

Relations Among Different Privacy Notions

Jun Zhao

junzhao@alumni.cmu.edu

Abstract—We present a comprehensive view of the relations among several privacy notions: differential privacy (DP) [1], Bayesian differential privacy (BDP) [2], semantic privacy (SP) [3], and membership privacy (MP) [4]. The results are organized into two parts. In part one, we extend the notion of semantic privacy (SP) to Bayesian semantic privacy (BSP) and show its essential equivalence with Bayesian differential privacy (BDP) in the quantitative sense. We prove the relations between BDP, BSP, and SP as follows: ϵ -BDP $\Leftarrow (\frac{1}{2} - \frac{1}{e^\epsilon + 1})$ -BSP, and ϵ -BDP $\Rightarrow (e^{2\epsilon} - 1)$ -BSP $\Rightarrow (e^{2\epsilon} - 1)$ -SP. In addition, we obtain a minor result ϵ -DP $\Leftarrow (\frac{1}{2} - \frac{1}{e^\epsilon + 1})$ -SP, which improves the result of Kasiviswanathan and Smith [3] stating ϵ -DP $\Leftarrow \epsilon/6$ -SP for $\epsilon \leq 1.35$. In part two, we establish the relations between BDP and MP. First, ϵ -BDP $\Rightarrow \epsilon$ -MP. Second, for a family of distributions that are downward scalable in the sense of Li *et al.* [4], it is shown that ϵ -BDP $\Leftarrow \epsilon$ -MP.

Keywords—Differential privacy, Bayesian differential privacy, semantic privacy, membership privacy.

I. INTRODUCTION

Differential privacy (DP). Differential privacy by Dwork *et al.* [1], [5] is a robust privacy standard that has been successfully applied to a range of data analysis tasks, since it provides a rigorous foundation for defining and preserving privacy. Differential privacy has received considerable attention in the literature [6]–[15]. Apple has incorporated differential privacy into its mobile operating system iOS 10 [16]. Google has implemented a differentially private tool called RAPPOR in the Chrome browser to collect information about clients [17]. A randomized algorithm Y satisfies ϵ -differential privacy if for all adjacent databases x, x' and any event E , it holds that $\mathbb{P}[Y(x) \in E] \leq e^\epsilon \mathbb{P}[Y(x') \in E]$, where $\mathbb{P}[\cdot]$ denotes the probability throughout this paper. Intuitively, under differential privacy, an adversary given access to the output do not have much confidence to determine whether it was sampled from the probability distribution generated by the algorithm when the database is x or when the database is x' .

Bayesian differential privacy (BDP). Yang *et al.* [2] introduce the notion of Bayesian differential privacy as follows. Bayesian differential privacy broadens the application scenarios of differential privacy when data records have dependencies. For a database x with n tuples, let $i \in \{1, 2, \dots, n\}$ be a tuple index

in the database and $S \subseteq \{1, 2, \dots, n\} \setminus i$ be a tuple index set. An adversary denoted by $A(i, S)$ knows the values of all tuples in S (denoted by x_S) and attempts to attack the value of tuple i (denoted by x_i). For a randomized perturbation mechanism $Y = \mathbb{P}[y \in \mathcal{Y} \mid x]$ on database x , the Bayesian differential privacy leakage (BDPL) of Y with respect to the adversary $A(i, S)$ is $\text{BDPL}_A(Y) = \sup_{x_i, x'_i, x_S, \mathcal{Y}} \ln \frac{\mathbb{P}[y \in \mathcal{Y} \mid x_i, x_S]}{\mathbb{P}[y \in \mathcal{Y} \mid x'_i, x_S]}$. The mechanism Y satisfies ϵ -Bayesian differential privacy if $\sup_A \text{BDPL}_A(Y) \leq \epsilon$.

Semantic privacy (SP). Kasiviswanathan and Smith [3] propose a Bayesian formulation of semantic privacy, inspired by the following interpretation of differential privacy explained in [1]: *Regardless of external knowledge, an adversary with access to the sanitized database draws the same conclusions whether or not any individual data is included in the original database.* The phrases “external knowledge” and “drawing conclusions” are formulated as follows in [3]. The external knowledge is modeled by a prior probability distribution b on \mathcal{D}^n , where b is short for “belief”, and databases are assumed to be vectors in \mathcal{D}^n for some domain \mathcal{D} . Conclusions are captured via the corresponding posterior distribution: given a transcript y , the adversary updates his belief b about the database x using Bayes’ rule to obtain a posterior \bar{b} : $\bar{b}[x|y] = \frac{\mathbb{P}[Y(x)=y]b[x]}{\sum_z \mathbb{P}[Y(z)=y]b[z]}$.

For the database x , Kasiviswanathan and Smith [3] further define x_{-i} to be the same vector except that the record at position i has been replaced by some fixed, default value \perp in \mathcal{D} .

Kasiviswanathan and Smith [3] define $n + 1$ related games, numbered 0 through n . In Game 0, the adversary interacts with $Y(x)$. This is the interaction that actually takes place between the adversary and the randomized mechanism Y . Hence, the distribution \bar{b}_0 is just the distribution \bar{b} as defined in (1); i.e., $\bar{b}_0[x|y] = \bar{b}[x|y] = \frac{\mathbb{P}[Y(x)=y]b[x]}{\sum_z \mathbb{P}[Y(z)=y]b[z]}$.

In Game i (for $1 \leq i \leq n$), the adversary interacts with $Y(x_{-i})$. Game i describes the hypothetical scenario where person i ’s record is not used. In Game i (for $1 \leq i \leq n$), given a transcript y , the adversary updates his belief b about database x again using Bayes’ rule to obtain a posterior \bar{b}_i as follows: $\bar{b}_i[x|y] = \frac{\mathbb{P}[Y(x_{-i})=y]b[x]}{\sum_z \mathbb{P}[Y(z_{-i})=y]b[z]}$.

Given a transcript y , Kasiviswanathan and Smith [3] say that privacy has been breached if the adversary would draw different conclusions about the world and, in particular, about a person i , depending on whether or not i ’s data was used. To this end, Kasiviswanathan and Smith [3] formally define ϵ -semantic privacy below, where the statistical difference $\text{SD}(X, Y)$ between random variables X and Y on the same discrete space D is defined

The author Jun Zhao obtained his PhD from Carnegie Mellon University, Pittsburgh, PA 15213, USA, where he was with the Cybersecurity Lab (CyLab). He was a postdoctoral scholar with Arizona State University, Tempe, AZ 85281, USA. He is now a research fellow at Nanyang Technological University in Singapore. Email: junzhao@alumni.cmu.edu

by $\text{SD}(X, Y) = \max_{S \subseteq D} |\mathbb{P}[X \in S] - \mathbb{P}[Y \in S]|$. A randomized mechanism Y is said to be ϵ -semantically private if for all belief distributions b on \mathcal{D}^n , for all possible transcripts y , and for all $i = 1, \dots, n$, it holds that $\text{SD}(\bar{b}_0[\cdot|y], \bar{b}_i[\cdot|y]) \leq \epsilon$.

