

# Heavy-Tailed Privacy: The Symmetric alpha-Stable Privacy Mechanism

**Christopher C. Zawacki**

CZAWACKI@UMD.EDU

*Dept. of Electrical and Computer Engineering  
University of Maryland  
College Park, MD 20742 USA*

**Eyad H. Abed**

ABED@UMD.EDU

*Dept. of Electrical and Computer Engineering  
University of Maryland  
College Park, MD 20742 USA*

**Editor:** Christopher C. Zawacki and Eyad H. Abed

## Abstract

With the rapid growth of digital platforms, there is increasing apprehension about how personal data is collected, stored, and used by various entities. These concerns arise from the increasing frequency of data breaches, cyber-attacks, and misuse of personal information for targeted advertising and surveillance. To address these matters, Differential Privacy (DP) has emerged as a prominent tool for quantifying a digital system's level of protection. The Gaussian mechanism is commonly used because the Gaussian density is closed under convolution, and is a common method utilized when aggregating datasets. However, the Gaussian mechanism only satisfies an approximate form of Differential Privacy. In this work, we present and analyze of the Symmetric alpha-Stable (SaS) mechanism. We prove that the mechanism achieves pure differential privacy while remaining closed under convolution. Additionally, we study the nuanced relationship between the level of privacy achieved and the parameters of the density. Lastly, we compare the expected error introduced to dataset queries by the Gaussian and SaS mechanisms. From our analysis, we believe the SaS Mechanism is an appealing choice for privacy-focused applications.

**Keywords:** Differential Privacy, Stable distributions, Data Privacy, Heavy Tails, Federated Learning

## 1 INTRODUCTION

Privacy is fundamental to individual autonomy and personal safety. It protects individuals from harassment and discrimination, fosters trust in institutions, and encourages free speech and innovation. As the world becomes increasingly digital, we have seen in Id-Theft-Center (2022) a corresponding increase in data breaches targeting the growing number of individual databases that hold client information. In recent years, the public and private sectors have begun to act. Political leaders are taking action to ensure the privacy of their citizens, for example the Internet Freedom Act USA (2011) and the General Data Protection Regulation EU (2016), and consumers are putting pressure on companies to adopt settings and methods that focus on the privacy of their customers, as discussed in Koetsier (2021) and Minto and Haller (2021).

Introduced by Dwork (2006); Dwork et al. (2006a), one common approach to protecting client data is known as Differential Privacy (DP). Differentially private systems inject carefully constructed noise into a dataset to obfuscate the specific participants, while retaining general trends in machine learning datasets. The Differential Privacy framework has been applied in various domains, from large-scale data analysis to machine learning; see Abadi et al. (2016). More recently, Differential Privacy has received renewed attention within the field of federated learning (FL), a privacy focused branch of machine learning introduced by McMahan et al. (2016). The objective of differentially private FL methods are to enhance privacy preservation while collaboratively training machine learning models across multiple decentralized devices or servers Wei et al. (2020). In Li et al. (2019), the authors use differentially private federated learning methods to train a machine learning model that segments images of brain tumors.

The work here extends the initial results presented in Zawacki and Abed (2024) with additional insight into the how the expected error increases as the level of noise increases. Related to the theme of this work, other groups have begun to examine the benefits of using heavy-tailed distributions within the differential privacy framework. Ito et al. (2021) use heavy-tailed distributions to mask contributions by outliers in the of filter/controller design for control systems. In Asi et al. (2024), the authors determine optimal rates of convergence (up to a logarithm) in the context of private convex optimization with heavy tailed distributions. Şimşekli et al. (2024) show that under broad conditions, the use of heavy-tailed distributions in differentially private stochastic gradient descent (SGD) eliminates the need for a projection step, decreasing the computational complexity. Our results differ in the level of privacy guaranteed by the privacy mechanism and we additionally provide a deeper analysis on the relationship between the parameters of the density and their effect on the level of protection.

The contributions of this work are threefold. First we introduce our new privacy mechanism which is based on the use of stable densities and prove that this mechanism is  $\epsilon$ -differentially private. Second, we show that the level of privacy scales inversely with the level of injected noise; aligning its behavior with existing privacy mechanisms. Lastly, we compare the expected distortion of our privacy mechanism against other commonly utilized privacy mechanisms.

The rest of the paper is organized as follows. Section 2 summarizes the basics of Differential Privacy. Section 3 introduces the definition of the Symmetric alpha-Stable Mechanism. Section 4 proves the privacy guarantee of the new mechanism. Section 5 studies how the privacy scales with the level of noise. Section 6 provides a measure of error for the mechanism introduces. Section 7 summarizes the results and provides comments on active related research efforts.

## 2 BACKGROUND

In this section we outline the background material required to derive our results.

### 2.1 Differential Privacy

Differential Privacy operates on a collaboratively constructed dataset, which we denote by  $\mathcal{D}$ . Conceptually, we can think of such a dataset as a table of records, where each row

represents a set of client data. Denote by  $f$  a function that operates on this dataset and produces a vector of  $m$  numerical values. For instance:

- How many clients have blue eyes?
- What is the average income of all clients?
- What are the optimized parameters of a given machine learning model across all clients?

Differential Privacy primarily prevents *passive* adversaries from obtaining information about a target client. A passive adversary is one that observes communications or interactions without actively tampering with them. For example, a passive adversary may eavesdrop on communication between a client and the server or combine publicly available datasets in a linking attack to re-identify anonymized client data. This is in contrast with an active adversary which seeks to directly disrupt model training.

By a slight abuse of notation, we use the symbol  $f$  for the query, despite the composition of the dataset,

**Definition 1** (*Query*) *A function  $f$  is termed a query if it takes a dataset  $\mathcal{D}$  as input and outputs a vector in  $\mathbb{R}^m$ :*

$$f : \mathcal{D} \rightarrow \mathbb{R}^m. \tag{1}$$



The types of queries commonly employed in Differential Privacy methods are those exhibiting finite  $\ell_p$ -sensitivity Dwork (2006); Dwork and Roth (2014):

**Definition 2** ( *$\ell_p$ -Sensitivity of Query*) *The  $\ell_p$ -sensitivity of a query  $f$ , denoted  $\Delta_p f$ , is defined to be a maximum of a  $p$ -norm over the domain of  $f$ ,  $\text{dom}(f)$ :*

$$\Delta_p f := \max_{\mathcal{D}_1 \simeq \mathcal{D}_2} \|f(\mathcal{D}_1) - f(\mathcal{D}_2)\|_p, \tag{2}$$

for all  $\mathcal{D}_1, \mathcal{D}_2 \in \text{dom}(f)$ .



From Definitions 1 and 2 above, it is clear that when the sensitivity of  $f$  is bounded, the range of the query is also bounded. While focusing on finite queries here, ongoing research aims to extend differentially private methods to handle queries with unbounded ranges, see Durfee (2024).

Now, we recall the definition of a privacy mechanism, which introduces stochastic noise to the result of a query.

**Definition 3** (*Privacy mechanism*) *A privacy mechanism for the query  $f$ , denoted  $\mathcal{M}_f$ , is defined to be a randomized algorithm that returns the result of the query perturbed by a vector of i.i.d. noise sampled from pre-selected densities  $Y_i$ ,*

$$\mathcal{M}_f(\mathcal{D}) = f(\mathcal{D}) + (Y_1, Y_2, \dots, Y_m)^T. \tag{3}$$



To simplify notation, we denote the resulting vector,  $\mathcal{M}_f(\mathcal{D})$ , as  $\mathbf{x} \in \mathbb{R}^m$ . Note that the noise variables,  $Y_i$ , induce a density, which we occasionally denote  $p = p(\mathbf{x})$  for  $\mathcal{M}_f$ , on a given dataset  $\mathcal{D}$ . Although not strictly necessary, we assume that the injected density is symmetric about the origin, simplifying the analysis.

The privacy mechanism aims to hinder an adversary from conclusively ascertaining the presence of a specific client within the dataset.

**Definition 4** (*Neighboring Datasets*) Two datasets, denoted  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , are known as neighboring datasets if they differ in the presence or absence of exactly one client record. We denote this relation as  $\mathcal{D}_1 \simeq \mathcal{D}_2$ . ◀

This concept is visualized in Figure 1, which depicts two datasets, one that contains the red client, and one that does not. Let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  represent these two scenarios respectively. To proceed, let us assume the red client has allowed their data to be included in the set and that  $\mathcal{D}_1$  is the *true* dataset. Denote a realization of a mechanism as  $x \sim \mathcal{M}_f(\mathcal{D}_a)$ . Informally, the mechanism  $\mathcal{M}_f$  is said to be differentially private if the inclusion or exclusion of a single individual in the dataset, illustrated in red in the figure, results in *essentially* the same distribution over the realized outputs  $x$ ,

$$\Pr[\mathcal{M}_f(\mathcal{D}_1) = x] \approx \Pr[\mathcal{M}_f(\mathcal{D}_2) = x]. \tag{4}$$

Differential Privacy then quantifies what *essentially* means mathematically:

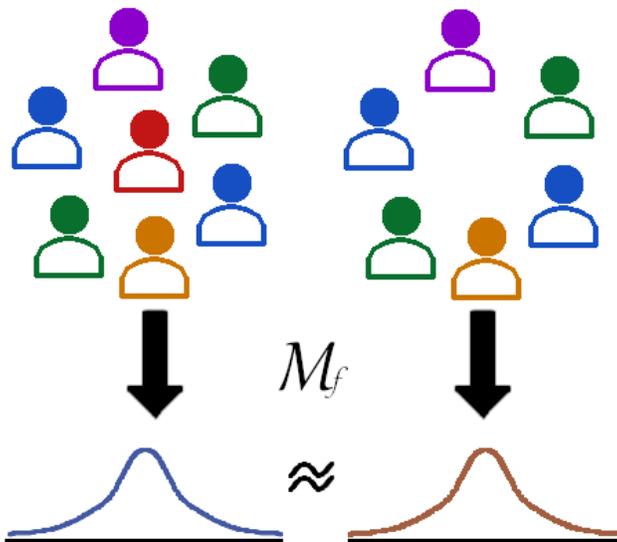


Figure 1: In order to protect client identity, Differential Privacy injects noise into the output of a query  $f$  on a dataset. This induces a probability density over possible outcomes. A mechanism,  $\mathcal{M}_f$ , is considered private, if the resulting distributions are *essentially* the same regardless of the inclusion or exclusion of a single client, shown here in red. Differential Privacy quantifies how much information an adversary is able to gain about the red client.

**Definition 5** (*Pure Differential Privacy*) Let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be any neighboring datasets. Given a query  $f$  that operates on  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , a privacy mechanism  $\mathcal{M}_f$  is said to be

$\epsilon$ -Differentially Private ( $\epsilon$ -DP or pure-DP) if it satisfies

$$\Pr[\mathcal{M}_f(\mathcal{D}_1) \in \mathcal{X}] \leq e^\epsilon \Pr[\mathcal{M}_f(\mathcal{D}_2) \in \mathcal{X}] \quad (5)$$

for some  $\epsilon > 0$  and any subset of outputs  $\mathcal{X} \subseteq \mathcal{R}(\mathcal{M}_f(\mathcal{D}_1))$ . The mechanism is defined to have no privacy ( $\epsilon = \infty$ ) if, upon its application to each dataset, the supports of the resulting densities are not equal, viz.  $\mathcal{R}(\mathcal{M}_f(\mathcal{D}_1)) \neq \mathcal{R}(\mathcal{M}_f(\mathcal{D}_2))$ . ◀

The parameter  $\epsilon$  is also referred to as the privacy budget. Smaller values of  $\epsilon$  are associated with stronger privacy. We remark that when  $\epsilon = 0$ , the definition yields perfect privacy. However, in that case, adding more client data results in no new information.

