Differential Privacy from Axioms

Guy Blanc
Stanford University, CA, USA

William Pires
Columbia University, New York, NY, USA

Toniann Pitassi
Columbia University, New York, NY, USA

Abstract

Differential privacy (DP) is the de facto notion of privacy both in theory and in practice. However, despite its popularity, DP imposes strict requirements which guard against strong worst-case scenarios. For example, it guards against seemingly unrealistic scenarios where an attacker has full information about all but one point in the data set, and still nothing can be learned about the remaining point. While preventing such a strong attack is desirable, many works have explored whether average-case relaxations of DP are easier to satisfy [17, 28, 5, 22].

In this work, we are motivated by the question of whether alternate, weaker notions of privacy are possible: can a weakened privacy notion still guarantee some basic level of privacy, and on the other hand, achieve privacy more efficiently and/or for a substantially broader set of tasks? Our main result shows the answer is no: even in the statistical setting, any reasonable measure of privacy satisfying nontrivial composition is equivalent to DP. To prove this, we identify a core set of four axioms or desiderata: pre-processing invariance, prohibition of blatant non-privacy, strong composition, and linear scalability. Our main theorem shows that any privacy measure satisfying our axioms is equivalent to DP, up to polynomial factors in sample complexity. We complement this result by showing our axioms are minimal: removing any one of our axioms enables ill-behaved measures of privacy.

2012 ACM Subject Classification Theory of computation → Theory of database privacy and security; Theory of computation → Machine learning theory

Keywords and phrases Differential Privacy, Privacy Amplification, Composition

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.21

Related Version *Full Version*: `https://arxiv.org/abs/2511.21876` [6]

Funding *Guy Blanc*: Supported by NSF awards 1942123, 2211237, and 2224246 and a Jane Street Graduate Research Fellowship.
William Pires: Supported by NSF awards 2106429, 2107187.
Toniann Pitassi: Supported by the Simons Foundation Collaboration on the Theory of Algorithmic Fairness.

Acknowledgements The authors thank the anonymous ITCS reviewers for their helpful feedback.

1 Introduction

Differential privacy has won. It is the de facto formalization of privacy both in theory (see, e.g., the textbooks [13, 27, 24]) and in practice (see, e.g., its use in the U.S. Census [3] and by various technology companies [1, 29, 9]).

▶ **Definition 1** ((ε, δ)-Differential Privacy, [12, 11]). *A randomized algorithm $\mathcal{M} : X^n \to Y$ is (ε, δ)-DP if, for every $S, S' \in X^n$ differing in only one of the n coordinates and $Y' \subseteq Y$,*

$$\Pr[\mathcal{M}(S) \in Y'] \leq e^{\varepsilon} \cdot \Pr[\mathcal{M}(S') \in Y'] + \delta.$$

A large part of the reason that differential privacy (DP) has been so successful is the extensive toolkit of DP algorithms for a variety of basic primitives [13]. This toolkit can then be combined with *strong composition*: The sequential combination of k-many of these primitives has a privacy loss (ε in Definition 1) that scales sublinearly in k [15, 19]. This allows for efficient and simple construction of DP algorithms for a variety of tasks (see e.g. [2] for how strong composition enables differentially private deep learning). This work is motivated by the following question.

▶ **Question 1.** How inevitable was Definition 1? Is it possible to construct a materially different formulation of privacy that still satisfies strong composition?

A natural reason to suspect alternative definitions of privacy may be useful is that Definition 1 guards against an incredibly strong, and in some cases unrealistic, attack. Even if the attacker is able to freely manipulate all but one point in the dataset, corresponding to the $n-1$ points S and S' agree on, they must still learn almost nothing about the one unknown point. In statistical settings, we model the entire dataset as being drawn from some unknown distribution $\boldsymbol{S} \sim \mathcal{D}^n$, in which case the attacker is not nearly as strong as Definition 1 suggests. That observation has motivated a number of relaxations of DP in which privacy must only be preserved on more "typical" datasets [17, 28, 5, 22].

Our main result shows that we may as well use the worst-case definition of differential privacy.

> *Even in the statistical setting, any reasonable measure of privacy that satisfies strong composition is equivalent to Definition 1 up to polynomial factors in the sample complexity.*

To formalize this, we define the following four privacy axioms that we posit should be satisfied by any measure of privacy that is both reasonable and useful.[1]

1. *Preprocessing:* Privacy is preserved under preprocessing. Specifically, privacy should hold regardless of the ordering of the dataset, and regardless of the ordering of the domain.
2. *Prohibits blatant non-privacy:* a private algorithm should not reveal almost all of the dataset.
3. *Strong composition:* the privacy measure should grow sublinearly under composition. I.e., the composition of ℓ-many ε-private algorithms should be $O(\varepsilon\ell^{\delta})$-private, for some $\delta < 1$.
4. *Linear scalability:* the privacy measure should decrease linearly with the number of samples.

See Section 2 for a more detailed description of these axioms, and justification for why we view these axioms as both reasonable and usable.

With these axioms in place, our main results are captured by the following three theorems. The first and most important theorem states that any algorithm satisfying our axioms is also differentially private:

▶ **Theorem 2** (Our axioms imply differential privacy)**.** *Let $\mathcal{P}$ be any privacy measure that satisfies Axioms 1–4 and $\mathcal{M} : X^n \to Y$ be any algorithm that is $\mathcal{P}$-private. For any $\varepsilon, \delta > 0$ and $m \coloneqq \mathrm{poly}(n, 1/\varepsilon, \log(1/\delta))$ there is an (ε,δ)-DP algorithm $\mathcal{M}' : X^m \to Y$ that is equivalent (as in Definition 6) to $\mathcal{M}$.*

[1] By a privacy measure, we mean a scalar quantity $\mathcal{P}(\mathcal{M})$ associated with an algorithm $\mathcal{M}$, and we say that $\mathcal{M}$ is $\mathcal{P}$-private if $\mathcal{P}(\mathcal{M})$ is at most 1.