Membership privacy (MP). Li *et al.* [4] propose membership privacy (MP) in consideration of the adversary's prior beliefs. Let the adversary's prior beliefs about the dataset be captured by a distribution \mathcal{D} . From the adversary's point of view, the dataset is a random variable drawn according to the distribution \mathcal{D} . With \bar{x}_i denoting the event that record x_i is not in the database, Li *et al.* [4] define membership privacy as follows. A mechanism Y achieves ϵ -membership privacy under a family \mathbb{D} of distributions, i.e., $\langle \mathbb{D}, \epsilon \rangle$ -MP, if and only if for any distribution $\mathcal{D} \in \mathbb{D}$ and for any record x_i , any possible set \mathcal{Y} for the output, we have¹ $\mathbb{P}_{\mathcal{D}, Y}[x_i | \mathcal{Y}] \leq e^\epsilon \mathbb{P}_{\mathcal{D}}[x_i]$ and $\mathbb{P}_{\mathcal{D}, Y}[\bar{x}_i | \mathcal{Y}] \geq e^{-\epsilon} \mathbb{P}_{\mathcal{D}}[\bar{x}_i]$.

The rest of the paper is organized as follows. Section II presents the results on the relations among several privacy notions: differential privacy (DP), Bayesian differential privacy (BDP), semantic privacy (SP), and membership privacy (MP). We elaborate their proofs in Sections III. Section IV surveys related work, and Section V concludes the paper.

II. THE RESULTS

Kasiviswanathan and Smith [3] introduce semantic privacy (SP) and show its essential equivalence with differential privacy (DP) in the quantitative sense (the notion of essential equivalence means ϵ -DP $\Leftarrow f(\epsilon)$ -SP and ϵ -DP $\Rightarrow g(\epsilon)$ -SP for some functions f and g). We extend their notion to Bayesian semantic privacy (BSP) and show its essential equivalence with Bayesian differential privacy (BDP) also in the quantitative sense. We prove the relations between BDP, BSP, and SP as follows:

- (i) ϵ -BDP $\Leftarrow (\frac{1}{2} - \frac{1}{e^{\epsilon+1}})$ -BSP.
- (ii) ϵ -BDP $\Rightarrow (e^{2\epsilon} - 1)$ -BSP $\Rightarrow (e^{2\epsilon} - 1)$ -SP.

We prove results (i) and (ii) in Section III-A, where we also obtain a minor result ϵ -DP $\Leftarrow (\frac{1}{2} - \frac{1}{e^{\epsilon+1}})$ -SP, which improves the result of Kasiviswanathan and Smith [3] stating ϵ -DP $\Leftarrow \epsilon/6$ -SP for $\epsilon \leq 1.35$.

Li *et al.* [4] propose membership privacy (MP), which is applicable to Bayesian data, in contrast to DP. However, no general algorithm has been proposed for this framework. We present the following relations between BDP and MP:

- (iii) ϵ -BDP $\Rightarrow \epsilon$ -MP.
- (iv) For a family of distributions that are downward scalable in the sense of Li *et al.* [4], ϵ -BDP $\Leftarrow \epsilon$ -MP (See [4] for the meaning of “downward scalable” distributions).

We prove results (iii) and (iv) in Section III-B.

¹ $\langle \mathbb{D}, \epsilon \rangle$ -membership privacy actually corresponds to $\langle \mathbb{D}, e^\epsilon \rangle$ -positive membership privacy in [4]. Li *et al.* [4] use γ and γ^{-1} instead of e^ϵ and $e^{-\epsilon}$ in (21) and (22) to define $\langle \mathbb{D}, \gamma \rangle$ -membership privacy. We use e^ϵ and $e^{-\epsilon}$ here for better comparison between membership privacy and Bayesian differential privacy. Also, by membership privacy, we mean positive membership privacy of [4]. We do not discuss negative membership privacy of [4].

III. PROOFS

A. Relations between our Bayesian differential privacy and Kasiviswanathan and Smith's semantic privacy [3]

We extend the work of Kasiviswanathan and Smith [3] on semantic privacy to tackle the case of correlated tuples. Specifically, we will present Bayesian semantic privacy and prove that the notions of Bayesian differential privacy and Bayesian semantic privacy are essentially (i.e., quantitatively) equivalent (see Theorem 1 below). Our result resembles [3, Theorem 2.2], which shows that differential privacy and semantic privacy are essentially equivalent.

Theorem 1. *ϵ -Bayesian differential privacy implies $(e^{2\epsilon} - 1)$ -Bayesian semantic privacy, and is implied by $(\frac{1}{2} - \frac{1}{e^{\epsilon+1}})$ -Bayesian semantic privacy.*

Theorem 2 (Improving the result of Kasiviswanathan and Smith [3]). *ϵ -Differential privacy implies $(e^{2\epsilon} - 1)$ -semantic privacy, and is implied by $(\frac{1}{2} - \frac{1}{e^{\epsilon+1}})$ -semantic privacy.*

Theorem 1 is one of our novel results. The first part of Theorem 2 is obtained by Kasiviswanathan and Smith [3]. The second part of Theorem 2 improves the corresponding result of Kasiviswanathan and Smith [3], which states that ϵ -differential privacy is implied by $\epsilon/6$ -semantic privacy for $\epsilon \leq 1.35$. The improvement can be seen from $\frac{1}{2} - \frac{1}{e^{\epsilon+1}} > \epsilon/6$ for $\epsilon \leq 1.35$.

The rest of the discussion is organized as follows. We review semantic privacy and define Bayesian semantic privacy in Section III-A1. In Section III-A2, we recall Bayesian differential privacy. Finally, we prove the above Theorem 1 in Section III-A3. The proof of Theorem 2 is similar to that of Theorem 1.

1) *Reviewing semantic privacy and defining Bayesian semantic privacy:* In this section, we first review semantic privacy from Kasiviswanathan and Smith [3], before presenting Bayesian semantic privacy, which extends the notion of semantic privacy to address correlated tuples.

A review of Kasiviswanathan and Smith [3] for semantic privacy:

Kasiviswanathan and Smith [3] propose a Bayesian formulation of semantic privacy, inspired by the following interpretation of differential privacy explained in [1]: *Regardless of external knowledge, an adversary with access to the sanitized database draws the same conclusions whether or not any individual data is included in the original database.* The phrases “external knowledge” and “drawing conclusions” are formulated as follows in [3]. The external knowledge is modeled by a prior probability distribution b on \mathcal{D}^n , where b is short for “belief,” and databases are assumed to be vectors in \mathcal{D}^n for some domain \mathcal{D} . Conclusions are captured via the corresponding posterior distribution: given a transcript y , the adversary updates his belief b about the database

x using Bayes' rule to obtain a posterior \bar{b} :²

$$\bar{b}[x|y] = \frac{\mathbb{P}[Y(x) = y] b[x]}{\sum_z \mathbb{P}[Y(z) = y] b[z]}. \quad (1)$$

For the database x , Kasiviswanathan and Smith [3] further define x_{-i} to be the same vector except that position i has been replaced by some fixed, default value in \mathcal{D} . Any valid value in \mathcal{D} will do for the default value. In addition, the default value can be understood as a special value \perp (e.g., “no data”); see [3, Page 3–Footnote 2] for details. We will use \perp whenever it is necessary to explicitly write out the default value.