Note that Eq. 5 holds for each element when the density of the distributions is considered in Dwork and Roth (2014):

**Theorem 6** (*Privacy as Densities*) Let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be neighboring datasets and  $f$  be a query that operates on them. Denote by  $p_1$  and  $p_2$  the densities of the privacy mechanism  $\mathcal{M}_f$  when applied to  $\mathcal{D}_1$  and  $\mathcal{D}_2$  respectively. Then, a privacy mechanism  $\mathcal{M}_f$  is  $\epsilon$ -Differentially Private if

$$p_1(x) \leq e^\epsilon p_2(x), \quad \forall x \in \mathcal{R}(\mathcal{M}_f(\mathcal{D}_1)) \quad (6)$$

for all  $\mathcal{D}_1 \simeq \mathcal{D}_2$ . ◀

Next, we give a brief proof for this known fact.

**Proof** Begin by writing condition (5) in terms of the generated densities,

$$\int_{\mathcal{X}} p_1(x) dx \leq \int_{\mathcal{X}} e^\epsilon p_2(x) dx. \quad (7)$$

Equation (7) can be rewritten as

$$0 \leq \int_{\mathcal{X}} e^\epsilon p_2(x) - p_1(x) dx. \quad (8)$$

Noting that (6) enforces the integrand in (8) to be non-negative, implying that (5) is satisfied. ■

Next, we recall a metric for evaluating the loss of privacy experienced by a participating client under a given privacy mechanism.

**Definition 7** (*Privacy Loss*) The privacy loss of an outcome  $x$  is defined to be the log-ratio of the densities when the mechanism is applied to  $\mathcal{D}_1$  and  $\mathcal{D}_2$  at  $x$  Dwork and Roth (2014):

$$\mathcal{L}_{\mathcal{D}_1 \parallel \mathcal{D}_2}(x) := \ln \frac{p_1(x)}{p_2(x)}. \quad (9)$$

By (6), it is evident that  $\epsilon$ -Differential Privacy (5) is equivalent to

$$|\mathcal{L}_{\mathcal{D}_1 \parallel \mathcal{D}_2}(x)| \leq \epsilon, \quad \forall x \in \mathcal{R}(\mathcal{M}_f(\mathcal{D}_1)) \quad (10)$$

for all neighboring datasets  $\mathcal{D}_1$  and  $\mathcal{D}_2$ . ◀

For mechanisms that are purely differential private, the privacy budget  $\varepsilon$  is the maximum over all observations  $x$ ,

$$\varepsilon = \max_{x \in \mathbb{R}} \mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}(x). \quad (11)$$

Due to its beneficial mathematical properties, the Gaussian density is a commonly chosen density for Differential Privacy. However, the Gaussian mechanisms fails to satisfy condition (5). To accommodate this, the condition can be relaxed through the inclusion of an additive constant  $\delta > 0$ , as in the following definition:

**Definition 8** (*Approximate Differential Privacy*) *Let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be any neighboring datasets. Given a query  $f$  that operates on  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , a privacy mechanism  $\mathcal{M}_f$  is said to be  $(\varepsilon, \delta)$ -Differentially Private if it satisfies*

$$\Pr[\mathcal{M}_f(\mathcal{D}_1) \in \mathcal{X}] \leq e^\varepsilon \Pr[\mathcal{M}_f(\mathcal{D}_2) \in \mathcal{X}] + \delta. \quad (12)$$

*This is known as approximate-Differential Privacy.* ◀

The accepted error term  $\delta$  is the probability that the result of the query provides more information to the adversary than expected from the bound  $\varepsilon$ .

One common modification relates to who applies the privacy mechanism. Up to this point, we have considered a mechanism in relation to a query over the entire dataset  $\mathcal{D}$ . It is then understood that the mechanism is applied by a *trusted aggregator*, who collects the clients' data prior to obfuscation. However, there does not always exist such a trusted central authority. For example, in a Federated Learning framework, the server is assumed untrustworthy by default. Another situation where clients may wish to apply noise locally is if they lack secure communication protocols. In this case, the clients' sensitive information could be leaked to an adversary during the transmission between the clients and server. To this end, a mechanism  $\mathcal{M}_f$  is said to be Locally Differentially Private (LDP) if the mechanism can be applied locally by the clients prior to transmission to the server.

**Definition 9** (*Local Differential Privacy*) *Let a client apply the privacy mechanism  $\mathcal{M}_f^{loc}$  to their local dataset  $\mathcal{D}$ . The mechanism  $\mathcal{M}_f^{loc}$  is said to be locally differentially privacy if, for any pair of data points  $v_1, v_2 \in \mathcal{D}$ , it satisfies the following Kasiviswanathan et al. (2011):*

$$\Pr[\mathcal{M}_f^{loc}(v_1) \in \mathcal{X}] \leq e^\varepsilon \Pr[\mathcal{M}_f^{loc}(v_2) \in \mathcal{X}] + \delta, \quad (13)$$

*for all  $\mathcal{X} \in \mathcal{R}(\mathcal{M}_f^{loc})$ .* ◀

The mechanism is called  $\varepsilon$ -LDP if  $\delta = 0$  and  $(\varepsilon, \delta)$ -LDP otherwise.

## 2.2 Selecting a Level of Privacy

Wasserman and Zhou (2009); Geng and Viswanath (2015) describe a useful connection between Differential Privacy and hypothesis testing. Their analysis considers the problem of client privacy from the perspective of an adversary deciding between two hypotheses. Denote by  $\mathcal{D}_1$  and  $\mathcal{D}_2$  two neighboring datasets. Let one of the following hypotheses hold:

- $H_0$  (The null hypothesis): the true dataset is  $\mathcal{D}_1$ .

- $H_1$  (The alternative hypothesis): the true dataset is  $\mathcal{D}_2$ .

The objective of the adversary is to determine, based on the output of a privacy mechanism  $\mathcal{M}_f$ , which hypothesis is true. Denote by  $p$  the probability of a false positive, that is, the adversary chooses  $H_1$  when  $H_0$  is true. Then, denote by  $q$  the probability of a false negative, i.e.,  $H_0$  is chosen when  $H_1$  is true. The authors show that if a mechanism  $\mathcal{M}_f$  is  $\varepsilon$ -differentially private, then the following two statements must hold:

$$p + e^\varepsilon q \geq 1 \text{ and } e^\varepsilon p + q \geq 1. \tag{14}$$

Combining the inequalities in (14) yields

$$p + q \geq \frac{2}{1 + e^\varepsilon}. \tag{15}$$

Consider that when  $\varepsilon \ll 1$ , which equates to high privacy, the adversary cannot achieve both low false positive and low false negative rates simultaneously. Often, it is more convenient to specify lower bounds for  $p$  and  $q$  and to use (15) to determine  $\varepsilon$  than it is to state the privacy budget directly. When a mechanism that satisfies pure Differential Privacy is employed, the maximum information an adversary may learn is strictly bounded. Figure 2 depicts the upper and lower bounds for a given privacy budget  $\varepsilon$  and initial probability  $p_1$  on a simple dataset. When a mechanism that only achieves approximate-Differential

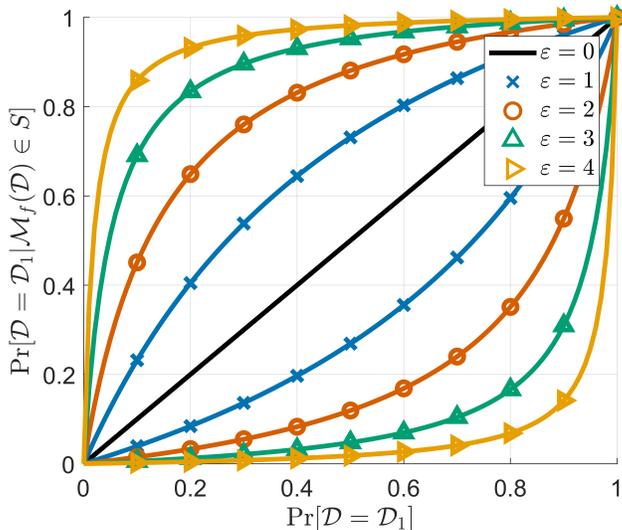


Figure 2: Pure-Differential Privacy limits the amount of information an adversary can gain from the outcome of private query. Based on the adversary’s initial estimate of the alternative hypothesis,  $\Pr[\mathcal{D} = \mathcal{D}_1]$ , a Differentially Private mechanism bounds the conditional probability given the outcome of the query. Each pair of matching curves represents the lower and upper bound for an adversary’s estimate of the alternative hypothesis after observing the outcome of the privacy mechanism. As the privacy budget  $\varepsilon$  is increased, the bound of the adversary’s updated estimate is increased.

Privacy is deployed, these bounds become probabilistic.

### 2.3 Common Privacy Mechanisms

The choice of probability density from which the noise is drawn significantly impacts the level of privacy achieved. To prevent bias from being introduced into the query, all mechanisms considered here are chosen to have zero mean.

Dwork et al. (2006b) and Dwork and Roth (2014) introduce the Laplace mechanism, one of the most commonly used mechanisms, which samples noise from the Laplace density:

$$p_{\text{Lap}}(x) = \frac{1}{2b} e^{-\frac{|x|}{b}}. \quad (16)$$

Here,  $b$  determines the spread of the distribution. The Laplace mechanism has been shown to satisfy pure Differential Privacy, equation (5) Dwork et al. (2006b); Dwork and Roth (2014). Unfortunately, the Laplace density does not trivially extend to local Differential Privacy, limiting its application in methods such as Federated Learning. This property additionally makes the Laplace mechanism challenging to use for training Neural Networks, which rely heavily on the repeated compositions of a mechanism, so it is not commonly employed for deep learning.

Another frequently employed mechanism is the Gaussian mechanism, studied in Dwork et al. (2006b) and Dwork and Roth (2014). This mechanism injects noise, drawn from a normal density with a mean of zero, into the output of a query:

$$p_{\text{Gaus}}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x}{\sigma}\right)^2}. \quad (17)$$

Since the Gaussian density is closed under convolution, the mechanism naturally extends to environments that require local application of the mechanism. However, the Gaussian mechanism is only approximately Differentially Private, i.e., it requires  $\delta > 0$  in equation (12). Traditionally, this has not been seen as an issue because, as shown in Dwork and Roth (2014), the repeated composition of approximate-DP methods scales better than simple composition of pure-DP methods.

The Exponential mechanism is another noteworthy approach introduced by McSherry and Talwar (2007). This mechanism selects outputs from a set probabilistically, weighting them according to their utility scores. By carefully choosing the scoring function, it provides a way to balance privacy and utility effectively. The noise added by the Exponential mechanism is drawn from the exponential density:

$$p_{\text{Exp}}(x) = \lambda e^{-\lambda x}. \quad (18)$$

Here,  $\lambda$  controls the rate of decay of the distribution. However, designing appropriate scoring functions that accurately capture utility while ensuring privacy remains a significant challenge in practical implementations.

With these common mechanisms in mind, we next proceed to define the Symmetric alpha-Stable mechanism and present novel analysis of its properties.

## 3 The Symmetric alpha-Stable Mechanism

To begin, it is essential to note that the Gaussian density belongs to a broader family of distributions called the Lévy alpha-Stable densities, all of which exhibit closure under

convolutions; see Lévy (1925). However, it is shown by Dwork et al. (2006a) and Dwork and Roth (2014) that, in the realm of Differential Privacy, the Gaussian mechanism only adheres to condition (12), approximate Differential Privacy. This section delves into the characteristics of a privacy mechanism drawn from a subset of the Lévy alpha-stable family. We refer to such mechanisms as Symmetric alpha-Stable mechanisms and provide a proof that they adhere to condition (5), pure Differential Privacy.

