Schrödinger Mechanisms: Optimal Differential Privacy Mechanisms for Small Sensitivity

Wael Alghamdi*[1], Shahab Asoodeh[2], Flavio P. Calmon[1], Juan Felipe Gomez[1], Oliver Kosut[3] and Lalitha Sankar[3]
[1] School of Engineering and Applied Sciences, Harvard University
(emails: alghamdi@g.harvard.edu, flavio@seas.harvard.edu, juangomez@g.harvard.edu)
[2] Department of Computing and Software, McMaster University (email: asoodehs@mcmaster.ca)
[3] School of Electrical, Computer, and Energy Engineering, Arizona State University (emails: {okosut, lsankar}@asu.edu)

2023 IEEE International Symposium on Information Theory (ISIT) | 978-1-6654-7554-9/23/$31.00 ©2023 IEEE | DOI: 10.1109/ISIT54713.2023.10206616

***Abstract*—We consider the problem of designing optimal differential privacy mechanisms with a favorable privacy-utility tradeoff in the limit of a large number n of compositions (i.e., sequential queries). Here, utility is measured by the average distance between the mechanism's input and output, evaluated by a cost function c. We show that if n is sufficiently large and the sensitivities of all queries are small, then the optimal additive noise mechanism has probability density function fully characterized by the ground-state eigenfunction of the Schrödinger operator with potential c. This leads to a family of optimal mechanisms, dubbed the Schrödinger mechanisms, depending on the choice of the cost function. Instantiating this result, we demonstrate that for $c(x) = x^2$ the Gaussian mechanism is optimal, and for $c(x) = |x|$, the optimal mechanism is obtained by the Airy function, thereby leading to the Airy mechanism.**
The full proofs can be found in the extended version at [1].

I. Introduction

Differential privacy (DP) [2] provides provable privacy guarantees for queries computed over sensitive data. Currently, DP is the standard definition used in privacy-preserving machine learning (ML) deployed in practice by, for example, Google [3], Apple [4], and Facebook [5]. The parameters of these mechanisms are determined by the desired level of privacy and the query's sensitivity, denoted by s. When incorporating DP into ML algorithms, one fundamental challenge is to accurately characterize the privacy loss in iterative algorithms. To address this challenge, numerous composition results have been proved in the literature, e.g., [6]–[16].

In this paper, we view composition problems from a different angle: Instead of assuming access to constituent mechanisms, we seek to *construct* a DP mechanism whose n-fold composition has the *optimal* privacy guarantee among all possible mechanisms. We investigate this problem under two assumptions: (1) large number of compositions n, and (2) small values of the query sensitivity s. The first assumption is inspired by iterative training procedures for ML models such as stochastic gradient descent, where a dataset is queried many times (often in the thousands) in order to update model parameters (e.g., weights of a neural network). Thus, it is a natural assumption for the privacy analysis of private ML algorithms (see, e.g., [12], [16], [17]). The second assumption holds, for example, when we are interested in counting queries over large datasets, because in this case the sensitivity is inversely proportional to the size of the dataset.

Optimal DP mechanisms under the first assumption (i.e., in the large composition regime) have been recently characterized in [17]. The main technical result validating the approach of [17] is that the privacy guarantee of the n-fold composition of a mechanism $P_{Y|X}$ scales as [16]

$$n \cdot \sup_{|x-x'|\le s} D(P_{Y|X=x} \,\|\, P_{Y|X=x'}), \tag{1}$$

where $D(\cdot\,\|\,\cdot)$ denotes the KL-divergence. It follows from this observation that the optimal mechanism is the one that solves the following optimization problem:

$$\inf_{\mathbb{E}[c(Y-x)|X=x]\le C,\ \forall x\in\mathbb{R}} \ \sup_{|x-x'|\le s} D(P_{Y|X=x} \,\|\, P_{Y|X=x'}), \tag{2}$$

where the outer infimum is taken over all mechanisms $P_{Y|X}$ that satisfy the prescribed cost constraint dictated by a cost function c and a cost bound C (e.g., a bounded variance). Not surprisingly, the optimal mechanism is additive and continuous (see [17, Theorem 1]), thus (2) reduces to the following optimization over probability density functions (PDFs) p:

$$\inf_{p:\ \mathbb{E}_p[c]\le C} \ \sup_{|a|\le s} D(p \,\|\, T_a p), \tag{3}$$

where $(T_a p)(x) := p(x-a)$ denotes the shift operator. The so-called *cactus mechanisms*, were shown by [17] to achieve the optimal value in (3) to arbitrary accuracy and for *fixed* sensitivity $s > 0$. In this paper, we seek to solve (3) with vanishing sensitivity, i.e., $s \to 0^+$. We achieve this goal by a sequence of reductions: from KL-divergence to Fisher information and then to the Schrödinger equation. Thus, we name the ensuing family of optimal mechanisms the *Schrödinger mechanisms*.