Kasiviswanathan and Smith [3] define $n + 1$ related games, numbered 0 through n . In Game 0, the adversary interacts with $Y(x)$. This is the interaction that actually takes place between the adversary and the randomized mechanism Y . Hence, the distribution \bar{b}_0 is just the distribution \bar{b} as defined in (1); i.e.,

$$\bar{b}_0[x|y] = \bar{b}[x|y] = \frac{\mathbb{P}[Y(x) = y] b[x]}{\sum_z \mathbb{P}[Y(z) = y] b[z]}. \quad (2)$$

In Game i (for $1 \leq i \leq n$), the adversary interacts with $Y(x_{-i})$. Game i describes the hypothetical scenario where person i 's record is not used. In Game i (for $1 \leq i \leq n$), given a transcript y , the adversary updates his belief b about database x again using Bayes' rule to obtain a posterior \bar{b}_i as follows:

$$\bar{b}_i[x|y] = \frac{\mathbb{P}[Y(x_{-i}) = y] b[x]}{\sum_z \mathbb{P}[Y(z_{-i}) = y] b[z]}. \quad (3)$$

Given a transcript y , Kasiviswanathan and Smith [3] say that privacy has been breached if the adversary would draw different conclusions about the world and, in particular, about a person i , depending on whether or not i 's data was used. To this end, Kasiviswanathan and Smith [3] formally define ϵ -semantic privacy below, where the statistical difference $\text{SD}(X, Y)$ between probability distributions (or random variables) X and Y on a discrete space D is defined by

$$\text{SD}(X, Y) = \max_{S \subseteq D} |\mathbb{P}[X \in S] - \mathbb{P}[Y \in S]|.$$

Definition 1 (ϵ -Semantical Privacy by [3, Definition 2.1]). A randomized mechanism Y is said to be ϵ -semantically private if for all belief distributions b on \mathcal{D}^n , for all possible transcripts y , and for all $i = 1, \dots, n$:

$$\text{SD}(\bar{b}_0[\cdot|y], \bar{b}_i[\cdot|y]) \leq \epsilon. \quad (4)$$

From (3) and (4), the above definition of ϵ -semantic privacy requires the use of x_{-i} , where x_{-i} is obtained after we replace position i at x by the default value \perp . If the tuples are correlated, changing position i at x might also result in changing other positions at x . Hence, ϵ -semantic privacy may not work well under correlated tuples. Given this, we next extend ϵ -semantic privacy to address correlated tuples and present ϵ -Bayesian semantic privacy.

²For simplicity, only discrete probability distributions are discussed. The results can be readily extended to the continuous case.

Extending semantic privacy to Bayesian semantic privacy to address correlated tuples:

As will become clear, our extension of semantic privacy to Bayesian semantic privacy is similar to the extension of differential privacy to Bayesian differential privacy.

We let a statistical database be $[X_1, X_2, \dots, X_n]$, where X_j for each $j \in \{1, 2, \dots, n\}$ is a *random variable*. We also let \mathcal{N} be $\{1, 2, \dots, n\}$. Then we consider the databases x and z used in (1)–(3) above to be

$$x = [X_1 = x_1, X_2 = x_2, \dots, X_n = x_n] = [X_j = x_j : j \in \mathcal{N}], \quad (5)$$

and

$$z = [X_1 = z_1, X_2 = z_2, \dots, X_n = z_n] = [X_j = z_j : j \in \mathcal{N}]. \quad (6)$$

When the data tuples are correlated, the adversary may gain more advantage in inferring x_i by using random variables $X_j|_{j \in \mathcal{S}}$'s instantiations $x_j|_{j \in \mathcal{S}}$, and random variables $X_j|_{j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}}$ for computation instead of using instantiations $x_j|_{j \in \mathcal{N} \setminus \{i\}}$ only, where $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$ (note that \mathcal{S} can be an arbitrary subset of $\mathcal{N} \setminus \{i\}$). For notation convenience, we define $x_{i+\mathcal{S}}$ and $z_{i+\mathcal{S}}$ by

$$x_{i+\mathcal{S}} = [X_i = x_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}], \quad (7)$$

and

$$z_{i+\mathcal{S}} = [X_i = z_i, X_j = z_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}]. \quad (8)$$

From (5)–(8), if $\mathcal{S} = \mathcal{N} \setminus \{i\}$, then $x_{i+\mathcal{S}}$ and $z_{i+\mathcal{S}}$ reduce to databases x and z , respectively.

Similar to the previous subsection, here we also let the adversary play $n + 1$ related games with the randomized mechanism Y , and define $\bar{b}, \bar{b}_0, \bar{b}_i|_{i=1, \dots, n}$ as detailed below. In Game 0, the adversary interacts with $Y(x_{i+\mathcal{S}})$. We generalize x and z in (2) to $x_{i+\mathcal{S}}$ and $z_{i+\mathcal{S}}$, so that (2) becomes

$$\bar{b}_0[x_{i+\mathcal{S}}|y] = \bar{b}[x_{i+\mathcal{S}}|y] = \frac{\mathbb{P}[Y(x_{i+\mathcal{S}}) = y] b[x_{i+\mathcal{S}}]}{\sum_{z_{i+\mathcal{S}}} \mathbb{P}[Y(z_{i+\mathcal{S}}) = y] b[z_{i+\mathcal{S}}]}. \quad (9)$$

For clarity, we explain the beliefs in (9). From (7) and (8), $b[x_{i+\mathcal{S}}]$ and $b[z_{i+\mathcal{S}}]$ in (9) are given by

$$\begin{aligned} b[x_{i+\mathcal{S}}] &= b[X_i = x_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}] \\ &= b[X_i = x_i, X_j = x_j : j \in \mathcal{S}], \end{aligned} \quad (10)$$

and

$$\begin{aligned} b[z_{i+\mathcal{S}}] &= b[X_i = z_i, X_j = z_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}] \\ &= b[X_i = z_i, X_j = z_j : j \in \mathcal{S}]. \end{aligned} \quad (11)$$

Similar to (10), from (7), $\bar{b}_0[x_{i+S}|y]$ is given by

$$\begin{aligned}\bar{b}_0[x_{i+S}|y] &= \bar{b}_0[X_i = x_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}|y] \\ &= \bar{b}_0[X_i = x_i, X_j = x_j : j \in \mathcal{S}|y].\end{aligned}\quad (12)$$

In Game i (for $1 \leq i \leq n$), we change position i at x_{i+S} to the default value \perp to obtain x_{-i+S} defined below; specifically, recalling x_{i+S} given by (7), we set x_{-i+S} by

$$x_{-i+S} = [X_i = \perp, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}]. \quad (13)$$

Similarly, we change position i at z_{i+S} to the default value \perp to obtain z_{-i+S} defined below; specifically, recalling z_{i+S} given by (8), we set z_{-i+S} by

$$z_{-i+S} = [X_i = \perp, X_j = z_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}]. \quad (14)$$

As x_{i+S} and z_{i+S} generalize x and z in (3), clearly x_{-i+S} and z_{-i+S} also generalize x_{-i} and z_{-i} in (3). In Game i (for $1 \leq i \leq n$), the adversary interacts with $Y(x_{-i+S})$. Then replacing x , z , x_{-i} and z_{-i} in (3) by x_{i+S} , z_{i+S} , x_{-i+S} and z_{-i+S} , respectively, we obtain

$$\bar{b}_i[x_{i+S}|y] = \frac{\mathbb{P}[Y(x_{-i+S}) = y] b[x_{i+S}]}{\sum_{z_{i+S}} \mathbb{P}[Y(z_{-i+S}) = y] b[z_{i+S}]}.\quad (15)$$

The beliefs $b[x_{i+S}]$ and $b[z_{i+S}]$ in (15) are already interpreted as (10) and (11). For clarity, we further explain $\bar{b}_i[x_{i+S}|y]$ in (15). Similar to (12), from (7), $\bar{b}_i[x_{i+S}|y]$ is given by

$$\begin{aligned}\bar{b}_i[x_{i+S}|y] &= \bar{b}_i[X_i = x_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}|y] \\ &= \bar{b}_i[X_i = x_i, X_j = x_j : j \in \mathcal{S}|y].\end{aligned}\quad (16)$$

With the above notation, we define ϵ -Bayesian semantical privacy below, in a way similar to that of ϵ -semantical privacy in Definition 1.