The exact polynomial in Theorem 2 depends on the constant c in our our strong composition axiom (Axiom 3). The best known constant for strong-composition is $c = 1/2$, in which case the sample-size in Theorem 2 would be $m \approx n^2$, provided the domain X is not too small.[2] We refer the reader to the full version of this work [6] for the formal version of Theorem 2. We note that given the known equivalences between DP, replicability, and various notions of stability, Theorem 2 shows that these other notions are also implied by our axioms.

Second, we show that differential privacy satisfies our axioms.

▶ **Theorem 3** (Informal). *Approximate differential privacy (Definition 1) satisfies Axioms 1–4.*

Lastly, we show that removal of any one of our axioms would allow for measures of privacy that do not intuitively align with any reasonable notion of privacy. The reader can find a brief overview of what those nonsensical privacy measures are in Section 3.3. For an in-depth discussion we refer the reader to the full version of this work [6]. We do have the following simple implication of those results.

▶ **Theorem 4** (Minimality of our axioms). *Theorem 2 does not hold if $\mathcal{P}$ is allowed to not satisfy any one of Axioms 1–4.*

Organization of Paper. In Section 2, we present our framework and our axiomatic formulation of privacy; in Section 3, we give high level overviews of the proofs of our main theorems, and in Section 4 we discuss related work and several open problems raised by our framework. The formal proofs of Theorems 2–4 can be found in the full version of this work [6].

2 Our Framework

All of our equivalences will hold with respect to algorithms that solve statistical tasks.

▶ **Definition 5** (Statistical task, [16, 7]). *A statistical task is defined by a set of distributions $\mathscr{D}$ over data domain X, an output space Y and a mapping $\mathcal{T}$ from distributions $\mathcal{D} \in \mathscr{D}$ to valid responses $\mathcal{T}(\mathcal{D}) \subseteq Y$. An algorithm $\mathcal{M} : X^n \to Y$ solves $\mathcal{T}$ with failure probability β if, for all $\mathcal{D} \in \mathscr{D}$,*

$$\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})] \geq 1 - \beta.$$

Statistical tasks capture essentially any setting where the algorithm is learning from i.i.d. data. We note that in many such tasks, there is an error parameter ε. This parameter is implicit in Definition 5 as we can define $\mathcal{T}(\mathcal{D})$ to only consist of outputs that are "ε-good." For example, if we aim to capture realizable PAC learning of a concept class $\mathcal{C}$ to error $1 - \varepsilon$, then $\mathscr{D}$ would consist of all distributions over labeled pairs $(\boldsymbol{x}, \boldsymbol{y})$ where $\boldsymbol{y} = f(\boldsymbol{x})$ for some single $f \in \mathcal{C}$ with probability 1. The valid responses $\mathcal{T}(\mathcal{D})$ would be any hypothesis h satisfying $\Pr_{\boldsymbol{x},\boldsymbol{y} \sim \mathcal{D}}[h(\boldsymbol{x}) = \boldsymbol{y}] \geq 1 - \varepsilon$. Our notion of equivalence will be agnostic to the particular statistical task an algorithm wishes to solve, and hence, automatically applies to all goals and error parameters.

▶ **Definition 6** (Equivalent algorithm). *We say an algorithm $\mathcal{M}' : X^m \to Y$ is (β, β')-equivalent to $\mathcal{M} : X^n \to Y$ if, any statistical task that $\mathcal{M}$ solves with failure probability β, $\mathcal{M}'$ solves with failure probability β'.*

[2] Our actual analysis case splits on the size of the domain, and gets a worse polynomial on very small domains.

2.1 Privacy measures and our axioms

To formalize Theorems 2–4 we define a series of axioms that any reasonable and useful *privacy measure* should satisfy.

▶ **Definition 7** (Privacy measure)**.** *A* privacy measure *is a mapping $\mathcal{P}$ from (possibly randomized) algorithms $\mathcal{M} : X^n \to Y$ to their level of privacy, parametrized as a number on $\mathbb{R}_{\geq 0}$. We adopt the convention that a lower values for $\mathcal{P}(\mathcal{M})$ indicate that $\mathcal{M}$ is more private. It will often be useful to succinctly say that $\mathcal{M}$ is $\mathcal{P}$-private if $\mathcal{P}(\mathcal{M}) \leq 1$.*

We remark upon a few basic properties about Definition 7. First, as is typical of previous definitions of privacy, a single privacy measure $\mathcal{P}$ must provide privacy levels for algorithms taking in samples of all sizes $n \in \mathbb{N}$. Later, our scaling axiom (Axiom 4) will enforce some amount of consistency between how $\mathcal{P}$ behaves on different sample sizes.

Second, Definition 7 is a single parameter definition of privacy, in contrast to the two-parameters of DP (Definition 1). This single parameter was a deliberate choice. A guiding philosophy in the development of our axioms was to not directly enforce specific meaning to the privacy value $\mathcal{P}(\mathcal{M})$, as we did not want our axioms to be biased by the meaning of ε and δ in DP. If we had a two (or more) parameter definition of privacy, we would need our axioms to somehow encode the distinction between those parameters, contradicting that guiding philosophy.