We use the folklore expansion [18, Section 2.6]

$$D(p \,\|\, T_a p) = \frac{a^2}{2} I(p) + o(a^2) \quad \text{as } a \to 0, \tag{4}$$

where $I(p)$ is the Fisher information. Consequently, the min-

*Corresponding author, and the remaining authors are in alphabetical order.

This material is based upon work supported by the National Science Foundation under Grant Nos. CAREER-1845852, FAI-2040880, CIF-1900750, SCH-2205080, CIF-1922971, CIF-1901243, and CIF-2007688; by NSERC Canada; and by the U.S. Department of Energy Award No. DE-SC0022158.

imax optimization of KL-divergence in (3) reduces to finding the *unique* minimizer of $I(p)$ over all PDFs p satisfying the utility constraint. This reduced formulation reveals a remarkable characterization of the optimizer $p^\star$: it is the square of the ground-state eigenfunction of a Schrödinger operator (Theorem 3). This general characterization provides a powerful tool to identify closed-form DP mechanisms with the optimal privacy-utility trade-off where the utility is measured via the cost function c. In particular, we show that $p^\star$ is the Gaussian PDF for the L^2 cost function (Proposition 3), thereby proving that the Gaussian mechanism is optimal in this sense in the small-sensitivity regime. Our results also show that $p^\star$ for the L^1 cost is given by the Airy function, leading to the introduction of a new optimal DP mechanism, which we call the *Airy mechanism* (see Definition 4).

A. Related Work and Contributions

Several optimal mechanisms in DP settings are known, e.g., stair-case mechanism [19]–[21], geometric mechanism [22], discrete Laplace mechanism [23], truncated Laplace mechanism [24], and uniform mechanism [25], to name a few. All these works assume a query with a given sensitivity in a single-shot setting (i.e., no compositions). Unlike these works, we focus on characterizing optimal mechanisms under large composition when the query's sensitivity is rather small. Although optimal mechanisms for the large-composition regime are treated in [17], the work therein considers *fixed* sensitivity.

Compared to existing literature on the problem of minimizing the Fisher information, we:

1) work with a larger class of cost functions,
2) do not restrict the support of the PDFs we optimize over,
3) do not require any regularity assumptions whatsoever on the PDFs we optimize over.

We go beyond existing literature by introducing a novel proof technique that does not depend on the calculus of variation, and also by deriving an estimate of the logarithmic derivative of the ground-state eigenfunction of the Schrödinger operator.

The statistics literature is rife with results on Fisher-information-minimizing distributions. The Cramér-Rao bound implies that Gaussian measures are optimal for a given variance. The minimizer over compactly-supported distributions or over those supported on $\mathbb{R}^+$ were characterized in [26] and [27], respectively. Kagan [28] studied the same problem for densities on $\mathbb{R}$ with fixed first and second moments, which was later extended to other moments by Ernst [29]. A connection between minimizing Fisher information and the Schrödinger equation has been established in [30, Example 5.1]. Formulating a privacy problem in terms of minimizing Fisher information has appeared in [31], [32], but not in a DP sense; rather, the analyses therein pertain to privacy-preserving battery charging methods to obfuscate household information, and the Fisher information itself is proposed as a privacy metric. Fisher information minimization in [31] is done for PDFs of compact support, and that is extended to unbounded support in [32] but for only a quadratic cost. Further, the PDFs considered in [31], [32] are assumed *a priori* to be twice continuously differentiable. Therefore, none of these previous works has a setup encompassing ours, namely, they minimize Fisher information: over PDFs supported over a compact set [26], [31] or over $\mathbb{R}^+$ [27]; assuming regularity of the PDFs [30]–[32]; or under a strictly smaller or different class of constraint functions [28], [29], [32].

We discuss in more detail how our work differs from the existing literature [29]–[32] in [1, Appendix A].

B. Notation and Assumptions

We let λ denote the Lebesgue measure on $\mathbb{R}$. The set of all probability density functions (PDFs) on $\mathbb{R}$ is denoted by $\mathcal{P}(\mathbb{R})$. For $p \in \mathcal{P}(\mathbb{R})$ and $c : \mathbb{R} \to \mathbb{R}$, the expectation is denoted by $\mathbb{E}_p[c] := \int_{\mathbb{R}} c(x)p(x)\,dx$. The shift operator is denoted by T_x, i.e., $(T_x r)(A) := r(A - x)$.