Definition 2 (ϵ -Bayesian Semantical Privacy). *A randomized mechanism Y is said to have ϵ -Bayesian semantical privacy if for all belief distributions b on \mathcal{D}^n , for all possible transcripts y , for all $i = 1, \dots, n$, and for all x_{i+S} and z_{i+S} defined in (7) and (8) with $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$:*

$$\text{SD}(\bar{b}_0[x_{i+S}|y], \bar{b}_i[x_{i+S}|y]) \leq \epsilon. \quad (17)$$

To understand the beliefs $\bar{b}_0[x_{i+S}|y]$ and $\bar{b}_i[x_{i+S}|y]$ in (17), we use their interpretations in (12) and (16). In Definition 2 for ϵ -Bayesian semantical privacy, we consider all possible $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$. In the hypothetical scenario where we consider \mathcal{S} only as $\mathcal{N} \setminus \{i\}$ in Definition 2, Definition 2 would reduce to Definition 1 for ϵ -semantical privacy.

2) *Recalling Bayesian differential privacy:* In this section, we recall Bayesian differential privacy and express its definition using some new notation.

With x_{i+S} defined in (7) (i.e., $x_{i+S} = [X_i = x_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}]$), for notation convenience, we further define x'_{i+S} by

$$x'_{i+S} = [X_i = x'_i, X_j = x_j : j \in \mathcal{S}, X_j : j \in \mathcal{N} \setminus \{i\} \setminus \mathcal{S}]. \quad (18)$$

Note that the only difference between x_{i+S} and x'_{i+S} is that the former has $X_i = x_i$, while the latter enforces $X_i = x'_i$. Then ϵ -Bayesian differential privacy means

$$\frac{\mathbb{P}[Y(x_{i+S}) = y]}{\mathbb{P}[Y(x'_{i+S}) = y]} \leq e^\epsilon. \quad (19)$$

3) *Proving Theorem 1 on the relations between Bayesian differential privacy and Bayesian semantic privacy:* Our Theorem 1 restated below presents the relations between Bayesian differential privacy and Bayesian semantic privacy.

Theorem 1 (Restated). *ϵ -Bayesian differential privacy implies $(e^{2\epsilon} - 1)$ -Bayesian semantic privacy, and is implied by $(\frac{1}{2} - \frac{1}{e^\epsilon + 1})$ -Bayesian semantic privacy.*

Theorem 1 shows that the notions of Bayesian differential privacy and Bayesian semantic privacy are essentially equivalent (of course, the parameters should be set appropriately). The proof of Theorem 1 below is just an extension of the reasoning by Kasiviswanathan and Smith [3].

Proof of Theorem 1. We show Theorem 1 in two parts below. We will use the following definition of point-wise $(\epsilon, 0)$ -indistinguishability from [3, Definition 3.2]: Two discrete random variables X and Y are point-wise $(\epsilon, 0)$ -indistinguishable if it holds for a drawn from either X or Y that $e^{-\epsilon} \mathbb{P}[Y = a] \leq \mathbb{P}[X = a] \leq e^\epsilon \mathbb{P}[Y = a]$.

Proving ϵ -Bayesian differential privacy $\implies (e^{2\epsilon} - 1)$ -Bayesian semantic privacy: To prove this part, we consider any database $x \in \mathcal{D}^n$. Let Y be an $\epsilon/2$ -Bayesian differentially private algorithm. Consider any belief distribution b . Let the posterior distributions $\bar{b}_0[x_{i+S}|y]$ and $\bar{b}_i[x_{i+S}|y]$ for some fixed i , \mathcal{S} and y be defined in (9) and (15). From (19), ϵ -Bayesian differential privacy implies that for every z_{i+S} ,

$$e^{-\epsilon} \mathbb{P}[Y(z_{-i+S}) = y] \leq \mathbb{P}[Y(z_{i+S}) = y] \leq e^\epsilon \mathbb{P}[Y(z_{-i+S}) = y].$$

These inequalities imply that the ratio of $\bar{b}_0[x_{i+S}|y]$ and $\bar{b}_i[x_{i+S}|y]$ (defined in (9) and (15)) is within $e^{\pm 2\epsilon}$. Since these inequalities hold for every x_{i+S} , we get:

$$e^{-2\epsilon} \bar{b}_i[x_{i+S}|y] \leq \bar{b}_0[x_{i+S}|y] \leq e^{2\epsilon} \bar{b}_i[x_{i+S}|y], \quad \forall x_{i+S}.$$

This implies that the random variables $\bar{b}_0[x_{i+S}|y]$ and $\bar{b}_i[x_{i+S}|y]$ are point-wise $(2\epsilon, 0)$ -indistinguishable. Applying [3, Lemma 3.3-Property 5], we obtain $\text{SD}(\bar{b}_0[x_{i+S}|y], \bar{b}_i[x_{i+S}|y]) \leq (e^{2\epsilon} - 1)$. Repeating the above arguments for every belief distribution, for every i , and for every y , we thus show that the mechanism Y is $(e^{2\epsilon} - 1)$ -Bayesian semantic private.

Proving $(\frac{1}{2} - \frac{1}{e^\epsilon + 1})$ -Bayesian semantic privacy $\implies \epsilon$ -Bayesian differential privacy: To prove this part, we consider a belief distribution b which is uniform over

$$x_{i+S} = [X_i = x_i, X_j = x_j : j \in S, X_j : j \in \mathcal{N} \setminus \{i\} \setminus S]$$

and

$$x'_{i+S} = [X_i = x'_i, X_j = x_j : j \in S, X_j : j \in \mathcal{N} \setminus \{i\} \setminus S];$$

i.e.,

$$b[x_{i+S}] = b[X_i = x_i, X_j = x_j : j \in S] = \frac{1}{2}$$

and

$$b[x'_{i+S}] = b[X_i = x'_i, X_j = x_j : j \in S] = \frac{1}{2}.$$

Fix a transcript y . The distribution $\bar{b}_i[\cdot|y]$ will be uniform over x_{i+S} and x'_{i+S} since they induce the same distribution on transcripts in Game i . This means that $\bar{b}_0[\cdot|y]$ will assign probabilities in the interval $[\frac{1}{2} - (\frac{1}{2} - \frac{1}{e^\epsilon + 1}), \frac{1}{2} + (\frac{1}{2} - \frac{1}{e^\epsilon + 1})]$ to each of x_{i+S} and x'_{i+S} (by Definition 1). Working through Bayes' rule shows that (note that $b[x_{i+S}] = b[x'_{i+S}]$)

$$\begin{aligned} & \frac{\mathbb{P}[Y(x_{i+S}) = y]}{\mathbb{P}[Y(x'_{i+S}) = y]} \\ &= \frac{\bar{b}_0[x_{i+S}|y]}{\bar{b}_0[x'_{i+S}|y]} \leq \frac{\frac{1}{2} + (\frac{1}{2} - \frac{1}{e^\epsilon + 1})}{\frac{1}{2} - (\frac{1}{2} - \frac{1}{e^\epsilon + 1})} = e^\epsilon. \end{aligned} \quad (20)$$

Since the bound in (20) holds for every y , $Y(x_{i+S})$ and $Y(x'_{i+S})$ are point-wise $(\epsilon, 0)$ -indistinguishable. From [3, Lemma 3.3-Property 5], $Y(x_{i+S})$ and $Y(x'_{i+S})$ are $(\epsilon, 0)$ -indistinguishable. Since this relation holds for every pair of x_{i+S} and x'_{i+S} , the mechanism Y is ϵ -Bayesian differentially private. ■

B. Relations between Bayesian differential privacy and membership privacy

The adversary may have prior beliefs about what the dataset is; this is captured by a distribution \mathcal{D} . From the adversary's point of view, the dataset is a random variable drawn according to the distribution \mathcal{D} . With \bar{x}_i denoting the event that record x_i is not in the database, Li *et al.* [4] define membership privacy as follows, where we reuse some notation of Li *et al.* [4].