Furthermore, despite DP having two parameters, they are not of equal importance. Typical applications of DP simply set δ small enough to ignore and focus on ε. Indeed, following the intuition that one only needs δ "small enough," we show in Section 3.2 how to collapse Definition 1 into a single parameter in a way that respects all of our axioms.

2.1.1 Axioms any reasonable definition of privacy should satisfy

We now proceed to define our axioms, beginning with those that any "reasonable" definition of privacy should satisfy. The first axioms encodes some basic operations that should maintain privacy.

▶ **Axiom 1** (Preprocessing maintains privacy)**.** We say a privacy measure $\mathcal{P}$ satisfies the *preprocessing axiom* if the following is true.

1. **Reordering the input maintains privacy:** For any algorithm $\mathcal{M} : X^n \to Y$ and permutation $\pi : [n] \to [n]$, defining

   $$\mathcal{M} \circ \pi(S) \coloneqq \mathcal{M}(S_{\pi(1)}, \ldots, S_{\pi(n)}),$$

   we have that $\mathcal{P}(\mathcal{M} \circ \pi) \leq \mathcal{P}(\mathcal{M})$.
2. **Remapping the domain maintains privacy:** For any mapping $\sigma : X \to X$ and algorithm $\mathcal{M} : X^n \to Y$, defining

   $$\mathcal{M} \circ \sigma(S) \coloneqq \mathcal{M}(\sigma(S_1), \ldots, \sigma(S_n)),$$

   we have that $\mathcal{P}(\mathcal{M} \circ \sigma) \leq \mathcal{P}(\mathcal{M})$.

The first criteria, that reordering the input maintains privacy, says that under $\mathcal{P}$ it is equally bad to leak information about the i^{th} point and j^{th} point for any $i, j \in [n]$. The second criteria similarly says that it is equally bad to leak information about users x and x' for any $x, x' \in X$.

While both of these criteria are intuitively reasonable, we also provide more formal justification for their inclusion as axioms. In the full version, we show show that removing any one of our four axioms would allow for ill-behaved privacy measures, illustrating why these axioms are necessary (see Section 3.3 for a briefer overview). Since Axiom 1 has two criteria, we will furthermore show that removing either of them would similarly result in ill-behaved privacy measures, helping to justify why both are necessary.

Our second axiom requires private mechanisms to not reveal (essentially) the entire dataset. This is the only axiom that directly enforces that $\mathcal{P}$ measures some notion of privacy.

▶ **Definition 8** (Blatantly non-private). *A mechanism $\mathcal{M} : X^n \to Y$ is* blatantly non-private *if there is a "high-entropy" distribution $\mathcal{D}$ (formally $\mathcal{D}(x) \leq 1/(100n^2)$ for all $x \in X$) and adversary A mapping mechanism outputs $y \in Y$ to datasets $S' \in X^n$ for which*[3]

$$\mathop{\mathbb{E}}_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow A(\mathcal{M}(\boldsymbol{S}))}} \left[\sum_{x\in\boldsymbol{S}} \mathbb{1}[x \in \boldsymbol{S}']\right] \geq 0.9n.$$

The "high-entropy" requirement of Definition 8 is designed to ensure the adversary's task is not too easy. In particular, it means that if the adversary were not able to see the $\mathcal{M}$'s output, it would not even be able to guess a single point in $\boldsymbol{S}$. This stands in sharp contrast to the adversary being able to guess nearly all of $\boldsymbol{S}$ upon seeing $\mathcal{M}$'s output.

▶ **Axiom 2** (Prohibits blatant non-privacy). We say a privacy measure $\mathcal{P}$ satisfies the *prohibits blatant non-privacy* axiom if any $\mathcal{P}$-private algorithm is not blatantly non-private.

2.1.2 Strong-composition axioms

While the first two axioms were meant to be minimal requirements of any privacy definition to capture some reasonable notion of privacy, our next two axioms together formalize the notion of *strong composition*. As discussed earlier, the fact that the privacy costs of Definition 1 scale sublinearly with composition is crucial to the widespread adoption of differential privacy. Our next axiom encodes that the composition of ℓ many algorithms each of which have privacy level ε results in an algorithm with privacy level $\varepsilon' := \varepsilon \cdot \ell^c$. We will state this in the minimal form we need: In particular, we only need that the composed algorithm is $\mathcal{P}$-private whenever $\varepsilon' \leq 1$.

▶ **Axiom 3** (Strong composition). For $c < 1$, we say a privacy measure $\mathcal{P}$ satisfies *c-strong composition* if for any algorithms $\mathcal{M}^1, \ldots, \mathcal{M}^\ell : X^n \to Y$ all satisfying $\mathcal{P}(\mathcal{M}^i) \leq \varepsilon$ and

$$\varepsilon' := \tilde{O}(\varepsilon \cdot \ell^c) = O(\varepsilon \cdot \ell^c \cdot \mathrm{polylog}(n)),$$

if $\varepsilon' \leq 1$, then the composed algorithm $\mathcal{M}' : X^n \to Y^\ell$ that takes in a sample $S \in X^n$ and outputs the ℓ responses $(\mathcal{M}^1(S), \ldots, \mathcal{M}^\ell(S))$ is $\mathcal{P}$-private.

Interestingly, we are able to define Axiom 3 to be qualitatively weaker than the strong composition DP satisfies. DP satisfies *adaptive* strong composition, where the choice of $\mathcal{M}_i$ may depend adaptively on the outputs of $\mathcal{M}_1, \ldots, \mathcal{M}_{i-1}$. In contrast, Axiom 3 only requires strong composition to hold when $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$ are fixed in advance. Yet, we are still able to show that our axioms imply DP. This shows, in some sense, that non-adaptive strong composition is enough to derive adaptive strong composition.