The Fisher information of $p \in \mathcal{P}(\mathbb{R})$ is denoted by $I(p)$, i.e., if p is absolutely continuous then

$$I(p) := \int_{\{x\in\mathbb{R}\,;\,p(x)>0\}} \frac{p'(x)^2}{p(x)}\,dx, \tag{5}$$

and $I(p) = \infty$ otherwise. The KL-divergence is denoted by $D(p\,\|\,q)$ if $p, q \in \mathcal{P}(\mathbb{R})$. The variance of the information density is denoted by (for $p, q \in \mathcal{P}(\mathbb{R})$)

$$\mathrm{V}(p \parallel q) := \mathbb{E}_p\left[\left(\log\frac{p}{q} - D(p\|q)\right)^2\right]. \tag{6}$$

It is well-known (see, e.g., [18, Section 2.6]) that, under mild regularity conditions on a PDF p, one has the expansion in (4). Define the subset of PDFs $\mathcal{F} \subset \mathcal{P}(\mathbb{R})$ by

$$\mathcal{F} := \left\{ p \in \mathcal{P}(\mathbb{R}) : (4) \text{ holds}, \sup_{|a|\le s} D(p\|T_a p) < \infty \text{ and } \sup_{|a|\le s} \mathrm{V}(p\|T_a p) < \infty \text{ for some } s > 0 \right\}. \tag{7}$$

The minimization problem we solve for Fisher information is global, i.e., over all of $\mathcal{P}(\mathbb{R})$, while the DP optimization we solve will be over the set $\mathcal{F}$ defined in (7).

The results of this paper hold for the following class of cost functions c. We note that this class includes functions such as $c(x) = \beta|x|^\alpha$ and $c(x) = \beta\log(|x|+1)^\alpha$ for any $\alpha, \beta > 0$.

Assumption 1. *The cost function $c : \mathbb{R} \to \mathbb{R}$ satisfies:*

1) Positivity: $c(x) \ge 0$ *for all* $x \in \mathbb{R}$, *and* $c(0) = 0$.
2) Symmetry: $c(x) = c(-x)$ *for all* $x \in \mathbb{R}$.
3) Monotonicity: $c(x_1) \le c(x_2)$ *if* $|x_1| \le |x_2|$.
4) Continuity: c *is continuous over* $\mathbb{R}$.
5) Unbounded: $c(x) \to \infty$ *as* $x \to \infty$,
6) Controlled derivative: $c'(x) = o\big(c(x)^{3/2}\big)$ *as* $x \to \infty$,
7) Tail regularity: $\int_{x_0}^\infty |c'|^2/|c|^{5/2}, \int_{x_0}^\infty |c''|/|c|^{3/2} < \infty$ *for some* $x_0 \in \mathbb{R}$,
8) Moderate growth: $x \mapsto \sqrt{c(x)}/\exp(\gamma\int_0^{|x|}\sqrt{c(t)}\,dt)$ *is integrable for all* $\gamma > 0$,
9) Additive/Multiplicative regularity: *there is a locally bounded strictly positive function* ρ *on* $\mathbb{R}$ *such that* $c(x-t), c(tx) \le \rho(t)(c(x)+1)$ *for all* $x, t \in \mathbb{R}$.

In the assumptions involving c' or c'', it is to be understood that c is required to be differentiable (or twice differentiable) *only* at large enough values.

II. From DP to KL-Divergence to Fisher Information

Let $\mathcal{D}$ be the collection of datasets, each of which contains sensitive data of several individuals, and $f : \mathcal{D} \to \mathbb{R}$ be a query function. The quantity of interest is $x = f(d)$, which is the outcome of the query f on the dataset $d \in \mathcal{D}$ (e.g., $f(d)$ could be the percentage of individuals in d falling inside a certain income bracket). To protect the privacy of individuals against membership and inference attacks, a typical approach is to perturb $f(d)$ using a channel (or mechanism) $P_{Y|X=f(d)}$ so that Y cannot be used to distinguish d from a neighboring dataset d' that differs from d in one entry. This approach, known as *differential privacy* [2], is formalized as follows. Given $\varepsilon \geq 0$ and $\delta \in [0,1]$, a mechanism $P_{Y|X}$ is said to be (ε, δ)-differentially private (or (ε, δ)-DP for short) if

$$\sup_{d \sim d'} \sup_{A \subset \mathcal{Y}} \left[P_{Y|X=f(d)}(A) - e^{\varepsilon} P_{Y|X=f(d')}(A) \right] \leq \delta, \quad (8)$$

where the outer supremum is taken over all pairs of neighboring datasets d and d', denoted by $d \sim d'$, and the inner supremum is taken over all measurable subsets A of the support $\mathcal{Y}$ of Y. If a mechanism $P_{Y|X}$ is (ε, δ)-DP for sufficiently small ε and δ, then an adversary observing Y cannot accurately distinguish d from an arbitrary neighboring d', thus providing a tunable privacy guarantee for each individual in d. A popular family of such DP mechanisms includes *additive* ones, that is, $Y = f(d) + Z$ where $Z \sim P$ is a noise variable drawn from a distribution P.[1]

We note that the DP definition in (8) can be more compactly expressed using the E_γ-divergence [33] defined for $\gamma \geq 0$ as

$$\mathsf{E}_\gamma(P \,\|\, Q) := \sup_{A \text{ Borel}} P(A) - \gamma Q(A). \quad (9)$$

With this definition at hand, we can say $P_{Y|X}$ is (ε, δ)-DP if

$$\sup_{|x - x'| \leq s} \mathsf{E}_{e^\varepsilon}(P_{Y|X=x} \parallel P_{Y|X=x'}) \leq \delta, \quad (10)$$

where s denote the sensitivity of the query f, defined as $s := \sup_{d \sim d'} |f(d) - f(d')|$.

