data was appropriately preprocessed, reducing bias and enhancing trust.*"

Understanding Utility: Hyperparameters clarify the privacy-utility trade-off. P06 explained, "*Including these details would allow a better understanding of how much signal you are getting, and how much of it is clipped, which in turn impacts the utility.*"

4.1.6 Deployment Model

The deployment model determines how data is processed, stored, and accessed, influencing privacy and security. It can involve local processing on a user's device (as in local DP), centralized processing by a data collector (as in central DP), or decentralized models like federated systems. Each model presents different privacy and security trade-offs, making it crucial to assess the privacy protections at various stages.

Significant Impact on Threat Models: The deployment model directly affects the threat model and associated security and privacy guarantees in a deployment. As P05 noted, "*It implies significantly different threat models, which impact data security concerns.*" Local DP avoids the need for a trusted data collector and reduces the risk of data breaches, while central DP increases the risk of a data breach but offers better utility.

4.1.7 Empirical Privacy Metrics

Empirical privacy metrics measure privacy experimentally, via reconstruction, re-identification, membership inference, or other attacks, or via experimental auditing procedures [11, 37, 51]. The results can provide a lower bound or average-case approximation of the privacy protection provided by a data release. One or many empirical metrics can be provided based on the concerns for a specific dataset.

Practical Risk Assessment: Empirical privacy metrics help assess actual attack risks. As P06 stated, *"I think empirical results are just always good to have for anyone because that's the thing that is tangible for people and people can assess their risks."* Empirical metrics are broadly applicable to evaluate datasets with or without DP, and may be easier for non-experts to understand, so they may be particularly useful for end-users, data users, and policymakers.

4.1.8 Privacy Interpretation or Semantics

Privacy interpretation or semantics communicate privacy by directly stating the semantics of the guarantee, often as a bound on the trade-off between Type 1 and Type 2 error in a hypothesis test.

Practical Risk Assessment: The interpretations are important for conveying privacy risks clearly, particularly for decision-makers both experts and non-experts who are familiar with concepts like type 1 & 2 errors better grasp privacy risks in practice. As P11 noted, *"I think it helps ... to have something like a type 1, type 2 error interpretation that is closer to a plain language statement..."*

4.1.9 Other Aspects or Parameters

Experts highlighted additional aspects or parameters that, while mentioned less frequently, could reflect nuanced dimensions of DP to improve transparency, facilitate interpretation, and support informed decision-making.

These include the **data creation pipeline/contextual integrity**, **data release metadata**, **goals of the release**, **DP software tool used** (e.g., OpenDP) **implementation and vetting details**, (e.g., compliance with DP principles and verifying correctness), **average-case guarantees**, **guidance for normal ranges of parameters**, **group privacy information** and **public information used**. Notably, P09 emphasized the importance of documenting the data creation pipeline and release goals stating, *"Error in DP often comes from data governance rather than aggregation... Ensuring clear documentation and attribution of all steps is critical."* This perspective is derived from this expert's prior experience with DP deployments.

Although these aspects were mentioned less frequently, they may become increasingly important as DP deployments grow in scale and complexity. In particular, many of these suggestions reflect the principle of *contextual integrity* [54, 55]—the idea that privacy depends not only on data protection mechanisms but also on preserving appropriate information flows relative to the context in which data is collected, processed, and shared. Capturing metadata about the release context, tool choices, and intended use aligns with this framework and supports transparency grounded in the expectations and norms of the data environment.

4.1.10 Overlooked Parameters

Many experts noted that discussions on DP tend to overemphasize epsilon and neglect other critical parameters. As P01 remarked,

"Everything besides epsilon is overlooked... but there's so much more that is important."

Seven experts highlighted the **unit of privacy** as another often-overlooked parameter in DP. P08 observed, *"In discussions, I feel like the unit of privacy is not so much discussed, at least not as much as privacy loss guarantees."* Similarly, P05 remarked, *"I don't think it's overlooked in the DP community, but I think that it gets overlooked in documentation around DP deployments."*

Four experts identified **algorithm hyperparameters** and **utility metrics** as underestimated parameters that significantly affect DP effectiveness and usability. P06 explained, *"When training with DP, hyperparameters like learning rates and batch sizes become nonlinear and more complex, which is often overlooked."* P03 emphasized, *"All the little algorithm quirks and data properties that impact utility are overlooked... Things that affect any machine learning algorithm probably get overlooked here too."*

One expert noted **sub-budgets** as often overlooked. P11 pointed out, *"The difference between per-attribute guarantees and the global budget is significant, yet these sub-budgets are rarely discussed."* Similarly, one expert identified **data universe assumptions** which define the set of possible records in a DP framework, as critical yet often overlooked. P12 remarked *"The definition of universes is probably one of the most overlooked but critical aspects."*

Experts emphasized the need for greater attention to these overlooked parameters to adopt a holistic approach to DP and improve transparency. This would enable stakeholders to fully understand the scope of privacy guarantees and associated trade-offs.

4.2 Challenges of Communicating Parameters

While many DP parameters are essential for defining privacy guarantees, experts in our study identified significant challenges in making these parameters understandable and useful—particularly for non-experts. We synthesized these overarching challenges in communicating DP parameters below.

Too Technical and Difficult to Understand: Experts emphasized that many DP parameters—particularly (ϵ) , (δ) and advanced constructs like Rényi DP—are too technical for non-technical audiences. P09 described this as DP's *"main Achilles heel... you have to have some mathy understanding of it to begin with."* P01 stated that for rho, *"It's overly technical, and you wouldn't want to explain this to a general audience."* This theme extended to the unit of privacy, deployment models, mechanisms, privacy semantic interpretations, and algorithm hyperparameters, which several experts considered too intricate to be broadly interpretable.

Five experts noted the difficulty in interpreting and comparing privacy units, particularly in complex settings like DPML. P09 explained, *"... If you're doing DPSGD, the privacy unit is harder to describe in human understandable terms."*

Similarly, experts noted that the privacy semantic interpretations are too complex to understand, making them not useful in practice, as many users are unlikely to invest the time to fully grasp them. Furthermore, experts also cautioned that mechanism details are too complex for a non-technical audience. P08 added, *"It makes sense for engineers but not if you're sharing it with everybody's grandma."* This complexity risks alienating general audiences.

For algorithm hyperparameters, experts warned that too many technical details can overwhelm non-experts. As P06 said, *"If they are ML like LLM practitioners but are not laypersons... they would care because just the size of the gradient tells you how much signal you're getting."* In addition, explaining deployment models to non-technical audiences can be challenging, as the nuances between models may be hard to convey clearly. Experts pointed out concerns about information overload when presenting multiple models without prioritization, making the disclosure technical and difficult to understand.

Risk of Misinterpretation and Overemphasis on Numbers: Several experts raised concerns about overly focusing on privacy parameters and numeric values, which can lead to misleading comparisons. P07 explained, *"It's much easier to compare numbers... and people would do that and get maybe wrong conclusions."* P05 added that *"Discussions around DP just became about epsilon... foreclosing all those other issues."* Without clear interpretive context, these values may obscure rather than clarify privacy risks.

Information Overload and Redundancy: Experts warned that including too many technical details could overwhelm users. P07 suggested that advanced parameters be included only in supplementary materials: *"That information benefits only experts, who can check the white paper for details."* P05 observed that parameters like delta are often already set by consensus and may not require user-facing disclosure. *"There's more consensus around delta, so maybe including it is just giving someone more information than they need."* Other experts highlighted the redundancy of presenting algorithm hyperparameters when the overall privacy guarantee is already disclosed. As P08 put it, *"There is already enough information through the privacy guarantee itself, so disclosing clipping bounds again might be redundant."*

Limited Relevance and Potential for Harm: Utility information presented its own challenges. Some experts argued it has limited value for privacy-conscious users or policymakers. P09 remarked, *"I don't think it's that important for someone who cares about their privacy."* P03 warned that revealing data thresholds could discourage participation: stating, *"If I'm rare enough, I might not fill this out at all because it already told me it's pointless."* Other experts cautioned that organizations could misuse utility metrics to justify weaker privacy protections. P05 noted, *"I wonder if organizations would use it to manipulate people... saying, 'Well, yeah, you have to give up your privacy for better utility."*

Lack of Standards and Complexity of Implementation: Experts considered utility information and empirical privacy metrics difficult to standardize or measure. P10 stated that utility information *"is difficult to characterize and will vary greatly from dataset to dataset."* Experts noted that although it is nice to have empirical privacy metrics to provide a sense of practicality, there is no standard measure available to do so that addresses all user concerns from every perspective. P11 expressed concern about 'privacy theater', since empirical metrics often focus on average-case performance, potentially downplaying worst-case attack scenarios. *"Even if all of your per-attribute privacy budgets are big, every theoretical guarantee you can provide is almost trivial. Then you need empirical attacks to show that you did something useful."* This highlights the need for careful interpretation of results. Experts also noted that it is infeasible to

implement and describe all possible attacks in a way that caters to concerns of all the stakeholders. Metrics that are not on the list can still show up in a way that becomes a source of concern for privacy. *Uncertain Practical Utility:* Some experts questioned the real-world utility of privacy parameters. P11 remarked, *"In practice, it's close to useless... we end up with very large values... for which the guarantees are also not comforting."* This highlights the gap between theoretical guarantees and practical privacy risks.