[3] This constant of 0.9 could be replaced with any $c < 1$.

Axiom 3 on its own is not enough to enforce any reasonable notion of strong composition because it does not enforce any notion of scaling. For example, suppose we had some privacy measure $\mathcal{P}$ that only satisfied linear composition[4] (Axiom 3 with $c = 1$). Then, we could simply define a new privacy measure $\mathcal{P}'$ as $\mathcal{P}'(\mathcal{M}) \coloneqq \sqrt{\mathcal{P}(\mathcal{M})}$. This new measure would satisfy Axiom 3 with $c = 1/2$. Our last axiom rectifies this.

▶ **Axiom 4** (Linear scalability)**.** We say a privacy measure $\mathcal{P}$ satisfies *linear scaling*, if for some polynomial $p : \mathbb{R}^2 \to \mathbb{R}$, any $\mathcal{P}$-private algorithm $\mathcal{M} : X^n \to Y$, any failure probability $\beta > 0$, and any large enough $k \geq p(n, 1/\beta)$, there exists a $(\beta, \beta' \coloneqq O(\beta))$-equivalent algorithm $\mathcal{M}'$ taking in $m \coloneqq kn$ samples that satisfies $\mathcal{P}(\mathcal{M}') \leq O(1/k)$.

Roughly speaking, linear scalability says that the privacy level can be improved by a factor of $1/k$ by increasing the sample size by a factor of k. For example, one common way to amplify privacy is *subsampling*, meaning $\mathcal{M}'$ is the randomized algorithm which runs $\mathcal{M}$ on a uniform size-n subsample of its size-m input dataset. Indeed, for Definition 1, subsampling an (ε, δ)-DP algorithm by a factor of k leads to an $(\varepsilon/k, \delta/k)$-DP algorithm, though we will need a slightly more complicated amplification algorithm after we collapse ε and δ to a single parameter (see Lemma 7.1 of the full version [6]).

Axioms 3 and 4 are best viewed as together enforcing the following notion of strong composition. If the goal is to do a sequence of ℓ operations that each require a sample of size n to perform privately, then only need a single sample size of $n \cdot \ell^{1-\Omega(1)}$. That is, we require some non-trivial improvement over a strategy that, for example, uses n separate samples for each of the ℓ operations. We prefer this definition of strong composition in terms of the sample size required for ℓ many operations over explicit definitions that enforce a particular meaning to the value of $\mathcal{P}(\mathcal{M})$ in lieu of Axiom 4.

3 Technical Overview

3.1 Overview of Theorem 2: Our axioms imply DP

Given any privacy measure $\mathcal{P}$ satisfying our axioms and $\mathcal{P}$-private algorithm $\mathcal{M}$, we wish to construct an equivalent $\mathcal{M}'$ that is (ε, δ)-DP. To do so, we use the following intermediate notion of stability.

▶ **Definition 9** (TV-Stability, also called TV-indistinguishability by [20])**.** *The* TV-stability *of an algorithm $\mathcal{M} : X^n \to Y$ under distribution $\mathcal{D}$ is defined as*

$$\mathrm{stab}_{\mathrm{tv}}(\mathcal{M}, \mathcal{D}) \coloneqq \mathop{\mathbb{E}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \mathcal{D}^n}[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{S}), \mathcal{M}(\boldsymbol{S}'))].$$

We simply say $\mathcal{M}$ is ρ-TV-stable *if* $\mathrm{stab}_{\mathrm{tv}}(\mathcal{M}, \mathcal{D}) \leq \rho$ *for all distributions $\mathcal{D}$ over X.*

This definition is useful because (slight modifications) of the results of [7] allow us to convert any TV-stable algorithm into an equivalent DP algorithm (see Lemma 6.1 of the full version [6] for a formal statement of that conversion). Most of our effort goes into converting a $\mathcal{P}$-private algorithm into a TV-stable algorithm.

▶ **Theorem 10** (Our privacy axioms imply TV-stability)**.** *Let $\mathcal{P}$ be any privacy measure that satisfies Axioms 1–4 and $\mathcal{M} : X^n \to Y$ be any algorithm that is $\mathcal{P}$-private (that is, $\mathcal{P}(\mathcal{M}) \leq 1$.) For any constant $\rho > 0$ and $m \coloneqq \mathrm{poly}_\rho(n)$, there is a TV-stable algorithm $\mathcal{M}' : X^m \to Y$ that is equivalent to $\mathcal{M}$.*

[4] As in the case for the average-case variants of DP defined in [17, 28, 22]

To prove Theorem 10, we show, roughly speaking, that for any non-TV-stable algorithm $\mathcal{M} : X^m \to Y$, there exists algorithms $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$ for $\ell \approx m$ satisfying,

1. Each $\mathcal{M}_i$ can be formed by preprocessing $\mathcal{M}$, and therefore, by the Axiom 1 (preprocessing), should have the same privacy.
2. The composed algorithm $\mathcal{M}_{\text{comp}}$ that takes as input S and outputs the tuple $(\mathcal{M}_1(S), \ldots, \mathcal{M}_\ell(S))$ is blatantly non-private.

By Axiom 2 (prohibition of blatant non-privacy), we can conclude that $\mathcal{M}_{\text{comp}}$ is not $\mathcal{P}$-private. Then, Axiom 3 (strong composition) says that at least one $\mathcal{M}_i$ must satisfy $\mathcal{P}(\mathcal{M}_i) \geq \tilde{\Omega}(\ell^{-c})$. By Axiom 1 (preprocessing) this in fact means that $\mathcal{P}(\mathcal{M}) \geq \tilde{\Omega}(\ell^{-c}) = \tilde{\Omega}(m^{-c})$.