Definition 3 (Li *et al.* [4]). A mechanism Y achieves ϵ -membership privacy under a family \mathbb{D} of distributions, i.e., $\langle \mathbb{D}, \epsilon \rangle$ -membership privacy, if and only if for any distribution $\mathcal{D} \in \mathbb{D}$ and for any record x_i , any possible set \mathcal{Y} for the output, we have³

$$\mathbb{P}_{\mathcal{D}, Y}[x_i | \mathcal{Y}] \leq e^\epsilon \mathbb{P}_{\mathcal{D}}[x_i] \quad (21)$$

and

$$\mathbb{P}_{\mathcal{D}, Y}[\bar{x}_i | \mathcal{Y}] \geq e^{-\epsilon} \mathbb{P}_{\mathcal{D}}[\bar{x}_i]. \quad (22)$$

³ $\langle \mathbb{D}, \epsilon \rangle$ -membership privacy actually corresponds to $\langle \mathbb{D}, e^\epsilon \rangle$ -membership privacy in [4]. Li *et al.* [4] use γ and γ^{-1} instead of e^ϵ and $e^{-\epsilon}$ in (21) and (22) to define $\langle \mathbb{D}, \gamma \rangle$ -membership privacy. We use e^ϵ and $e^{-\epsilon}$ here for better comparison between membership privacy and Bayesian differential privacy.

We discuss the adversary model considered here. Let $\mathcal{D}_{i,K}$ denote a distribution where $\mathbb{P}[x_i, x_K] = p$ and $\mathbb{P}[x'_i, x_K] = 1 - p$ for some p . Define $\mathbb{D}_* \stackrel{\text{def}}{=} \bigcup_{\substack{i \in \{1, \dots, n\}, \\ K \subseteq \{1, \dots, n\} \setminus \{i\}}} \mathcal{D}_{i,K}$. The adversary model will be captured by the family \mathbb{D}_* of distributions. For simplicity, we will refer to $\langle \mathbb{D}_*, \epsilon \rangle$ -membership privacy as ϵ -membership privacy.

1) From Bayesian differential privacy to membership privacy:

Theorem 3. ϵ -Bayesian differential privacy implies ϵ -membership privacy.

Lemma 1. A mechanism Y achieves ϵ -membership privacy under a family \mathbb{D} of distributions, i.e., $\langle \mathbb{D}, \epsilon \rangle$ -MP, if and only if it holds for any distribution $\mathcal{D} \in \mathbb{D}$ that⁴

$$\frac{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | \bar{x}_i]} \leq \begin{cases} \frac{1 - \mathbb{P}_{\mathcal{D}}[x_i]}{e^{-\epsilon} - \mathbb{P}_{\mathcal{D}}[x_i]}, & \text{if } 0 \leq \mathbb{P}_{\mathcal{D}}[x_i] \leq \frac{1}{1+e^\epsilon}, \quad (23a) \\ \frac{e^\epsilon - 1 + \mathbb{P}_{\mathcal{D}}[x_i]}{\mathbb{P}_{\mathcal{D}}[x_i]}, & \text{if } \frac{1}{1+e^\epsilon} < \mathbb{P}_{\mathcal{D}}[x_i] \leq 1. \quad (23b) \end{cases}$$

We will explain that Lemma 1 implies the following corollary, which will be used to show Theorem 3.

Corollary 1. A mechanism Y achieves ϵ -membership privacy under a family \mathbb{D} of distributions, i.e., $\langle \mathbb{D}, \epsilon \rangle$ -MP, if it holds for any distribution $\mathcal{D} \in \mathbb{D}$ that $\frac{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | \bar{x}_i]} \leq e^\epsilon$.

Proof of Theorem 3 using Corollary 1. Under distribution $\mathcal{D}_{i,K}$ where $\mathbb{P}[x_i, x_K] = p$ and $\mathbb{P}[x'_i, x_K] = 1 - p$ for some p , we have

$$\mathbb{P}_{\mathcal{D}_{i,K}, Y}[\mathcal{Y} | x_i] = \mathbb{P}[Y(x_i, x_K, X_{\bar{K}}) \in \mathcal{Y}], \quad (24)$$

and

$$\mathbb{P}_{\mathcal{D}_{i,K}, Y}[\mathcal{Y} | \bar{x}_i] = \mathbb{P}[Y(x'_i, x_K, X_{\bar{K}}) \in \mathcal{Y}]. \quad (25)$$

Under ϵ -Bayesian differential privacy, we have $\frac{\mathbb{P}[Y(x_i, x_K, X_{\bar{K}}) \in \mathcal{Y}]}{\mathbb{P}[Y(x'_i, x_K, X_{\bar{K}}) \in \mathcal{Y}]} \leq e^\epsilon$, which along with (24) and (25) yields $\frac{\mathbb{P}_{\mathcal{D}_{i,K}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}_{i,K}, Y}[\mathcal{Y} | \bar{x}_i]} \leq e^\epsilon$, then $\langle \mathbb{D}_*, \epsilon \rangle$ -membership privacy (i.e., ϵ -membership privacy) follows for $\mathbb{D}_* \stackrel{\text{def}}{=} \bigcup_{\substack{i \in \{1, \dots, n\}, \\ K \subseteq \{1, \dots, n\} \setminus \{i\}}} \mathcal{D}_{i,K}$. ■

Proof of Corollary 1 using Lemma 1. Note that (23a) and (23b) in Lemma 1 can be written as $\frac{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | \bar{x}_i]} \leq g(\mathbb{P}_{\mathcal{D}}[x_i])$, where $g(b)$ is a function defined as follows:

$$g(b) \stackrel{\text{def}}{=} \begin{cases} \frac{1-b}{e^{-\epsilon}-b}, & \text{if } 0 \leq b \leq \frac{1}{1+e^\epsilon}, \\ \frac{e^\epsilon-1+b}{b}, & \text{if } \frac{1}{1+e^\epsilon} < b \leq 1. \end{cases} \quad (26)$$

The function $g(b)$ increases as b increases for $0 \leq b \leq \frac{1}{1+e^\epsilon}$ and decreases as b increases for $\frac{1}{1+e^\epsilon} < b \leq 1$. Hence, at $b = 0$ or $b = 1$, $g(b)$ takes its minimum $g(0) = g(1) = e^\epsilon$. Then $\frac{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | \bar{x}_i]} \leq e^\epsilon$ implies $\frac{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}, Y}[\mathcal{Y} | \bar{x}_i]} \leq g(\mathbb{P}_{\mathcal{D}}[x_i])$ for any $\mathbb{P}_{\mathcal{D}}[x_i]$. In view this, we obtain Corollary 1 from Lemma 1.

⁴We let $\frac{0}{0} = 1$ and $\frac{\text{non-zero}}{0} = \infty$ to address the degenerate cases.