4.3 Typical Ranges for Parameters

This subsection summarizes expert perspectives on typical or context-appropriate ranges for DP parameters, when applicable. Experts emphasized that parameter values vary widely depending on factors like data sensitivity, dataset size, and use case, and that many parameters do not have a standard range. For a subset of parameters, they offered ranges supported by community norms. Because each parameter involves different trade-offs, we present ranges on a parameter-by-parameter basis, focusing only on those for which experts provided concrete numerical guidance. Parameters such as the deployment model or mechanism are not included here, as they do not have ranges. Table 3 summarizes expert-suggested value ranges for key DP parameters.

Epsilon: Experts agreed that ϵ should be context-dependent, with common values ranging from **0.001 to 20**. Many (P01, P04, P05, P09, P11) recommended $\epsilon \leq 4$ as ideal, though higher values are common in industry. P04 explained, *"Guidance suggests 1 or 0.5, but we've seen higher in industry."* P03 and P06 viewed values up to 10 as reasonable; P06 noted, *"Below 1 means high privacy... above 10 is basically no privacy."* For high-dimensional data, P08 found values up to 20 acceptable.

Delta: Experts agreed δ should be **very small**, often 10^{-5} to 10^{-6} (P01). P03 suggested $1/\sqrt{n}$; P08 preferred $1/(\log n \cdot n)$. P05 favored extremely small values to ensure strong guarantees. P01 also noted that larger datasets may demand even smaller deltas for meaningful privacy protection.

Unit of Privacy: Most experts supported user-level or individual-level units (P04, P08, P09, P10). P08 noted, *"A normal range is anything that is individual level or a proxy for an individual."* Others (P03, P05) emphasized the need for stakeholder-specific definitions. P09 proposed time-based units like *"one day to one month or a quarter."*

Utility Information: Utility needs varied by use case. P03 linked it to group sizes: *"100 is reasonable, 50 is good, 10 is problematic."* P08 stressed that utility depends on context. P09 recommended relative error $\leq 5\text{--}6\%$ to avoid misleading data.

Algorithm Hyperparameters: Experts had diverse views. P04 emphasized balancing bias and variance; P06 cited clipping norms from 0.1 to 1. P09 highlighted output thresholds for large datasets.

Empirical Privacy Metrics: P06 recommended percentiles and averages for risk assessment: *"Gives a clearer picture."* But cautioned that averages may miss individual risks.

Privacy Interpretations or Semantics: For privacy interpretations, experts referenced Type 1 error rates (0.01, 0.05, 0.1) and Type 2 power ranges (0.1 to 0.7). These offer intuitive, bounded guarantees, but may not apply when data is already public or low-risk.

Epsilon: Experts agreed that epsilon values vary by application, data sensitivity, and risk tolerance, with suggested thresholds from

Parameter	Typical Range / Guidance
Epsilon	0.001 – 20 (ideal ≤ 4)
Delta	10^{-5} to 10^{-6} ; or $\frac{1}{\sqrt{n}}$
Unit of Privacy	User-level; time-based (day-quarter)
Utility Info	Group size $\geq 50 – 100$; error $\leq 5 – 6\%$
Algorithm Hyperparameters	Clip norm: 0.1–1; tuned noise scale
Empirical Metrics	Percentiles, std, min/max risk
Privacy Semantics	Type I error: 0.01–0.1; Power: 0.1–0.7

Table 3: Examples of expert-recommended typical ranges for DP parameters.

0.001 to 20. Five experts (P01, P04, P05, P09, P11) considered $\epsilon \leq 4$ common or ideal, though industry settings often use larger values. P04 explained, "...general guidance suggests keeping epsilon at most 1 or 0.5, but it varies with data sensitivity...we've seen higher values in industry." For flexible applications, two experts (P03, P06) found values up to 10 reasonable. P06 noted, "Below 1 means high privacy, 1 to 10 is medium, but above 10 is basically no privacy." For high-dimensional or synthetic data, P08 suggested thresholds up to 20 stating, "For low-dimensional data, anything above 1 or 2 is too much, but for high-dimensional data...even up to 20 might be acceptable." These variations highlight the challenge of a universal epsilon range and the need for context-sensitive evaluations.

Delta: Experts agreed that delta should be **very small**, though specific values varied. P01 suggested 10^{-5} to 10^{-6} as common practice, stating, "We just want delta to be really small." P03 and P08 emphasized its dependence on dataset size. P03 cited $1/\sqrt{n}$ as a practical guideline, while P08 proposed $1/(\log n \cdot n)$. P05 favored extremely small values (e.g., ".000X") to maintain privacy while improving accuracy. P01 also noted that larger databases may require even smaller deltas for meaningful privacy protection.

Unit of Privacy: Experts emphasized tailoring the unit of privacy to specific contexts. Four experts (P04, P08, P09, P10) favored a user-level focus, with P08 stating, "A normal range is anything that is individual level or a proxy for an individual." P03 and P05 highlighted stakeholder-driven, context-dependent ranges. P09 proposed time-based units, suggesting bounds from "one day to one month or a quarter."

Utility Information: Experts emphasized that utility metrics vary based on data characteristics and intended use. P03 phrased utility in terms of the sizes of groups which can be reliably measured in a DP data release, with group sizes of 100 corresponding to good utility, and 50 still acceptable but 10 problematic: "If you hit 100, you're reasonable. If you hit fifty, you're good." P08 stressed its dependency on use cases, stating, "For utility, it depends on how the data's going to be used... some may accept lower utility, while for others, high utility is critical." P09 suggested a median relative error of $\leq 5\text{--}6\%$, warning against excessive hallucinated data.

Algorithm Hyperparameters: Experts have varied views on normal ranges. P04 emphasized a good balance between bias and variance for noise scale, while P06 mentioned heuristic ranges for clipping

norms, typically between 0.1 and 1. P09 stressed output thresholds' role in maintaining data quality, especially for large datasets.

Empirical Privacy Metrics: P06 suggested using percentiles, minimums, maximums, and standard deviations to quantify risk, stating, "it helps to quantify the risk in terms of percentiles... giving people an average gives a clearer picture." They also cautioned that averages may not reflect individual experiences.

Privacy Interpretations or Semantics: For frequentist interpretations, the normal range for Type 1 errors is conventionally set at values like 0.01, 0.05, or 0.1. Power (Type 2 error) is considered meaningful between 0.4 and 0.7, and strong between 0.1 and 0.3. The range provides interpretable guarantees, since statisticians are used to thinking of uncertainty in terms of Type 1 and Type 2 error.

4.4 Target Audiences

To understand for whom each parameter matters, we asked experts whether they considered each parameter to be "very important," "somewhat important," or "not important" for the general public, and then repeated the question for technical users and data analysts. We also asked them to explain why.

These questions enabled us to capture expert opinions on how different stakeholders, both general and technical audiences, could benefit from seeing each parameter on a potential DP label. We then synthesized their responses and areas of consensus and disagreement on the appropriate target audience(s) for each parameter. We organize our synthesis into three categories of audience relevance: (1) parameters considered relevant to both general and technical audiences; (2) parameters viewed as primarily relevant to technical users; and (3) parameters with mixed or contested relevance. To simplify reporting, we grouped expert responses into two categories: responses of "very important" and "somewhat important" were combined to indicate that a parameter was considered important; responses of "not important" were kept distinct. This approach allowed us to summarize consensus on whether a parameter was viewed as worth communicating to a particular audience.

Table 4 summarizes expert recommendations on the relevance of each parameter for different audiences. Each icon reflects the level of consensus among experts, and mention counts are provided in parentheses for context.

Parameter	General Audience	Technical Audience
Epsilon	✓ (8 yes, 2 no)	✓ (9 yes, 1 no)
Delta	✗ (1 yes, 2 no)	✓ (3 yes)
Unit of Privacy	✓ (6 yes, 2 no)	✓ (7 yes)
Utility Information	★ (2 yes, 3 no)	✓ (5 yes)
Mechanism	✗ (4 no)	✓ (4 yes)
Algorithm Hyperparameters	✗ (4 no)	✓ (4 yes)
Deployment Model	✓ (3 yes)	✓ (3 yes)
Empirical Metrics	✓ (2 yes)	✓ (2 yes)
Privacy Semantics	✗ (1 no)	✓ (1 yes)

Table 4: Expert recommendations on the relevance of each parameter for different audiences. (✓ indicates "relevant"; ✗ indicates "not relevant"; ★ indicates mixed views on relevance)

4.4.1 Parameters Relevant to Both Audiences

Epsilon: For the general public, 4 experts (P01, P05, P08, P12) viewed epsilon as very important, with P01 stating, "*epsilon would be very important, it controls what is at the crux of the privacy guarantees.*" Another 4 experts (P04, P06, P07, P09) considered it somewhat important, and 2 experts (P03, P11) said it was not important, citing its technical complexity. For data analysts and technical audience; 7 experts (P01, P03, P04, P05, P06, P08, P12) considered epsilon very important, with P05 noting, "*Epsilon is crucial for analysts as it's their surface of interaction with DP, helping them allocate it across queries.*" Two other experts (P07, P09) viewed it as somewhat important, and one expert (P11) noted that it was not important.

