Differential privacy from axioms

Guy Blanc
Stanford University

William Pires
Columbia University

Toniann Pitassi
Columbia University

December 1, 2025

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 [HWR13, WLF16, BF16, LWX23].

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.

Contents

1   Introduction

Differential privacy has won. It is the de facto formalization of privacy both in theory (see, e.g., the textbooks [DR14, Vad17, NH⁺21]) and in practice (see, e.g., its use in the U.S. Census [AACM⁺22] and by various technology companies [App, XZA⁺23, DKY17]).

Definition 1 ((ε, δ)-Differential Privacy, [DMNS06, DKM⁺06])**.** *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 $Y' \subseteq Y$,*

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

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 [DR14]. 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 [DRV10, KOV15]. This allows for efficient and simple construction of DP algorithms for a variety of tasks (see e.g. [ACG⁺16] 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 [HWR13, WLF16, BF16, LWX23].

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 $\delta < 1$.

---

[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.

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 1 (Our axioms imply differential privacy). *Let $\mathcal{P}$ be any privacy measure that satisfies Axioms 1 to 4 and $\mathcal{M} : X^n \to Y$ be any algorithm that is $\mathcal{P}$-private. For any $\varepsilon, \delta > 0$ and $m := \mathrm{poly}(n, 1/\varepsilon, \log(1/\delta))$ there is an (ε,δ)-DP algorithm $\mathcal{M}' : X^m \to Y$ that is equivalent (as in Definition 3) to $\mathcal{M}$.*

　　The exact polynomial in Theorem 1 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 1 would be $m \approx n^2$, provided the domain X is not too small.[2] We refer the reader to Section 6 for the full version of Theorem 1. We note that given the known equivalences between DP, replicability, and various notions of stability, Theorem 1 shows that these other notions are also implied by our axioms.

　　Second, we show that differential privacy satisfies our axioms.

Theorem 2 (Informal). *Approximate differential privacy (Definition 1) satisfies Axioms 1 to 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. To keep this overview succinct, we defer an in-depth discussion of what those nonsensical privacy measures are to Section 8 (with a briefer overview in Section 3.3). We do have the following simple implication of those results.

Theorem 3 (Minimality of our axioms). *Theorem 1 does not hold if $\mathcal{P}$ is allowed to not satisfy any one of Axioms 1 to 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. After a brief preliminaries (Section 5), the remaining sections (Sections 6 to 8) give formal proofs of Theorems 1 to 3 respectively.

2　Our Framework

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

Definition 2 (Statistical task, [Fel17, BGH$^+$23]). *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 $\mathcal{D} \in \mathscr{D}$,*

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

---

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

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 2 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 3 (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'.*

2.1 Privacy measures and our axioms

To formalize Theorems 1 to 3 we define a series of axioms that any reasonable and useful *privacy measure* should satisfy.

Definition 4 (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 4. 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 4 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

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

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

   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 Section 8, we will 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 5 (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 5 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 $\varepsilon' \leq 1$.

---

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

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.

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 $\mathcal{P}'(\mathcal{M}) := \sqrt{\mathcal{P}(\mathcal{M})}$. 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 $(\beta, \beta' := O(\beta))$-equivalent algorithm $\mathcal{M}'$ taking in $m := kn$ 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).

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 1: 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.

---

[4]As in the case for the average-case variants of DP defined in [HWR13, WLF16, LWX23]

Definition 6 (TV-Stability, also called TV-indistinguishability by [KKMV23]))**.** *The* TV-stability *of an algorithm* $\mathcal{M} : X^n \to Y$ *under distribution* $\mathcal{D}$ *is defined as*

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

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 [BGH⁺23] allow us to convert any TV-stable algorithm into an equivalent DP algorithm (see Lemma 6.1 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 4 (Our privacy axioms imply TV-stability)**.** *Let* $\mathcal{P}$ *be any privacy measure that satisfies Axioms 1 to 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* $\mathcal{M}$.

To prove Theorem 4, 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 $\mathcal{M}_{\mathrm{comp}}$ 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 $\mathcal{M}_{\mathrm{comp}}$ 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 4 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 4 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.

---

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

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

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

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

where $\mathrm{dist}(S, S') := m - |S \cap S'|$ *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 3.1 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 $\mathrm{dist}(S, S') \geq 0.01n$. Since the number of possible datasets S is $\binom{|X|}{m}$, it is possible to determine a dataset close to S by observing $\mathcal{M} \circ \sigma_1, \dots, \mathcal{M} \circ \sigma_\ell$ for $\ell := O(\log \binom{|X|}{m}) = O(m \log |X|)$. 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 3.1 is constructing the following random walk.

Lemma 3.2 (Random walk to disjoint samples)**.** *For any* $S, S' \in \binom{X}{m}$*, setting* $d := \mathrm{dist}(S, S')$ *and* $k := \lceil m/d \rceil$*, 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* $\boldsymbol{U}, \boldsymbol{U}' \sim \mathrm{Unif}(\binom{X}{m})$ *conditioned on* $|\boldsymbol{U} \cap \boldsymbol{U}'| = 0$*.*

The intuition behind Lemma 3.2 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 3.2 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 $\mathrm{dist}(\boldsymbol{T}^0, \boldsymbol{T}^k) < m$.

Nonetheless, we show with a careful construction that Lemma 3.2 holds.

3.2 Overview of Theorem 2: 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,

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

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 2. We first state the formal version of Theorem 2.

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

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

satisfies Axioms 1 to 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 [DR14])**.**
For all $\varepsilon, \delta \geq 0$, if $\mathcal{M}^1, \ldots, \mathcal{M}^\ell$ are (ε, δ)-differentially private, then the composed algorithm $(\mathcal{M}^1, \ldots, \mathcal{M}^\ell)$ 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 [BBG18] 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 [BGH+23] to prove $\mathcal{P}_{\mathrm{DP}}$ satisfies linear scalability. See Lemma 7.1 for this amplification procedure and the surrounding discussion for comparison to [BGH+23]'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 3: 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 3 will build on these ill-behaved privacy measures by showing that they allow algorithms solving statistical tasks that no differentially private algorithm can solve. (see Section 8 for details).

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 ($\mathcal{P}_{\mathrm{half}}$ from Definition 16) 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 ($\mathcal{P}_{\text{heavy}}$ in Definition 17) 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, $\mathcal{P}_{\text{heavy}}(\mathcal{M}) = 0$, indicating that $\mathcal{M}$ should have "perfect" privacy. We further observe (see Remark 1) that $\mathcal{P}_{\text{heavy}}$ 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 $\mathcal{P}_{\text{all}}$ that deems *all* algorithms perfectly private, i.e. $\mathcal{P}_{\text{all}}(\mathcal{M}) = 0$ 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, $\mathcal{P}_{\text{junta}}$ (see Definition 20) with the following behavior: The algorithm $\mathcal{M} : X^n \to X^k$ which outputs the first k points in its dataset satisfies $\mathcal{P}_{\text{junta}}(\mathcal{M}) = \frac{k}{2n}$. For example, an algorithm which outputs the first half of its dataset is still $\mathcal{P}_{\text{junta}}$-private.

If we remove Axiom 4 (linear scaling), then there is a rescaling, $\mathcal{P}_{\sqrt{\text{junta}}}$, 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 $\mathcal{P}_{\sqrt{\text{junta}}}(\mathcal{M}) = \sqrt{\frac{k}{2n}}$. 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 $\mathcal{P}_{\sqrt{\text{junta}}}$-private.

4 Discussion and Related Work

Computational efficiency. In Theorem 1, 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 2 as we utilize the following (computationally inefficient) procedure to prove that DP fits our axioms:

Theorem (DP-Amplification, [BGH+23].) *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* $m := 1/\varepsilon' \cdot \text{poly}(n, \log 1/\delta')$ *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 [BBG18]. 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 in Section 8, removing Axiom 3 allows for a degenerate privacy measure which is much weaker than pure DP, so a different approach is needed.

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 [Mir17, DR16, BS16]. 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 [Ste22] for an excellent overview.

Given this, it's natural to expect these variants would play nicely with our axioms. Indeed, we show in Section B 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 2 (see Section B.2). In the other direction, we show a variant of Theorem 1, that our axioms imply Réyni DP. One distinction between that statement (Theorem 7) and Theorem 1 is that the Réyni DP algorithm has a sample size that depends on $\log \log |Y|$, which we show is necessary in Lemma B.9.

Related Work. Perhaps most in the spirit of our results is recent work on reproducibility [ILPS22], and in particular the followup paper of Bun et al. [BGH+23] (see also [KKMV23]). 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 [BGH+23] 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 [KL10, Su24], which both investigate what measures of distance satisfy other reasonable axioms. Also in this spirit is the central limit theorem of [DRS22]. 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.