Proof of Lemma 1. For simplicity, we define

$$A \stackrel{\text{def}}{=} \frac{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid x_i]}{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i]}, \quad (27)$$

Then the goal of Lemma 1 is to show the combination of (21) and (22) is equivalent to $A \leq g(\mathbb{P}_{\mathcal{D}}[x_i])$. Hence, we will establish Lemma 1 once proving the following three results:

$$(21) \iff \left\{ 1 - \mathbb{P}_{\mathcal{D}}[x_i] \geq A(e^{-\epsilon} - \mathbb{P}_{\mathcal{D}}[x_i]) \right\}, \quad (28)$$

$$(22) \iff \left\{ A \times \mathbb{P}_{\mathcal{D}}[x_i] + 1 - \mathbb{P}_{\mathcal{D}}[x_i] \leq e^\epsilon \right\}, \quad (29)$$

and

$$\left. \begin{aligned} &1 - \mathbb{P}_{\mathcal{D}}[x_i] \geq A(e^{-\epsilon} - \mathbb{P}_{\mathcal{D}}[x_i]), \\ &\text{and } A \times \mathbb{P}_{\mathcal{D}}[x_i] + 1 - \mathbb{P}_{\mathcal{D}}[x_i] \leq e^\epsilon \end{aligned} \right\} \iff A \leq g(\mathbb{P}_{\mathcal{D}}[x_i]). \quad (30)$$

Below we demonstrate (28) (29) and (30), respectively.

Proving (28):

By Bayes' theorem, it holds that

$$\mathbb{P}_{\mathcal{D},Y}[x_i \mid \mathcal{Y}] = \frac{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid x_i] \mathbb{P}_{\mathcal{D}}[x_i]}{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}]}. \quad (31)$$

Given (31), we have

$$(21) \iff \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid x_i] \leq e^\epsilon \times \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}]. \quad (32)$$

To prove (32), we express $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}]$ by the law of total probability, and find

$$\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}] = \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid x_i] \mathbb{P}_{\mathcal{D}}[x_i] + \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i] \mathbb{P}_{\mathcal{D}}[\bar{x}_i]. \quad (33)$$

Applying (27) to (33), we obtain

$$\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}] = \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid x_i] \times \left\{ \mathbb{P}_{\mathcal{D}}[x_i] + A^{-1} \times \mathbb{P}_{\mathcal{D}}[\bar{x}_i] \right\}. \quad (34)$$

Then it follows from (32) and (34) that

$$\begin{aligned} (21) &\iff \mathbb{P}_{\mathcal{D}}[x_i] + A^{-1} \times \mathbb{P}_{\mathcal{D}}[\bar{x}_i] \geq e^{-\epsilon} \\ &\iff 1 - \mathbb{P}_{\mathcal{D}}[x_i] \geq A(e^{-\epsilon} - \mathbb{P}_{\mathcal{D}}[x_i]); \end{aligned}$$

i.e., (28) is established.

Proving (29):

By Bayes' theorem, it holds that

$$\mathbb{P}_{\mathcal{D},Y}[\bar{x}_i \mid \mathcal{Y}] = \frac{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i] \mathbb{P}_{\mathcal{D}}[\bar{x}_i]}{\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}]}. \quad (35)$$

Given (35), we have

$$(22) \iff \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i] \geq e^{-\epsilon} \times \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i]. \quad (36)$$

We recall (34). Applying (27) to (34), we obtain

$$\begin{aligned} &\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y}] \\ &= A \times \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} \mid \bar{x}_i] \times \left\{ \mathbb{P}_{\mathcal{D}}[x_i] + A^{-1} \times \mathbb{P}_{\mathcal{D}}[\bar{x}_i] \right\}. \end{aligned} \quad (37)$$

Then it follows from (36) and (37) that

$$\begin{aligned} (22) &\iff A \times \left\{ \mathbb{P}_{\mathcal{D}}[x_i] + A^{-1} \times \mathbb{P}_{\mathcal{D}}[\bar{x}_i] \right\} \leq e^\epsilon \\ &\iff A \times \mathbb{P}_{\mathcal{D}}[x_i] + 1 - \mathbb{P}_{\mathcal{D}}[x_i] \leq e^\epsilon; \end{aligned}$$

i.e., (29) is established.

Proving (30):

With $\mathbb{P}_{\mathcal{D}}[x_i]$ replaced by real $x \in [0, 1]$, (30) will follow once we show for $x \in [0, 1]$ that

$$\left. \begin{aligned} &1 - x \geq A(e^{-\epsilon} - x), \\ &\text{and } A \times x + 1 - x \leq e^\epsilon \end{aligned} \right\} \iff A \leq g(x). \quad (38)$$

We first prove the “ \implies ” part in (38). If $0 \leq x < e^{-\epsilon}$, we obtain from $1 - x \geq A(e^{-\epsilon} - x)$ that $A \leq \frac{1-x}{e^{-\epsilon}-x}$. If $0 < x \leq 1$, we obtain from $A \times x + 1 - x \leq e^\epsilon$ that $A \leq \frac{e^\epsilon - 1 + x}{x}$. With $g_1(x)$ denoting $\frac{1-x}{e^{-\epsilon}-x}$ for $0 \leq x < e^{-\epsilon}$ and $g_2(x)$ denoting $\frac{e^\epsilon - 1 + x}{x}$ for $0 < x \leq 1$, we see that $g(x)$ equals $g_1(x)$ if $0 \leq x \leq \frac{x}{1+e^\epsilon}$, and equals $g_2(x)$ if $\frac{1}{1+e^\epsilon} < x \leq 1$. Given the above, if $0 \leq x \leq \frac{1}{1+e^\epsilon}$, we have $A \leq g_1(x) = g(x)$, and if $\frac{1}{1+e^\epsilon} < x \leq 1$, we have $A \leq g_2(x) = g(x)$. Hence, the “ \implies ” part in (38) immediately follows.

We then prove the “ \impliedby ” part in (38). For any $x \in [0, 1]$, we will establish i) $1 - x \geq A(e^{-\epsilon} - x)$, and ii) $A \times x + 1 - x \leq e^\epsilon$, respectively. We still use $g_1(x)$ and $g_2(x)$ defined above. Note that $g_1(x)$ is only defined for $0 \leq x < e^{-\epsilon}$ and $g_2(x)$ is only defined for $0 < x \leq 1$. It is straightforward to show $g_1(x) \leq g_2(x)$ if $0 < x \leq \frac{1}{1+e^\epsilon}$, and $g_1(x) \geq g_2(x)$ if $\frac{1}{1+e^\epsilon} < x < e^{-\epsilon}$.

- i) If $0 \leq x \leq \frac{1}{1+e^\epsilon}$, we obtain from $A \leq g(x) = g_1(x)$ that $A \leq \frac{1-x}{e^{-\epsilon}-x}$, implying $1 - x \geq A(e^{-\epsilon} - x)$. If $\frac{1}{1+e^\epsilon} < x < e^{-\epsilon}$, we obtain from $A \leq g(x) = g_2(x) \leq g_1(x)$ that $A \leq \frac{1-x}{e^{-\epsilon}-x}$, yielding $1 - x \geq A(e^{-\epsilon} - x)$. If $e^{-\epsilon} \leq x \leq 1$, it holds that $1 - x \geq 0 \geq A(e^{-\epsilon} - x)$. To summarize, for any $x \in [0, 1]$, it follows that $1 - x \geq A(e^{-\epsilon} - x)$.
- ii) If $\frac{1}{1+e^\epsilon} < x \leq 1$, we obtain from $A \leq g(x) = g_2(x)$ that $A \leq \frac{e^\epsilon - 1 + x}{x}$, implying $A \times x + 1 - x \leq e^\epsilon$. If $0 < x \leq \frac{1}{1+e^\epsilon}$, we obtain from $A \leq g(x) = g_1(x) \leq g_2(x)$ that $A \leq \frac{e^\epsilon - 1 + x}{x}$, yielding $A \times x + 1 - x \leq e^\epsilon$. If $x = 0$, we have $A \times x + 1 - x = 1 \leq e^\epsilon$. To summarize, for any $x \in [0, 1]$, it follows that $A \times x + 1 - x \leq e^\epsilon$.