Unit of Privacy: This parameter received the most consensus, with most experts agreeing it matters for both audiences. For the general public; 6 experts (P04, P05, P07, P08, P09, P10) considered the unit of privacy very important, with P07 stating, "*It is the one technical thing I would like the general public to understand.*" Experts (P03, P12) considered it not important, arguing it is too complex. For data analysts and technical audience; 7 experts (P03, P04, P05, P07, P08, P10, P12) found it very important, with P07 stating that "*for data analysts and technical users, what we are protecting is very important.*"

Deployment Model : Three experts (P01, P03, P05) identified deployment models as very important for both general and technical audiences. P01 emphasized its significance for public understanding of data flows and privacy risks stating that "*It really changes things like who can see your data at various stages.*" The same experts also highlighted its relevance for technical audiences in assessing utility, accuracy, and threat models.

Empirical Privacy Metrics: Despite fewer mentions, those who addressed this parameter saw it as meaningful for both groups. For the general public, 2 experts (P06, P11) viewed these metrics as very important, citing that real-world attacks are more understandable than abstract guarantees. P11 noted, "*Empirical attacks are easier to explain... they feel more concrete because they involve real data rather than abstract mathematical concepts.*" For data analysts and technical audiences, P06 and P11 described empirical privacy metrics as very important, emphasizing that even technical audiences need empirical validation.

4.4.2 Parameters Primarily Relevant to Technical Audience

Delta: For the general public, only one expert (P01) considered delta somewhat important noting, "*it's not really driving things, but it's sort of like important that it isn't in the bad case.*" Two experts (P03, P05) said it was not important, with P05 stating, "*If there's agreement around what it should be and it's set very small, then there's no need to convey it.*" For technical audiences, 3 experts (P01, P03, P05) considered delta very important, especially for data curators.

Mechanism: For the general public; 4 experts (P01, P08, P10, P12) agreed mechanisms are not important to the general public, with P01 stating, "*This is more technical content than they are able to internalize and process, I think it will add more confusion than it clarifies.*" For data analysts and technical audiences; Some experts

considered it very (P10 & P12) or somewhat important (P01 & P08). P12 emphasized, "*The actual mechanism being used would be particularly important to them... for example, the scale term for Gaussian or Laplace noise.*"

Algorithm Hyperparameters: Four experts (P04, P06, P08, P09) considered algorithm hyperparameters very important for analysts and technical audiences but not relevant for the general public. As P04 noted, "*the public doesn't need details on noise scale or distribution,*" while P06 emphasized that such parameters "*help you debug.*"

Privacy Semantics: P11 viewed privacy semantics as very important for analysts and technical audiences—especially policymakers and privacy professionals—but not important for the general public, explaining that, "*There are a lot of limitations on what we can explain about privacy guarantees to the general public in a way that's useful,*" but added that, "*frequentist interpretations are very important... especially for the privacy community.*"

4.4.3 Parameters with Mixed or Contested Relevance

Utility Information: Expert opinions on the relevance of utility information varied widely. For the general public, two experts (P03, P05) considered it very important, emphasizing that people care about accuracy. P05 stated, "*People care about accuracy too.*" In contrast, three experts (P08, P09, P10) found it not important, with P10 remarking that, "*Technical users will care a lot about the utility.*" For technical users, four experts (P03, P05, P09, P10) rated utility as very important, while P08 considered it somewhat important. P05 explained, "*Accuracy implications are very important for analysts... They have an interest in knowing to what extent the results of their analysis will helpfully inform that real-world purpose.*"

4.5 Parameter-Specific Presentation Formats

We asked experts how each parameter could be more effectively presented within a DP label. Their feedback included specific suggestions for visual encoding—such as text, numbers, or range indicators—as well as preferences for placing parameters in the primary or secondary layer of a two-layer label. These expert insights were based on structured interview questions about where each parameter should appear and how it should be displayed.

To synthesize these insights, we categorized suggestions into three types of presentation formats; text explanations (e.g., plain-language descriptions), numbers (e.g., raw values like $\epsilon = 1.0$), and range indicators (e.g., color-coded or high/medium/low ranges). Table 5 summarizes expert suggestions by parameter, indicating the recommended formats, how many experts supported each, and the preferred layer of presentation.

Importantly, despite offering these concrete suggestions, many experts acknowledged the uncertainty about how best to present each parameter—underscoring the complexity of DP communication and the need for continued research to refine how each parameter is best presented. As P01 reflected "*There may not necessarily be one single unified explanation for each of these parameters, but I do think it should include some type of explanation so that users can interpret it.*" This section does not prescribe a single best format

but instead highlights areas of consensus and divergence that can inform future design iterations and empirical testing.

Text: Experts recommended text explanations for abstract parameters like unit of privacy, mechanism, deployment model, and empirical privacy metrics. P03 stated, “*a simple text explanation of what aspect of the data is being protected*” is best for the unit of privacy. P01 suggested written descriptions for mechanisms stating “*researchers can understand more, you can just tell them something about the algorithm.*” For the deployment model, P01 and P03 also recommended plain-language text to describe where data resides or flows, e.g., “*Does the data leave your device?*” P11 supported using text to present empirical metrics, noting that empirical privacy losses are easier to communicate in natural language than with formulas or numeric ranges.

Numbers: Numeric formats—such as displaying raw values like $\epsilon = 1.0$ —were primarily suggested for epsilon, delta, and privacy semantics. Three experts (P01, P05, P11) explicitly recommended numeric presentations. For epsilon, P01 and P05 suggested numeric values be shown, especially for technical users. However, P01 noted that numerical values alone may lack meaning without accompanying textual explanations. For delta, P01 indicated that a numerical value could be helpful when placed in context. P11 recommended numeric interpretations for privacy semantics to aid expert audiences, particularly for explaining frequentist interpretations and policy-relevant guarantees.

Indication of range: Color-coded or categorical range indicators (e.g., “high,” “medium,” “low”) were also recommended by experts for parameters like epsilon, delta, and utility information. Four experts (P01, P03, P06, P07) suggested that abstracting numeric values into ranges or visual symbols could help users interpret privacy strength with ranges calibrated to different models or data releases. P06 recommended a scheme “*like the energy consumption ABC kind of thing*” to make privacy more intuitive. P01 proposed applying categorical color labels to delta values—“*You can do colors like high, medium, low delta or something.*” For utility information, P03 inferred that users might better understand data usability if given an approximate group size, e.g., “small,” “medium,” or “large,” perhaps aided by graphic elements.

Beyond format types, a few experts shared interface-level suggestions for improving how parameters are accessed and prioritized. P05 and P09 recommended interactive display features—such as drop-downs or clickable links—to let users expand or drill down into detailed information. P04 emphasized that the *order* of parameter presentation also matters, suggesting that “*unit of privacy and privacy loss parameters should appear at the top, with relaxations listed below.*”

4.6 Expert Recommendations for the DP Label

We sought expert recommendations on the overall design of our proposed DP label. In general, we asked open-ended questions to solicit independent suggestions for how to design a label, with one exception: we proposed a two-layer DP label structure inspired by [24] during the interview, to solicit expert feedback on it; experts did not independently suggest a two-layer design. Wherever

Parameter	Text	Numbers	Range Indication	Primary Layer	Second Layer
Epsilon	•
Delta	•	•	•	•••	—
Unit of Privacy	...	—	—	•••••	—
Mechanism	•	—	—	•	•••
Deployment Model	..	—	—	•••	—
Utility Info	•	—	•	•••	•••
Empirical Metrics	•	—	—	••	•
Privacy Semantics	•	•	—	—	•
Algorithm Hyperparameters	—	—	—	—	•••

Table 5: Expert Suggestions for Presentation Format and Layer Placement (by Parameter). One dot (•) represents one expert suggestion.

possible, we clarify whether a suggestion aligned with or extended beyond our proposed label structure.