Next, consider a typical composition setting where a dataset d is queried n times sequentially with query functions f_j, $1 \leq j \leq n$, and a mechanism $P_{Y|X}$ is used n times to generate the n-tuple $Y^n = (Y_1, \ldots, Y_n)$ as a private version of the n-tuple $(f_1(d), \cdots, f_n(d))$. For simplicity, we assume that each f_j has the same sensitivity s. Therefore, this n-fold composition $P_{Y|X}^{\circ n}$ is $(\varepsilon, \delta_{P_{Y|X}^{\circ n}, s}(\varepsilon))$-DP for any $\varepsilon \geq 0$, where[2]

$$\delta_{P_{Y|X}^{\circ n}, s}(\varepsilon) := \sup_{\substack{|u_j - v_j| \leq s \\ 1 \leq j \leq n}} \mathsf{E}_{e^\varepsilon} \left(\prod_{j=1}^n P_{Y|X=u_j} \middle\| \prod_{j=1}^n P_{Y|X=v_j} \right). \quad (11)$$

Equivalently, $P_{Y|X}^{\circ n}$ is $(\varepsilon_{P_{Y|X}^{\circ n}, s}(\delta), \delta)$-DP for $\delta \in [0,1]$, where

$$\varepsilon_{P_{Y|X}^{\circ n}, s}(\delta) := \inf \Big\{ \varepsilon \geq 0 \ : \ \delta_{P_{Y|X}^{\circ n}, s}(\varepsilon) \leq \delta \Big\}. \quad (12)$$

Since additive continuous channels were shown to be optimal in [17], we henceforth consider only channels of the form $P_{Y|X=x} = T_x P$ with P being absolutely continuous with respect to the Lebesgue measure λ, for which we use the simplified notation $\varepsilon_{p^{\circ n}, s}(\delta)$ where $p = dP/d\lambda$ is the PDF. We derive the following asymptotic formula for $\varepsilon_{P_{Y|X}^{\circ n}, s}(\delta)$.[3]

Theorem 1. *For any PDF $p \in \mathcal{P}(\mathbb{R})$ and $s > 0$ satisfying $\sup_{|a| \leq s} \mathrm{V}(p \| T_a p) < \infty$, and for any $\delta \in (0, 1/2)$, we have*

$$\lim_{n \to \infty} \frac{\varepsilon_{p^{\circ n}, s}(\delta)}{n} = \sup_{|a| \leq s} D(p \| T_a p). \quad (13)$$

According to this theorem, characterizing $\varepsilon_{p^{\circ n}, s}(\delta)$ for sufficiently large n boils down to computing the maximum of $D(p \parallel T_a p)$ over all $|a| \leq s$.

Analogous to [17], we address the utility of the mechanism $P_{Y|X}$ by imposing the bound $\mathbb{E}[c(Y - x) \mid X = x] \leq C$ for all $x \in \mathbb{R}$ and a given $C \geq 0$, where $c : \mathbb{R} \to \mathbb{R}^+$ is a measurable cost function. Notice that for additive mechanisms $P_{Y|X=x} = T_x P$, this utility constraint reduces to $\mathbb{E}_P[c] \leq C$. Motivated by the asymptotic given in Theorem 1, we consider the following optimality in the small sensitivity regime.

Definition 1. Let $\mathcal{F} \subset \mathcal{P}(\mathbb{R})$ be as defined in (7).[4] We say that a PDF $p \in \mathcal{F}$ is *optimal in the small-sensitivity regime* for the cost function c and the cost bound C if $\mathbb{E}_p[c] \leq C$, and for every other PDF $q \in \mathcal{F}$ (i.e., $\lambda(\{p = q\}) = 0$) satisfying $\mathbb{E}_q[c] \leq C$ there is a constant $s(q) > 0$ such that $0 < s < s(q)$ implies

$$\sup_{0 < \delta < \frac{1}{2}} \lim_{n \to \infty} \frac{\varepsilon_{p^{\circ n}, s}(\delta)}{\varepsilon_{q^{\circ n}, s}(\delta)} < 1. \quad (14)$$

An immediate corollary of Theorem 1 is that the unique minimizer of the Fisher information is automatically the optimal PDF in the small-sensitivity regime.

Corollary 1. *If $p \in \mathcal{F}$ is the unique minimizer*

$$p = \operatorname*{argmin}_{\substack{q \in \mathcal{F} \\ \mathbb{E}_q[c] \leq C}} I(q), \quad (15)$$

[1] Alternatively, one can express additive mechanisms by $P_{Y|X=x} = T_x P$, where T_x denotes the shift operator defined as $(T_x P)(A) := P(A - x)$.