By contrapositive, this allows us to prove something just short of our goal: Any $\mathcal{M}$ satisfying sufficiently strong privacy, $\mathcal{P}(\mathcal{M}) \leq \tilde{O}(m^{-c})$, then $\mathcal{M}$ itself must be TV-stable.[5] In contrast, Theorem 10 only assume that $\mathcal{M}$ is $\mathcal{P}$-private. Here, we can exploit linear scalability: Using Axiom 4, we can convert any $\mathcal{M} : X^n \to Y$ that is $\mathcal{P}$-private to an $\mathcal{M}' : X^m \to Y$ satisfying $\mathcal{P}(\mathcal{M}') \leq (1/m^{-c})$ with only a polynomial increase in the sample size. This is the step where we crucially utilize the combined power of linear scalability and strong composition: Ultimately, we want to convert any $\mathcal{P}$-stable algorithm using n samples into one using $O(m)$ samples with the additional property that it can be composed m times and still be $\mathcal{P}$-stable. Axioms 3 and 4 together allow us to do this.

3.1.1 Exploiting TV-unstable algorithms

The key step in proving Theorem 10 is to show that if we compose $\approx m$ many preprocessed copies of a non-TV-stable algorithm $\mathcal{M} : X^m \to Y$, we will obtain a blatantly non-private algorithm. To prove this, we show a single random preprocessing reveals much information about the sample. It will be most convenient to state this lemma in terms of algorithms that take as input an unordered size-m set as input, and we will use $\binom{X}{m}$ to denote all such sets.

▶ **Lemma 11** (Key lemma, uniform permutations distinguish far samples). *For any $\mathcal{M} : \binom{X}{m} \to Y$ where $|X| \geq 2m$, define*

$$\rho \coloneqq \mathop{\mathbb{E}}_{\boldsymbol{S},\boldsymbol{S}' \sim \text{Unif}\left(\binom{X}{m}\right)} \left[d_{\text{TV}}(\mathcal{M}(\boldsymbol{S}), \mathcal{M}(\boldsymbol{S}')) \,\middle|\, |\boldsymbol{S} \cap \boldsymbol{S}'| = 0 \right]. \tag{1}$$

Then, for any $S, S' \in \binom{X}{m}$ and $\boldsymbol{\sigma} : X \to X$ a uniform permutation,

$$\mathbb{E}[d_{\text{TV}}(\mathcal{M} \circ \boldsymbol{\sigma}(S), \mathcal{M} \circ \boldsymbol{\sigma}(S'))] \geq \frac{\rho}{2} \cdot \text{dist}(S, S')/m,$$

where $\text{dist}(S, S') \coloneqq m - |S \cap S'|$ is the number of points S and S' differ on.

Since we start with a $\mathcal{M}$ that is not TV-stable, the quantity ρ in Equation (1) is promised to be somewhat large. Lemma 11 says that, if we draw just one $\boldsymbol{\sigma}$, the algorithm $\mathcal{M} \circ \boldsymbol{\sigma}$ provides roughly "$\Omega(1)$ bit" of useful information in distinguishing any S and S' that are somewhat far, satisfying $\text{dist}(S, S') \geq 0.01n$. Since the number of possible datasets S is $\binom{|X|}{m}$, it is possible to determine a dataset close to S by observing $\mathcal{M} \circ \sigma_1, \ldots, \mathcal{M} \circ \sigma_\ell$ for $\ell \coloneqq O(\log \binom{|X|}{m}) = O(m \log |X|)$. We furthermore show in the body of the paper how to reduce to the case where $|X| = O(m^2)$, in which case $\ell = O(m \log m)$ suffices.

The key step in proving Lemma 11 is constructing the following random walk.

[5] We note that there are some caveats to this statement: Briefly, it only holds for *symmetric* algorithms, those whose output does not depend on the order of its input, and assumes the domain is not too small. Both details are handled in the body.

▶ **Lemma 12** (Random walk to disjoint samples). *For any $S, S' \in \binom{X}{m}$, setting $d \coloneqq \operatorname{dist}(S, S')$ and $k \coloneqq \lceil m/d \rceil$, there exists random variables $\boldsymbol{T}^0, \dots, \boldsymbol{T}^k$ with the following properties:*

1. *For any $i \in [k]$ the marginal distribution of $(\boldsymbol{T}^{i-1}, \boldsymbol{T}^i)$ is equal to the distribution of $(\boldsymbol{\sigma}(S), \boldsymbol{\sigma}(S'))$ when $\boldsymbol{\sigma} : X \to X$ is a uniform permutation.*
2. *The marginal distribution of $(\boldsymbol{T}^0, \boldsymbol{T}^k)$ is equal to the distribution of $\boldsymbol{U}, \boldsymbol{U}' \sim \operatorname{Unif}\left(\binom{X}{m}\right)$ conditioned on $|\boldsymbol{U} \cap \boldsymbol{U}'| = 0$.*

The intuition behind Lemma 12 is that $\boldsymbol{T}^i$ can be formed by "rerandomizing" exactly d many of the elements in $\boldsymbol{T}^{i-1}$. As long as we have at least m/d steps, we can ensure all elements get rerandomized. The actual proof of Lemma 12 is a bit precise. In particular we need to use a non-Markovian walk (in that the distribution of $\boldsymbol{T}^i$ is not independent of $\boldsymbol{T}^1, \dots, \boldsymbol{T}^{i-2}$ conditioned on $\boldsymbol{T}^{i-1}$) for the following reasons:

1. In order to ensure all elements get rerandomized, the steps of the random walk cannot be independent. Instead, we enforce that the elements rerandomized in each step are different, while still ensuring that all the pairwise marginal $(\boldsymbol{T}^{i-1}, \boldsymbol{T}^i)$ have the right distribution.
2. When m/d is not exactly an integer some elements will be rerandomized twice. In this case, we need to ensure that no element accidentally gets rerandomized back into an element appearing in $\boldsymbol{T}^0$ as that would cause $\operatorname{dist}(\boldsymbol{T}^0, \boldsymbol{T}^k) < m$.