Standardized Contents and Formats: Eleven experts praised our idea of a DP label with **standardized contents** for DP communication. As P08 noted, “*Every data release provides different parameters, having a standard would make people’s lives so much easier.*” Five experts (P03, P06, P09, P11, P12) emphasized the importance of **standardized formats** aligned with existing tools and workflows. P09 recommended using Markdown format for broader accessibility because “*markdown is easily readable in a GitHub or GitLab repository*”. P06 suggested integrating DP labels for machine learning data releases into existing practices, like model cards on the Hugging Face platform, to “*ensure consistency and visibility*”. P12 highlighted the importance of uniformity, suggesting the label should be “*recognizable across different formats... so people know right away it’s the formal DP label.*”

Layered Structure: Experts unanimously supported our proposed layered structure for DP labels. Ten experts favored the two-layer structure proposed by our team: The primary layer provides high-level summaries for non-experts, while the secondary layer offers technical details for advanced DP users. P05 emphasized the value of this approach, stating, “*The primary layer helps people understand the space of things they might need to know... researchers can click to access more details.*” Two experts (P01, P11) proposed adding a **third, intermediary layer** to bridge the gap between technical users and the general public (i.e., semi-technical audience). P11 noted the need for differentiation, explaining, “*There’s a significant difference between a demographer and the general public... an intermediate category could address this.*” Two experts (P05, P09) desired an **interactive interface** to easily navigate different layers. P05 described a structure with “*a high-level primary layer and the ability*

to drill down for more details." There is an expert consensus that layered structure can balance the DP label's accessibility and depth. *Placement for Accessibility, Visibility, and Relevance:* When asked about where to place the DP labels, experts offered various recommendations. Seven experts (P03, P06, P08, P09, P10, P11, P12) favored placing the label **alongside data releases**. P10 noted, "*It makes sense to stick the label on the website where the dataset lives and describes the release.*" Another seven experts (P01, P03, P04, P05, P09, P10, P12) advocated **embedding the label directly into apps or services** where user interactions happen. P05 recommended placing the label in privacy centers within apps, stating that "*a [DP] label could be the first summary users see.*" P04 suggested interface areas like app permissions or URL bars to increase visibility. Three experts (P05, P06, P12) recommended housing labels in **dedicated databases or registries** to provide centralized access. P12 noted, "*Having like a secondary database too, we could look these things up.*" Furthermore, experts recommended including DP labels in **blog posts, press releases, or documentation** about specific data releases to improve transparency (P01, P07, P11), and integrating DP labels for machine learning data releases into existing formats like **model cards** (P06, P09) for relevance. These recommendations reflect a shared view that the placement of DP labels should maximize accessibility, visibility, and relevance.

Balance Between Simplicity and Education: P04 believed that a simple label can help adoption, advising "*Keep it simple so that it's easy to adopt.*" Experts (P01, P03, P11) also suggested that visual elements can simplify complex information, with P01 recommending "*something picture-based or color-based, but not overly technical.*" Other experts (P10, P12) felt that DP labels should provide adequate DP education to assist user comprehension. P10 commented, "*Education is key to making the label understandable.*" P06 wanted the DP label to prevent potential misinterpretations, suggesting "*warnings... may be needed for cases where the label doesn't account for issues like data correlation.*" Overall, experts expressed different opinions on how to maximize the understandability of DP labels while keeping the label information accurate and rigorous, surfacing the need to balance between simplicity and education.

As a starting point for future work on DP label design and validation, we include our initial design of the expert-informed DP label in Appendix E.

5 Discussion

Effective communication of differential privacy (DP) remains a significant challenge. Previous research focused on how to better explain obviously important parameters (e.g. epsilon (ϵ)) to non-expert audiences [50], without considering the fact that explaining one or two DP parameters in isolation cannot adequately convey the real-world privacy protection of DP deployments. Our study with DP experts addresses this gap by focusing on **what** parameters are essential to accurately describe the privacy guarantees of DP deployments. It also partially tackles **how** to properly explain DP parameters through the initial design of a standardized DP label. We discuss our contributions in relation to prior work, as well as open questions in improving DP communication as follows.

Towards Comprehensive DP Communication Standards: Unlike existing work explaining isolated parameters [50], our interview study

with DP experts seeks to holistically communicate the privacy guarantees of DP deployments. As detailed in 4.1, we identified a comprehensive list of parameters necessary to accurately convey the real-world privacy protection of a given DP deployment. Particularly, this list contains several critical, yet often overlooked parameters in 4.1.10. We also articulated the specific considerations for each of these parameters, including their importance in determining privacy guarantees, possible normal ranges, target audiences, and methods of presentation. Our work provides the DP community with guidance on what should be included to describe the privacy guarantees of DP deployments, laying the foundation for future standardization efforts to improve DP communication.

Privacy Labels to Improve DP Communication: Inspired by the approach of privacy labels [39, 40] to standardize privacy information communication, we designed a specialized privacy label for DP by synthesizing expert feedback from this study (Appendix E). Our initial label design draws from prior work [5, 17] to incorporate multiple effective communication techniques (e.g. text explanations, visual elements) to improve the clarity and accessibility of DP information. This DP label includes two layers similar to the proposed privacy label for IoT [23]: simple summaries for the general public and detailed explanations for more technical users; and interactive interface to help users navigate between layers. This two-layer design particularly addresses the needs of our target audience – technical users who require technical details to comprehensively assess the privacy implications of a given DP deployment.

Bridging the Gap Between Theory and DP Deployments: Many existing DP explanations rely on theoretical guarantees, often abstracted in ways that make them difficult to interpret in real-world scenarios [22, 58]. For example, risk assessments often omit key factors such as dataset size, adversarial capabilities, or the impact of repeated queries. Practitioners find it challenging to gauge actual privacy risks. Our study bridges this gap by emphasizing empirical privacy metrics and other contextually relevant parameters in our DP label, helping practitioners assess DP performance in practical deployments. This enhances DP communication in operational settings, ensuring that privacy guarantees are both theoretically sound and practically meaningful.

Enhancing Transparency and Trust in DP Deployments: Technical transparency is crucial for building trust and promoting DP adoption. Current projects like the Oblivious Privacy Deployments Registry [59] and Wikimedia Foundation's DP dataset documentation [28] primarily record privacy parameters. However, these documents often lack clear explanations about key aspects like privacy units and post-processing effects. Desfontaines' informal list [19] adds more properties to increase transparency, but still misses some important parameters. As Cummings et al. [18] point out, transparency without clear communication can backfire, leading to misunderstandings and decreased trust. We build on these works by focusing on both transparency and interpretability. Our proposed DP label records essential DP parameters and presents them in an accessible way. This approach connects theoretical DP guarantees with practical deployments and supports responsible and informed adoption of privacy-preserving technologies.

Open Question: Setting Parameter Ranges: The question of how to set ϵ is as old as DP itself [21], and experts in our study identified

similar issues with many other parameters. Appropriate settings for many parameters depend on context, making the problem even more challenging. The development of evidence-based guidance for how to set these parameters is an open research question. Approaches based on theory alone [45] have not yielded satisfying answers, and we believe that the most promising avenues would consider parameter settings in the social context of the associated deployment—for example, by taking into account utility requirements and adversary models specific to the deployment [16].

Open Question: How to Communicate: This study focuses on *what* to communicate (rather than *how*). In an attempt to answer both questions at once, we sought experts' recommendations for communicating DP parameters to both technical and general audiences. Unfortunately, the major point of consensus among the experts we interviewed was: *how to communicate with the general audience* remains a major open question. A significant amount of prior work has studied how to communicate key parameters to end users [5, 9, 29, 38, 44, 50, 69, 73–75]; we hope our work will motivate similar studies for other important DP parameters.

6 Limitations and Future Work

Stakeholders in DP deployments. Our study focused on a single group of stakeholders: DP experts who may want to use a label to evaluate the technical strength of the DP guarantee associated with a deployment. Typical DP deployments include many different stakeholders with different goals and levels of expertise, including data subjects, data curators, data analysts, system engineers, compliance and legal officers, regulators, policy and decision makers, and auditing or watchdog groups. Since each group of stakeholders represents a different population, the insights from our study likely do not generalize to all of the other groups. Further research is needed to identify and address the needs of each group, and to inform methods for ensuring transparency with each of them.

Study design and recruitment. Our qualitative study design and recruitment strategy produced rich insights about transparency in DP deployments, but the context and design of our study also has inherent limitations. Our choice of a qualitative study was motivated by the desire to enumerate important parameters for transparency in DP guarantees—since no such list previously existed, a structured, quantitative study on the relative importance of these parameters was not possible. In addition, recruiting DP experts for any type of study is challenging, since the pool of experts in this specialized field is relatively small; recruiting enough experts to obtain meaningful quantitative results is thus not feasible. Our purposive and snowball sampling may have introduced selection bias, but was necessitated by the challenges of recruiting DP experts.

Context of DP deployments. Full transparency in DP deployments requires more than communicating important parameters of the DP guarantee. For example, central-model deployments may be subject to data breaches, so real-world privacy rests on security of the stored data in addition to the parameters of the DP guarantee. The blurred line between DP parameters and other system properties affecting real-world privacy was reflected in a focus on threat modeling and deployment models in our study. A complete approach for *context-inclusive* transparency in DP deployments will likely need to integrate with frameworks like like contextual integrity [54, 55] and

human-centric approaches for enumerating and addressing privacy threats like the NIST Privacy Risk Assessment Methodology [60].