[2] While the sensitivity s is usually suppressed from the notation of δ and ε in the literature, we include it here since we consider a variable sensitivity.

[3] A similar result appears in [16] under additional third-moment constraints, and also under the assumption of existence of "worst-case shifts." Thus, our result can be seen as a generalization of the asymptotic formula in [16].

[4] For the Gaussian density φ^σ, we have $D(\varphi^\sigma \| T_a \varphi^\sigma) = a^2/(2\sigma^2)$. Thus, if one insists that the PDF p satisfy $D(p \| T_a p) \leq D(\varphi^\sigma \| T_a \varphi^\sigma)$ for all small a, then the mapping $a \mapsto D(p \| T_a p)$ is necessarily differentiable at $a = 0$ with vanishing derivative. In particular, one reasonably expects that desirable PDFs for the small-sensitivity regime to belong to $\mathcal{F}$.

then p is the optimal PDF in the small-sensitivity regime for the cost function c and the cost bound C.

We derive in the next section minimizers of Fisher information over *all* PDFs $\mathcal{P}(\mathbb{R})$, then we also show that such minimizers in fact fall within the set $\mathcal{F}$.

III. From Fisher Information to the Schrödinger Equation

Solving the Fisher information minimization problem reveals a bridge between DP and the Schrödinger operator. This connection enables us to show that the global minimizers of Fisher information are fully characterized by the minimal-eigenvalue eigenfunctions of the Schrödinger operator (see Theorem 2) with the potential given by the cost function c.

A. The Schrödinger Equation

We begin by recalling the setup of the Schrödinger operator eigen-problem and some of its known properties.

Definition 2 ([34, Section 2.4])**.** For a measurable $v : \mathbb{R} \to \mathbb{R}$, the Schrödinger operator $\mathcal{H}_v$ on $L^2(\mathbb{R})$ with potential v is defined as[5]

$$\mathcal{H}_v(y) := -y'' + vy. \tag{16}$$

We say that $y \in L^2(\mathbb{R})$ is an eigenfunction of $\mathcal{H}_v$ if y is differentiable, y' is absolutely continuous, and there exists a constant E such that $\mathcal{H}_v(y) = Ey$ holds a.e.

The spectrum of $\mathcal{H}_v$ is discrete: if v is locally bounded and $\lim_{|x|\to\infty} v(x) = \infty$ then $L^2(\mathbb{R})$ has an orthonormal complete set consisting of eigenfunctions of $\mathcal{H}_v$ with eigenvalues $\{E_k\}_{k\in\mathbb{N}}$ such that $E_k \to \infty$ (see [34, Chapter 2, Theorem 3.1]). Moreover, one may order the E_k in an increasing fashion, and then the eigenfunction associated to E_k has exactly k zeros (see [34, Chapter 2, Theorem 3.5]). We are interested in the smallest eigenvalue E_0 and the associated eigenfunction, i.e., the ground-state eigenfunction. An easy consequence of known properties of the ground-state eigenfunction is as follows.

Lemma 1. *For any $\theta > 0$, there exists a unique unit-L^2-norm eigenfunction $y_{\theta,c}$ of $\mathcal{H}_{\theta c}$ satisfying $y_{\theta,c}(x) > 0$ for all $x \in \mathbb{R}$. Further, $y_{\theta,c}$ is even, and its eigenvalue is the smallest eigenvalue of $\mathcal{H}_{\theta c}$.*

The notation $y_{\theta,c}$ as given by Lemma 1 will be used in the remainder of the paper.

B. Global Minimization of Fisher Information

Recall the recipe we provide in Section II for finding optimal DP mechanisms in the small-sensitivity regime:

1) globally minimize Fisher information (i.e., over $\mathcal{P}(\mathbb{R})$),
2) show that the solution in fact falls within $\mathcal{F}$,
3) use Theorem 1 to conclude that the Fisher information global minimizer is the optimal DP mechanism.

We carry out step 1 in Theorem 2 below, where we show that $y^2_{\theta,c}$ is the unique global minimizer of the Fisher information. After that, we complete our general derivations in Proposition 2 by showing that step 2 holds, i.e., $y^2_{\theta,c} \in \mathcal{F}$.

Theorem 2. *Suppose c satisfies Assumption 1, fix $\theta > 0$, consider the PDF $p = y^2_{\theta,c}$, and set $C = \mathbb{E}_p[c]$. Then, the PDF p uniquely minimizes the Fisher information among all PDFs $q \in \mathcal{P}(\mathbb{R})$ that satisfy $\mathbb{E}_q[c] \le C$, i.e.,*

$$p = \operatorname*{argmin}_{\substack{q\in\mathcal{P}(\mathbb{R})\\ \mathbb{E}_q[c]\le C}} I(q). \tag{17}$$

Since Theorem 2 gives a general unconditional result, our work can be seen as a way to fill the gaps in [29]–[32]. In the next section, we also provide a new *explicit* solution for the absolute-value cost case. Our method of proof deviates from those in [29]–[32], where we borrow results from the quantum mechanics literature (such as [34]) to show that the needed properties for p can be derived instead of assumed. For instance, we show that the unique eigenfunction $y_{\theta,c}$ as given by Lemma 1 satisfies the following bound.