Nonetheless, we show with a careful construction that Lemma 12 holds.

3.2 Overview of Theorem 3: DP satisfies the axioms

Since Definition 1 has two parameters, ε and δ, we must collapse them to one parameter for our framework. We do this by defining,

$$\mathcal{P}_{\mathrm{DP}}(\mathcal{M}) \coloneqq \arg\min_{v} \Big\{ \mathcal{M} : X^n \to Y \text{ is } (\varepsilon = \Theta(v^{4/5}), \delta = \Theta(v^{8/5}/n^3)\text{-DP} \Big\}. \tag{2}$$

There is just one of many ways to collapse ε and δ into a single parameter in a way that respects our axioms. The exponents $4/5$ and $8/5$ could be replaced with α and 2α for any $\alpha \in (0.5, 1)$. Furthermore, the n^3 factor could be replaced with any n^β for $\beta > 3$. With this privacy measure, we can state the formal version of Theorem 3. We first state the formal version of Theorem 3.

▶ **Theorem 2** (DP implies our axioms, formal version). *The privacy measure*

$$\mathcal{P}_{\mathrm{DP}}(\mathcal{M}) \coloneqq \arg\min_{v} \Big\{ \mathcal{M} : X^n \to Y \ \textit{is} \ (\varepsilon = \Theta(v^{4/5}), \delta = \Theta(v^{8/5}/n^3)\textit{-DP} \Big\}$$

satisfies Axioms 1–4.

The reason we want δ to be much smaller than ε is because that's the regime in which differential privacy satisfies strong composition. The following well-known theorem shows that DP has strong composition.

▶ **Theorem** (DP-Strong-Composition, Theorem 3.20 in [13]). *For all $\varepsilon, \delta \geq 0$, if $\mathcal{M}^1, \dots, \mathcal{M}^\ell$ are (ε, δ)-differentially private, then the composed algorithm $(\mathcal{M}^1, \dots, \mathcal{M}^\ell)$ is (ε', δ')-differentially private where $\delta' \coloneqq 2\ell\delta$ and*

$$\varepsilon' \coloneqq \varepsilon\sqrt{\ell \ln(1/(\ell\delta))} + \ell\varepsilon(e^\varepsilon - 1).$$

Given that we have forced δ to be small, we cannot simply use subsampling [4] to ensure that $\mathcal{P}_{\mathrm{DP}}$ satisfies linear scalability, as subsampling's effect on δ is too mild. Instead, we use (a small modification of) a recent result of [7] to prove $\mathcal{P}_{\mathrm{DP}}$ satisfies linear scalability. See Lemma 7.1 of the full version [6] for this amplification procedure and the surrounding discussion for comparison to [7]'s result. Given the well-known strong-composition theorem for DP and this amplification procedure, showing that $\mathcal{P}_{\mathrm{DP}}$ satisfies all our axioms is straightforward.

3.3 Overview of Theorem 4: Necessity of our axioms

Here, we explain why all four of our axioms are necessary. For each axiom, we exhibit ill-behaved notions of privacy that would be allowed if we removed the axiom. In the case of Axiom 1, we even show this is true if only one of the two parts of it are removed, and in the case of Axiom 3, we will show it is true even if we replace strong composition with linear composition (i.e. setting $c = 1$). The proof of Theorem 4, given in the full version, will build on these ill-behaved privacy measures by showing that they allow algorithms solving statistical tasks that no differentially private algorithm can solve..

If we remove just the first part of Axiom 1, that reordering the input maintains privacy, then there is a privacy measure satisfying the remaining axioms, that we call $\mathcal{P}_{\mathrm{half}}$, which deems the algorithm $\mathcal{M} : X^n \to X^{\lfloor n/2 \rfloor}$ that outputs the first half of its dataset perfectly private, satisfying $\mathcal{P}_{\mathrm{half}}(\mathcal{M}) = 0$.

If we remove the second half of Axiom 1, that remapping the domain maintains privacy, then there is a privacy measure satisfying the remaining axioms, that we call $\mathcal{P}_{\mathrm{heavy}}$, that deems the following algorithm perfectly private: Let $\mathcal{M} : X^n \to X^n \cup \{\varnothing\}$ be the algorithm with the following behavior:

$$\mathcal{M}(S) = \begin{cases} S & \text{if there is some } x \text{ appearing at least } 0.6n \text{ times in } S. \\ \varnothing & \text{otherwise.} \end{cases}$$

Essentially, $\mathcal{M}$ is allowed to leak the entire dataset if there is any element appearing frequently enough. Despite this leakage, $\mathcal{P}_{\mathrm{heavy}}(\mathcal{M}) = 0$, indicating that $\mathcal{M}$ should have "perfect" privacy. We further observe that $\mathcal{P}_{\mathrm{heavy}}$ still satisfies that *permuting* the domain maintains privacy. This shows that we could not have replaced the arbitrary *mappings* $\sigma : X \to X$ in Axiom 1 with arbitrary *permutations* without allowing this ill-behaved notion of privacy.