5 Preliminaries

Throughout this work, we will assume all domains are finite. For a domain X, we denote,

1. All ordered tuples of n elements as X^n.

2. All ordered tuples of n distinct elements as $X^{(n)}$.

3. All unordered sets of n distinct elements as $\binom{X}{n}$.

4. All permutations $\sigma : X \to X$ as $\mathfrak{S}(X)$.

For a natural number n, we use $[n]$ as shorthand for $\{1, \dots, n\}$. We use **boldface** letters, e.g. $\boldsymbol{x}, \boldsymbol{S}$. For brevity, we will sometimes use $\boldsymbol{x} \sim S$ as shorthand for $\boldsymbol{x} \sim \mathrm{Unif}(S)$. For example, $\boldsymbol{\sigma} \sim \mathfrak{S}(X)$ indicates that $\boldsymbol{\sigma}$ is a uniform permutation mapping X to itself.

It will sometimes be useful to convert any algorithm into its symmetric counterpart.

Definition 7 (Symmetrization of an algorithm)**.** *For a randomized algorithm $\mathcal{M} : X^n \to Y$, its symmetrization, which we denote $\widetilde{\mathcal{M}} : X^n \to Y$, is defined as follows. On input $S \in X^n$, $\widetilde{\mathcal{M}}(S)$ first draws a uniform permutation $\boldsymbol{\pi} \sim \mathfrak{S}([n])$ and then outputs $\mathcal{M}(S_{\boldsymbol{\pi}(1)}, \dots, S_{\boldsymbol{\pi}(n)})$.*

In particular, it's easy to see that the symmetrized algorithm $\widetilde{\mathcal{M}}$ doesn't depend on the order of its input. As such we will often abuse notation and view $\widetilde{\mathcal{M}}$ as both taking ordered sets and unordered sets as input. We also observe that for statistical tasks, permuting the input uniformly at random doesn't impact the correctness of the algorithm.

Fact 5.1. *Let $\mathcal{M} : X^n \to Y$ be an algorithm with symmetrized version $\widetilde{\mathcal{M}}$. For any $1 \geq \beta \geq 0$, we have $\widetilde{\mathcal{M}}$ is (β, β)-equivalent to $\mathcal{M}$.*

Proof. Note that for any set $S \in X^n$ and permutation $\pi : [n] \to [n]$ and distribution $\mathcal{D}$ over X we have: $\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\boldsymbol{S} = S] = \Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\boldsymbol{S} = \pi(S)]$. So, let $\mathcal{T}$ be a statistical task that $\mathcal{M}$ solves with failure probability at most β, recalling that $\mathcal{T}(\mathcal{D})$ is the set of correct answers to the task under distribution $\mathcal{D}$, we have:

$$
\begin{aligned}
\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\widetilde{\mathcal{M}}(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})] &= \Pr_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{\pi} \sim \mathfrak{S}(n)}}[\mathcal{M}(\boldsymbol{\pi}(\boldsymbol{S})) \in \mathcal{T}(\mathcal{D})] \\
&= \Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})] \\
&\geq 1 - \beta.
\end{aligned}
$$

So $\widetilde{\mathcal{M}}$ also solves task $\mathcal{T}$ with failure probability at most β. $\qquad\square$

We also recall the definition of TV-distance.

Definition 8. *Let $\mathcal{D}, \mathcal{D}'$ be two distributions over Y. The TV-distance between $\mathcal{D}$ and $\mathcal{D}'$ is defined as:*

$$d_{\mathrm{TV}}(\mathcal{D}, \mathcal{D}') := \max_{Y' \subseteq Y} \Big| \Pr_{\boldsymbol{y} \sim \mathcal{D}}[\boldsymbol{y} \in Y] - \Pr_{\boldsymbol{y} \sim \mathcal{D}'}[\boldsymbol{y} \in Y] \Big|.$$

It will sometimes be convenient to work with the following equivalent definition:

$$d_{\mathrm{TV}}(\mathcal{D}, \mathcal{D}') := \frac{1}{2} \|\mathcal{D} - \mathcal{D}'\|_1 = \frac{1}{2} \sum_{y \in Y} |\mathcal{D}(y) - \mathcal{D}'(y)|.$$

Finally, we recall the postprocessing property of differential privacy.

Proposition 5.2 (Postprocessing, see Proposition 2.1 in [DR14]). *Let $\mathcal{M} : X^n \to Y$ and $\mathcal{A} : Y \to Z$ be algorithms. If $\mathcal{M}$ is (ε, δ)-differentially private then the composed algorithm $\mathcal{A} \circ \mathcal{M} : X^n \to Z$ is (ε, δ)-differentially private.*

Definition 9 (Distance between two sets). *Let $S, S' \in X^n$ (or alternatively $S, S' \in \binom{X}{n}$), we define $\mathrm{dist}(S, S') := \sum_{i \in [n]} \mathbb{1}[S_i \notin S']$.*

6 Proof of Theorem 1: Our axioms imply DP

In this section, we prove that for any privacy measure $\mathcal{P}$ satisfying our axioms, any $\mathcal{P}$-private algorithm can be converted to a differential private one with only a polynomial increase in the sample size. We begin with a formal version of Theorem 1.

Theorem 1 (Our axioms imply DP, formal version). *Let $\mathcal{P}$ be any privacy measure satisfying Axioms 1 to 4 and $\mathcal{M} : X^n \to Y$ be any $\mathcal{P}$-private algorithm. For any $\varepsilon, \delta, \beta > 0$ and, c the constant in Axiom 3 and p the polynomial in Axiom 4, define*

$$m' := \tilde{O}\left(\frac{r^2 \cdot n^2 \cdot \log(1/\delta)}{\beta^2 \cdot \varepsilon} \right) \text{ where } r = \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}} \right).$$

Then there is an (ε, δ)-DP algorithm $\mathcal{M}'$ using m' samples that is $(\beta, \beta' := O(\beta))$-equivalent to $\mathcal{M}$.

This conversion relies on the two following lemmas. The first allows us to go from our axioms to TV-Stability. While the later one allows us to go from TV-Stability to differential privacy.

Theorem 4 (Our privacy axioms imply TV-Stability, formal version). *Let $\rho > 0$ be a constant, $\mathcal{P}$ be any privacy measure satisfying Axioms 1 to 4 and $\mathcal{M} : X^n \to Y$ be any $\mathcal{P}$-private algorithm. Let be c the constant in Axiom 3 and p the polynomial in Axiom 4, define*

$$m' := \tilde{O}\left(\frac{r^2 \cdot n^2}{\beta^2} \right) \text{ where } r = \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}} \right).$$

Then there is a ρ-TV stable $\mathcal{M}'$ using m' samples that is $(\beta, \beta' := O(\beta))$-equivalent to $\mathcal{M}$.

Lemma 6.1. *There is a universal constant $1 > \rho^\star > 0$ such that if $\mathcal{M} : X^n \to Y$ is an $\rho^\star$-TV-Stable algorithm, then there exists a $(\beta, 5\beta)$-equivalent algorithm $\mathcal{M}' : X^m \to Y$ which is (ε, δ)-differentially private using*

$$m = n \cdot O\left(\log(1/\beta) \cdot \frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon} \right) \text{ samples.}$$

With the above lemmas in mind, the proof of Theorem 1 follows easily.

Proof of Theorem 1. We first apply Theorem 4 to $\mathcal{M}$ to get a (β, β')-equivalent, $\beta' = O(\beta)$, and $\rho^\star$-TV-Stable algorithm $\mathcal{M}' : X^{m'} \to Y$ algorithm and using

$$m' = \tilde{O}(m^2 \cdot n^2 / \beta^2) \text{ where } m = n \cdot \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}}\right) \text{ samples.}$$

By applying Lemma 6.1 to $\mathcal{M}'$, we have that there exists a $(\beta', 5\beta')$-equivalent algorithm $\mathcal{M}^* : X^{m^*} \to Y$ which is (ε, δ)-differentially private and using

$$m^* = m' \cdot O\left(\log(1/\beta) \frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon}\right) = \tilde{O}\left(\frac{m^2 \cdot n^2 \cdot \log(1/\delta)}{\beta^2 \cdot \varepsilon}\right) \text{ samples.}$$

In particular, $\mathcal{M}$ and $\mathcal{M}'$ are $(\beta, O(\beta))$-equivalent. □

We prove Lemma 6.1 in Section A. In particular Lemma 6.1 follows almost identically from the work of [BGH+23] who gave a transformation from perfect generalization (which is related to our notion of TV-Stability) to differential privacy. The main difference is that the transformation of [BGH+23] makes the error probability of the algorithm go from β to $O(\log(1/\beta)\beta)$. While our error only grows to $O(\beta)$, this also comes at the cost of a worse dependency on $\log(1/\beta)$ in sample complexity compared to the result of [BGH+23].