Proposition 1. *For c satisfying Assumption 1 and any $\theta > 0$, we have the bound*

$$\limsup_{|x|\to\infty} \left| \frac{y'_{\theta,c}(x)}{y_{\theta,c}(x)\sqrt{c(x)}} \right| \le \sqrt{\theta}. \tag{18}$$

Finally, we show in the following result that the PDF $y^2_{\theta,c}$ falls within the set $\mathcal{F}$ defined in (7).

Proposition 2. *For any c satisfying Assumption 1 and any $\theta > 0$, we have that $y^2_{\theta,c} \in \mathcal{F}$.*

Next, we combine Theorems 1–2 and Proposition 2 to show in Theorem 3 that the PDF $y^2_{\theta,c}$ is the optimal DP mechanism in the sense of Definition 1.

C. The Schrödinger Mechanism

Since $y_{\theta,c}$ is a Borel function satisfying $\|y_{\theta,c}\|_2 = 1$, we get that $y^2_{\theta,c}$ is a PDF. We call $y^2_{\theta,c}$ the Schrödinger mechanism.

Definition 3. The *Schrödinger mechanism* given the cost function c and parameter $\theta > 0$ is defined by $Y = X + Z$ for Z having the PDF $y^2_{\theta,c}$ where $y = y_{\theta,c}$ is the unique unit-L^2-norm and strictly positive solution to the Schrödinger equation

$$y'' = (\theta c - E)y, \tag{19}$$

with E an arbitrary constant.[6]

Combining our results, we get that the Schrödinger mechanism is optimal in the small-sensitivity regime.

Theorem 3. *If a cost satisfies Assumption 1, the Schrödinger mechanism is optimal in the small-sensitivity regime in the sense of Definition 1.*

[5] One may define $\mathcal{H}_v$ initially on compactly-supported C^∞ functions, then show that its closure is self-adjoint if v satisfies mild conditions (see [34, Chapter 2, Theorem 1.1]). In particular, this extension goes through if v is nonnegative (and measurable).

[6] By Lemma 1, there is a unique E for which the ODE (19) is solvable with the prescribed properties for the solution y, and the solution then is $y = y_{\theta,c}$.

Remark 1. For the two examples we discuss in the next section, we give a reversing procedure producing θ given C that takes the form $\theta = aC^{-b}$ for absolute constants a and b.

IV. From the Schrödinger Equation to the Gaussian and Airy Mechanisms

Next, we instantiate Theorem 3 for two different cost functions, namely the quadratic and absolute-value cost functions.

A. Quadratic cost: optimality of Gaussian

Consider first the quadratic cost function $c(x) = x^2$. By particularizing Theorem 3 to this case, we show that the Gaussian distribution is optimal in the small-sensitivity regime in the sense of Definition 1. This is a direct consequence of the Cramér-Rao bound, but we derive it here using Theorem 3. The Schrödinger equation to be solved becomes

$$y''(x) = \big(\theta x^2 - E\big)y(x). \tag{20}$$

Proposition 3. *For a quadratic cost $c(x) = x^2$, the Gaussian mechanism is optimal in the small-sensitivity regime.*

B. Absolute value cost: optimality of Airy

We next consider the absolute-value cost $c(x) = |x|$. In this case, the eigenvalue problem $\mathcal{H}_{\theta c}(y) = Ey$ becomes

$$y''(x) = (\theta|x| - E)y(x), \tag{21}$$

for some $\theta > 0$. It will be useful to recall the definition of the Airy functions. The differential equation

$$y''(x) = xy(x) \tag{22}$$

has two linearly independent solutions, called the Airy functions [35, Chapter 9]. They are denoted by Ai and Bi, where Ai is the solution such that $\mathsf{Ai}(x) \to 0$ as $x \to \infty$. This function can be expressed by the improper Riemann integral

$$\mathsf{Ai}(x) = \frac{1}{\pi} \lim_{N\to\infty} \int_0^N \cos\left(\frac{t^3}{3} + xt\right) dt. \tag{23}$$

This function is analytic, and there are countably many zeros of Ai and Ai$'$ all falling on the negative half-line. As is customary, the zeros of Ai and Ai$'$ are denoted by $a_1 > a_2 > \cdots$ and $a_1' > a_2' > \cdots$, respectively. It is known that approximately

$$a_1 = -2.33810, \quad a_1' = -1.01879, \quad \text{and} \quad \mathsf{Ai}(a_1') = 0.53565.$$

In particular, the function Ai is strictly positive and strictly decreasing over $[a_1', \infty)$. We use the Airy function to construct the following density, which we show afterwards to be optimal.