Label design and validation. Our study sought to identify the important aspects for transparency in communicating DP guarantees, and can serve as a starting point for designing DP privacy labels for this purpose. We plan to employ the qualitative participatory design method [67] with a representative sample of target users of a DP privacy label, including experts, data users, end users, and other stakeholders discussed earlier. We plan to further iterate the DP label design through large-scale evaluation studies with diverse user groups, with the goal to accommodate the communication needs of the various DP stakeholders.

7 Conclusion

Through an in-depth interview study with 12 DP experts, we identified the important parameters to communicate privacy guarantees of DP deployments, articulated the key considerations for each parameter, and discussed approaches to improving DP communication such as our proposed DP label. Our findings highlight the need for standards in DP disclosures and caution against over-reliance on single parameters like ϵ . This study lays the groundwork for a standardized approach to communicating the privacy guarantees of DP deployments, fostering greater transparency and trust among DP stakeholders.

Acknowledgments

The authors would like to thank the anonymous reviewers at PETS and SOUPS for their helpful suggestions. This material is based upon work supported by the National Science Foundation under Grant No. 2238442 and No. 2336550. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

References

- [1] John M Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2867–2867, 2018.
- [2] John M Abowd, Robert Ashmead, Ryan Cummings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, et al. The 2020 census disclosure avoidance system top-down algorithm. *Harvard Data Science Review*, 2, 2022.
- [3] Apple. Apple: Differential Privacy Overview, 2023. https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf.
- [4] D Apple. Learning with privacy at scale. *Apple Machine Learning Journal*, 1(8):71, 2017.
- [5] Narges Ashena, Oana Inel, Badrie L Persaud, and Abraham Bernstein. Casual users and rational choices within differential privacy. In *IEEE Symposium on Security and Privacy. Proceedings*, pages 87–87. Institute of Electrical and Electronics Engineers, 2024.
- [6] Aleix Bassolas, Hugo Barbosa-Filho, Brian Dickinson, Xerxes Dotiwalla, Paul Eastham, Riccardo Gallotti, Gourab Ghoshal, Bryant Gipson, Surendra A Hazarie, Henry Kautz, et al. Hierarchical organization of urban mobility and its connection with city livability. *Nature communications*, 10(1):4817, 2019.
- [7] Aleix Bassolas, Hugo Barbosa-Filho, Brian Dickinson, Xerxes Dotiwalla, Paul Eastham, Riccardo Gallotti, Gourab Ghoshal, Bryant Gipson, Surendra A Hazarie, Henry Kautz, et al. Reply to: On the difficulty of achieving differential privacy in practice: user-level guarantees in aggregate location data. *Nature Communications*, 13(1):30, 2022.
- [8] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- [9] Brooke Bullek, Stephanie Garboski, Darakhshan J Mir, and Evan M Peck. Towards understanding differential privacy: When do people trust randomized response

- technique? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3833–3837, 2017.
- [10] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of cryptography conference*, pages 635–658. Springer, 2016.
- [11] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer. Membership inference attacks from first principles. In *2022 IEEE symposium on security and privacy (SP)*, pages 1897–1914. IEEE, 2022.
- [12] Andrew Chin and Anne Klinefelter. Differential privacy as a response to the reidentification threat: The facebook advertiser case study. *NCL Rev.*, 90:1417, 2011.
- [13] Aloni Cohen, Moon Duchin, J Matthews, and Bhushan Suwal. Private numbers in public policy: Census, differential privacy, and redistricting. *Harvard Data Science Review*, (Special Issue 2), 2022.
- [14] Aloni Cohen and Kobbi Nissim. Towards formalizing the gdpr’s notion of singling out. *Proceedings of the National Academy of Sciences*, 117(15):8344–8352, 2020.
- [15] Lorrie Faith Cranor. Mobile-app privacy nutrition labels missing key ingredients for success. *Communications of the ACM*, 65(11):26–28, 2022.
- [16] Rachel Cummings, Shlomi Hod, Jayshree Sarathy, and Marika Swanberg. Attaxonomy: Unpacking differential privacy guarantees against practical adversaries. *arXiv preprint arXiv:2405.01716*, 2024.
- [17] Rachel Cummings, Gabriel Kapchuk, and Elissa M Redmiles. “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*, pages 3037–3052, 2021.
- [18] Rachel Cummings and Jayshree Sarathy. Centering policy and practice: Research gaps around usable differential privacy. In *2023 5th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pages 122–135. IEEE, 2023.
- [19] Damien Desfontaines. A list of real-world uses of differential privacy. <https://desfontain.es/blog/real-world-differential-privacy.html>, 2021. Originally published: 2021-10-01, Last updated: 2024-09-13, Accessed: 2024-09-17.
- [20] Cynthia Dwork, Nitin Kohli, and Deirdre Mulligan. Differential privacy in practice: Expose your epsilons! *Journal of Privacy and Confidentiality*, 9(2), 2019.
- [21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4–7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.
- [22] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [23] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. Ask the experts: What should be on an iot privacy and security label? In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 447–464. IEEE, 2020.
- [24] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. An informative security and privacy “nutrition” label for internet of things devices. *IEEE Security & Privacy*, 20(2):31–39, 2021.
- [25] Georgina Evans, Gary King, Margaret Schwengsfeier, and Abhradeep Thakurta. Statistically valid inferences from privacy-protected data. *American Political Science Review*, 117(4):1275–1290, 2023.
- [26] Giulia Fanti, Vasyl Pihur, and Úlfar Erlingsson. Building a rapport with the unknown: Privacy-preserving learning of associations and data dictionaries. *arXiv preprint arXiv:1503.01214*, 2015.
- [27] Jennifer Fereday and Eimear Muir-Cochrane. Demonstrating rigor using thematic analysis: A hybrid approach of inductive and deductive coding and theme development. *International journal of qualitative methods*, 5(1):80–92, 2006.
- [28] Wikimedia Foundation. Pageviews differential privacy – current. https://analytics.wikimedia.org/published/datasets/country_project_page/00_README.html, 2024. Accessed: 2024-09-17.
- [29] Daniel Franzen, Saskia Nuñez von Voigt, Peter Sörries, Florian Tschorsch, and Claudia Müller-Birn. 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*, pages 1125–1139, 2022.
- [30] Gonzalo Munilla Garrido, Xiaoyuan Liu, Florian Matthes, and Dawn Song. Lessons learned: Surveying the practicality of differential privacy in the industry. *arXiv preprint arXiv:2211.03898*, 2022.
- [31] Google. Google: Differentially private heatmaps, 2023. <https://blog.research.google/2023/04/differentially-private-heatmaps.html>.
- [32] Michael Hay, Ashwin Machanavajjhala, Jerome Miklau, Yan Chen, and Dan Zhang. Principled evaluation of differentially private algorithms using dpbench. In *Proceedings of the 2016 International Conference on Management of Data*, pages 139–154, 2016.
- [33] Shlomi Hod and Ran Canetti. Differentially private release of israel’s national registry of live births. *arXiv preprint arXiv:2405.00267*, 2024.
- [34] Sarah Holland, Ahmed Hosny, Sarah Newman, Joshua Joseph, and Kasia Chmielinski. The dataset nutrition label. *Data Protection and Privacy*, 12(12):1, 2020.
- [35] Julian Holzel. Differential privacy and the gdpr. *Eur. Data Prot. L. Rev.*, 5:184, 2019.
- [36] Florimond Houssiau, Luc Rocher, and Yves-Alexandre de Montjoye. On the difficulty of achieving differential privacy in practice: user-level guarantees in aggregate location data. *Nature communications*, 13(1):29, 2022.
- [37] Matthew Jagielski, Jonathan Ullman, and Alina Oprea. Auditing differentially private machine learning: How private is private sgd? *Advances in Neural Information Processing Systems*, 33:22205–22216, 2020.
- [38] Farzaneh Karegar, Alaa Alaqra, and Simone Fischer-Hübner. Exploring {User-Suitable} metaphors for differentially private data analyses. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 175–193, 2022.
- [39] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A “nutrition label” for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, 2009.
- [40] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: an online study of the nutrition label approach. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pages 1573–1582, 2010.
- [41] Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 us census. *Science advances*, 7(41):eabk3283, 2021.
- [42] Krishnaram Kenthapadi and Thanh TL Tran. Pripearl: A framework for privacy-preserving analytics and reporting at linkedin. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2183–2191, 2018.
- [43] Konrad Kollnig, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. Goodbye tracking? impact of ios app tracking transparency and privacy labels. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 508–520, 2022.
- [44] Patrick Kühtreiber, Viktoriya Pak, and Delphine Reinhardt. Replication: the effect of differential privacy communication on german users’ comprehension and data sharing attitudes. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 117–134, 2022.
- [45] Jaewoo Lee and Chris Clifton. How much is enough? choosing ϵ for differential privacy. In *Information Security: 14th International Conference, ISC 2011, Xi’an, China, October 26–29, 2011. Proceedings 14*, pages 325–340. Springer, 2011.
- [46] Microsoft. Microsoft AI: Differential Privacy, 2023. <https://www.microsoft.com/en-us/ai/ai-lab-differential-privacy>.
- [47] Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017.
- [48] J Tom Mueller and Alexia R Santos-Lozada. The 2020 us census differential privacy method introduces disproportionate discrepancies for rural and non-white populations. *Population research and policy review*, 41(4):1417–1430, 2022.
- [49] Priyanka Nanayakkara, Joes Bater, Xi He, Jessica Hullman, and Jennie Rogers. Visualizing privacy-utility trade-offs in differentially private data releases. *arXiv preprint arXiv:2201.05964*, 2022.
- [50] Priyanka Nanayakkara, Mary Anne Smart, Rachel Cummings, Gabriel Kapchuk, and Elissa M Redmiles. What are the chances? explaining the epsilon parameter in differential privacy. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1613–1630, 2023.
- [51] Milad Nasr, Reza Shokri, and Amir Houmansadr. Machine learning with membership privacy using adversarial regularization. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 634–646, 2018.
- [52] Joseph P Near and Chiké Abuah. Programming differential privacy. *URL: https://uvm*, 2021.
- [53] Joseph P Near, David Darais, Naomi Lefkowitz, Gary Howarth, et al. Guidelines for evaluating differential privacy guarantees. *National Institute of Standards and Technology, Tech. Rep.*, 2023.
- [54] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.
- [55] Helen Nissenbaum. Privacy in context: Technology, policy, and the integrity of social life. In *Privacy in context*. Stanford University Press, 2009.
- [56] Kobbi Nissim and Alexandra Wood. Is privacy privacy? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128):20170358, 2018.
- [57] Briony J Oates, Marie Griffiths, and Rachel McLean. *Researching information systems and computing*. Sage, 2022.
- [58] Jonathan A Obar and Anne Oeldorf-Hirsch. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society*, 23(1):128–147, 2020.
- [59] Oblivious. The privacy deployments registry. <https://registry.oblivious.com/>, 2024. Accessed: 2024-09-17.
- [60] National Institute of Standards and Technology. Nist privacy framework: a tool for improving privacy through enterprise risk management, 2020. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.01162020.pdf>.
- [61] Kevin Proudfoot. Inductive/deductive hybrid thematic analysis in mixed methods research. *Journal of mixed methods research*, 17(3):308–326, 2023.