The concept of stable densities, extensively explored in Lévy (1925), refers to a particular set of probability distributions. These distributions possess a notable property; the convolution of two distributions from the family is also a member of the family; this is otherwise known as closure under convolution.

**Definition 10** (*The Stable Family*) *A probability density function  $Y$  is termed stable if, for any constants  $a, b > 0$ , there exist constants  $c(a, b) > 0$  and  $d(a, b) \in \mathbb{R}$  such that the following holds for two independent and identically distributed random variables  $Y_1$  and  $Y_2$ :*

$$aY_1 + bY_2 = cY + d. \tag{19}$$

*If  $d$  equals 0, the distribution is termed strictly stable.*



Nolan (2020) shows that, except for certain special cases, there is no known closed-form expression for the density of a general stable distribution. Nonetheless, several parameterizations of the characteristic function of a stable density are documented; see Nolan (2020). One common representation of the characteristic function is as follows:

$$\varphi(t; \alpha, \beta, \gamma, \mu) = \exp(it\mu - |\gamma t|^\alpha + i\beta \operatorname{sgn}(t)\Phi(t)), \tag{20}$$

where

$$\Phi(t) = \begin{cases} \tan(\frac{\pi\alpha}{2}) & \alpha \neq 1 \\ -\frac{2}{\pi} \log |t| & \alpha = 1. \end{cases} \tag{21}$$

The density function is then given by the integral:

$$p(x; \alpha, \beta, \gamma, \mu) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \varphi(t; \alpha, \beta, \gamma, \mu) e^{-ixt} dt. \tag{22}$$

In Figure 3, we present three examples of the symmetric form  $\beta = 0$ :  $\alpha = 1$  (blue),  $\alpha = 1.5$  (orange), and  $\alpha = 2$  (green). Each of the three depicted densities has zero for the location parameter ( $\mu = 0$ ) and unit scale ( $\gamma = 1$ ). The Symmetric alpha-Stable densities with  $\alpha = 1$  and  $\alpha = 2$  are the only two densities with support on the whole real line that have a known closed form. When  $\alpha = 1$ , the density is known as the Cauchy density. When  $\alpha = 2$ , we recover the Gaussian density.

Working with the family of stable distributions presents a significant challenge due to the absence of a closed-form solution for the general density. This difficulty arises because the value at any given point is determined by integrating an infinitely oscillating function.

Denote the real part of the integrand in Equation (22) by  $q(t; x)$ . A visualization of this function for  $x = 10$  is illustrated in Figure 4.

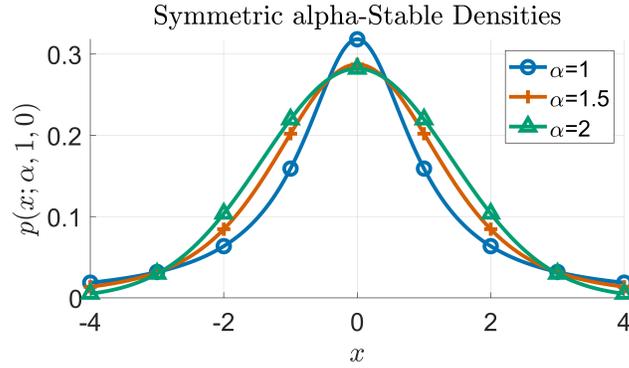


Figure 3: The family of Symmetric alpha-Stable densities consists of bell shaped densities with varying tail weights determined by the stability parameter  $\alpha$ . This family of densities is unique because it is the only set of densities that are closed under convolution. When  $\alpha = 1$ , shown in **blue**  $\circ$ , the density is known as the Cauchy. When  $\alpha = 2$ , shown in **green**  $\triangle$ , the density is known as the Gaussian. No other values of alpha (for the symmetric case) have a known closed form solution, for example  $\alpha = 1.5$ , shown in **orange**  $+$ .

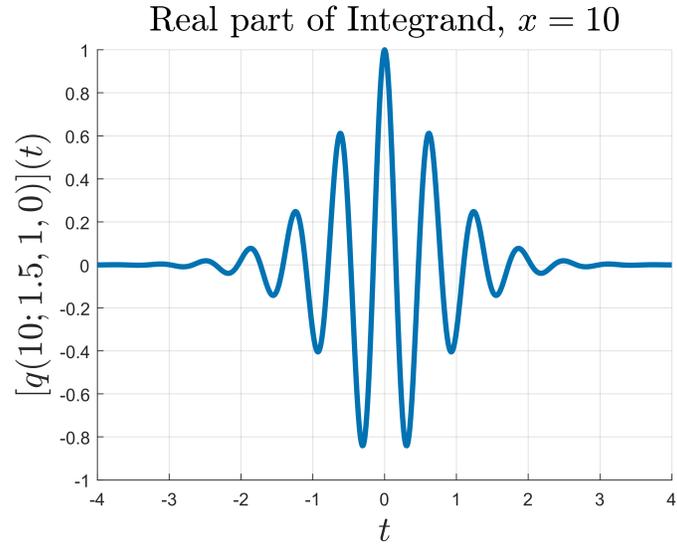


Figure 4: The real part of the integrand of (22) for  $\alpha = 1.5$ ,  $\gamma = 1$ , and  $\mu = 0$  is an infinitely oscillating function. The value of the stable density with these parameters at the point  $x = 10$  is the integral of this function on the real line.

In order for Equation (22) to represent a valid probability density, the parameter  $\alpha$  must fall within the interval  $(0, 2]$ . The value of  $\alpha$  dictates the rate of decay of the density's tail. The expected value of the density is only defined in the range  $\alpha \in (1, 2]$  and is not defined for  $\alpha \leq 1$ . Furthermore, the density exhibits infinite variance for  $\alpha \in (0, 2)$  and finite variance only when  $\alpha = 2$ . In this study, we confine  $\alpha$  to  $(1, 2]$ , deferring median or mode estimators for future exploration. The parameter  $\beta$ , constrained to  $(-1, 1)$ , serves as a measure linked to skewness, noting that the strict definition of skewness lacks meaning for  $\alpha < 2$ . Our focus lies specifically on symmetric alpha-stable (SaS) densities, where  $\beta = 0$ :

$$\begin{aligned}
 p_{\text{SaS}}(x; \alpha, \gamma, \mu) &:= \\
 p(x; \alpha, 0, \gamma, \mu) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha - it(x-\mu)} dt.
 \end{aligned} \tag{23}$$

SaS densities have a known closed form for two values of the parameter  $\alpha$ : the Cauchy, for  $\alpha = 1$ , and the Gaussian, for  $\alpha = 2$ . The last two parameters,  $\gamma > 0$  and  $\mu \in \mathbb{R}$ , are the scale and location parameter's respectively.

**Remark 11** *For stable densities, it is common for the location parameter to be denoted  $\delta$  rather than  $\mu$ , to signify that it is not always equal to the expected value. In our context, we choose to use  $\mu$  and reserve  $\delta$  for the definition of approximate Differential Privacy (12) as is common in the DP literature. Because we are restricting the domain of interest to  $\alpha \in (1, 2]$ , where the mean is well-defined, we do not believe this notation will be cause for confusion.*

We are now prepared to define the Symmetric alpha-Stable mechanism:

**Definition 12** *(The Symmetric alpha-Stable mechanism) For a given dataset  $\mathcal{D}$  and a query function  $f$ , we define a privacy mechanism  $\mathcal{M}_f$  to be a Symmetric alpha-Stable (SaS) mechanism if each element of the vector of injected values,  $Y_i$  for  $i \in \{1, \dots, m\}$ , is drawn independently from a SaS density*

$$\begin{aligned}
 p_{\text{SaS}}(x; \alpha, \gamma) &:= \\
 p(x; \alpha, 0, \gamma, 0) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha - itx} dt.
 \end{aligned} \tag{24}$$

◀

While this family of mechanisms is closely related to the Gaussian, the Gaussian mechanism only satisfies approximate Differential Privacy. We show in the next section that the heaviness of the SaS density's tail allows the privacy mechanism to satisfy pure-Differential Privacy (when  $\alpha < 2$ ).

## 4 Pure-Differential Privacy of SaS Mechanism

In this section, we establish that the SaS mechanism, when  $\alpha \in [1, 2)$ , satisfies (5), providing  $\epsilon$ -Differential Privacy. A significant challenge in working with stable densities, excluding the Cauchy and Gaussian, arises from their lack of a closed-form expression for the density. To ensure that the stable distribution covers the entire real number line, it's crucial to demonstrate that the privacy loss remains finite within a compact set. To prevent the denominator of the privacy loss in equation (7) from becoming zero, resulting in infinite privacy loss, we first provide a lemma that states that the density is nonzero over the whole real line.

**Lemma 13** *(Support of SaS Density) The support of the symmetric alpha-stable density (23) is  $\mathbb{R}$ .*

**Proof** See (Nolan, 2020, Lemma 1.1). ■

Additionally, we recall a partial sum expansion, as described in Bergström (1952), wherein the remainder term possesses a smaller order of magnitude (for large  $|x|$ ) than the final term in the series.

**Lemma 14** (*Finite Series Expansion of SaS Distribution*) *The symmetric alpha-stable density (23), with  $\alpha \in (1, 2]$  and  $\gamma = 1$ , admits the following finite series expansion:*

$$p_{SaS}(x; \alpha, 1, 0) = -\frac{1}{\pi} \sum_{k=1}^n (-1)^k \frac{\Gamma(\alpha k + 1)}{(x)^{\alpha k + 1}} \sin\left(\frac{k\alpha\pi}{2}\right) + O\left(x^{-\alpha(n+1)-1}\right), \quad (25)$$

as  $|x| \rightarrow \infty$ .

**Proof** Bergström (1952) offers an expanded form of (25) that is valid for the complete range  $\beta \in (-1, 1)$ . However, because we have restricted the parameter set to  $\beta = 0$ , we only require the form provided for our purposes. ■

We use the foregoing lemma to argue that the privacy loss remains bounded as the observation  $|x|$  tends to infinity. However, Eq. (25) is stated for  $\gamma = 1$ . The next lemma states that the asymptotic behavior of the privacy loss as  $|x| \rightarrow \infty$  is independent of  $\gamma$ .

**Lemma 15** (*No Scale Dependence in the Limit*) *Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be two neighboring datasets. Denote by  $\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x; \gamma)$  the privacy loss of observation  $x$  for a bounded query  $f$  perturbed by a SaS mechanism  $\mathcal{M}_f$  with scale parameter  $\gamma$ . In the limit as  $|x|$  tends to  $\infty$ , the behavior of the privacy loss is indistinguishably asymptotic for distinct choices of  $\gamma$ :*

$$\lim_{|x| \rightarrow \infty} \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x; \gamma_1) = \lim_{|x| \rightarrow \infty} \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x; \gamma_2), \quad (26)$$

for  $\gamma_1 \neq \gamma_2$ .

**Proof**

$$p(x; \alpha, \gamma, \mu) = \frac{1}{2\gamma\pi} \int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i(\hat{x} - \mu_i)\hat{t}} d\hat{t} = p(\hat{x}; \alpha, 1, \mu). \quad (27)$$

Substituting (27) into the privacy loss function (9) gives

$$\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x; \gamma) = \ln \frac{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i(\hat{x} - f(\mathcal{D}_1))\hat{t}} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i(\hat{x} - f(\mathcal{D}_2))\hat{t}} d\hat{t}} = \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(\hat{x}; 1). \quad (28)$$

Observing that  $|\hat{x}|$  tends to  $\infty$  as  $|x|$  is driven to  $\infty$ , we have

$$\lim_{|x| \rightarrow \infty} \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x; \gamma) = \lim_{|\hat{x}| \rightarrow \infty} \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(\hat{x}; 1), \quad \forall \gamma. \quad (29)$$

In the limit, as  $|x|$  tends to infinity, the shift and scale of  $\hat{x}_1$  and  $\hat{x}_2$  are irrelevant. ■

With the results above, we are now in a position to state and prove our main contribution, namely, that for  $\alpha \in [1, 2)$ , the privacy loss of the SaS mechanism is bounded, i.e. the SaS mechanism is  $\varepsilon$ -differentially private.