The rest of this section is dedicated to the proof of Theorem 4.

Useful notation: Note that throughout this section, we often use γ to denote a pair of (σ, π) where $\sigma \in \mathfrak{S}(X)$ and $\pi \in \mathfrak{S}([n])$. For ease of notation, we will use $\boldsymbol{\gamma} \sim \mathfrak{S}(X, n)$ to denote drawing such a pair uniformly at random. Finally, given $\mathcal{M} : X^n \to Y$, we denote by $\mathcal{M} \circ \gamma$, the algorithm such that $\mathcal{M} \circ \gamma(S) = \mathcal{M} \circ \sigma \circ \pi(S)$ (i.e. the algorithm reorders the elements according to π and then remaps them according to σ and then runs $\mathcal{M}$).

6.1 A first step toward Theorem 4

We begin with a simplified version of Theorem 4. While Theorem 4 guarantees TV-Stability on *all distributions*, this simplified setting will only guarantee we obtain an algorithm that is TV-stable on uniform distributions with small support. It will also assume $\mathcal{M}$ satisfies a stronger privacy guarantee than in Theorem 4. We will do away with these assumptions in Section 6.2.

Lemma 6.2 (TV-Stability in a simplified setting). *Fix any constant $\rho > 0$. For any domain $X^\star$ of size at least $\frac{100}{\rho} \cdot n^2$, privacy measure $\mathcal{P}$ satisfying Axioms 1 to 3, and algorithm $\mathcal{M} : (X^\star)^n \to Y$. There exists constant $\alpha > 0$ and $t > 0$ such that if*

$$\mathcal{P}(\mathcal{M}) \leq \frac{1}{\alpha \cdot n^c \log^t(n)},$$

where c is the constant in Axiom 3, and t depends on the polylog factors in Axiom 3, then the "symmetrized" version of $\mathcal{M}$ (See Definition 7) denoted $\widetilde{\mathcal{M}}$

satisfies $\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}}, \mathcal{D} := \mathrm{Unif}(X))) \leq \rho$ *for every $X \subseteq X^\star$ of size $\lceil \frac{100}{\rho} \rceil \cdot n^2$.*

The domain size of $|X| \geq \Omega(n^2)$ ensures that $\boldsymbol{S} \sim \mathrm{Unif}(X)^n$ is very likely to take on n unique values. Furthermore, we work with the symmetrized version of $\mathcal{M}$ because the output behavior of $\widetilde{\mathcal{M}}(S)$ does not depend on the order of elements in S. The proof of Lemma 6.2 proceeds by contradiction: We show that if $\widetilde{\mathcal{M}}$ is *not* TV-stable under $\mathrm{Unif}(X)$, we can compose many "preprocessed" copies of $\widetilde{\mathcal{M}}$ to create an algorithm $\mathcal{M}'$ that should be $\mathcal{P}$-private according to Axioms 1 and 3. However, using Lemma 3.1 we will show, in the following subsections, the existence of an adversary which given the output of $\mathcal{M}'(\boldsymbol{S})$ for $\boldsymbol{S} \sim \mathrm{Unif}(X)^n$ can guess more than 90% of the dataset. In particular, this implies $\mathcal{M}'$ is blatantly non-private, violating Axiom 2.

6.1.1 Exploiting Lemma 3.1 and hypothesis selection

Throughout this section we fix an algorithm $\mathcal{M} : X^n \to Y$, where the reader should think that $|X| = \Theta(n^2)$. We first show how to reduce the adversary's task to classical *hypothesis selection*.

Fact 6.3 (Hypothesis selection, Ch. 6 of [DL01] or [ABS24]). *For any distributions $\mathcal{D}_1, \ldots, \mathcal{D}_M$ there is an algorithm $\mathcal{H}$ with the following guarantee: Given*

$$\ell := O\left(\frac{\log(M/\delta)}{\varepsilon^2}\right)$$

i.i.d. samples from some unknown $\mathcal{D}_i$, $\mathcal{H}$ outputs a j satisfying $d_{\mathrm{TV}}(\mathcal{D}_i, \mathcal{D}_j) \leq \varepsilon$ with probability at least $1 - \delta$.

Having fixed the algorithm $\mathcal{M}$, we will be interested in the following set of distributions.

Definition 10. *For $S \in X^{(n)}$, we define the distribution $\mathcal{D}_S$ has the distribution $(\boldsymbol{\gamma}, \mathcal{M} \circ \boldsymbol{\gamma}(S))$ for $\boldsymbol{\gamma} \sim \mathfrak{S}(X, n)$. We denote by $\mathcal{F}$ the set of all such distribution $\mathcal{D}_S$.*

As sketched in the introduction, we will use Lemma 3.1 on $\widetilde{\mathcal{M}}$ to argue that for any $S, S' \in X^{(n)}$, if S and S' differ on $\Omega(n)$ coordinates, the distributions $\mathcal{D}_S$ and $\mathcal{D}_{S'}$ will be far. We present this idea formally in the next lemma. Since the number of possible datasets S is at most $|X|^n$, the adversary, using hypothesis selection, can find a dataset close to an unknown set $S \in X^{(n)}$ by observing $O(n \log |X|)$ many samples from $\mathcal{D}_S$.

Lemma 6.4. *Fix $\mathcal{M} : X^n \to Y$ with $|X| \geq 2n$, if for its "symmetrized" version $\widetilde{\mathcal{M}}$ we have:*

$$\rho := \mathop{\mathbb{E}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \binom{X}{n}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}}(\boldsymbol{S}), \widetilde{\mathcal{M}}(\boldsymbol{S}')) \,\middle|\, |\boldsymbol{S} \cap \boldsymbol{S}'| = 0 \right].$$

Then, for any $S, S' \in X^{(n)}$ we have

$$d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{S'}) \geq \frac{1}{2} \cdot \mathrm{dist}(S, S')/n \cdot \rho.$$

Proof. Let $S, S' \in X^{(n)}$ and let $d = \mathrm{dist}(S, S')$. We let $\widetilde{S}, \widetilde{S}'$ denote the unordered version of S, S'. Note that $d = \mathrm{dist}(\widetilde{S}, \widetilde{S}')$. By viewing $\widetilde{\mathcal{M}}$ as taking unordered sets as input and using Lemma 3.1, we have:

$$\frac{\rho \cdot d}{2n} \leq \mathop{\mathbb{E}}_{\boldsymbol{\sigma}} \left[d_{\mathrm{TV}}(\widetilde{M} \circ \boldsymbol{\sigma}(\widetilde{S}), \widetilde{M} \circ \boldsymbol{\sigma}(\widetilde{S}')) \right]$$

$$= \frac{1}{2} \mathop{\mathbb{E}}_{\boldsymbol{\sigma}} \left[\|\widetilde{M} \circ \boldsymbol{\sigma}(\widetilde{S}) - \widetilde{M} \circ \boldsymbol{\sigma}(\widetilde{S}')\|_1 \right].$$

By the definition we have $\widetilde{\mathcal{M}}(S) = \sum_\pi \Pr[\pi]\mathcal{M}(S)$, hence:

$$\frac{\rho \cdot d}{n} = \mathop{\mathbb{E}}_{\boldsymbol{\sigma}}\left[\|\sum_\pi \Pr[\pi]\mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S) - \sum_\pi \Pr[\pi]\mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S')\|_1\right]$$

$$= \mathop{\mathbb{E}}_{\boldsymbol{\sigma}}\left[\|\sum_\pi \Pr[\pi]\left(\mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S) - \mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S')\right)\|_1\right]$$

$$\leq \mathop{\mathbb{E}}_{\boldsymbol{\sigma}}\left[\sum_\pi \Pr[\pi]\|\left(\mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S) - \mathcal{M} \circ \pi \circ \boldsymbol{\sigma}(S')\right)\|_1\right]$$

$$= \mathop{\mathbb{E}}_{\boldsymbol{\pi},\boldsymbol{\sigma}}\left[\|\left(\mathcal{M} \circ \boldsymbol{\pi} \circ \boldsymbol{\sigma}(S) - \mathcal{M} \circ \boldsymbol{\pi} \circ \boldsymbol{\sigma}(S')\right)\|_1\right]$$

$$= \mathop{\mathbb{E}}_{\boldsymbol{\gamma}=(\boldsymbol{\sigma},\boldsymbol{\pi})}\left[\|\left(\mathcal{M} \circ \boldsymbol{\gamma}(S) - \mathcal{M} \circ \boldsymbol{\gamma}(S')\right)\|_1\right].$$