- [62] Lucas Rosenblatt, Bernease Herman, Anastasia Holovenko, Wonkwon Lee, Joshua Loftus, Elizabeth McKinnie, Taras Rumezhak, Andrii Stadnik, Bill Howe, and Julia Stoyanovich. Epistemic parity: Reproducibility as an evaluation metric for differential privacy. *Proceedings of the VLDB Endowment*, 16(11), 2023.
- [63] Steven Ruggles, Catherine Fitch, Diana Magnuson, and Jonathan Schroeder. Differential privacy and census data: Implications for social and economic research. In *AEA papers and proceedings*, volume 109, pages 403–408. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2019.
- [64] Jayshree Sarathy, Sophia Song, Audrey Haque, Tania Schlatter, and Salil Vadhan. Don't look at the data! how differential privacy reconfigures the practices of data science. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2023.
- [65] Mary Anne Smart, Priyanka Nanayakkara, Rachel Cummings, Gabriel Kaptschuk, and Elissa Redmiles. Models matter: Setting accurate privacy expectations for local and central differential privacy. *arXiv preprint arXiv:2408.08475*, 2024.
- [66] Ana Sokolovska and Ljupco Kocarev. Integrating technical and legal concepts of privacy. *Ieee Access*, 6:26543–26557, 2018.
- [67] Clay Spinuzzi. The methodology of participatory design. *Technical communication*, 52(2):163–174, 2005.
- [68] U.S. Census Bureau. Why the Census Bureau Chose Differential Privacy, 2023. <https://www.census.gov/library/publications/2023/decennial/c2020br-03.html>.
- [69] Zikai Alex Wen, Jingyu Jia, Hongyang Yan, Yaxing Yao, Zheli Liu, and Changyu Dong. The influence of explanation designs on user understanding differential privacy and making data-sharing decision. *Information Sciences*, 642:118799, 2023.
- [70] Aaron R Williams, Joshua Snoke, C Bowen, and Andrés Felipe Barrientos. Disclosing economists' privacy perspectives: A survey of american economic association members' views on differential privacy and the usability of noise-infused data. *Harvard Data Science Review*, 2024.
- [71] Alexandra Wood, Micah Altman, Aaron Bembenek, Mark Bun, Marco Gaboardi, James Honaker, Kobbi Nissim, David R O'Brien, Thomas Steinke, and Salil Vadhan. Differential privacy: A primer for a non-technical audience. *Vand. J. Ent. & Tech. L.*, 21:209, 2018.
- [72] Aiping Xiong. Effect of facts box on users' comprehension of differential privacy: A preliminary study. In *Proceedings of the Human Factors and Ergonomics Society 2020 Annual Meeting*, 2020.
- [73] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. Towards effective differential privacy communication for users' data sharing decision and comprehension. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 392–410. IEEE, 2020.
- [74] Aiping Xiong, Chuhao Wu, Tianhao Wang, Robert W Proctor, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Using illustrations to communicate differential privacy trust models: an investigation of users' comprehension, perception, and data sharing decision. *arXiv preprint arXiv:2202.10014*, 2022.
- [75] Aiping Xiong, Chuhao Wu, Tianhao Wang, Robert W Proctor, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Exploring use of explanatory illustrations to communicate differential privacy models. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 67, pages 226–232. SAGE Publications Sage CA: Los Angeles, CA, 2023.
- [76] Shikun Zhang, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh. How usable are ios app privacy labels? *Proceedings on Privacy Enhancing Technologies*, 2022.

Appendix

Appendix A: Eligibility Survey

After displaying the IRB-approved consent form, the following eligibility questions were used to assess participants' qualifications for the study, collect informed consent, and gather demographic information. For questions that tested participants' understanding of DP concepts, the correct answers are highlighted in bold.

Eligibility Survey Questions

Consent/Basic Eligibility

I have read and understood the information above.

No/Yes

I want to proceed to complete the eligibility survey for this research study.

No/Yes

Are you at least 18 years old?

No/Yes

Do you currently reside in the United States?

No/Yes

Do you have at least five (5) years of research or practical experience in the field of differential privacy (DP)?

No/Yes

Are you willing to participate in a 1-hour remote usability study via Microsoft Teams to help design a differential privacy nutrition label?

No/Yes

DP Background

How many years of experience do you have working or researching in the field of Differential Privacy?

(a) 0-5 (b) 5-10 (c) More than 10

Which of the following best describes your current field of expertise?

(a) Academic Research (b) Industry Practice (c) Policy Making (d) Other

What is your current job role?

(a) Professor/Researcher in a University setting (b) Data Privacy Engineer (c) Policy Advisor/Consultant in Data Privacy (d) Data Scientist with a focus on Privacy Technologies (e) Other (please describe)

Releasing two differentially private statistics, one with $\epsilon_1 = 0.1$ and the other with $\epsilon_2 = 0.5$, results in a total privacy loss of:

(a) 0.1 (b) 0.5 (c) **0.6** (d) 0.05 (e) I don't know

In differential privacy, which value of the privacy parameter ϵ provides stronger privacy?

(a) **0.1** (b) 1.0 (c) I don't know

Demographics

What is your age?

(a) 18-24 (b) 25-29 (c) 30-34 (d) 35-39 (e) 40-44 (f) 45-49 (g) 50-54 (h) 55-59 (i) I prefer not to answer

What is your gender?

(a) Male (b) Female (c) Nonbinary (d) I prefer not to say

Do you hold a PhD in Computer Science with a focus on Differential Privacy or a related area?

Yes/No

Appendix B: Interview Guide

The following section outlines the interview guidelines, which are organized into three focus areas: 1) Identifying Key Differential Privacy (DP) Parameters, 2) Assessing the key considerations including - reason for parameter inclusion, typical/normal ranges, and relevance of each DP parameter to various audiences, and 3) Gathering recommendations for the presentation and placement of the DP Label.

Interview Questions

Focus 1: Identifying Key DP Parameters

Can you list the DP parameters that you think are essential for inclusion in a DP Nutrition Label?

In your opinion, is there a parameter whose importance is often underestimated or overlooked in discussions about DP? If so, can you elaborate on this?

Focus 2: Assessment of key considerations

Could you explain why you believe [parameter #i] should be included in the DP Nutrition Label?