**Theorem 16** *(The SaS mechanism is  $\epsilon$ -DP) Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be two neighboring datasets and let  $f$  be a bounded query that operates on them. Consider the SaS mechanism, which we denote by  $\mathcal{M}_f$ , with stability parameter  $\alpha$  in the reduced range  $\alpha \in [1, 2)$ . Then, the mechanism  $\mathcal{M}_f$  satisfies (5), pure Differential Privacy.*

**Proof** Each element of the mechanism's output is the perturbation of the query's response by an independent sample from the uni-variate density in (23). Thus, the joint density is equal to the product of the individual densities. As a result, we can express the privacy loss for a given observation vector  $\mathbf{x}$  as

$$\mathcal{L}_{\mathcal{D}_1 \parallel \mathcal{D}_2}^{SaS}(\mathbf{x}) = \ln \frac{\prod_{i=1}^m p_{SaS}(x_i; \alpha, \gamma, f(\mathcal{D}_1)_i)}{\prod_{i=1}^m p_{SaS}(x_i; \alpha, \gamma, f(\mathcal{D}_2)_i)}. \quad (30)$$

This can be written as the sum of the log-ratios of the individual elements:

$$\mathcal{L}_{\mathcal{D}_1, \mathcal{D}_2}^{SaS}(\mathbf{x}) = \sum_{i=1}^m \ln \frac{p_{SaS}(x_i; \alpha, \gamma, f(\mathcal{D}_1)_i)}{p_{SaS}(x_i; \alpha, \gamma, f(\mathcal{D}_2)_i)}. \quad (31)$$

Without loss of generality, let this sum be written in decreasing order of magnitudes of the terms, i.e. the first term,  $i = 1$ , has the largest magnitude. We now have the following bound:

$$|\mathcal{L}_{\mathcal{D}_1 \parallel \mathcal{D}_2}^{SaS}(\mathbf{x})| \leq m \left| \ln \frac{p_{SaS}(x_1; \alpha, \gamma, f(\mathcal{D}_1)_1)}{p_{SaS}(x_1; \alpha, \gamma, f(\mathcal{D}_2)_1)} \right|. \quad (32)$$

Our objective is to prove that the right side of (32) is bounded as function of  $x_1$ , which will imply, by Theorem 6, that the mechanism is  $\varepsilon$ -differentially private. We do so by first proving that the privacy loss is bounded on any compact set. Note that this is not immediate, since we are dealing with the log of a ratio and have no assurance that the numerator or denominator ever vanishes. Then, we show that in the limit as  $|x|$  tends to infinity, the privacy loss tends to 0, and thus does not diverge.

Initially, let  $x_1$  be an element in a compact set  $[a, b] \subset \mathbb{R}$ . The log-ratio of the densities could become unbounded within a finite interval in two ways: the argument vanishes or diverges. Consider first the case where one of the densities vanishes within the interval. By Lemma 13, an SaS density has support on the entire real line,  $\mathbb{R}$ . Therefore, the density is strictly positive over all compact sets  $[a, b] \subset \mathbb{R}$ .

Then, we consider if the numerator or denominator of (32) could be unbounded within the interval  $[a, b]$ . For simplicity, let  $\mu = 0$  and apply the substitution  $e^{-ix_1} = \cos(tx_1) - i \sin(tx_1)$  to the representation of the SaS density (23):

$$p_{SaS}(x_1; \alpha, \gamma, 0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha} (\cos(tx_1) - i \sin(tx_1)) dt. \quad (33)$$

Splitting the integral we have

$$\begin{aligned} p_{SaS}(x_1; \alpha, \gamma, 0) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha} \cos(tx_1) dt \\ &\quad - i \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha} \sin(tx_1) dt. \end{aligned} \quad (34)$$

Since  $\sin(tx_1)$  is an odd function the second integral vanishes:

$$p_{SaS}(x_1; \alpha, \gamma, 0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha} \cos(tx_1) dt. \quad (35)$$

As  $\cos(tx_1)$  is bounded above by 1, the density is bounded above:

$$p_{SaS}(x_1; \alpha, \gamma, 0) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha} dt. \quad (36)$$

Observe that the integrand in (36) is symmetric about  $t = 0$ , so we can remove the absolute value by adjusting the limits of integration:

$$p_{SaS}(x_1; \alpha, \gamma, 0) \leq \frac{1}{\pi} \int_0^{\infty} e^{-(\gamma t)^\alpha} dt. \quad (37)$$

Letting  $\hat{t} = (\gamma t)^\alpha$ , substitute  $\hat{t}$  into the inequality (37):

$$\begin{aligned} p_{SaS}(x_1; \alpha, \gamma, 0) &\leq \frac{1}{\alpha\gamma\pi} \int_0^{\infty} \hat{t}^{\frac{1}{\alpha}-1} e^{-\hat{t}} d\hat{t} \\ &= \frac{\Gamma(\frac{1}{\alpha})}{\alpha\gamma\pi}, \end{aligned} \quad (38)$$

where  $\Gamma$  is the standard Gamma function. Note that the Gamma function is finite on the interval  $1/\alpha \in (1/2, 1)$ ; see OEIS Foundation Inc. (2023). Equation (38) states that the density  $p_{SaS}$  is bounded over the real line. It is therefore bounded on the compact subset  $[a, b]$ . We proceed to prove that the privacy loss remains bounded in the limit as  $|x_1|$  tends to infinity.

Recall the series expansion presented in Lemma 14, for scale  $\gamma = 1$ . Truncate the series to a single term by taking  $n = 1$  and consider the privacy loss after substitution in (32):

$$|\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(\mathbf{x})| \leq m \left| \ln \frac{(x_1 - f(\mathcal{D}_1))^{-\alpha-1} + O(x_1^{-2\alpha-1})}{(x_1 - f(\mathcal{D}_2))^{-\alpha-1} + O(x_1^{-2\alpha-1})} \right|. \quad (39)$$

In the limit, as  $|x_1|$  tends infinity, the error terms in the numerator and denominator are dominated by the first terms:

$$\begin{aligned} \lim_{\|\mathbf{x}\| \rightarrow \infty} |\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(\mathbf{x})| &\leq \\ \lim_{|x_1| \rightarrow \infty} m \left| \ln \frac{(x_1 - f(\mathcal{D}_1))^{-\alpha-1} + O(x_1^{-2\alpha-1})}{(x_1 - f(\mathcal{D}_2))^{-\alpha-1} + O(x_1^{-2\alpha-1})} \right| &= \\ \lim_{|x_1| \rightarrow \infty} m \left| \ln \frac{(x_1 - f(\mathcal{D}_1))^{-\alpha-1}}{(x_1 - f(\mathcal{D}_2))^{-\alpha-1}} \right| &= 0, \end{aligned} \quad (40)$$

and the privacy loss converges to 0. By Lemma 15, the choice of  $\gamma$  does not impact the asymptotic behavior. Since this result holds for any value of  $\mathbf{x} \in \mathcal{R}(\mathcal{M}_f)$ , by Theorem 6, we have proved that the SaS mechanism is  $\varepsilon$ -differentially private.  $\blacksquare$

Although Theorem 16 establishes that the SaS mechanism is purely-Differentially Private, it does not offer a connection between the sensitivity of the query  $\Delta f$ , the scale of the noise distribution  $\gamma$ , and the achieved level of privacy  $\varepsilon$ . This limitation stems from the absence of a known closed-form expression for the density  $p_{SaS}$ . Before pursuing further details on these relationships, we revisit the Differential Privacy characteristics of two widely used privacy mechanisms to facilitate the subsequent comparison.

## 5 Privacy Scaling with Noise

In this section, we recall the characteristics of two common privacy mechanisms put forth in Dwork (2006); Dwork and Roth (2014): the Laplace mechanism and the Gaussian mechanism. After discussing these mechanisms, we proceed to study the relation between the privacy budget  $\varepsilon$  and the scale  $\gamma$  of the SaS mechanism and to provide related numerical results. We proceed to argue that the privacy budget of the SaS mechanism scales with the same order as the Laplace and Gaussian mechanisms, i.e., we wish to show that

$$\varepsilon_{SaS} \stackrel{?}{\propto} \frac{\Delta_1 f}{\gamma} \tag{41}$$

which is similar to

$$\varepsilon_{Lap} = \frac{\Delta_1 f}{b} \quad \text{and} \quad \varepsilon_{Gau} \propto \frac{\Delta_2 f}{\sigma}. \tag{42}$$

### 5.1 Level of privacy afforded by the SaS mechanism

For a given problem, there are three factors to consider when setting the parameters of a mechanism: the sensitivity of the query  $\Delta f$ , the scale of the noise  $\gamma$ , and the privacy budget  $\varepsilon$ . In this section, we study the relationship between these three quantities for the SaS mechanism.

Theorem 16 bounds the privacy loss by considering the largest component of the  $m$ -dimensional response of a query (see Eq. (32)). This motivates us to focus on the case  $m = 1$ , i.e., we now restrict to real-valued queries,  $f(\mathcal{D}) \in [a, b] \subset \mathbb{R}$ . Furthermore, in this section, when referring to the sensitivity of query  $f$  we exclusively use the  $\ell_1$ -sensitivity and denote it by  $\Delta_1$ .

We begin by establishing the linear relation between sensitivity and scale. To do so, we first prove that the extremes of the privacy loss, for a given privacy budget  $\varepsilon$ , occur when the query over datasets  $\mathcal{D}_1 \simeq \mathcal{D}_2$  returns values in the boundary of the range,  $\mathcal{R}(f)$ . For instance, when  $f(\mathcal{D}_1) = b$  and  $f(\mathcal{D}_2) = a$  (or vice versa,  $f(\mathcal{D}_1) = a$  and  $f(\mathcal{D}_2) = b$ , by symmetry of the absolute value of the query).

In order to prove that the privacy loss is maximized at the boundary of the query's range, we first establish that the density is monotonic on each semi-infinite interval to the left and right of the location parameter  $\mu$ . We give a proof for the generic symmetric stable

density using the fact that the density is *bell-shaped*, the definition of which is recalled next from Kwaśnicki (2020).

**Definition 17** (*Bell-Shaped Function*) *A continuous real-valued function is said to be bell-shaped if the  $n^{\text{th}}$  derivative,  $f^{(n)}$  for each  $n \in \mathbb{N}_0$ , changes sign exactly  $n$  times over its domain.* ◀

**Lemma 18** (*Monotonic First Derivative*) *The symmetric alpha-stable density (23) with location parameter  $\mu$  is monotonically increasing from  $-\infty$  to  $\mu$  and monotonically decreasing from  $\mu$  to  $\infty$ .*

**Proof** See the proof of (Kwaśnicki, 2020, Cor. 1.3) which asserts that all stable distributions are bell-shaped densities. Taking  $n = 1$  in Def. 17 implies that the first derivative of the density,  $f'$ , changes sign exactly once. Because the density is symmetric, the change in sign must occur at the axis of symmetry and the density must then decrease monotonically to 0 in the limit as  $|x| \rightarrow \infty$ . ■

We now utilize Lemma 18 to prove that, out of all neighboring datasets  $\mathcal{D}_1 \simeq \mathcal{D}_2$ , the maximum of the privacy loss occurs at the boundary of the query's range  $[a, b]$ . Recall that the SaS mechanism involves injecting noise with a location parameter of 0. Thus, the location parameter is the result of the query,  $\mu_i = f(\mathcal{D}_i)$ , and is itself bounded by the range of the query. By Theorem 16, the privacy loss of the SaS mechanism is bounded. As a result, we denote by  $x^*(\mu_1, \mu_2)$  the point at which the maximum privacy loss occurs as a function of the location parameters  $\mu_1$  and  $\mu_2$  generated by datasets  $\mathcal{D}_1$  and  $\mathcal{D}_2$  respectively.