(38) is proved since its “ \implies ” and “ \impliedby ” both hold. \blacksquare

2) From membership privacy to Bayesian differential privacy:

Theorem 4. For a family of distributions that are downward scalable in the sense of Li et al. [4], ϵ -membership privacy implies ϵ -Bayesian differential privacy.

Proof of Theorem 4. The proof is similar to that of [4, Theorem 3.6]. For completeness, we still present the details below.

Assume, for the sake of contradiction, that mechanism Y achieves ϵ -membership privacy yet does not satisfy ϵ -Bayesian

differential privacy. Then there exists a distribution \mathcal{D} and entity x_i such that $0 < \mathbb{P}_{\mathcal{D}}[x_i] < 1$ and $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | x_i] > e^\epsilon \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i]$. We discuss two cases below.

Case one: $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i] = 0$ and $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | x_i] > 0$. Since \mathbb{D} is downward scalable, by definition \mathbb{D} contains some \mathbb{D}' which is x_i -scaled from \mathbb{D} such that $\mathbb{P}_{\mathcal{D}'}[x_i] < e^{-\epsilon}$. From [4, Lemma 3.4], we have $\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] = \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i]$, which with the case condition $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i] = 0$ means $\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] = 0$, further yielding $\mathbb{P}_{\mathcal{D}',Y}[x_i | \mathcal{Y}] = 1$. Therefore, $\mathbb{P}_{\mathcal{D}',Y}[x_i | \mathcal{Y}] = 1 > e^\epsilon \mathbb{P}_{\mathcal{D}'}[x_i]$, which contradicts the fact that Y achieves ϵ -membership privacy.

Case two: $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | x_i] = \alpha \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i]$, where $\alpha > e^\epsilon$. Since \mathbb{D} is downward scalable, by definition \mathbb{D} contains some \mathbb{D}' which is x_i -scaled from \mathbb{D} such that $\mathbb{P}_{\mathcal{D}'}[x_i] = q$ for an arbitrarily small q (see [4] for the meaning of “ x_i -scaled”). From [4, Lemma 3.4], we have $\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | x_i] = \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | x_i]$ and $\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] = \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i]$. These with the case condition $\mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | x_i] = \alpha \mathbb{P}_{\mathcal{D},Y}[\mathcal{Y} | \bar{x}_i]$ gives $\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | x_i] = \alpha \mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i]$. Then, under \mathcal{D}' , we have

$$\begin{aligned} \frac{\mathbb{P}_{\mathcal{D}',Y}[x_i | \mathcal{Y}]}{\mathbb{P}_{\mathcal{D}'}[x_i]} &= \frac{\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y}]} \\ &= \frac{\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | x_i]}{\mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | x_i] \mathbb{P}_{\mathcal{D}'}[x_i] + \mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] \mathbb{P}_{\mathcal{D}'}[\bar{x}_i]} \\ &= \frac{\alpha \mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i]}{\alpha \mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] \cdot q + \mathbb{P}_{\mathcal{D}',Y}[\mathcal{Y} | \bar{x}_i] \cdot (1 - q)} \\ &= \frac{\alpha}{\alpha q + 1 - q}. \end{aligned} \quad (39)$$

The above ratio $\frac{\alpha}{\alpha q + 1 - q}$ is greater than e^ϵ given $\alpha > e^\epsilon$, once we ensure $q < \frac{\alpha - e^\epsilon}{e^\epsilon(\alpha - 1)}$. This will give $\mathbb{P}_{\mathcal{D}',Y}[x_i | \mathcal{Y}] > e^\epsilon \mathbb{P}_{\mathcal{D}'}[x_i]$, which contradicts the fact that Y achieves ϵ -membership privacy.

Summarizing the above two cases, we have proved the desired result. \blacksquare

IV. RELATED WORK

The notion of *differential privacy* (DP) [1], [5] provides a rigorous foundation for privacy protection. Intuitively, DP implies that changing one entry in the database does not significantly change the query output, so that an adversary, seeing the query output and knowing all records except the one to be inferred, draws almost the same conclusion on whether or not a record is in the database. Differential privacy has received considerable interest in the literature [17]–[26]. Yang *et al.* [2] and Liu *et al.* [27] propose Bayesian differential privacy and dependent differential privacy respectively to generalize differential privacy for correlated data. Kasiviswanathan and Smith [3] propose a Bayesian formulation of semantic privacy, inspired by the following interpretation of differential privacy explained in [1]: *Regardless of external knowledge, an adversary with access to the sanitized database draws the same conclusions whether or not any individual data is included in the original database.* To present the notion of semantic privacy, Kasiviswanathan and Smith model the external knowledge via a prior probability distribution, and model conclusions via the corresponding posterior distribution.

Li *et al.* [4] introduce membership privacy (MP) in consideration of the adversary’s prior beliefs.

Dwork and Rothblum [28] recently proposed the notion of *concentrated differential privacy*, a relaxation of differential privacy achieving better accuracy than differential privacy without compromising on cumulative privacy cost over multiple computations. Motivated by [28], Bun and Steinke [29] suggest a relaxation of concentrated differential privacy. Instead of treating the privacy loss as a subgaussian random variable as [28] does, Bun and Steinke [29] instead formulate the problem in terms of Renyi entropy, giving a relaxation of concentrated differential privacy. Jorgensen *et al.* [30] introduce a new privacy definition called personalized differential privacy, a generalization of differential privacy in which users specify a personal privacy level for their data. They show that by accepting that not all users demand the same level of privacy, a higher level of utility can often be obtained by not providing excess privacy budget to those who do not need it. They present a mechanism for achieving personalized differential privacy, inspired by the well-known exponential mechanism of differential privacy. Hall *et al.* [31] introduce additional randomness to extend differential privacy to the notion of random differential privacy. Compared with differential privacy, Lee and Clifton [32] give an alternate formulation, differential identifiability, parameterized by the probability of individual identification. Their notion provides the strong privacy guarantees of differential privacy, while allowing policy makers to set parameters based on the privacy concept of individual identifiability.

Bohli and Andreas [33] discuss the relations among several privacy definitions, but the discussion does not cover differential privacy. Li *et al.* [34] present the relation between k -anonymization and differential privacy, where the k -anonymity notion by [35], [36] means that when only quasi-identifiers are considered, each record in a k -anonymized dataset should appear at least k times. Wang *et al.* [37] analyze the relation between differential privacy, mutual-information privacy, and identifiability. Mironov *et al.* [38] present several relaxations of differential privacy by requiring privacy guarantees to hold only against computationally bounded adversaries. They establish various relations among these notions, and show that the notions exhibit close connection with the theory of pseudodense sets [39].