On the other hand, we also have:

$$\|\mathcal{D}_S - \mathcal{D}_{S'}\|_1 = \sum_{\gamma,y}\|D_S(\gamma,y) - \mathcal{D}_{S'}(\gamma,y)\|_1$$

$$= \sum_\gamma \Pr[\gamma] \cdot \sum_y \left|\left(\mathcal{M} \circ \gamma(S)\right)(y) - \left(\mathcal{M} \circ \gamma(S)\right)(y)\right|$$

$$= \mathop{\mathbb{E}}_\gamma\left[\sum_y \left|\left(\mathcal{M} \circ \gamma(S)\right)(y) - \left(\mathcal{M} \circ \gamma(S)\right)(y)\right|\right]$$

$$= \mathop{\mathbb{E}}_\gamma\left[\|\left(\mathcal{M} \circ \gamma(S) - \mathcal{M} \circ \gamma(S')\right)\|_1\right]$$

Combining the above yields, $d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{S'}) = \frac{1}{2}\|D_S - \mathcal{D}_{S'}\|_1 \geq \frac{\rho \cdot d}{2n}$ as claimed. □

6.1.2 The adversary and proof of Lemma 6.2

We consider $\mathcal{M} : X^n \to Y$ where $|X| = \Theta(n^2)$. Fix some unknown set $S \in X^{(n)}$. In the above subsection, we explained how in this setting the adversary, upon seeing $\ell = O(n \log(|X|))$ samples from $\mathcal{D}_S$ could use hypothesis selection to output a set S' that has many elements in common with S. In particular, consider an adversary that first samples $\overline{\gamma} = (\gamma_1, \ldots, \gamma_\ell) \sim \mathfrak{S}(X, n)^\ell$, and then runs

$$\mathcal{M}_{\overline{\gamma}} := (\mathcal{M} \circ \gamma_1(S), \ldots, \mathcal{M} \circ \gamma_\ell(S)). \tag{3}$$

Using the output of the algorithm, they build samples $(\gamma_i, \mathcal{M} \circ \gamma_i(S))_{i \in [\ell]}$ (which are distributed iid. from $\mathcal{D}_S$) and then run hypothesis selection. It's easy to see that such an adversary would succeed with high probability in finding S' which overlaps a lot with S. So why are we not done already ? Axiom 1 only gives us guarantees on the privacy of $\mathcal{M} \circ \gamma$ for a fixed γ; we have no guarantee on an algorithm that samples first γ and then runs $\mathcal{M} \circ \gamma(S)$. This is why we need to "derandomize" the above strategy to argue there is a fixed choice of $\overline{\gamma}^\star$ such that the adversary succeeds with high probability when seeing output of $\mathcal{M}_{\overline{\gamma}^\star}$.

Lemma 6.5. *Let $\rho > 0$ be a constant. Assume that for every $S, S' \in X^{(n)}$ we have*

$$d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{S'}) \geq \frac{\rho}{2} \cdot \mathrm{dist}(S, S')/n.$$

Then, there exists $\overline{\gamma} \in \mathfrak{S}(X, n)^{\ell}$ where $\ell = O(n \log(|X|))$ and adversary $\mathcal{A} : Y^{\ell} \to X^{(n)}$ such that:

$$\underset{\boldsymbol{S} \sim X^{(n)}}{\mathbb{E}} \left[\underset{\boldsymbol{S}' \leftarrow \mathcal{A}(\mathcal{M}_{\overline{\gamma}}(S))}{\mathbb{E}} \left[\mathrm{dist}(S, \boldsymbol{S}') \right] \right] \leq 0.02n$$

Proof. We denote by $\mathcal{H}$ the hypothesis selection algorithm of Fact 6.3 which we will aim to run on the set of candidates distribution $\mathcal{F}$. In particular, we let $\ell = O(\log(|\mathcal{F}|)) = O(n \log(|X|))$ be large enough so that for every $\mathcal{D}_S \in \mathcal{F}$, with probability at least 0.99 we have that $\mathcal{H}$ returns a distribution $\mathcal{D}_{S'}$ with

$$d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{S'}) < \frac{1}{100} \cdot \frac{\rho}{2}.$$

For every $\overline{\gamma} \in \mathfrak{S}(X, n)^{\ell}$, we define an adversary $\mathcal{A}_{\overline{\gamma}}$ described in Figure 1 who runs the algorithm $\mathcal{M}_{\overline{\gamma}}(S)$. We will prove the following:

$$\underset{\overline{\gamma} \sim \mathfrak{S}(X,n)^{\ell}}{\mathbb{E}} \left[\underset{\substack{\boldsymbol{S} \sim X^{(n)}, \\ \boldsymbol{S}' \leftarrow \mathcal{A}_{\overline{\gamma}}(\mathcal{M}_{\overline{\gamma}}(\boldsymbol{S}))}}{\mathbb{E}} \left[\mathrm{dist}(\boldsymbol{S}, \boldsymbol{S}') \right] \right] \leq 0.02n. \tag{4}$$

In particular, this implies there exists a fixed choice of $\overline{\gamma}^{\star}$, such that that

$$\underset{\substack{\boldsymbol{S} \sim X^{(n)}, \\ \boldsymbol{S}' \leftarrow \mathcal{A}_{\overline{\gamma}^{\star}}(\mathcal{M}_{\overline{\gamma}^{\star}}(\boldsymbol{S}))}}{\mathbb{E}} \left[\mathrm{dist}(\boldsymbol{S}, \boldsymbol{S}') \right] \leq 0.02n,$$

proving the lemma. We now turn to the prove of Equation (4), we will simply show that for every $S \in X^{(n)}$ (unknown to the adversary), we have

$$\underset{\overline{\gamma} \sim \mathfrak{S}(X,n)^{\ell}}{\mathbb{E}} \left[\underset{\boldsymbol{S}' \leftarrow \mathcal{A}_{\overline{\gamma}}(\mathcal{M}_{\overline{\gamma}}(S))}{\mathbb{E}} \left[\mathrm{dist}(S, \boldsymbol{S}') \right] \right] \leq 0.02n.$$

Fix $S \in X^{(n)}$, by our choice of ℓ and the properties of $\mathcal{H}$ we have:

$$\underset{\substack{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_{\ell} \sim \mathcal{D}_S \\ \boldsymbol{S}' \leftarrow \mathcal{H}(\boldsymbol{z}_1, \ldots, \boldsymbol{z}_{\ell})}}{\mathrm{Pr}} \left[d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{\boldsymbol{S}'}) \geq 0.01\rho/2 \right] \leq 0.01$$

By assumption, we have that for any S' with $\mathrm{dist}(S, S') \geq 0.01n$ we have $d_{\mathrm{TV}}(\mathcal{D}_S, \mathcal{D}_{S'}) \geq \frac{1}{100}\frac{\rho}{2}$, so we can conclude that:

$$\underset{\substack{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_{\ell} \sim \mathcal{D}_S \\ \boldsymbol{S}' \leftarrow \mathcal{H}(\boldsymbol{z}_1, \ldots, \boldsymbol{z}_{\ell})}}{\mathbb{E}} \left[\mathrm{dist}(S, \boldsymbol{S}') \right] \leq n \cdot 0.01 + 0.01n \cdot (1 - 0.01) \leq 0.02n$$

Finally note that the output of $\boldsymbol{S}' \leftarrow \mathcal{M}_{\overline{\gamma}}$ is distributed as follows: First run $\mathcal{M}_{\overline{\gamma}}(S)$ to get

$(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\ell)$ and then run $\mathcal{H}((\boldsymbol{\gamma}_1, \boldsymbol{y}_1), \ldots, (\boldsymbol{\gamma}_\ell, \boldsymbol{y}_\ell))$ to get $\boldsymbol{S}'$. With this in mind, we have:

$$
\mathop{\mathbb{E}}_{\overline{\gamma} \sim \mathfrak{S}(X,n)^\ell} \left[\mathop{\mathbb{E}}_{\boldsymbol{S}' \leftarrow \mathcal{A}_{\overline{\gamma}}(\mathcal{M}_{\overline{\gamma}}(S))} \left[\operatorname{dist}(S, \boldsymbol{S}') \right] \right] = \mathop{\mathbb{E}}_{\overline{\gamma} \sim \mathfrak{S}(X,n)^\ell} \left[\mathop{\mathbb{E}}_{\substack{(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\ell) \leftarrow \mathcal{M}_{\overline{\gamma}}(S) \\ \mathcal{D}_{\boldsymbol{S}'} \leftarrow \mathcal{H}((\boldsymbol{\gamma}_1, \boldsymbol{y}_1), \ldots, (\boldsymbol{\gamma}_\ell, \boldsymbol{y}_\ell))}} \left[\operatorname{dist}(S, \boldsymbol{S}') \right] \right]
$$