**Theorem 19** (*Privacy Loss Maximization Occurs at Boundary*) *Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be neighboring datasets and denote by  $f$  a bounded query that operates on them and returns values in the compact set  $[a, b] \subset \mathbb{R}$ . Denote the SaS mechanism's privacy loss for an observation  $x$  by  $\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(x)$ . Let the location parameters of the two densities be  $\mu_1 = f(\mathcal{D}_1)$  and  $\mu_2 = f(\mathcal{D}_2)$ , with  $\mu_1 \neq \mu_2$ . Then*

$$\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(x^*(\mu_1, \mu_2)) \leq \mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(x^*(b, a)). \quad (43)$$

**Proof** Without loss of generality, take  $\mu_1 > \mu_2$ . Recall that the privacy loss, (9), is given by the log-ratio of two densities. Consider Figure 5 and let  $p(x; \mu_1)$ , in blue, and  $p(x; \mu_2)$ , in orange, represent the numerator and denominator of the privacy loss respectively. Let  $\epsilon$  be a value in  $[0, b - \mu_1]$ . First, we show that if the privacy loss achieves a maximum  $x^*(\mu_1, \mu_2)$ , then  $\mu_1 \leq x^*(\mu_1, \mu_2)$ . Observe that, by construction,  $p(x = \mu_1; \mu_1) \geq p(x = \mu_1; \mu_2)$ . Consider a point to the left of  $\mu_1$ . By the symmetry of SaS densities,  $p(\mu_1 - \epsilon; \mu_1) = p(\mu_1 + \epsilon; \mu_1)$  and because the first derivative is negative, Lemma 18,  $p(\mu_1 - \epsilon; \mu_2) \geq p(\mu_1 + \epsilon; \mu_2)$ . Thus,

$$\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(\mu_1 - \epsilon) \leq \mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(\mu_1 + \epsilon). \quad (44)$$

Next, let  $\mu_1 < b$ . Then, observe that  $p(x^*(\mu_1, \mu_2); \mu_1) = p(x^*(\mu_1, \mu_2) + \epsilon; \mu_1 + \epsilon)$ , illustrated by the upper two marked points in Figure 5. Similarly, by Lemma 18,  $p(x^*(\mu_1, \mu_2) + \epsilon; \mu_2) \leq p(x^*(\mu_1, \mu_2); \mu_2)$ , marked by the two lower points. Thus,  $\mathcal{L}$  can only be made larger

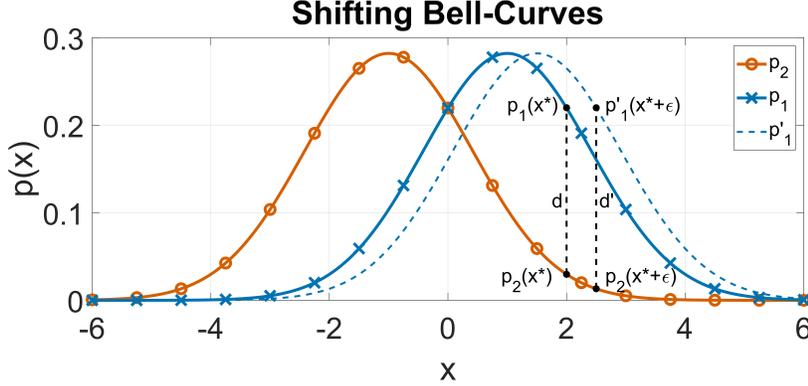


Figure 5: Consider two bell curves, shown here as  $p_1$  in blue and  $p_2$  in orange, with location parameters  $\mu_1 > \mu_2$  respectively. Given a point  $x^* > \mu_1$ , denote by  $d$  the distance between the curves at  $x^*$ :  $d := p_1(x^*) - p_2(x^*)$ . Shifting the distribution  $p_1$  to the right by some positive value  $\epsilon$ , gives the curve  $p'_1$  shown as a dotted line. By Lemma 18, the distance  $d' := p'_1(x^* + \epsilon) - p_2(x^* + \epsilon)$  is necessarily larger than  $d$ .

by increasing  $\mu_1$  in the direction of the bound  $b$ . Likewise, if  $\mu_1 = b$ , then shifting the distribution to the left can only decrease the maximum. A similar argument shows that the log-ratio cannot be decreased by shifting  $p(x; \mu_2)$  towards  $p(x; a)$ , which completes the proof.  $\blacksquare$

Because the maximum of the privacy loss is invariant to translation, Theorem 19 additionally implies the following corollary.

**Corollary 20** (*Relative Location Parameter*) *Let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be any neighboring datasets. Consider the privacy loss of the SaS mechanism, with  $\alpha \in (1, 2)$ , for a one-dimensional query  $f$ , with bounded range  $\mathcal{R}(f) = [a, b]$ . Denote by  $\Delta_1$  the  $\ell_1$ -sensitivity of  $f$ . Then, for a given  $\alpha \in (1, 2)$  and scale  $\gamma$ ,*

$$\max_{\mathcal{D}_1 \simeq \mathcal{D}_2} \max_{x \in \mathbb{R}} \mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{\text{SaS}}(x) = \max_{x \in \mathbb{R}} \ln \frac{p_{\text{SaS}}(x; \alpha, \gamma, \Delta_1)}{p_{\text{SaS}}(x; \alpha, \gamma, 0)}. \quad (45)$$

**Proof** The result follows directly from Theorem 19 and observing that the maximum of the privacy loss is invariant under translation.  $\blacksquare$

We can now assert that there is a linear relation between the sensitivity of the query  $\Delta_1$  and the scale of the density  $\gamma$ .

**Theorem 21** (*Linearity of Scale and Query's Sensitivity*) *Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be neighboring datasets and  $f$  be a one-dimensional query with bounded range  $\mathcal{R}(f) = [a, b]$ . Denote by  $\Delta_1$  the  $\ell_1$ -sensitivity of  $f$ . Let  $p_{\text{SaS}}$  be the SaS density as described in equation (23). Then, the level of privacy  $\epsilon$  remains the same if the sensitivity  $\Delta_1$  and the scale  $\gamma$  are both scaled by the same constant  $c > 0$ :*

$$\max_{x' \in \mathbb{R}} \ln \frac{p_{\text{SaS}}(x'; \alpha, c\gamma, c\Delta_1)}{p_{\text{SaS}}(x'; \alpha, c\gamma, 0)} = \max_{x \in \mathbb{R}} \ln \frac{p_{\text{SaS}}(x; \alpha, \gamma, \Delta_1)}{p_{\text{SaS}}(x; \alpha, \gamma, 0)}. \quad (46)$$

**Proof** We proceed by contradiction. Denote by  $x^*$  optimal argument on the right side of (46). Consider the left side of (46) in terms of the expression (23):

$$\max_{x' \in \mathbb{R}} \ln \frac{\int_{-\infty}^{\infty} e^{-|c\gamma t|^\alpha - it(x' - c\Delta_1)} dt}{\int_{-\infty}^{\infty} e^{-|c\gamma t|^\alpha - itx'} dt}. \quad (47)$$

The change of variables  $\hat{t} = ct$  results in the equivalent expression

$$\max_{x' \in \mathbb{R}} \ln \frac{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}(\frac{x'}{c} - \Delta_1)} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}\frac{x'}{c}} d\hat{t}}. \quad (48)$$

Denote by  $x'^*$  the location of the maximum in (48) and assume that it is not equal to  $cx^*$ . This leads to the following contradiction

$$\begin{aligned} \max_{cx' \in \mathbb{R}} \ln \frac{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}(x' - \Delta_1)} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}x'} d\hat{t}} &\neq \\ \max_{x \in \mathbb{R}} \ln \frac{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}(x - \Delta_1)} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\gamma \hat{t}|^\alpha - i\hat{t}x} d\hat{t}} & \end{aligned} \quad (49)$$

which is equivalent to

$$\max_{cx' \in \mathbb{R}} \ln \frac{p_{SaS}(x'; \alpha, \gamma, \Delta_1)}{p_{SaS}(x'; \alpha, \gamma, 0)} \neq \max_{x \in \mathbb{R}} \ln \frac{p_{SaS}(x; \alpha, \gamma, \Delta_1)}{p_{SaS}(x; \alpha, \gamma, 0)}. \quad (50)$$

■

**Remark 22** (Normalized Form) Because the scale and sensitivity are related linearly, we can combine  $\gamma$  and  $\Delta_1$  into a single parameter  $\hat{\gamma} = \gamma/\Delta_1$  by taking  $c = 1/\Delta_1$ :

$$\max_{x \in \mathbb{R}} \ln \frac{p_{SaS}(x; \alpha, \gamma, \Delta_1)}{p_{SaS}(x; \alpha, \gamma, 0)} = \max_{x' \in \mathbb{R}} \ln \frac{p_{SaS}(x'; \alpha, \hat{\gamma}, 1)}{p_{SaS}(x'; \alpha, \hat{\gamma}, 0)}. \quad (51)$$

We use the normalized form on the right side of Eq. (51) to gain an intuitive understanding of how the maximum of the privacy loss behaves as  $\alpha$  and  $\gamma$  are allowed to vary. Figure 6 fixes  $\gamma = \Delta_1 = 1$  and illustrates how the privacy loss approaches a straight line as  $\alpha$  tends to 2. Note that when  $\alpha = 2$ , corresponding to the privacy loss of the Gaussian mechanism, the loss is unbounded, illustrating that the Gaussian mechanism is not purely

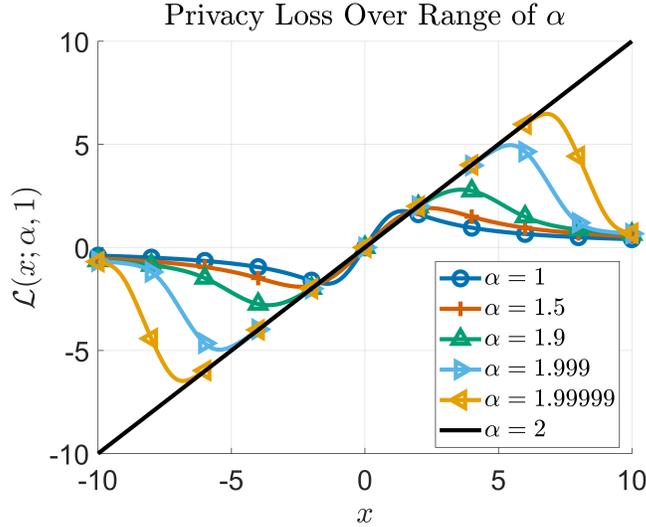


Figure 6: Denote by  $\mathcal{L}(x; \alpha, \gamma = 1)$  the privacy loss of the SaS mechanism with unit scale over observations  $x$ . Without loss of generality, let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be neighboring datasets such that the privacy loss of the Gaussian mechanism is linear:  $\mathcal{L}(x; 2, 1) = x$  shown in **black**. As the stability parameter  $\alpha$  is reduced, we observe that the privacy loss becomes bounded, reaching a peak before converging to the  $x$ -axis.