V. CONCLUSION

In this paper, we present a comprehensive view of the relations among different privacy notions: differential privacy (DP), Bayesian differential privacy (BDP), semantic privacy (SP), and membership privacy (MP). In particular, we extend the notion of semantic privacy (SP) to Bayesian semantic privacy (BSP) and prove its essential equivalence with Bayesian differential privacy (BDP) in the quantitative sense. We show the relations between BDP, BSP, and SP as follows: ϵ -BDP $\Leftarrow (\frac{1}{2} - \frac{1}{e^\epsilon + 1})$ -BSP, and ϵ -BDP $\Rightarrow (e^{2\epsilon} - 1)$ -BSP $\Rightarrow (e^{2\epsilon} - 1)$ -SP. Moreover, we derive the following relations between BDP and MP. First, ϵ -BDP $\Rightarrow \epsilon$ -MP. Second, For a family of

distributions that are downward scalable in the sense of Li *et al.* [4], it holds that ϵ -BDP \Leftarrow ϵ -MP.

REFERENCES

- [1] C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [2] B. Yang, I. Sato, and H. Nakagawa, “Bayesian differential privacy on correlated data,” in *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2015, pp. 747–762.
- [3] S. Kasiviswanathan and A. Smith, “On the ‘semantics’ of differential privacy: A Bayesian formulation,” *Journal of Privacy and Confidentiality*, vol. 6, no. 1, pp. 1–16, 2014.
- [4] N. Li, W. Qardaji, D. Su, Y. Wu, and W. Yang, “Membership privacy: A unifying framework for privacy definitions,” in *ACM Conference on Computer and Communications Security (CCS)*, 2013, pp. 889–900.
- [5] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of Cryptography Conference (TCC)*, 2006, pp. 265–284.
- [6] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, “Geo-indistinguishability: Differential privacy for location-based systems,” in *ACM Conference on Computer and Communications Security (CCS)*, 2013, pp. 901–914.
- [7] J. Zhang, X. Xiao, and X. Xie, “PrivTree: A differentially private algorithm for hierarchical decompositions,” in *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2016, pp. 155–170.
- [8] J. Blocki, A. Datta, and J. Bonneau, “Differentially private password frequency lists,” in *Network and Distributed System Security (NDSS) Symposium*, 2016.
- [9] Y. Xiao and L. Xiong, “Protecting locations with differential privacy under temporal correlations,” in *ACM Conference on Computer and Communications Security (CCS)*, 2015, pp. 1298–1309.
- [10] X. Lou, R. Tan, D. K. Yau, and P. Cheng, “Cost of differential privacy in demand reporting for smart grid economic dispatch,” in *IEEE Conference on Computer Communications (INFOCOM)*, 2017.
- [11] R. Shokri and V. Shmatikov, “Privacy-preserving deep learning,” in *ACM Conference on Computer and Communications Security (CCS)*, 2015, pp. 1310–1321.
- [12] S. Song and K. Chaudhuri, “Composition properties of inferential privacy for time-series data,” in *Allerton Conference on Communication, Control, and Computing*, October 2017.
- [13] S. Song, Y. Wang, and K. Chaudhuri, “Pufferfish privacy mechanisms for correlated data,” in *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2017, pp. 1291–1306.
- [14] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *IEEE Symposium on Security and Privacy*, 2017, pp. 3–18.
- [15] N. Phan, X. Wu, H. Hu, and D. Dou, “Adaptive Laplace mechanism: Differential privacy preservation in deep learning,” in *IEEE International Conference on Data Mining series (ICDM)*, 2017.
- [16] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang, “Privacy loss in Apple’s implementation of differential privacy on MacOS 10.12,” *arXiv preprint arXiv:1709.02753*, 2017.
- [17] Ú. Erlingsson, V. Pihur, and A. Korolova, “RAPPOR: Randomized aggregatable privacy-preserving ordinal response,” in *ACM Conference on Computer and Communications Security (CCS)*, 2014, pp. 1054–1067.
- [18] N. Wang, X. Xiao, Y. Yang, Z. Zhang, Y. Gu, and G. Yu, “PrivSuper: A superset-first approach to frequent itemset mining under differential privacy,” in *IEEE International Conference on Data Engineering (ICDE)*, 2017, pp. 809–820.
- [19] R. Bassily, A. Groce, J. Katz, and A. Smith, “Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy,” in *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013, pp. 439–448.
- [20] J. Zhang, G. Cormode, C. M. Procopiuc, D. Srivastava, and X. Xiao, “Private release of graph statistics using ladder functions,” in *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2015, pp. 731–745.
- [21] F. Tramèr, Z. Huang, J.-P. Hubaux, and E. Ayday, “Differential privacy with bounded priors: Reconciling utility and privacy in genome-wide association studies,” in *ACM Conference on Computer and Communications Security (CCS)*, 2015, pp. 1286–1297.
- [22] Z. Qin, Y. Yang, T. Yu, I. Khalil, X. Xiao, and K. Ren, “Heavy hitter estimation over set-valued data with local differential privacy,” in *ACM Conference on Computer and Communications Security (CCS)*, 2016, pp. 192–203.
- [23] K. Jiang, D. Shao, S. Bressan, T. Kister, and K.-L. Tan, “Publishing trajectories with differential privacy guarantees,” in *International Conference on Scientific and Statistical Database Management (SSDBM)*, 2013, pp. 12:1–12:12.
- [24] G. Acs, L. Melis, C. Castelluccia, and E. De Cristofaro, “Differentially private mixture of generative neural networks,” in *IEEE International Conference on Data Mining series (ICDM)*, 2017.
- [25] J. Zhao and J. Zhang, “Preserving privacy enables “coexistence equilibrium” of competitive diffusion in social networks,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 282–297, June 2017.
- [26] P. Mohassel and Y. Zhang, “SecureML: A system for scalable privacy-preserving machine learning,” in *IEEE Symposium on Security and Privacy*, 2017, pp. 19–38.
- [27] C. Liu, S. Chakraborty, and P. Mittal, “Dependence makes you vulnerable: Differential privacy under dependent tuples,” in *Network and Distributed System Security (NDSS) Symposium*, 2016.
- [28] C. Dwork and G. N. Rothblum, “Concentrated differential privacy,” *arXiv preprint arXiv:1603.01887*, 2016.
- [29] M. Bun and T. Steinke, “Concentrated differential privacy: Simplifications, extensions, and lower bounds,” in *Theory of Cryptography Conference*. Springer, 2016, pp. 635–658.
- [30] Z. Jorgensen, T. Yu, and G. Cormode, “Conservative or liberal? Personalized differential privacy,” in *International Conference on Data Engineering (ICDE)*, 2015, pp. 1023–1034.
- [31] R. Hall, A. Rinaldo, and L. Wasserman, “Random differential privacy,” *Journal of Privacy and Confidentiality*, vol. 4, no. 2, pp. 43–59, 2012.
- [32] J. Lee and C. Clifton, “Differential identifiability,” in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2012, pp. 1041–1049.
- [33] J.-M. Bohli and A. Pashalidis, “Relations among privacy notions,” *ACM Transactions on Information and System Security (now ACM Transactions on Privacy and Security)*, vol. 14, no. 1, p. 4, 2011.
- [34] N. Li, W. Qardaji, and D. Su, “On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy,” in *ACM Symposium on Information, Computer and Communications Security (AsiaCCS)*, 2012, pp. 32–33.
- [35] L. Sweeney, “k-Anonymity: A model for protecting privacy,” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [36] P. Samarati, “Protecting respondents identities in microdata release,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 1010–1027, 2001.
- [37] W. Wang, L. Ying, and J. Zhang, “On the relation between identifiability, differential privacy, and mutual-information privacy,” *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5018–5029, 2016.
- [38] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan, “Computational differential privacy,” in *Advances in Cryptology (CRYPTO)*, 2009, pp. 126–142.
- [39] O. Reingold, L. Trevisan, M. Tulsiani, and S. Vadhan, “Dense subsets of pseudorandom sets,” in *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2008, pp. 76–85.