Could you give a potential argument against including the parameter?

What do you consider to be its 'normal range' in the context of differential privacy?

Why do you think this range is appropriate for this parameter?

Are there circumstances under which this 'normal range' might not be applicable or effective? If so, could you describe these situations or give examples?

For the general public, do you consider [parameter i] to be 'Very Important,' 'Somewhat Important,' or 'Not Important'? Why?

For data analysts and other technical users, do you consider [parameter #i] to be 'Very Important,' 'Somewhat Important,' or 'Not Important'?

Focus 3: Structure, Presentation Placement and Design Recommendations

We are considering a layered structure for the DP Nutrition Label, where the primary layer displays the most important information for the general public and non-technical users. Users can then click a link or follow a QR code to access a secondary layer with more detailed information for experts, researchers, and technical practitioners. What are your thoughts on this proposed design with the primary and secondary layers?

Considering the layered structure of the DP Nutrition Label, please indicate on which layer you think [parameter #i] should be presented and provide your reasons.

How do you think the parameter(s) should be displayed on the label?

Where should the DP Nutrition Labels be placed to ensure they are easily accessible and visible to users?

What are your thoughts on how the DP Nutrition Label should be designed, including its layout?

Do you have any other recommendations for the label that we haven't discussed?

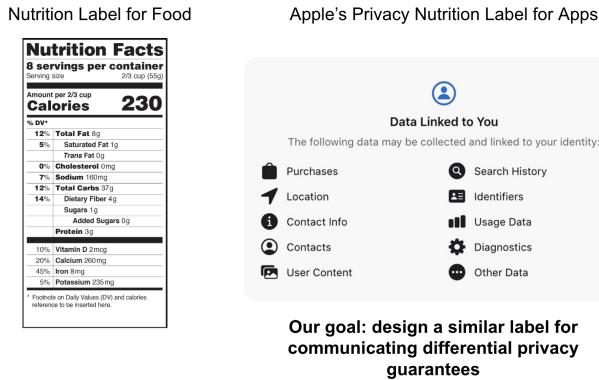


Figure 1: Example Labels

Appendix C: More Demographics

The following tables detail the expert demographics in terms of work experience, job roles, and gender, as well as the industries and countries they represent.

Gender		Role	
Male	7	Academic Researcher	4
Female	5	Data Privacy Engineer	4
Age		Country	
25-29	4	Data Scientist	3
35-39	3	Software Engineer	1
40-44	1	United States	11
55-59	1	Europe	1

Table 6: Expert demographics, $n = 12$

Appendix D: List of Parameters Mentioned by Experts

Here's the list of parameters as mentioned by the experts in our study. We use "/" for the same parameters but called different names by our expert participants.

Parameter Mentioned	Category	Description
Epsilon (ϵ)	Privacy Parameters	Core privacy parameter controlling the tradeoff between privacy and accuracy.
Delta (δ)	Privacy Parameters	Represents the probability of privacy failure in worst-case scenarios.
Privacy Loss Parameters (e.g., epsilon, delta)	Privacy Parameters	Additional parameters (e.g., rho for zCDP) used to quantify privacy loss.
Other parameters outside of epsilon, delta (e.g., rho for zCDP)	Privacy Parameters	Alternative privacy loss metrics like rho for zCDP or Rényi DP.
Composition (one query or many queries)	Privacy Parameters	Whether privacy guarantees apply to a single query or across multiple queries.
Query & Mechanism used / Process for DP being applied	Mechanism Used	Describes the method used to apply differential privacy (e.g., Laplacian or Gaussian noise).
Privacy Definition (e.g., Pure epsilon-DP vs zCDP)	Mechanism Used	Specifies the type of differential privacy applied (e.g., epsilon-DP or zCDP).
Deployment Model	Deployment Model	Describes how data is processed (centralized vs. decentralized), affecting privacy risks.
Unit of Privacy	Unit of Privacy	Defines the level of privacy protection, such as individual or group-level privacy.
Budget Replenishment / Budget Reset	Privacy Parameters	Describes whether the privacy budget is replenished or reset, impacting long-term privacy.
Time Frame of Guarantee (e.g., epsilon per day, per month)	Privacy Parameters	Defines the duration of the privacy guarantee (e.g., daily, monthly).
Parameters that affect utility (Noise distribution, Noise scale)	Utility Information	Defines the impact of noise and scale on the utility of the data after privacy is applied.
Utility Metrics	Utility Information	Quantifies how well the data retains quality after applying DP.

Parameter Mentioned	Category	Description	Parameter Mentioned	Category	Description
Worst-case vs average-case guarantee	Utility Information	Describes whether the privacy guarantees are based on worst-case scenarios or average cases.	Where is the Data Coming From / How is it Collected	Data Collection	Describes the origin and collection methods for the data.
What is considered public information vs not public information	Privacy Parameters	Defines what data is considered public versus private within the DP context.	What Data was Published / Who Was it Shared With	Data Sharing	Describes the published data and who has access to it.
Training Hyperparameters (Number of samples, etc.)	Algorithm Hyperparameters	Describes the parameters (e.g., number of samples, iterations) affecting privacy and utility.	What Software was Used / if Source Code is Available	Software and Implementation	Describes the software used for implementing DP and whether the source code is available.
Sensitivity /Noise multiplier / Clipping norm	Algorithm Hyperparameters	Key parameters that control how noise is applied and data is clipped during processing.	Publication Date / Version Number	Data Release	Describes when the data was released and the version number.
Utility Implications / Utility and/or bias	Utility Information	Describes how privacy parameters influence utility and potential bias in the data output.	Links to Additional Materials (e.g., whitepapers)	Additional Materials	Provides references to related materials or research papers that support the DP deployment.
Summary of Empirical Results from Privacy Attacks	Empirical Privacy Metrics	Provides results from attacks (e.g., MIA, reconstruction) assessing real-world privacy risks.	Goal of the Release	Data Release	Describes the intended and unintended uses of the data being released.
Non-DP Results and their impact	Empirical Privacy Metrics	Describes the impact of non-DP data on the privacy outcomes.			
Definition of Neighboring Datasets	Privacy Parameters	Clarifies what constitutes a neighboring dataset for differential privacy purposes.			
Universe of Possible Records	Privacy Parameters	Defines the set of potential records within the DP framework.			
Bounding User Contribution Across Releases	Privacy Parameters	Defines how user contributions are bounded across multiple data releases.			
Vetting of DP Implementation	Privacy Verification	Describes the process of verifying the correctness of a DP implementation.			
Interpretation via Frequentist or Bayesian Framework	Privacy Interpretation	Describes the statistical framework used for interpreting the privacy guarantees (frequentist vs Bayesian).			
Privacy for Natural Groups within the Data	Group Privacy	Ensures privacy protections apply to groups within the data, not just individuals.			

Table 7: Mentioned parameters

Appendix E: Prototype DP Label

Differential Privacy Deployment Facts	
Parameter	Value
Privacy Parameters	
 Privacy Budget Summary	2.0 ϵ total 
Description: Overall privacy budget allocation and usage.	
Remaining Budget: 1.5 ϵ	
Per Query Budget: 0.1 ϵ	
Non-Ideal Range: > 20	
Further reading: For a comprehensive overview of differential privacy concepts and guidance on setting privacy budgets, please refer to this primer: [link]	
 Epsilon (ϵ)	2.0 
 Delta (δ)	1e-7 
 Rho (ρ)	3.85 
 Unit of Privacy	User-Day 
 Utility Information	Relative Error 
 Mechanism	Laplace 
Algorithm Hyperparameters	
 Noise Multiplier	1.2 
 Clipping Norm	1.0 
 Batch Size	512 
 Iterations	5 
 Learning Rate	0.01 
 Deployment Model	Central 
 Empirical Privacy Metrics	None 
 Privacy Interpretation-Semantics	None 
 Other Parameters	Release Metadata 

Figure 2: An exploratory visualization of expert-informed DP parameters, based on publicly documented values from Wikimedia's deployment. This is included as a visual reference to support transparency-focused research.

This label is presented solely as an exploratory visualization to illustrate how expert-informed DP parameters might be communicated. It reflects the consensus from expert interviews but is not evaluated or proposed as a standard. We include it to support future research. The interactive version is available at:

<https://privacylevel4dp.github.io/Privacy-Label-for-Differential-Privacy/>
For illustrative purposes, the label uses:

- **Icons** to represent parameter types (e.g., lock = ϵ , gear = mechanism).
- **Color coding** to indicate parameter priority (e.g., green = core, blue = contextual).

Appendix F: Codebook

The codebook developed during the analysis of expert interviews includes detailed descriptions of each theme extracted from the interview data along with definitions.