$$
= \mathop{\mathbb{E}}_{\gamma_1, \ldots, \gamma_\ell \sim \mathfrak{S}(X,n)} \left[\mathop{\mathbb{E}}_{\substack{\boldsymbol{y}_1 \leftarrow \mathcal{M} \circ \gamma_1(S), \ldots, \boldsymbol{y}_\ell \leftarrow \mathcal{M} \circ \gamma_\ell(S) \\ \mathcal{D}_{\boldsymbol{S}'} \leftarrow \mathcal{H}((\boldsymbol{\gamma}_1, \boldsymbol{y}_1), \ldots, (\boldsymbol{\gamma}_\ell, \boldsymbol{y}_\ell))}} \left[\operatorname{dist}(S, \boldsymbol{S}') \right] \right]
$$

$$
= \mathop{\mathbb{E}}_{\substack{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_\ell \sim \mathcal{D}_S \\ \mathcal{D}_{\boldsymbol{S}'} \leftarrow \mathcal{H}(\boldsymbol{z}_1, \ldots, \boldsymbol{z}_\ell)}} \left[\operatorname{dist}(S, \boldsymbol{S}') \right]
$$

$$
\leq 0.02n. \qquad \square
$$

---

1. Run the algorithm $\mathcal{M}_{\overline{\gamma}} = (\mathcal{M} \circ \gamma_1, \ldots, \mathcal{M} \circ \gamma_\ell)$ on the unknown set S.
2. Let $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\ell)$ be the output of $\mathcal{M}_{\overline{\gamma}}$.
3. Build a set of samples $\boldsymbol{Z} = (\boldsymbol{z}_1, \ldots, \boldsymbol{z}_\ell)$ where $\boldsymbol{z}_i = (\gamma_i, \boldsymbol{y}_i)$.
4. Runs the hypothesis selection algorithm $\mathcal{H}$ on the sample set $\boldsymbol{Z}$ with the set of candidate distributions $\mathcal{F}$ (See Definition 10).
5. If $\mathcal{H}$ returns $\boldsymbol{S}'$, do the same.

Figure 1: The adversary $\mathcal{M}_{\overline{\gamma}}$ where $\overline{\gamma} = (\gamma_1, \ldots, \gamma_\ell) \in \mathfrak{S}(X,n)^\ell$.

In the above we assumed that the hidden set S was in $X^{(n)}$. By our choice of $|X| = \Theta(n^2)$, we have with high probability that $\boldsymbol{S} \sim \operatorname{Unif}(X)^n$ has no duplicates. As such, the above adversary can just ignore the (small) probability that $\boldsymbol{S} \sim \operatorname{Unif}(X)^n$ has duplicates, and will still output a set $\boldsymbol{S}'$ such that $\mathbb{E}_{\boldsymbol{S} \sim \operatorname{Unif}(X)^n}[\operatorname{dist}(\boldsymbol{S}, \boldsymbol{S}')] < 0.1n$, guaranteeing that $\mathcal{M}_{\overline{\gamma}}$ is blatantly non-private. With this in mind, we now present the proof of Lemma 6.2:

Proof of Lemma 6.2. Let $X \subseteq X^\star$ of the size $\lceil \frac{100}{\rho} \rceil \cdot n^2$. Assume for contradiction that $\operatorname{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}}, \operatorname{Unif}(X)) > \rho$. Let E be the event over the draw of $\boldsymbol{S}, \boldsymbol{S}' \sim X^n$ that $\boldsymbol{S}, \boldsymbol{S}'$ have no common elements nor any duplicates. We have:

$$
\mathop{\mathrm{Pr}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \operatorname{Unif}(X)^n}[E] \geq \prod_{i=1}^{2n}(1 - \rho \frac{i-1}{100 \cdot n^2}) \geq (1 - \rho \frac{2n}{100 \cdot n^2})^{2n} \geq 1 - \frac{\rho}{25}.
$$

In which case we have:

$$
\mathop{\mathbb{E}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \binom{X}{n}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}}(\boldsymbol{S}), \widetilde{\mathcal{M}}(\boldsymbol{S}')) \,\middle|\, |\boldsymbol{S} \cap \boldsymbol{S}'| = 0 \right] = \mathop{\mathbb{E}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \operatorname{Unif}(X)^n} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}}(\boldsymbol{S}), \widetilde{\mathcal{M}}(\boldsymbol{S}')) \,\middle|\, E \right]
$$

$$
\geq \mathop{\mathbb{E}}_{\boldsymbol{S}, \boldsymbol{S}' \sim \operatorname{Unif}(X)^n} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}}(\boldsymbol{S}), \widetilde{\mathcal{M}}(\boldsymbol{S}')) \right] - (1 - \mathrm{Pr}[E])
$$

$$
\geq \rho - \frac{\rho}{2}
$$

$$
\geq \rho/2.
$$

Now consider the algorithm $\mathcal{M}' : X^n \to Y$ which is just $\mathcal{M}$ restricted to inputs in X, we denote its symmetrized version by $\widetilde{\mathcal{M}'}$. It's easy to see that we also have:

$$\mathbb{E}_{\boldsymbol{S},\boldsymbol{S}'\sim\binom{X}{n}}\left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}'}(\boldsymbol{S}),\widetilde{\mathcal{M}'}(\boldsymbol{S}'))\,\Big|\,|\boldsymbol{S}\cap\boldsymbol{S}'|=0\right]\geq\rho/2.$$

Hence, by Lemma 6.4 we have that for any $S, S' \in X^{(n)}$ with $\mathrm{dist}(S,S')\geq 0.01n$ we have

$$d_{\mathrm{TV}}(\mathcal{D}_S,\mathcal{D}_{S'})\geq 0.01\rho/4. \tag{5}$$

By Lemma 6.5, we have there exists $\overline{\gamma}=(\gamma_1,\ldots,\gamma_\ell)\in\mathfrak{S}(X,n)^\ell$ where $\ell=O(n\log(|X|))=O(n\log(n))$ and adversary $\mathcal{A}:Y^\ell\to X^{(n)}$ such that:

$$\mathbb{E}_{\boldsymbol{S}\sim X^{(n)}}\left[\mathbb{E}_{\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}'_{\overline{\gamma}}(S))}\left[\mathrm{dist}(\boldsymbol{S},\boldsymbol{S}')\right]\right]\leq 0.02n,$$

where $\mathcal{M}'_{\overline{\gamma}}$ is as defined in Equation (3). Finally, it follows by our bound on the probability of E that whenever we draw $\boldsymbol{S}\sim\mathrm{Unif}(X)^n$, we have that $\boldsymbol{S}\in X^{(n)}$ (i.e. $\boldsymbol{S}$ has no duplicates) with probability at least $1-1/25\geq 0.96$.

As such we have:

$$\mathbb{E}_{\boldsymbol{S}\sim\mathrm{Unif}(X)^n}\left[\mathbb{E}_{\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}'_{\overline{\gamma}}(\boldsymbol{S}))}\left[\sum_{x\in\boldsymbol{S}}^n\mathbb{1}[x\in\boldsymbol{S}']\right]\right]\geq 0.96\,\mathbb{E}_{\boldsymbol{S}\sim\mathrm{Unif}\left(X^{(n)}\right)}\left[\mathbb{E}_{\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}'_{\overline{\gamma}}(\boldsymbol{S}))}\left[\sum_{x\in\boldsymbol{S}}^n\mathbb{1}[x\in\boldsymbol{S}']\right]\right]$$

$$=0.96\,\mathbb{E}_{\boldsymbol{S}\sim\mathrm{Unif}\left(X^{(n)}\right)}\left[\mathbb{E}_{\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}'_{\overline{\gamma}}(\boldsymbol{S}))}\left[n-\mathrm{dist}(\boldsymbol{S},\boldsymbol{S}')\right]\right]$$

$$\geq 0.94n.$$

Right now, this means there is an adversary which can leak most of the dataset when seeing the output of $\mathcal{M}'_{\overline{\gamma}}$, but we have no guarantee for the privacy of $\mathcal{M}'$. Hence, we now want to translate this to an adversary working for $\mathcal{M}$. For each $\gamma_i=(\sigma_i,\pi_i)\in\overline{\gamma}$ we define $\sigma_i^\star:X^\star\to X^\star$ as:

$$\sigma_i^\star(x)=\begin{cases}x\text{ if }x\in X^\star\setminus X\\\sigma_i(x)\text{ otherwise.}\end{cases}$$

We then let $\gamma_i^\star=(\sigma_i^\star,\pi_i)$ and define $\overline{\gamma}^\star=(\gamma_1^\star,\ldots,\gamma_\ell^\star)$. It's easy to see that that for any $S\in X^{(n)}$ we have the following distributions are equivalent

$$\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}'_{\overline{\gamma}}(S))\text{ and }\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}_{\overline{\gamma}^\star}(S)).$$

Hence,

$$\mathbb{E}_{\boldsymbol{S}\sim X^n}\left[\mathbb{E}_{\boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}_{\overline{\gamma}^\star}(\boldsymbol{S}))}\left[\sum_{x\in\boldsymbol{S}}^n\mathbb{1}[x\in\boldsymbol{S}']\right]\right]\geq 0.94n.$$

Recall that $\ell=O(n\log(n))$, hence by making $\alpha,t>0$ large enough we have that $\varepsilon:=\mathcal{P}(\mathcal{M})\leq\frac{1}{\alpha n\log^t(n)}$ is so that, in Axiom 3 (strong composition), we have $\varepsilon'=O(\varepsilon\cdot\ell^c\cdot\mathrm{polylog}(n)\leq 1$.