Definition 4. For $C > 0$, we define the *Airy distribution* with first absolute moment C as the probability measure whose PDF $p_{\mathsf{Ai},C}$ is given by

$$p_{\mathsf{Ai},C}(x) := \frac{1}{3C\mathsf{Ai}(a_1')^2} \mathsf{Ai}\left(\frac{-2a_1'}{3C}|x| + a_1'\right)^2. \tag{24}$$

Remark 2. It can be verified with some algebra that $p_{\mathsf{Ai},C}$ is a valid PDF and that its first absolute moment is indeed C.

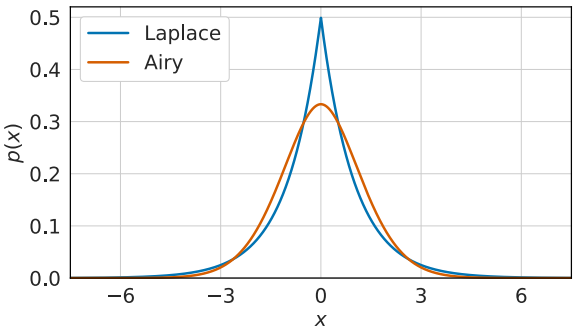

Fig. 1: The densities of the Laplace and Airy distributions ($p_{\mathsf{Ai},C}(x)$, introduced in Definition 4), with $C = \mathbb{E}[|X|] = 1$.

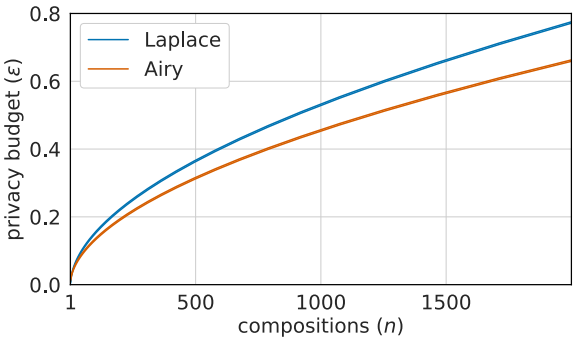

Fig. 2: The privacy budget ε versus the number of the compositions n, for the constraint $C = \mathbb{E}[|X|] = 2$, $s = 1$, fixed privacy parameter $\delta = 10^{-8}$, and subsampling rate $q = 0.01$.

Proposition 4. *For an absolute-value cost $c(x) = |x|$, the Airy mechanism is optimal in the small-sensitivity regime.*

In Figure 1, we illustrate the Airy distribution and compare it with the Laplace distribution. We note that the Airy distribution has a lighter tail than that of the Laplace distribution, where the exponential decay of the former is $e^{-2x^{3/2}/3}$ and that of the latter is e^{-x}.

Experiments: Finally, we demonstrate that the Airy mechanism can achieve better DP parameters than the Laplace mechanism for the same fixed absolute-value cost. In particular, we subsample both mechanisms, following standard practice in the DP machine learning community for amplifying privacy [8], [36], [37]. We also use the arbitrary-accuracy FFT-based numerical accountant introduced in [14] to compute tight privacy bounds for finite compositions. In Figure 2, we fix $\delta = 10^{-8}$ and estimate ε under a varying number of compositions. Under this construction, the accountant in [14] computes both upper and lower bounds on ε. We choose $\varepsilon_{\text{error}} = 0.002$ and $\delta_{\text{error}} = 10^{-10}$, effectively making the upper and lower bounds indistinguishable (they are both plotted in Figure 2). The Airy mechanism provides stronger privacy guarantees for all values of compositions ($1 \le n \le 2000$), and the gap increases with composition.

REFERENCES

[1] W. Alghamdi, S. Asoodeh, F. P. Calmon, J. F. Gomez, O. Kosut, and L. Sankar, "Schrödinger mechanisms: Optimal differential privacy mechanisms for small sensitivity," 2023. [Online]. Available: https://github.com/WaelAlghamdi/DP-Schrodinger

[2] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory of Cryptography (TCC)*, Berlin, Heidelberg, 2006, pp. 265–284.

[3] Ú. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. ACM, 2014, pp. 1054–1067.

[4] Differential privacy team Apple, "Learning with privacy at scale," 2017.

[5] D. Kifer, S. Messing, A. Roth, A. Thakurta, and D. Zhang, "Guidelines for implementing and auditing differentially private systems," *ArXiv*, vol. abs/2002.04049, 2020.

[6] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *51st Annual Symposium on Foundations of Computer Science*. IEEE, 2010, pp. 51–60.

[7] J. Murtagh and S. Vadhan, "The complexity of computing the optimal composition of differential privacy," in *Proc. Int. Conf. Theory of Cryptography*, 2016, pp. 157–175.

[8] 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*, 2016, pp. 308–318.