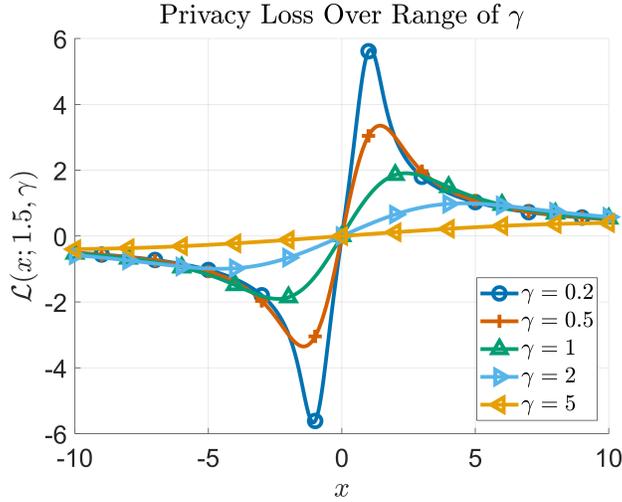


Figure 7: Denote by  $\mathcal{L}(x; \alpha = 1.5, \gamma)$  the privacy loss of the SaS mechanism with stability parameter  $\alpha = 1.5$  over observations  $x$ . Without loss of generality, let  $\mathcal{D}_1$  and  $\mathcal{D}_2$  be neighboring datasets such that the privacy loss is symmetric about the origin. As the scale  $\gamma$  of the density is increased, we observe that the increase in noise decreases the maximum possible privacy loss, compressing the curve toward the  $x$ -axis.

differentially private. In Figure 7, with  $\alpha$  fixed at  $3/2$ , we see that as the scale of the density,  $\gamma$ , increases, the level of privacy also increases (seen in the decreasing maximum  $\epsilon$  value; recall from 11 that  $\epsilon = \max_x \mathcal{L}(x)$ ).

Next, to derive the behavior of the privacy loss at observation  $x$  in terms of the scale, we use a special case of the second partial sum expansion discussed by in Bergström (1952).

**Lemma 23** (*A Second Finite Series Expansion*) *The symmetric alpha-stable density (23), with  $\alpha \in (1, 2)$  and  $\gamma = 1$ , has the following finite series expansion:*

$$p_{SaS}(x; \alpha, 1, 0) = \frac{1}{\pi} \sum_{k=0}^n (-1)^k \frac{\Gamma(\frac{k+1}{\alpha})}{k! \alpha} (x)^k \cos\left(\frac{k\pi}{2}\right) + O(|x|^{n+1}), \quad (52)$$

as  $|x| \rightarrow 0$ .

**Proof** The full form provided in Bergström (1952) states the result for the full range  $\beta \in (-1, 1)$ . In our work, we only require (52), so for brevity, we leave out the full form of the series.  $\blacksquare$

Below we make use of the following two elementary Taylor series expansions. For any  $c \neq 0$ :

$$\frac{1}{c+x} = \frac{1}{c} - \frac{x}{c^2} + \frac{x^2}{c^3} - \frac{x^3}{c^4} + O(x^4), \quad (53)$$

and

$$\ln \frac{c+x}{c} = \frac{x}{c} - \frac{x^2}{2c^2} + \frac{x^3}{3c^3} - \frac{x^4}{4c^4} + O(x^5). \quad (54)$$

Using Lemma 23, we now assert a relationship between the privacy loss for a given observation  $x$  and the scale of the SaS mechanism  $\gamma$ .

**Theorem 24** *Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be neighboring datasets and  $f$  a bounded query that operates on them. Denote by  $\Delta_1$  the  $\ell_1$ -sensitivity of the query  $f$ . Let  $\mathcal{M}_f$  be a SaS mechanism with  $\alpha \in (1, 2)$ . Let the observation  $x$  be fixed and take  $\gamma$  to be the independent variable. Then*

$$[\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{SaS}(x)](\gamma) = O\left(\frac{\Delta_1}{\gamma}\right) \text{ as } \gamma \rightarrow \infty. \quad (55)$$

( $\Delta_1$  is included in (55) in order to highlight the analogy with (42).)

**Proof** Fix the observation  $x$ , then, by Lemma 19, the maximum privacy loss for  $x$  over the datasets  $\mathcal{D}_1$  and  $\mathcal{D}_2$  is

$$[\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{\int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha - it(x - \Delta_1)} dt}{\int_{-\infty}^{\infty} e^{-|\gamma t|^\alpha - itx} dt}. \quad (56)$$

Let  $\hat{t} = \gamma t$ ,  $\hat{x} = x\Delta_1$ , and  $\hat{\gamma} = \gamma/\Delta_1$  and denote  $(\hat{x} - 1)/\hat{\gamma}$  and  $\hat{x}/\hat{\gamma}$  by  $x_1$  and  $x_2$ . The Eq. (56) becomes

$$[\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i\hat{t}x_1} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i\hat{t}x_2} d\hat{t}}. \quad (57)$$

We consider the numerator first, followed by the denominator. Expand the numerator in (57) using the partial sum expansion given in Lemma 23 with  $n = 0$ :

$$p_{SaS}(x_1; \alpha, 1, 0) = \frac{\Gamma(\frac{1}{\alpha})}{\alpha} + O(|x_1|), \quad |x_1| \rightarrow 0. \quad (58)$$

For simplicity we denote

$$a(\alpha) = \frac{\Gamma(\frac{1}{\alpha})}{\alpha}, \quad (59)$$

which gives

$$p_{SaS}(x_1; \alpha, 1, 0) = a(\alpha) + O(|x_1|), \quad |x_1| \rightarrow 0. \quad (60)$$

Thus, there exist positive constants  $C$  and  $x_0$  such that

$$|p_{SaS}(x_1; \alpha, 1, 0)| \leq a + C|x_1|, \quad \forall |x_1| \leq x_0. \quad (61)$$

Replace  $x_1$  by its definition in (61), first noting that the translations in  $x$  are described by the last parameter in the notation for the SaS density:

$$p_{SaS}\left(\frac{\hat{x} - 1}{\hat{\gamma}}; \alpha, 1, 0\right) = p_{SaS}\left(\frac{\hat{x}}{\hat{\gamma}}; \alpha, 1, \frac{1}{\hat{\gamma}}\right). \quad (62)$$

Then

$$\left|p_{SaS}\left(\frac{\hat{x}}{\hat{\gamma}}; \alpha, 1, \frac{1}{\hat{\gamma}}\right)\right| \leq a + C\left|\frac{\hat{x} - 1}{\hat{\gamma}}\right|, \quad \forall \left|\frac{\hat{x} - 1}{\hat{\gamma}}\right| \leq x_0. \quad (63)$$

Note that the range restriction in (63) is equivalent to  $\hat{\gamma} \geq |\hat{x} - 1|/x_0$ . Thus, denote  $|\hat{x} - 1|/x_0$  and  $C|\hat{x} - 1|$  by  $\gamma_0$  and  $\hat{C}$  respectively. Then

$$\left|p_{SaS}\left(\frac{\hat{x}}{\hat{\gamma}}; \alpha, 1, \frac{1}{\hat{\gamma}}\right)\right| \leq a + \hat{C} \cdot \frac{1}{\hat{\gamma}}, \quad \forall \hat{\gamma} \geq \gamma_0, \quad (64)$$

which can be represented in big O notation as

$$p_{SaS}\left(\frac{\hat{x}}{\hat{\gamma}}; \alpha, 1, \frac{1}{\hat{\gamma}}\right) = a + O\left(\frac{1}{\hat{\gamma}}\right), \quad \hat{\gamma} \rightarrow \infty. \quad (65)$$

Using the same logic, the denominator in (56) can be represented as

$$p_{SaS}\left(\frac{\hat{x}}{\hat{\gamma}}; \alpha, 1, 0\right) = a + O\left(\frac{1}{\hat{\gamma}}\right), \quad \hat{\gamma} \rightarrow \infty. \quad (66)$$

Combining (65) and (66), (57) can now be expressed for large  $\hat{\gamma}$  in the form

$$[\mathcal{L}_{\mathcal{D}_1 || \mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{a + O(\frac{1}{\hat{\gamma}})}{a + O(\frac{1}{\hat{\gamma}})}, \quad \hat{\gamma} \rightarrow \infty. \quad (67)$$

Using the elementary Taylor series (53), we rewrite the denominator as

$$\frac{1}{a + O(\frac{1}{\hat{\gamma}})} = \frac{1}{a} + O\left(\frac{1}{\hat{\gamma}}\right) = \frac{1 + O(\frac{1}{\hat{\gamma}})}{a}, \quad (68)$$

as  $\hat{\gamma} \rightarrow \infty$ . Substituting (68) into (67) gives

$$\begin{aligned} [\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) &= \ln \frac{a + O(\frac{1}{\hat{\gamma}})}{1} \cdot \frac{1 + O(\frac{1}{\hat{\gamma}})}{a} \\ &= \ln \frac{a + O(\frac{1}{\hat{\gamma}}) + O(\frac{1}{\hat{\gamma}^2})}{a}, \end{aligned} \quad (69)$$

as  $\hat{\gamma} \rightarrow \infty$ . The squared term is dominated in the limit and leaves

$$[\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{a + O(\frac{1}{\hat{\gamma}})}{a}, \quad \hat{\gamma} \rightarrow \infty. \quad (70)$$

Next, we use the elementary Taylor series (54) and expand to

$$[\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) = \frac{O(\frac{1}{\hat{\gamma}})}{a} + O\left(\frac{1}{\hat{\gamma}^2}\right), \quad \hat{\gamma} \rightarrow \infty. \quad (71)$$

Recalling that  $\hat{\gamma} = \gamma/\Delta_1$ , we complete the proof:

$$[\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) = O\left(\frac{\Delta_1}{\gamma}\right), \quad \gamma \rightarrow \infty. \quad (72)$$

■

Theorem 24 only guarantees that the privacy of a specific observation  $x$  scales as  $O(\Delta_1/\gamma)$  for large  $\gamma$ . Without additional information about the location of the maximum, which is difficult to attain due to the lack of a known closed form for the general SaS density, Theorem 24 does not allow us to conclude that the maximum over all observations scales in the same manor. Because of this, in Figure 8 we provide numerical results graphing the max privacy loss  $\varepsilon$  over a range of scale values  $\gamma$  (with  $\Delta_1 = 1$ ) for a selection of  $\alpha$  values. Note that because this is a log-log plot, a vertical shift, as seen with  $\alpha = 1.9$ , corresponds to a multiplicative scalar in the limiting behavior. We observe that as  $\gamma$  increases, the maximum privacy loss falls off at a rate similar to that for  $\alpha = 1$  (at least for practically useful values of  $\varepsilon$  and  $\gamma$ ). To this end, we take advantage of the closed form of the density when  $\alpha = 1$  to provide more concrete results for that case.