If we remove Axiom 2 (prohibition of blatant non-privacy), then a privacy measure $\mathcal{P}_{\mathrm{all}}$ that deems *all* algorithms perfectly private, i.e. $\mathcal{P}_{\mathrm{all}}(\mathcal{M}) = 0$ for all $\mathcal{M}$, satisfies the remaining axioms.

If we relax strong composition to linear composition, i.e. allow $c = 1$ in Axiom 3, then there is a privacy measure, that we call $\mathcal{P}_{\mathrm{junta}}$ with the following behavior: The algorithm $\mathcal{M} : X^n \to X^k$ which outputs the first k points in its dataset satisfies $\mathcal{P}_{\mathrm{junta}}(\mathcal{M}) = \frac{k}{2n}$. For example, an algorithm which outputs the first half of its dataset is still $\mathcal{P}_{\mathrm{junta}}$-private.

If we remove Axiom 4 (linear scaling), then there is a rescaling, $\mathcal{P}_{\sqrt{\mathrm{junta}}}$, of the above privacy measure that satisfies the remaining axioms. The algorithm $\mathcal{M} : X^n \to X^k$ which outputs the first k points in its dataset is satisfies $\mathcal{P}_{\sqrt{\mathrm{junta}}}(\mathcal{M}) = \sqrt{\frac{k}{2n}}$. This still has essentially the same consequences as if we weakened Axiom 3. For example, we still have that the algorithm which outputs the first half of its dataset is $\mathcal{P}_{\sqrt{\mathrm{junta}}}$-private.

4 Discussion and Related Work

Computational efficiency. In Theorem 2, we guarantee that any sample efficient $\mathcal{P}$-private algorithm $\mathcal{M}$ can be transformed into an equivalent DP algorithm $\mathcal{M}'$ with approximately the same sample complexity. While our transformation is constructive, it does not necessarily preserve computational efficiency. Part of the reason is that Axiom 4 does not require the scaling to preserve computational efficiency, and we utilize a scaled version of $\mathcal{M}$ to construct $\mathcal{M}'$. This choice to allow for non-computationally efficient amplification is crucial to Theorem 3 as we utilize the following (computationally inefficient) procedure to prove that DP fits our axioms:

▶ **Theorem** (DP-Amplification, [7]). *For any $(\varepsilon = O(1), \delta = O(1/n^3))$-DP algorithm $\mathcal{M} : X^n \to Y$, there exists an equivalent (ε', δ')-DP algorithm $\mathcal{M}' : X^m \to Y$ using $m := 1/\varepsilon' \cdot \mathrm{poly}(n, \log 1/\delta')$ samples.*

We remark that there is a computationally efficient way to amplify an (ε, δ)-DP algorithm to $(\varepsilon/k, \delta/k)$-DP at the cost of a $O(k)$ increase in the sample size, via subsampling [4]. While subsampling's linear amplification of ε is as good as DP-Amplification, the linear amplification of δ is not sufficient for our purposes, and so we need to utilize the computationally inefficient amplification of DP-Amplification.

As far as we are aware, despite it being of independent interest, it is unknown whether a computationally efficient analogue of DP-Amplification exists. More broadly, we leave open the possibility that it is possible to obtain a computationally efficient analogue of our results, possibly by adjusting the axioms appropriately.

Other formalizations of differentially privacy. We focused on the well-studied (ε, δ)-DP formulation of Definition 1 (often called *approximate DP*). One popular alternative, *pure* DP, is equivalent to Definition 1 where δ is fixed to be 0. We did not focus on pure-DP because it does not satisfy strong composition, which makes it more difficult to utilize in practice and also that it does not fit our axioms. That said, it would be interesting to come up with an alternative set of axioms that characterize pure DP in the same sense as our axioms characterize approximate DP. One tempting solution is to simply remove our strong composition axiom (Axiom 3). However, as we show, removing Axiom 3 allows for a degenerate privacy measure which is much weaker than pure DP, so a different approach is needed (see Section 8 in the full version [6]).

A second popular generalization of approximate DP follows from the simple observation that algorithms are not (ε, δ)-DP for a single fixed choice of ε and δ. Rather, for any algorithm $\mathcal{M}$, there is an entire "curve" $\varepsilon : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ for which $\mathcal{M}$ is $(\varepsilon(\delta), \delta)$-DP for all choices of $\delta > 0$. There are a variety of formulations of DP that bound the behavior of this curve (e.g. through bounding appropriately defined "moments") such as Réyni DP and concentrated DP [23, 14, 8]. These variations are popular precisely because they allow for easy (and often strong) composition, and in appropriate parameter regimes, also are amplified by subsampling. We refer the reader to [25] for an excellent overview.

Given this, it's natural to expect these variants would play nicely with our axioms. Indeed, we show in that the privacy measure that assigns to $\mathcal{M}$ the smallest privacy value v s.t. $\mathcal{M}$ is $(2, \sqrt{v})$-Réyni DP respects all our axioms with an even more straightforward analysis than the proof of Theorem 3 (For more details we refer the reader to appendix of the full version). In the other direction, we show a variant of Theorem 2, that our axioms imply Réyni DP. One distinction between that statement (This corresponds to Theorem 7 in the full version) and Theorem 2 is that the Réyni DP algorithm has a sample size that depends on $\log \log |Y|$, which we also show is necessary in the appendix of the full version.