[9] S. Asoodeh, J. Liao, F. P. Calmon, O. Kosut, and L. Sankar, "Three variants of differential privacy: Lossless conversion and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 208–222, 2021.

[10] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," in *Proceedings of the 32nd International Conference on Machine Learning*, F. Bach and D. Blei, Eds., vol. 37, 2015, pp. 1376–1385.

[11] S. Meiser and E. Mohammadi, "Tight on budget? tight bounds for r-fold approximate differential privacy," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18, 2018, pp. 247–264.

[12] J. Dong, A. Roth, and W. J. Su, "Gaussian differential privacy," *CoRR*, vol. abs/1905.02383, 2019. [Online]. Available: http://arxiv.org/abs/1905.02383

[13] A. Koskela, J. Jälkö, and A. Honkela, "Computing tight differential privacy guarantees using fft," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2560–2569.

[14] S. Gopi, Y. T. Lee, and L. Wutschitz, "Numerical composition of differential privacy," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[15] A. Koskela, J. Jälkö, L. Prediger, and A. Honkela, "Tight differential privacy for discrete-valued mechanisms and for the subsampled gaussian mechanism using fft," in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Banerjee and K. Fukumizu, Eds., vol. 130. PMLR, 13–15 Apr 2021, pp. 3358–3366. [Online]. Available: https://proceedings.mlr.press/v130/koskela21a.html

[16] W. Alghamdi, S. Asoodeh, F. P. Calmon, J. F. Gomez, O. Kosut, L. Sankar, and F. Wei, "The saddle-point accountant for differential privacy," *arXiv preprint arXiv:2208.09595*, 2022.

[17] W. Alghamdi, S. Asoodeh, F. P. Calmon, O. Kosut, L. Sankar, and F. Wei, "Cactus mechanisms: Optimal differential privacy mechanisms in the large-composition regime," in *2022 IEEE International Symposium on Information Theory (ISIT)*, 2022, pp. 1838–1843.

[18] S. Kullback, *Information Theory and Statistics*. Wiley, 1959.

[19] Q. Geng and P. Viswanath, "The optimal noise-adding mechanism in differential privacy," *IEEE Transactions on Information Theory*, vol. 62, no. 2, pp. 925–951, 2015.

[20] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1176–1184, 2015.

[21] J. Soria-Comas and J. Domingo-Ferrer, "Optimal data-independent noise for differential privacy," *Information Sciences*, vol. 250, pp. 200–214, 2013.

[22] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," *SIAM Journal on Computing*, vol. 41, no. 6, pp. 1673–1693, 2012.

[23] Q. Geng and P. Viswanath, "Optimal noise adding mechanisms for approximate differential privacy," *IEEE Transactions on Information Theory*, vol. 62, no. 2, pp. 952–969, 2016.

[24] Q. Geng, W. Ding, R. Guo, and S. Kumar, "Tight analysis of privacy and utility tradeoff in approximate differential privacy," in *Proc. International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, S. Chiappa and R. Calandra, Eds., vol. 108, 2020, pp. 89–99.

[25] ——, "Optimal noise-adding mechanism in additive differential privacy," in *Proc. International Conference on Artificial Intelligence and Statistics*, K. Chaudhuri and M. Sugiyama, Eds., vol. 89, 2019, pp. 11–20.

[26] E. Uhrmann-Klingen, "Minimal fisher information distributions with compact-supports," *Sankhyā: The Indian Journal of Statistics*, vol. 57, no. 3, pp. 360–374, 1995.

[27] J. F. Bercher and C. Vignat, "On minimum fisher information distributions with restricted support and fixed variance," *Inf. Sci.*, vol. 179, no. 22, pp. 3832–3842, 2009.

[28] A. M. Kagan, "Information property of exponential families," *Theory of Probability & Its Applications,*, vol. 30, no. 4, pp. 831–835, 1986.

[29] P. A. Ernst, "Minimizing fisher information with absolute moment constraints," *Statistics & Probability Letters*, vol. 129, pp. 167–170, 2017.

[30] P. J. Huber and E. M. Ronchetti, *Robust Statistics, Second Edition*. Wiley, 2009.

[31] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2018.

[32] ——, "Ensuring privacy with constrained additive noise by minimizing fisher information," *Automatica*, vol. 99, pp. 275–288, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109818304862

[33] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.

[34] F. A. Berezin and M. Shubin, *The Schrödinger Equation*. Dordrecht: Springer, 1991.

[35] "*NIST Digital Library of Mathematical Functions*," http://dlmf.nist.gov/, Release 1.1.4 of 2022-01-15, f. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. [Online]. Available: http://dlmf.nist.gov/

[36] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?" *SIAM J. Comput.*, vol. 40, no. 3, p. 793–826, jun 2011.

[37] A. Beimel, K. Nissim, and U. Stemmer, "Characterizing the sample complexity of private learners," in *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, 2013, p. 97–110.