**Theorem 25** *Let  $\mathcal{D}_1 \simeq \mathcal{D}_2$  be neighboring datasets and  $f$  a bounded query with  $\ell_1$ -sensitivity  $\Delta_1$  that operates on them. Take the stability parameter  $\alpha = 1$  for the SaS mechanism. Then, the privacy budget  $\varepsilon$  as a function of the scale  $\gamma$  is given by*

$$\varepsilon(\gamma) = \ln \frac{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} + 1}{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} - 1}. \quad (73)$$

**Proof** Consider the privacy budget  $\varepsilon$  for the SaS mechanism when  $\alpha = 1$ :

$$\varepsilon := \max_{x \in \mathbb{R}} \mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x). \quad (74)$$

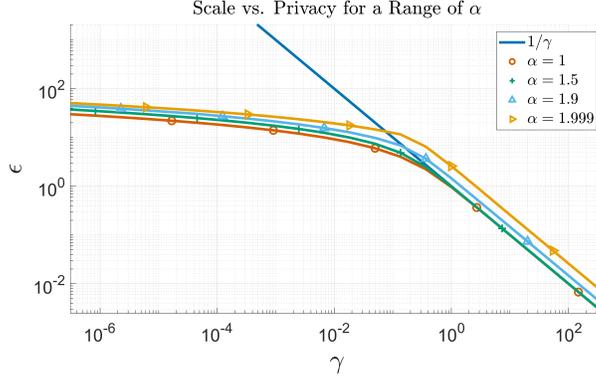


Figure 8: The maximum privacy loss of the Laplace mechanism is inversely related to the scale of the injected noise  $\gamma$ , show linearly on the log-log plot in **blue**. For large values of  $\gamma$ , the privacy loss of the SaS mechanism falls off at the same rate, see Corollary 26. However, for small scales, as  $\gamma$  is decreased, the SaS mechanism increases at a rate of  $O(\log(1/\gamma))$  as shown in Corollary 27 (as opposed to  $O(1/\gamma)$  for the Laplace). The equation for the Cauchy's privacy loss, shown here in **orange**, is explicitly given in Equation (73) with  $\Delta_1 = 1$ .

As in the proof of the foregoing theorem, let  $\hat{t} = \gamma t$ ,  $\hat{x} = x\Delta_1$ , and  $\hat{\gamma} = \gamma/\Delta_1$  in (56):

$$[\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i\hat{t}\frac{\hat{x}-1}{\hat{\gamma}}} d\hat{t}}{\int_{-\infty}^{\infty} e^{-|\hat{t}|^\alpha - i\hat{t}\frac{\hat{x}}{\hat{\gamma}}} d\hat{t}}. \quad (75)$$

Note that the SaS density, when  $\alpha = 1$ , takes the closed form

$$p_{SaS}(x; 1, \gamma, \mu) = \frac{1}{\pi\gamma(1 + (\frac{x-\mu}{\gamma})^2)}. \quad (76)$$

Substituting (76) into (75) gives

$$[\mathcal{L}_{\mathcal{D}_1||\mathcal{D}_2}^{SaS}(x)](\gamma) = \ln \frac{1 + (\frac{\hat{x}}{\hat{\gamma}})^2}{1 + (\frac{\hat{x}-1}{\hat{\gamma}})^2}. \quad (77)$$

To find the maximum, we take the derivative of the right side with respect to  $\hat{x}$ ,

$$\frac{d}{d\hat{x}} \ln \frac{1 + (\frac{\hat{x}}{\hat{\gamma}})^2}{1 + (\frac{\hat{x}-1}{\hat{\gamma}})^2} = \frac{-2(\hat{x}^2 - \hat{x} - \hat{\gamma}^2)}{(\hat{\gamma}^2 + (\hat{x}-1)^2)(\hat{\gamma}^2 + \hat{x}^2)}. \quad (78)$$

This equates to 0 when

$$\hat{x}^2 - \hat{x} - \hat{\gamma}^2 = 0. \quad (79)$$

There are two solutions:

$$\hat{x}^* = \frac{1}{2}(1 \pm \sqrt{1 + 4\hat{\gamma}^2}). \quad (80)$$

Since the privacy loss is symmetric, we take the positive solution without loss of generality. Substituting the positive maximum location into (77) gives

$$\varepsilon(\hat{\gamma}) = \ln \frac{1 + \frac{1 + \sqrt{1 + 4\hat{\gamma}^2}}{4\hat{\gamma}^2}}{1 + \frac{1 + 4\hat{\gamma}^2}{4\hat{\gamma}^2}}. \quad (81)$$

Recalling that  $\hat{\gamma} = \gamma/\Delta_1$ , equation (81) is equivalent to the following expression after simplification,

$$\varepsilon|_{\alpha=1}(\gamma) = \ln \frac{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} + 1}{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} - 1}. \quad (82)$$

■

To study the limiting behavior of the privacy loss, we again invoke three elementary Taylor series:

$$\sqrt{1+x^2} \pm x = 1 \pm x + \frac{x^2}{2} - \frac{x^4}{8} + O(x^6), \quad (83)$$

$$\ln \frac{(1+x)}{(1-x)} = 1 + 2x + \frac{2x^3}{3} + \frac{2x^5}{5} + O(x^7), \quad (84)$$

and

$$\sqrt{4x^2 + 1} + c = (c+1) + 2x^2 - 2x^4 + O(x^5). \quad (85)$$

**Corollary 26** (*Large scale approximation*) *In the limit as  $\gamma$  grows without bound, when  $\alpha = 1$  the privacy budget  $\varepsilon$ , falls off at the following rate:*

$$\varepsilon(\gamma)|_{\alpha=1} \approx \frac{\Delta_1}{\gamma}, \quad \text{as } \gamma \rightarrow \infty. \quad (86)$$

**Proof** The change of variables  $x = \Delta_1/(2\gamma)$  applied to Eq. (82) gives

$$\varepsilon(x)|_{\alpha=1} = \ln \frac{\sqrt{\frac{1}{x^2} + 1} + 1}{\sqrt{\frac{1}{x^2} + 1} - 1}. \quad (87)$$

Because  $x$  only equates to 0 in the limit of  $\gamma \rightarrow \infty$ , and we seek the dynamics when  $\gamma$  is large but finite, we can safely multiply the argument of the logarithm in (87) by  $x/x$  giving

$$\varepsilon(x)|_{\alpha=1} = \ln \frac{\sqrt{1+x^2} + x}{\sqrt{1+x^2} - x}. \quad (88)$$

Expand the numerator and denominator of (88) using the elementary Taylor series (83), giving the following expression for small  $x$  after eliminating the higher order terms:

$$\varepsilon(x)|_{\alpha=1} \approx \ln \frac{1+x}{1-x}, \quad \text{as } x \rightarrow 0. \quad (89)$$

This can be further simplified by appealing to the Taylor series expansion (84) yielding a first order approximation

$$\varepsilon(x)|_{\alpha=1} \approx 2x, \quad \text{as } x \rightarrow 0. \quad (90)$$

Recalling that  $x = \Delta_1/(2\gamma)$  now gives

$$\varepsilon(\gamma)|_{\alpha=1} \approx \frac{\Delta_1}{\gamma}, \quad \text{as } \gamma \rightarrow \infty. \quad (91)$$

■

**Corollary 27** (*Small scale approximation*) *In the limit as  $\gamma$  becomes vanishing small, for  $\alpha = 1$  the privacy budget  $\varepsilon$ , increases at the following rate:*

$$\varepsilon(\gamma)|_{\alpha=1} \approx 2 \ln \frac{\sqrt{2}\Delta_1}{\gamma}, \text{ as } \gamma \rightarrow 0. \quad (92)$$

**Proof** Begin by expanding the argument of the logarithm in (82) using the elementary Taylor series (84),

$$\ln \frac{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} + 1}{\sqrt{4(\frac{\gamma}{\Delta_1})^2 + 1} - 1} = \ln \frac{2 + 2(\frac{\gamma}{\Delta_1})^2 + O(\gamma^3)}{2(\frac{\gamma}{\Delta_1})^2 + O(\gamma^3)}, \quad \gamma \rightarrow 0. \quad (93)$$

As  $\gamma$  tends to 0, the higher order behavior is dominated by  $\gamma^2$  and we have

$$\varepsilon(\gamma)|_{\alpha=1} \approx \ln \frac{2}{(\frac{\gamma}{\Delta_1})^2}, \quad \gamma \rightarrow 0. \quad (94)$$

Equivalently, the expression in (94) gives

$$\varepsilon(\gamma)|_{\alpha=1} \approx 2 \ln \frac{\sqrt{2}\Delta_1}{\gamma}, \text{ as } \gamma \rightarrow 0. \quad (95)$$

■

In Figure 9, we supplement the numerical results with graphs of the limiting behavior derived in Corollaries 26 and 27. The figure numerically confirms that, for small  $\gamma$ , the SaS

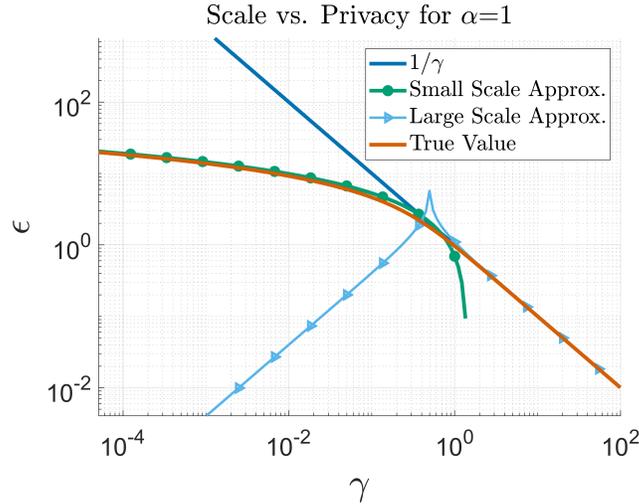


Figure 9: The privacy loss  $\varepsilon$  for the SaS mechanism, with  $\alpha = 1$  and  $\Delta_1 = 1$ , over a range of scale values  $\gamma$  described by (73) and shown here in **orange**. For small  $\gamma$ , the privacy loss is approximated by Eq. (95) in **green**, and, for large  $\gamma$ , the privacy loss is approximated by Eq. (91), shown in **light blue**. For comparison, the privacy loss of the Laplace mechanism is shown in **blue**.

mechanism, due the appearance of the logarithm in (95), scales better than the Laplace and Gaussian mechanisms as recalled in (42).

Now that we have shown that the SaS mechanism behaves in a similar manner to other common privacy mechanisms, we move on to describe the expected error that the SaS mechanisms introduces into the query's result by using any of these mechanisms.

## 6 Error Analysis

It is typical for methods to employ the  $\ell_2$ -norm when defining a measure of error. However, the moments of SaS densities are only defined up to  $\alpha$ , and since we consider  $\alpha < 2$ , the second moment lacks a clear definition Nolan (2020). In lieu of the  $\ell_2$ -norm, we opt for the mean absolute deviation (MAD), as used in Dwork and Roth (2014):

**Definition 28** (*Expected Privacy Distortion*) *Let  $\mathcal{D}$  be a dataset and denote by  $f(\mathcal{D})$  and  $\mathcal{M}_f(\mathcal{D})$  the response of a query and privacy mechanism respectively. Denote the density of the privacy mechanism by  $Y$ . The mean absolute deviation is*

$$E(f(\mathcal{D}), \mathcal{M}_f(\mathcal{D})) := \mathbb{E}|f(\mathcal{D}) - \mathcal{M}_f(\mathcal{D})|, \quad (96)$$

which is equivalent to the expectation of the absolute value of the injected noise  $Y$ :

$$E(f(\mathcal{D}), \mathcal{M}_f(\mathcal{D})) = \mathbb{E}|Y|. \quad (97)$$

◀

Before beginning the analysis of the error incurred by the SaS mechanism, we establish that the SaS mechanism adheres to strict stability.