We have that $\mathcal{M}_{\overline{\gamma}^\star} = (\mathcal{M} \circ \gamma_1^\star, \ldots, \mathcal{M} \circ \gamma_\ell^\star)$. In particular, for each $1 \leq i \leq \ell$, by Axiom 1 we have that $\mathcal{P}(\mathcal{M} \circ \gamma_i^\star) \leq \mathcal{P}(\mathcal{M})$. [6]

Hence by Axiom 3 we must have that $\mathcal{M}_{\overline{\gamma}^\star}$ is $\mathcal{P}$-private. But combined with existence of the adversary $\widetilde{\mathcal{A}}$ we have that Axiom 2 can't hold, reaching a contradiction. Hence, it must be that $\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}}, \mathrm{Unif}(X)) \leq \rho$. $\qquad\square$

6.2 Proof of Theorem 4: Getting TV-Stability on all distributions

Lemma 6.2 only guarantees that $\widetilde{\mathcal{M}}$ is ρ-TV-stable on distributions that are uniform over a small set elements. To achieve Theorem 4 we actually need to guarantee ρ-TV-Stability on *all* distributions. To get around this, we use two separate arguments depending on the size of the domain. In the first case, when $|X|$ is large we will show the following:

Claim 6.6 (TV-Stability on one distribution to TV-Stability on all distributions). *For any $\mathcal{M}$: $X^n \to Y$ where $X \geq 8n^2$, let $X' \subseteq X$ be any subset of size at least $8n^2$. For any privacy measure $\mathcal{P}$ satisfying Axiom 1 and distribution $\mathcal{D}$ on X, there exists an algorithm $\mathcal{M}' : X^n \to Y$ for which $\mathcal{P}(\mathcal{M}') \leq \mathcal{P}(\mathcal{M})$ satisfying,*

$$\frac{1}{2} \cdot \mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}}, \mathcal{D}) \leq \mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}'}, \mathrm{Unif}(X')). \tag{6}$$

where $\widetilde{\mathcal{M}}$ is the symmetrized version of $\mathcal{M}$ defined in Definition 7.

In the case where the domain is small, we take an entirely different approach. Instead, we show that any algorithm, without any assumption that it is $\mathcal{P}$-private, can be converted into a TV-stable algorithm.

Lemma 6.7 (Any algorithm can be made TV-stable over small domains). *Let $\mathcal{M} : X^n \to Y$ be an algorithm and $\rho > 0$ be a constant. There exists a $(\beta, 2\beta)$ equivalent algorithm $\mathcal{M}' : X^m \to Y$ such that $\mathcal{M}'$ is ρ-TV-stable and*

$$m = O\left(\frac{|X| + \log(1/\beta)}{(\beta/n)^2}\right).$$

We first prove Claim 6.6.

Proof of Claim 6.6. Let $\mathcal{D}$ be an arbitrary distribution over X and $X' \subseteq X$ with $|X'| \geq 8n^2$. Consider the random map $\boldsymbol{\sigma} : X \to X$ where $\boldsymbol{\sigma}(x) \sim \mathcal{D}$. Consider the two following distributions:

1. Draw $\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n$ and $\boldsymbol{\sigma}$ as above and output $(\boldsymbol{\sigma}(\boldsymbol{T}), \boldsymbol{\sigma}(\boldsymbol{T}'))$.

2. Draw $\boldsymbol{S}, \boldsymbol{S}' \sim \mathcal{D}^n$ and output $(\boldsymbol{S}, \boldsymbol{S}')$.

Let E denote the event where $\boldsymbol{T}, \boldsymbol{T}'$ contain no duplicates and no element in common. Observe that conditioned on E, both $\boldsymbol{\sigma}(\boldsymbol{T}), \boldsymbol{\sigma}(\boldsymbol{T}')$ and $(\boldsymbol{S}, \boldsymbol{S}')$ follow the same distribution.

We now have:

$$\Pr_{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n}[E] \geq \prod_{i=1}^{2n}\left(1 - \frac{i-1}{|X'|}\right) \geq \left(1 - \frac{2n}{|X'|}\right)^{2n} \geq 1/2.$$

---

[6] Recall that $\gamma_i^\star = (\sigma_i^\star, \pi_i)$ for some permutation $\pi : [n] \to [n]$ and $\sigma_i = X \to X$. So this follows by applying both points of the axioms to $\mathcal{M} \circ \gamma_i = \mathcal{M} \circ \sigma_i^\star \circ \pi_i$.

Where the last inequality follows by having $|X'| \geq 8n^2$. Hence, we have:

$$
\begin{aligned}
\mathrm{stab}_{\mathrm{tv}}(\widetilde{M}, \mathcal{D}) &= \underset{\boldsymbol{S}, \boldsymbol{S}' \sim \mathcal{D}^n}{\mathbb{E}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}}(\boldsymbol{S}), \tilde{\mathcal{M}}(\boldsymbol{S}')) \right] \\
&= \underset{\substack{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n \\ \boldsymbol{\sigma} \sim \mathfrak{S}(X)}}{\mathbb{E}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}} \circ \boldsymbol{\sigma}(\boldsymbol{T}), \widetilde{\mathcal{M}} \circ (\boldsymbol{T}') \mid E) \right] \\
&\leq \frac{1}{\Pr_{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n}[E]} \cdot \underset{\substack{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n \\ \boldsymbol{\sigma}}}{\mathbb{E}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}} \circ \boldsymbol{\sigma}(\boldsymbol{T}), \widetilde{\mathcal{M}} \circ (\boldsymbol{T}')) \right] \\
&\leq 2 \underset{\substack{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n \\ \boldsymbol{\sigma}}}{\mathbb{E}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}} \circ \boldsymbol{\sigma}(\boldsymbol{T}), \widetilde{\mathcal{M}} \circ (\boldsymbol{T}')) \right]. \\
&\leq 2 \underset{\boldsymbol{\sigma}}{\mathbb{E}} \left[\underset{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(X')^n}{\mathbb{E}} \left[d_{\mathrm{TV}}(\widetilde{\mathcal{M}} \circ \boldsymbol{\sigma}(\boldsymbol{T}), \widetilde{\mathcal{M}} \circ (\boldsymbol{T}')) \right] \right]. \\
&\leq 2 \underset{\boldsymbol{\sigma}}{\mathbb{E}} \left[\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}} \circ \boldsymbol{\sigma}, \mathrm{Unif}(X')) \right]
\end{aligned}
$$

Hence, there exists a choice of $\sigma^\star : X \to X$, such that $\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}} \circ \sigma^\star, \mathrm{Unif}(X')) \geq \frac{1}{2} \mathrm{stab}_{\mathrm{tv}}(\widetilde{M}, \mathcal{D})$. Now, let $\mathcal{M}' = \mathcal{M} \circ \sigma^\star$. It's easy to see that $\widetilde{\mathcal{M}'} = \widetilde{\mathcal{M}} \circ \sigma^\star$, meaning $\mathcal{M}'$ respects Equation (6). Furthermore, by Axiom 1 (preprocessing), we must have $\mathcal{P}(\mathcal{M}') \leq \mathcal{P}(\mathcal{M})$. $\qquad \square$

We now turn to the proof of Lemma 6.7. First, we will need the following well known fact, see for instance Theorem 1 of [Can20]:

Lemma 6.8. *Let X be a finite set. There exists an algorithm taking $O\left(\frac{|X| + \log(1/\delta)}{\varepsilon^2} \right)$ samples from an unknown distribution $\mathcal{D}$ over X and outputs a distribution Sim_D such that $d_{\mathrm{TV}}(\mathcal{D}, \mathsf{Sim}_{\mathcal{D}}) \leq \varepsilon$ with probability at least $1 - \delta$.*

Proof of Lemma 6.7. Let $\mathcal{D}$ be an arbitrary distribution over X We consider the algorithm $\mathcal{M}'$ from Figure 2, which uses the claimed bound on the number of samples from the distribution $\mathcal{D}$. It remains to show that $\mathcal{M}'$ outputs a correct answer with probability at least $1 - 2\beta$ and that it is TV-stable, meaning:

$$
\underset{\boldsymbol{S}^1, \boldsymbol{S}^2 \sim \mathcal{D}^m}{\mathbb{E}} \left[d_{\mathrm{TV}} \left(\mathcal{M}'(\boldsymbol{S}^1), \mathcal{M}'(\boldsymbol{S}^2) \right) \right] \leq \rho.
$$