Code	Description
1. Reason to Include	
Key Parameters	Critical parameters, such as epsilon (ϵ) and delta (δ), that directly impact both privacy guarantees and utility of a DP system, defining the trade-off between data protection and usability.
Interpretation	The ability to understand and compare privacy guarantees, ensuring precision in their definition and communication.
Scope of Protection	Defines the boundaries of data protection, specifying what aspects of the data (e.g., individual records, sensitive attributes) are safeguarded under a privacy model.
Transparency	Clear communication of privacy guarantees, enabling stakeholders to verify and trust privacy mechanisms through reproducibility, accuracy, and correctness of implementation, ultimately building confidence in the system's privacy assurances.
Utility	Represents the usefulness of data after applying privacy mechanisms, balancing privacy protection with the accuracy and practicality of the output.
Privacy Threats	Scenarios where certain parameters create an illusion of strong (robust) privacy protection but fail to provide meaningful privacy or offer little actual privacy, often leading to a false sense of trust in a system's guarantees and neglect of other essential privacy parameters.
Practical Risk Assessment	Involves evaluating risks tied to real-world harms, both positive and negative, to ensure privacy guarantees align with practical applications and realistic threat scenarios.
Understanding Utility	Entails assessing how much usability and quality a data release retains after applying DP mechanisms.
Incentive to Share	Parameters that encourage data sharing by demonstrating that privacy mechanisms protect sensitive information while preserving data utility.

Code	Description
2. Why not Include	
None	No valid reason exists not to disclose these parameters; they are essential for clarity and transparency.
Privacy Theatre	Parameters that mislead stakeholders by appearing to provide strong (robust) privacy protection but fail to provide meaningful privacy or offer little actual privacy, often leading to a false sense of trust in a system's guarantees and neglect of other critical privacy considerations.
Communication Challenge	Parameters that are too technical, complex, or lengthy to explain effectively, making them difficult for stakeholders to understand or interpret or better suited for detailed documentation elsewhere.
Doesn't Belong	Includes parameters deemed unnecessary to disclose due to consensus, redundancy with more expressive or critical parameters, irrelevance, or being better suited for disclosure in other contexts.
Privacy Leakage	Parameters that might inadvertently reveal sensitive or private information, undermining privacy goals.
Consensus	Parameters with no universal agreement or standardization, complicating comparisons between mechanisms and reducing their clarity or usefulness.
Practical Risk Assessment	Assesses risks tied to real-world harms but may be excluded if overly complex, impractical, or irrelevant to stakeholders.
Incentive to Share	Parameters that fail to significantly influence stakeholders' willingness to share data or are deemed less critical.
Implementation	Parameters that are challenging to implement due to technical complexity or resource constraints or feasibility challenges.
3. Normal Range	
Data/Problem or Algorithm Dependent	The appropriate range for parameters varies depending on the dataset, the problem being addressed, or the algorithm used. These ranges are inherently context-specific and influenced by the underlying scenario.
Risk Tolerance	The range reflects how much privacy risk stakeholders or data subjects are willing to accept in real-world applications, balancing privacy protection with acceptable exposure.
Specific Guidance	Clear recommendations or constraints provided by experts, such as using very small ranges, user-level settings, bounded/unbounded configurations, or time-based constraints to minimize risk.

Code	Description
Consensus	Denotes whether there is widespread agreement or standardization about the range. While consensus may exist, discrepancies in concrete values highlight the lack of absolute uniformity.
Theoretical Justification	Ranges grounded in formal definitions or mathematical principles, providing logical or theoretical support for their use.
4. Why the Range	
Theoretical Justification	The range is supported by established theoretical frameworks or formal definitions, ensuring mathematical soundness.
Consensus	A range commonly accepted or standardized within the field, facilitating comparisons and usability across implementations.
Data/Problem or Algorithm Dependent	The range is determined by the specific context, such as the dataset's characteristics, the problem's requirements, or the algorithm being employed.
Utility	The range ensures a balance between maintaining privacy and retaining practical usability of the data, optimizing both goals.
Stakeholder Feedback or Expectation	Ranges are influenced by stakeholder inputs, reflecting their expectations for balancing privacy guarantees with practical application needs.
Quantifies Risk	The range allows stakeholders to assess and quantify the potential privacy risks associated with specific parameter choices.
5. Circumstances the normal range is not applicable	
Composition / Complexity	In scenarios involving complex models or multiple composed queries, the normal range may not apply due to added layers of difficulty in maintaining consistency.
Utility	If adhering to the normal range significantly reduces data usability, the range becomes impractical for the intended purpose.
None	There is no valid reason not to apply the range.
Data/Problem or Algorithm Dependent	Variability in data, problem scope, or algorithm characteristics may render the normal range irrelevant or inapplicable.
Implementation	Challenges in implementing the range due to technical complexity or limitations or resource constraints may prevent its applicability.
Stakeholder Feedback or Expectation	Deviations from the normal range may be necessary if stakeholders require different balances of privacy and utility.
Threat model	The range may not apply if the threat model changes or introduces additional risks that require a different parameter setting.

Code	Description
Consensus	The lack of agreement or standardization in the field can make the normal range irrelevant or inapplicable.
6. Important for General Public	
Very Important	Indicates that the parameter is crucial for the general public to understand, as it significantly impacts their trust, informed decisions (data-sharing decisions) about data privacy and its implications, or perception of privacy guarantees.
Somewhat Important	Denotes that the parameter is moderately relevant to the general public, providing useful context but not essential for their understanding or decision-making processes.
Not Important	Denotes that the parameter holds little to no relevance to the general public and does not contribute meaningfully to their understanding or data-sharing decisions.
7. Important for Technical Users	
Very Important	Indicates that the parameter is essential for technical users to implement, evaluate, or optimize DP systems effectively.
Somewhat Important	Suggests that the parameter is moderately relevant for technical users, aiding their tasks but not being critical for successful implementation or evaluation.
Not Important	Implies that the parameter is not relevant or necessary for technical users in their roles or tasks.
8. Thoughts about layered structure	
Support	Reflects agreement or approval that a layered structure is effective for presenting parameters to audiences with different expertise levels, ensuring accessibility for all stakeholders.
Add Layer	Suggests enhancing the layered structure by introducing a third layer to tailor the presentation of parameters for different user groups, such as adding a deeper technical level or simplifying for non-technical users.
9. Layer to Present Parameters	
Primary	The top layer designed to present the most critical parameters, offering high-level, accessible information for a broad audience.
Secondary	A deeper layer intended for more detailed or technical explanations, tailored for users with advanced knowledge (technical users) or specific needs.

Code	Description
10. How to Present Parameters	
Full Technical Info	Present the complete and detailed information about the parameter, including numerical values, descriptions, and links to supporting documentation or research papers for technical users.
Color-coded	The use of a color scheme to visually differentiate between parameter categories, levels of risk, or importance for quick comprehension (e.g., green for safe, yellow for caution, red for high-risk).
High-Level Summary	A brief, non-technical explanation of parameters, designed to be understandable by non-expert audiences while conveying the core concepts.
11. Where to Place the Label	
With the Data Release	Attaching the DP label directly to the dataset being shared to ensure immediate access by users evaluating the data.
In the App	Embedding the DP label within the application or tool that interacts with the data, ensuring it's accessible during use.
In a Registry	Storing the DP label in a centralized registry where datasets and their privacy parameters are documented for easy reference.
In a Blog Post/-Press Release/-Docs	Including the DP label in public communications or documentation to provide transparency and educate broader audiences about privacy protections.
In a Model Card	Integrating the DP label into a structured document (model card) that outlines key details about the data, its use, and privacy guarantees.
12. Thoughts on design and layout	
Color-coding	The use of a color scheme to visually differentiate between parameter categories, levels of risk, or importance for quick comprehension (e.g., green for safe, yellow for caution, red for high-risk).
Consistency and Standardization	The DP label follows a uniform design and structure to improve usability and avoid confusion across datasets or tools.
Interactivity	The incorporation of interactive elements, such as tooltips or expandable sections, to allow users to explore detailed parameter information as needed.
Use of Visuals	Addition of visual elements like icons, charts, or diagrams to enhance the DP label's clarity and make complex information easier to understand.

Code	Description
Simple File Format	The DP label should be in a format that is easy to integrate and share, such as README.
Order Matters	Structuring the DP label to prioritize the most critical parameters ensuring important parameters are noticed first.
13. Recommendations	
Educational Resources	These are supplementary materials, such as tutorials or guides, to help users understand the parameters and their implications.
Mirror Existing Labels	Adopting design elements or structures from established labels (e.g., nutrition labels) to leverage familiarity and reduce the learning curve.
Checklist	A suggestion of having a checklist format, allowing users to verify that all critical privacy elements are addressed.
Seek Feedback	Engage with users, stakeholders, or experts to refine the label's design and content based on their needs and suggestions.
Simplicity	The DP label should be straightforward, avoiding unnecessary complexity to ensure it is accessible and user-friendly.
Highlight Potentially Misleading Aspects	Clearly identifying areas where the label might unintentionally cause misinterpretations in privacy guarantees.