**Lemma 29** (*SaS density is Strictly Stable*) *The SaS density (23) with location parameter  $\mu = 0$  is strictly stable.*

**Proof** Consider three independent and identically distributed SaS densities denoted by  $Y_1$ ,  $Y_2$ , and  $Y$  with  $\mu = 0$ . Let  $a$  and  $b$  represent two scalar values. Next, examine the density of the combined random variable  $aY_1 + bY_2$ . Since SaS densities are determined by their characteristic functions, we establish the following relation:

$$\varphi_{aY_1+bY_2}(t) = \varphi_{aY_1}(t)\varphi_{bY_2}(t). \quad (98)$$

Using the definition of a characteristic function, we bring the constants into the argument

$$\begin{aligned} \varphi_{aY_1}(t)\varphi_{bY_2}(t) &= \mathbb{E}[e^{itaY_1}]\mathbb{E}[e^{itbY_2}] \\ &= \varphi_{Y_1}(at)\varphi_{Y_2}(bt). \end{aligned} \quad (99)$$

Expand by substituting the expression for the characteristic function of a stable distribution with  $\mu = 0$  (20) into both functions on the right side,

$$\begin{aligned} \varphi_{Y_1}(at)\varphi_{Y_2}(bt) &= \exp(|\gamma at|^\alpha) \exp(|\gamma bt|^\alpha) \\ &= \exp|(a^\alpha + b^\alpha)^{1/\alpha} \gamma t|^\alpha. \end{aligned} \quad (100)$$

Setting  $c = (a^\alpha + b^\alpha)^{1/\alpha}$  gives  $aY_1 + bY_2 = cY$ . ■

We are now equipped to determine the expected error introduced into the query's response by the SaS mechanism.

**Theorem 30** (*Expected Distortion Due to SaS mechanism*) *Let  $f$  be a bounded query that operates on dataset  $\mathcal{D}$ . Denote by  $\mathcal{M}_f$  the SaS mechanism and take the stability parameter  $\alpha$  to be restricted to the range  $\alpha \in (1, 2)$ . Then, the mean absolute distortion is*

$$E(f(\mathcal{D}, \mathcal{M}_f(\mathcal{D}))) = \frac{2\gamma}{\pi} \Gamma\left(1 - \frac{1}{\alpha}\right). \quad (101)$$

**Proof** Note that, by Lemma 29, the noise injected by the SaS mechanism is strictly stable. In Nolan (2020), the proof of Corollary 3.5 includes a statement that if a density  $Y$  is strictly stable, then its mean absolute deviation is given by

$$\mathbb{E}[|Y|] = \frac{2\gamma}{\pi} \Gamma\left(1 - \frac{1}{\alpha}\right). \quad (102)$$

■

We now provide the expected distortions of the two most common privacy mechanisms from Dwork and Roth (2014); Dwork et al. (2006a): the Laplace and the Gaussian mechanisms, to show that each induces an error linear with the scale of the noise. The mean absolute deviation of the Laplace density is

$$\mathbb{E}[|Lap(0, b)|] = \mathbb{E}[Exp(b^{-1})] = b. \quad (103)$$

The mean absolute deviation Gaussian density is the expected value of the half-normal random variable

$$\mathbb{E}[|\mathcal{N}(0, \sigma^2)|] = \sqrt{\frac{2}{\pi}} \sigma. \quad (104)$$

Note that for each of the three densities, the error is related linearly to the density's respective scale. From (102) we recover the distortion of the Gaussian mechanism by taking  $\alpha = 2$  and  $\gamma = \sigma/\sqrt{2}$ . Next, we proceed to prove that the expected distortion is monotonic in  $\alpha$ , reaching a minimum when  $\alpha = 2$  and diverging as  $\alpha$  tends to 1.

**Remark 31** *We note that there is a linear relationship between the expected error of the SaS mechanism  $\mathbb{E}[|Y^{SaS}|]$  and the scale of the density  $\gamma$  in (102). Recall that the privacy budget  $\varepsilon$  is inversely proportional to large values of  $\gamma$  (Figure 8). This relationship is proven for  $\alpha = 1$  in Corollary 26. Therefore, small values of  $\varepsilon$ , while enhancing the client's privacy, necessarily increase the expected error induced by the mechanism. In other words, the level of privacy is inversely related to the accuracy of the query. We note that this relationship is shared by the other common mechanisms.*

**Corollary 32** (*Error is Monotonic in  $\alpha$* ) *The mean absolute distortion injected into a query by the SaS mechanism decreases monotonically as  $\alpha$  increases from 1 to 2.*

**Proof** Because the stability parameter  $\alpha$  is chosen from the bounded set  $(1, 2)$ , the argument of the Gamma function in (102) varies between  $(0, 1/2)$ . The Gamma function has an asymptote at  $x = 0$  and reaches a local minimum in the right plane at  $x \approx 1.462$ , see OEIS Foundation Inc. (2023). Thus, for a given  $\gamma$ , the distortion in Eq. (102) is minimized when  $\alpha$  tends to 2. ■

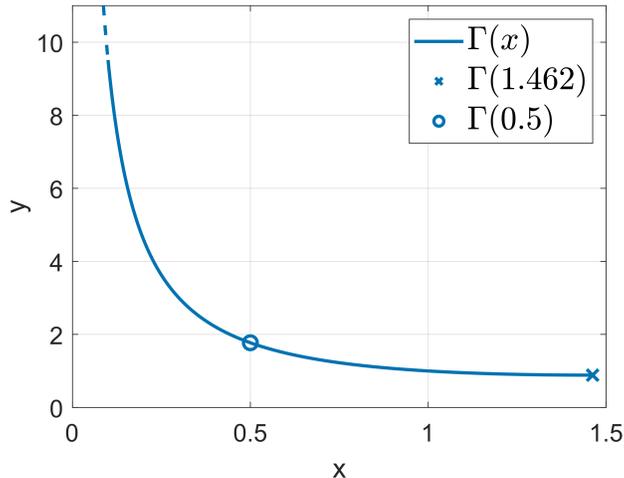


Figure 10: The Gamma function  $\Gamma(x)$  achieves a minimum value in the right hand plane at  $x \approx 1.462$ . When  $\alpha$  is bounded between  $[1, 2)$ , the Gamma component of the injected error takes input values in  $[0, 0.5)$ . The Gamma function is monotonically decreasing in this interval from  $\infty$  to 1.7725.

This minimum is proven in Corollary 32 and depicted in Figure 10. A naive first thought is thus that  $\alpha = 2$  is the *optimal* value for the parameter  $\alpha$  as the injected error achieves a minimum. However, this is not necessarily the case as choosing  $\alpha = 2$  increases the required scale  $\gamma$  necessary to achieve a given privacy budget  $\epsilon$ . Table 1 provides a list of expected distortions for a selection of stability values  $\alpha$ . By selecting  $\alpha$  close to 2 we can achieve an essentially equivalent expected distortion but provide better levels of privacy.

Table 1: The expected distortion for the Gaussian mechanism ( $\alpha = 2$ ) and a selection of SaS mechanisms ( $\alpha < 2$ ). The expected distortion is given as a multiple of the injected noise scale  $\gamma$ .

$\alpha$	Expected Distortion
2	$1.1284\gamma$
1.999	$1.1289\gamma$
1.99	$1.1340\gamma$
1.95	$1.1576\gamma$
1.9	$1.1903\gamma$
1.8	$1.2687\gamma$
1.0	$\infty$

In particular, note that the expected distortion between  $\alpha = 2$  and  $\alpha = 1.999$  differ only by only 0.044%. Thus, to achieve similar accuracy results to the Gaussian mechanism, we can choose to focus on  $\alpha$  close to 2, leaving a further exploration of an optimal choice of  $\alpha$  for future work. It should be noted that decreasing  $\alpha$  enhances the privacy experienced by the client, thus decreasing  $\epsilon$ . Therefore, for a fixed privacy budget  $\epsilon$ , one can compose an optimization that minimizes a weighted sum of privacy and injected noise. This optimization is problem specific since it deals with the particular sensitivity of the dataset in question, as

well as the desired weighting between error and privacy. We leave it for a more application focused examination.

## 7 Concluding Remarks

We have presented, the SaS mechanism represents and shown how it advances the field of Differential Privacy. This mechanism not only provides strong guarantees of privacy but also offers distinct advantages when compared to other common privacy mechanisms. We proved that the SaS mechanism achieves pure Differential Privacy, ensuring that individual data points remain protected even in the face of powerful adversaries. This is starkly contrasted with the Gaussian mechanism, which only achieves approximate Differential Privacy. Additionally, the SaS mechanism utilizes a stable density, allowing it to be used in local applications where the Laplace mechanisms is difficult to analyze.

We showed that the expected distortion introduced by the SaS mechanism into query results can be made essentially equivalent to that of the Gaussian mechanism. The expected distortion can additionally be formulated as a compromise between injected error and privacy guarantees. As a result, we conclude that there is little reason to use the Gaussian mechanism over the SaS mechanism.

## ACKNOWLEDGMENT

This work was supported in part by Grant ECCS-2024493 from the U.S. National Science Foundation.

## References

- M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- H. Asi, D. Liu, and K. Tian. Private stochastic convex optimization with heavy tails: Near-optimality from simple reductions. *arXiv preprint arXiv:2406.02789*, 2024.
- H. Bergström. On some expansions of stable distribution functions. *Arkiv för Matematik*, 2(4):375–378, 1952.
- D. Durfee. Unbounded differentially private quantile and maximum estimation. *Advances in Neural Information Processing Systems*, 36, 2024.
- C. Dwork. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer, 2006.
- C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic*

- Techniques, St. Petersburg, Russia, May 28-June 1, 2006. Proceedings 25*, pages 486–503. Springer, 2006a.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006b.
- EU. General data protection regulation, 2016.
- Q. Geng and P. Viswanath. Optimal noise adding mechanisms for approximate differential privacy. *IEEE Transactions on Information Theory*, 62(2):952–969, 2015.
- Id-Theft-Center. 2021 annual data breach report, Apr 2022. URL <https://www.idtheftcenter.org/publication/2021-annual-data-breach-report-2/>.
- K. Ito, Y. Kawano, and K. Kashima. Privacy protection with heavy-tailed noise for linear dynamical systems. *Automatica*, 131:109732, 2021.
- S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- J. Koetsier. Privacy checkup: Limit ad tracking up 216% on ios, but down 85% on android, Jul 2021. URL <https://www.singular.net/blog/limit-ad-tracking-privacy-checkup-in-2020/>.
- M. Kwaśnicki. A new class of bell-shaped functions. *Transactions of the American Mathematical Society*, 373(4):2255–2280, 2020.
- P. Lévy. *Calcul des probabilités*. Gauthier-Villars, 1925.
- W. Li, F. Milletari, D. Xu, N. Rieke, J. Hancox, W. Zhu, M. Baust, Y. Cheng, S. Ourselin, M. J. Cardoso, and A. Feng. Privacy-preserving federated brain tumour segmentation, 2019.
- H. B. McMahan, E. Moore, D. Ramage, and B. A. y Arcas. Federated learning of deep networks using model averaging. *CoRR*, abs/1602.05629, 2016. URL <http://arxiv.org/abs/1602.05629>.
- F. McSherry and K. Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103. IEEE, 2007.
- L. Minto and M. Haller. Using federated learning to improve brave’s on-device recommendations while protecting your privacy, Jun 2021. URL <https://brave.com/federated-learning/>.
- J. P. Nolan. *Univariate stable distributions*. Springer, 2020.
- OEIS Foundation Inc. The On-Line Encyclopedia of Integer Sequences, 2023. Published electronically at <http://oeis.org>.

- U. Şimşekli, M. Gürbüzbalaban, S. Yıldırım, and L. Zhu. Differential privacy of noisy (s)gd under heavy-tailed perturbations. *arXiv preprint arXiv:2403.02051*, 2024.
- USA. Internet Freedom 2011. *47 CFR § 8.1(a) 2011*, 2011.
- L. Wasserman and S. Zhou. A statistical framework for differential privacy, 2009.
- K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, S. Jin, T. Q. S. Quek, and H. V. Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.
- Christopher C. Zawacki and Eyad H. Abed. The symmetric alpha-stable privacy mechanism. In *2024 58th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6, 2024. doi: 10.1109/CISS59072.2024.10480198.