To do this, let $\mathcal{O}$ be the distribution over Y induced by sampling $\boldsymbol{S} \sim \mathcal{D}^n$ and then running $\mathcal{M}(\boldsymbol{S})$. We will show that the following equation holds:

$$
\underset{\boldsymbol{S} \sim \mathcal{D}^m}{\mathbb{E}}[d_{\mathrm{TV}}(\mathcal{M}'(\boldsymbol{S}), \mathcal{O})] \leq \frac{1}{2} \cdot \beta \cdot \rho \tag{7}
$$

TV-Stability then follows from the above by the triangle inequality. For the correctness of $\mathcal{M}'$, let $\mathcal{T}(\mathcal{D})$ be the set of good answer to the statistical task $\mathcal{T}$ under distribution $\mathcal{D}$. By the correctness of $\mathcal{M}$ we have that $\Pr_{\boldsymbol{A} \sim \mathcal{O}}[\boldsymbol{A} \in \mathcal{T}(\mathcal{D})] = \Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})] \geq 1 - \beta$. Hence, assuming Equation (7) holds, we have:

$$\Pr_{\boldsymbol{S}\sim\mathcal{D}^m}[\mathcal{M}'(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})] = \mathbb{E}_{\boldsymbol{S}\sim D^m}\Big[\Pr[\mathcal{M}'(\boldsymbol{S}) \in \mathcal{T}(\mathcal{D})]\Big]$$

$$\geq \mathbb{E}_{\boldsymbol{S}\sim\mathcal{D}^m}\Big[\Pr_{\boldsymbol{A}\sim\mathcal{O}}[\boldsymbol{A} \in \mathcal{T}(\mathcal{D})] - d_{\mathrm{TV}}(\mathcal{M}'(\boldsymbol{S}),\mathcal{O})\Big]$$

$$= \Pr_{\boldsymbol{A}\sim\mathcal{O}}[\boldsymbol{A} \in \mathcal{T}(\mathcal{D})] - \mathbb{E}_{\boldsymbol{S}\sim\mathcal{D}^n}[d_{\mathrm{TV}}(\mathcal{M}'(\boldsymbol{S}),\mathcal{O})]$$

$$\geq 1 - \beta - \beta.$$

So $\mathcal{M}'$ has failure probability at most 2β. We now prove Equation (7). Let $\tau = \frac{\beta\cdot\rho}{4}$. By setting the hidden constant in m to be large enough, we have by Lemma 6.8 that for any distribution $\mathcal{D}$, with failure probability $1 - \tau$ the distribution $\mathsf{Sim}_\mathcal{D}$ obtained by the algorithm is such that $d_{\mathrm{TV}}(\mathcal{D},\mathsf{Sim}_\mathcal{D}) \leq \tau/n$. Conditioned on $d_{\mathrm{TV}}(\mathcal{D},\mathsf{Sim}_\mathcal{D}) \leq \tau$ we have:

$$\tau \geq d_{\mathrm{TV}}\left(\mathsf{Sim}_\mathcal{D}^n,\mathcal{D}^n\right) \geq d_{\mathrm{TV}}\left(\mathcal{M}(\mathsf{Sim}_\mathcal{D}^n),\mathcal{M}(\mathcal{D}^n)\right) = d_{\mathrm{TV}}\left(\mathcal{M}(\mathsf{Sim}_\mathcal{D}^n),\mathcal{O}\right)$$

Hence, we have that:

$$\mathbb{E}_{\boldsymbol{S}\sim\mathcal{D}^n}\left[d_{\mathrm{TV}}\left(\mathcal{M}'(\boldsymbol{S}),\mathcal{O}\right)\right] \leq \tau + \mathbb{E}_{\boldsymbol{S}\sim\mathcal{D}^n}\left[d_{\mathrm{TV}}\left(\mathcal{M}'(\boldsymbol{S}),\mathcal{O}\right) \mid d_{\mathrm{TV}}\left(\mathsf{Sim}_\mathcal{D},D\right) \leq \tau/n\right]$$

$$= \tau + \mathbb{E}_{\boldsymbol{S}\sim\mathcal{D}^n}\left[d_{\mathrm{TV}}\left(\mathcal{M}(\mathsf{Sim}_\mathcal{D}^n),\mathcal{O}\right) \mid d_{\mathrm{TV}}(\mathsf{Sim}_\mathcal{D},D) \leq \tau/n\right]$$

$$\leq 2\tau.$$

By our choice of τ, this proves Equation (7) $\qquad\qquad\square$

Input: An algorithm $\mathcal{M} : X^n \to Y$, parameter $\beta > 0$. Sample access to an unknown distribution $\mathcal{D}$ over X.

1. Draw a set $\boldsymbol{S}$ of m samples $\mathcal{D}$ where

$$m = O\left(\frac{|X| + \log(1/\beta)}{(\beta/n)^2}\right).$$

2. Run the algorithm of Lemma 6.8 to get a distribution $\mathsf{Sim}_\mathcal{D}$ over X.
3. Sample $\boldsymbol{S}^\star \sim \mathsf{Sim}_\mathcal{D}^n$.
4. Output $\mathcal{M}(\boldsymbol{S}^\star)$.

Figure 2: A procedure to turn $\mathcal{M} : X^n \to Y$ into a TV-Stable algorithm $\mathcal{M}'$ by learning the distribution $\mathcal{D}$.

Finally, we turn to the proof of Theorem 4, which we recall bellow for convenience.

Theorem 4 (Our privacy axioms imply TV-Stability, formal version)**.** *Let $\rho > 0$ be a constant, $\mathcal{P}$ be any privacy measure satisfying Axioms 1 to 4 and $\mathcal{M} : X^n \to Y$ be any $\mathcal{P}$-private algorithm. Let be c the constant in Axiom 3 and p the polynomial in Axiom 4, define*

$$m' := \tilde{O}\left(\frac{r^2 \cdot n^2}{\beta^2}\right) \text{ where } r = \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}}\right).$$

Then there is a ρ-TV stable $\mathcal{M}'$ using m' samples that is $(\beta, \beta' \coloneqq O(\beta))$-equivalent to $\mathcal{M}$.

Proof. As mentioned in the overview, we will want to use Axiom 4 (linear scalability). Let t be the constant in Lemma 6.2. We will choose the smallest $k \geq p(n, 1/\beta)$ satisfying, for $m = kn$,

$$1/k \leq O\left(\frac{1}{m^c \log^t(m)}\right).$$

Ultimately, this will allow us to apply Lemma 6.2. Choosing any k at least $\tilde{O}(n^{c/(1-c)})$ suffices for the above expression to hold. In which case $m = \tilde{O}(n^{1/(1-c)})$.

We then set $m \coloneqq n \cdot k$ and case split on the size of the domain X.

Small domains: If $|X| \leq 100/\rho \cdot m^2$, we use Lemma 6.7 on $\mathcal{M}$ to get a ρ-TV-stable algorithm $\mathcal{M}' : X^{m'} \to Y$, which is $(\beta, 2\beta)$-equivalent to $\mathcal{M}$ and where

$$m' = O\left(\frac{m^2 + \log(1/\beta)}{\beta^2} \cdot n^2\right).$$

Note that above fits the claimed bound on the sample complexity of the lemma.

Large domains: If $|X| \geq 100/\rho \cdot m^2$. Since $k \geq p(n, 1/\beta)$, we can use Axiom 4 to obtain an algorithm to get an algorithm $\mathcal{M}^{\mathsf{A}} : X^m \to Y$ which is $(\beta, \beta' = O(\beta))$-equivalent to $\mathcal{M}$ using m samples satisfying

$$\mathcal{P}(\mathcal{M}^{\mathsf{A}}) \leq O(1/k) \leq O\left(\frac{1}{m^c \cdot \log^t(m)}\right).$$

We claim that the symmetrized version of $\mathcal{M}^{\mathsf{A}}$, $\widetilde{\mathcal{M}^{\mathsf{A}}}$ must be ρ-TV-Stable. Assume for contradiction it's not. Then, there exists a distribution $\mathcal{D}$ such that $\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}^{\mathsf{A}}}, D) > \rho$. We fix an arbitrary subset X' of size $100/\rho \cdot m^2$ of X. By Claim 6.6 we have that there exists an algorithm $\mathcal{M}' : X^m \to Y$ with $\rho/2 < \mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}'}, \mathrm{Unif}(X'))$ and $\mathcal{P}(\mathcal{M}') \leq \mathcal{P}(\mathcal{M}) = O(1/m^c \log^t(m))$.