Related Work. Perhaps most in the spirit of our results is recent work on reproducibility [18], and in particular the followup paper of Bun et al. [7] (see also [20]). That work examines the broader context of *algorithmic stability*, which are various ways of formalizing that an algorithms output does not depend too much on its input. They show that some of these measures of stability, replicability, max-information, and perfect generalization, are equivalent to differential privacy using the same formalization of equivalence as us. Measures of algorithmic stability and privacy share many of the same basic properties. In some sense, the only distinction between algorithmic stability and privacy is simply that measures of algorithmic stability were designed for applications other than privacy. Indeed, one could just as easily view our axioms as desirable properties for any measure of algorithmic stability. From this perspective, our work is a natural evolution of [7] as we show all measures of stability satisfying our axioms are equivalent to privacy. We also utilize some of their techniques to prove our results.

More broadly, there have been several works formalizing axioms that any "reasonable" definition of privacy should satisfy. Often this includes an axiom or assumption that privacy should be some measure of distance between the distributions $\mathcal{M}(S)$ and $\mathcal{M}(S')$ for worst-case S and S' (as in Definition 1). This includes [21, 26], which both investigate what measures of distance satisfy other reasonable axioms. Also in this spirit is the central limit theorem of [10]. Roughly speaking, it says that if we consider only privacy definitions based on some distance between $\mathcal{M}(S)$ and $\mathcal{M}(S')$, in the limit of many compositions, we may as well define "Gaussian differential privacy." The key distinction between all of these works and ours is that we aim to justify why the most successful privacy definitions are measures of distance between $\mathcal{M}(S)$ and $\mathcal{M}(S')$ for worst-case S and S', whereas previous works take that as an assumption or axiom.

References

1 Differential privacy: Technical overview. Apple Inc. White paper; documents Apple's local DP deployment and budgets. Accessed Aug 19, 2025. URL: `https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf`.

2 Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016. `doi:10.1145/2976749.2978318`.

3 John M Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, et al. The 2020 census disclosure avoidance system topdown algorithm. *Harvard Data Science Review*, 2, 2022.

4 Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. *Advances in neural information processing systems*, 31, 2018.

5 Raef Bassily and Yoav Freund. Typical stability. *arXiv preprint arXiv:1604.03336*, 2016.

6 Guy Blanc, William Pires, and Toniann Pitassi. Differential privacy from axioms, 2025. `arXiv:2511.21876`.

7 Mark Bun, Marco Gaboardi, Max Hopkins, Russell Impagliazzo, Rex Lei, Toniann Pitassi, Satchit Sivakumar, and Jessica Sorrell. Stability is stable: Connections between replicability, privacy, and adaptive generalization. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 520–527, New York, NY, USA, 2023. Association for Computing Machinery. `doi:10.1145/3564246.3585246`.

8 Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of cryptography conference*, pages 635–658. Springer, 2016. doi:10.1007/978-3-662-53641-4_24.

9 Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.

10 Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2022.

11 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 486–503. Springer, 2006. doi:10.1007/11761679_29.

12 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006. doi:10.1007/11681878_14.

13 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, 2014. doi:10.1561/0400000042.

14 Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016. arXiv:1603.01887.

15 Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st annual symposium on foundations of computer science*, pages 51–60. IEEE, 2010. doi:10.1109/FOCS.2010.12.

16 Vitaly Feldman. A general characterization of the statistical query complexity. In *Conference on learning theory*, pages 785–830. PMLR, 2017. URL: http://proceedings.mlr.press/v65/feldman17c.html.

17 Robert Hall, Larry Wasserman, and Alessandro Rinaldo. Random differential privacy. *Journal of Privacy and Confidentiality*, 4(2), 2013. doi:10.29012/JPC.V4I2.621.

18 Russell Impagliazzo, Rex Lei, Toniann Pitassi, and Jessica Sorrell. Reproducibility in learning. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, pages 818–831, 2022. doi:10.1145/3519935.3519973.

19 Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR, 2015. URL: http://proceedings.mlr.press/v37/kairouz15.html.

20 Alkis Kalavasis, Amin Karbasi, Shay Moran, and Grigoris Velegkas. Statistical indistinguishability of learning algorithms. In *International Conference on Machine Learning*, pages 15586–15622. PMLR, 2023. URL: https://proceedings.mlr.press/v202/kalavasis23a.html.

21 Daniel Kifer and Bing-Rong Lin. Towards an axiomatization of statistical privacy and utility. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 147–158, 2010. doi:10.1145/1807085.1807106.

22 Ao Liu, Yu-Xiang Wang, and Lirong Xia. Smoothed differential privacy. *Transactions on Machine Learning Research*, 2023.

23 Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017. doi:10.1109/CSF.2017.11.

24 Joseph P Near, Xi He, et al. Differential privacy for databases. *Foundations and Trends® in Databases*, 11(2):109–225, 2021. doi:10.1561/1900000066.

25 Thomas Steinke. Composition of differential privacy & privacy amplification by subsampling. *arXiv preprint arXiv:2210.00597*, 2022. doi:10.48550/arXiv.2210.00597.

26 Weijie J Su. A statistical viewpoint on differential privacy: Hypothesis testing, representation, and blackwell's theorem. *Annual Review of Statistics and Its Application*, 12, 2024.

27 Salil Vadhan. The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, pages 347–450. Springer, 2017. doi:10.1007/978-3-319-57048-8_7.

28 Yu-Xiang Wang, Jing Lei, and Stephen E Fienberg. On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms. In *International Conference on Privacy in Statistical Databases*, pages 121–134. Springer, 2016. doi:10.1007/978-3-319-45381-1_10.

29 Zheng Xu, Yanxiang Zhang, Galen Andrew, Christopher A Choquette-Choo, Peter Kairouz, H Brendan McMahan, Jesse Rosenstock, and Yuanbo Zhang. Federated learning of gboard language models with differential privacy. *arXiv preprint arXiv:2305.18465*, 2023. doi:10.48550/arXiv.2305.18465.