By making the hidden constant in the $O(\cdot)$ small enough and using $|X'| = \frac{100}{\rho} m^2$, we have by Lemma 6.2, that $\mathrm{stab}_{\mathrm{tv}}(\widetilde{\mathcal{M}'}, \mathrm{Unif}(X')) \leq \rho/2$. Hence, using Claim 6.6, it must have been the case that $\widetilde{\mathcal{M}^{\mathsf{A}}}$ was ρ-TV-Stable. Furthermore, $\widetilde{\mathcal{M}^{\mathsf{A}}}$ uses m samples, which clearly matches the claim bounds on sample complexity of the lemma. Finally, by Fact 5.1 we have that $\widetilde{\mathcal{M}^{\mathsf{A}}}$ is (β', β') equivalent to $\mathcal{M}^{\mathsf{A}}$, and thus is $(\beta, O(\beta))$ equivalent to $\mathcal{M}$. $\qquad\square$

6.3 The proof of Lemma 3.1

In this subsection, we prove Lemma 3.1. Which we recall for convenience.

Lemma 3.1 (Key lemma, uniform permutations distinguish far samples). *For any $\mathcal{M} : \binom{X}{m} \to Y$ where $|X| \geq 2m$, define*

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

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

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

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

22

As mentioned in our overview, the first step toward the proof of Lemma 3.1 is Lemma 3.2 which we recall bellow:

Lemma 3.2 (Random walk to disjoint samples). *For any $S, S' \in \binom{X}{m}$, setting $d := \mathrm{dist}(S, S')$ and $k := \lceil m/d \rceil$, there exists random variables $\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k$ 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 $\boldsymbol{U}, \boldsymbol{U}' \sim \mathrm{Unif}\left(\binom{X}{m}\right)$ conditioned on $|\boldsymbol{U} \cap \boldsymbol{U}'| = 0$.*

We first prove Lemma 3.1, and delay the proof of Lemma 3.2 to the end of this subsection.

Proof. Let $S, S' \in \binom{X}{m}$ and let $d = \mathrm{dist}(S, S')$. By Lemma 3.2 we have that there exists random variables $\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k$ with $k = \lceil m/d \rceil$. Such that:

1. For any $i \in [k]$ the marginal distribution of $(\boldsymbol{T}^{i-1}, \boldsymbol{T}^i)$ is the same as $(\boldsymbol{\sigma}(S), \boldsymbol{\sigma}(S'))$ where $\boldsymbol{\sigma} \sim \mathfrak{S}(X)$.

2. The marginal distribution of $(\boldsymbol{T}^0, \boldsymbol{T}^k)$ is the same as drawing $\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}(\binom{X}{m})$ conditioned on $|\boldsymbol{T} \cap \boldsymbol{T}'| = 0$.

We thus have:

$$
\begin{aligned}
\rho &= \mathop{\mathbb{E}}_{\boldsymbol{T}, \boldsymbol{T}' \sim \mathrm{Unif}\left(\binom{X}{m}\right)} \left[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{T}), \mathcal{M}(\boldsymbol{T}')) \,\Big|\, |\boldsymbol{T} \cap \boldsymbol{T}'| = 0 \right] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k} \left[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{T}^0), \mathcal{M}(\boldsymbol{T}^k)) \right] \\
&\leq \mathop{\mathbb{E}}_{\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k} \left[\sum_{i=1}^{k} d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{T}^{i-1}), \mathcal{M}(\boldsymbol{T}^i)) \right] && \text{(Triangle inequality)} \\
&= \sum_{i=1}^{k} \mathop{\mathbb{E}}_{\boldsymbol{T}^{i-1}, \boldsymbol{T}^i} \left[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{T}^{i-1}), \mathcal{M}(\boldsymbol{T}^i)) \right] \\
&= \sum_{i=1}^{k} \mathop{\mathbb{E}}_{\boldsymbol{\sigma} \sim \mathfrak{S}(X)} \left[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{\sigma}(\boldsymbol{S})), \mathcal{M}(\boldsymbol{\sigma}(\boldsymbol{S}'))) \right] \\
&= k \mathop{\mathbb{E}}_{\boldsymbol{\sigma} \sim \mathfrak{S}(X)} \left[d_{\mathrm{TV}}(\mathcal{M} \circ \boldsymbol{\sigma}(\boldsymbol{S}), \mathcal{M} \circ \boldsymbol{\sigma}(\boldsymbol{S}')) \right].
\end{aligned}
$$

Finally, since $k \leq 2m/d$ we have,

$$
\mathop{\mathbb{E}}_{\boldsymbol{\sigma} \sim \mathfrak{S}(X)} \left[d_{\mathrm{TV}}(\mathcal{M} \circ \boldsymbol{\sigma}(\boldsymbol{S}), \mathcal{M} \circ \boldsymbol{\sigma}(\boldsymbol{S}')) \right] \geq \frac{\rho}{2} \cdot \frac{\mathrm{dist}(S, S')}{m}. \qquad \square
$$

Proof of Lemma 3.2. To better understand the distributions of $(\boldsymbol{\sigma}(S), \boldsymbol{\sigma}(S'))$ when $\boldsymbol{\sigma} \sim \mathfrak{S}(X)$ we will work with the following distributions:

Definition 11 (Joint distribution at distance d). *For any distance $d \in [0, m]$ define $\mathcal{J}_d$ to be the uniform distribution over all $T, T' \in \binom{X}{m}$ satisfying $\mathrm{dist}(T, T') = d$.*

In particular, we have the following:

Proposition 6.9 (The joint distribution of uniform permutations)**.** *For any* $S, S' \in \binom{X}{m}$ *satisfying* $\mathrm{dist}(S, S') = d$, *if we draw* $\boldsymbol{\sigma} \sim \mathfrak{S}(X)$, *the joint distribution of* $(\boldsymbol{\sigma}(S), \boldsymbol{\sigma}(S'))$ *is exactly* $\mathcal{J}_d$.

Proof. We observe that, for any permutation $\sigma : X \to X$ and set $T = \sigma(S)$, $T' = \sigma(S')$ then,

$$\mathrm{dist}(T, T') = m - |\sigma(S) \cap \sigma(S')| = m - |S \cap S'| = \mathrm{dist}(S, S').$$

Furthermore, if we choose $\boldsymbol{\sigma}$ uniformly then, by symmetry, every choice of $\boldsymbol{T}$ and $\boldsymbol{T}'$ with distance d are equally likely, giving the desired distribution. □

We now turn to the proof of Lemma 3.2.

Proof of Lemma 3.2. We start by constructing $\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k$. To begin, we pick $\boldsymbol{T}^0 \sim \mathrm{Unif}(\binom{X}{m})$ uniformly among all size-n samples. Then, for each $i \in [k]$ we will use the following procedure to form $\boldsymbol{T}^i$:

1. We choose d many elements to remove from $\boldsymbol{T}^{i-1}$. Here, there are two cases:

   (a) For the first $k-1$ steps (i.e. when $i < k$), there will be at least d many elements remaining from $\boldsymbol{T}^0$, and we choose the elements to remove uniformly from them:
   $$\boldsymbol{R}^i \sim \binom{\boldsymbol{T}^0 \cap \boldsymbol{T}^{i-1}}{d}.$$

   (b) In the last step, when $i = k$, the number of remaining elements in $\boldsymbol{T}^0 \cap \boldsymbol{T}^{i-1}$ will be $m - (k-1)d$. In particular, if $kd \neq m$ (which happens whenever n/d is not exactly an integer), the number of remaining elements will be strictly less than d. In this case, we will remove all $m - (k-1)d$ elements plus $kd - m$ other uniform elements from $\boldsymbol{T}^{i-1} \setminus \boldsymbol{T}^0$,
   $$\boldsymbol{R}^k = (\boldsymbol{T}^0 \cap \boldsymbol{T}^{i-1}) \cup \boldsymbol{E} \qquad \text{where } \boldsymbol{E} \sim \binom{\boldsymbol{T}^{i-1} \setminus \boldsymbol{T}^0}{kd - n}$$

2. We choose d many elements to add to $\boldsymbol{T}^{i-1}$. Here, we simply choose uniformly among all elements that have not appeared in the process yet,
   $$\boldsymbol{A}^i \sim \binom{X \setminus (\boldsymbol{T}^0 \cup \cdots \boldsymbol{T}^{i-1})}{d}.$$

3. We then construct $\boldsymbol{T}^i$ as:
   $$\boldsymbol{T}^i = (\boldsymbol{T}^{i-1} \setminus \boldsymbol{E}^i) \cup \boldsymbol{A}^i.$$

At each step, we construct $\boldsymbol{T}^i$ by swapping exactly d elements from $\boldsymbol{T}^{i-1}$, so $\mathrm{dist}(\boldsymbol{T}^{i-1}, \boldsymbol{T}^i) = d$ with probability 1. Furthermore, we eventually swap every element that started in $\boldsymbol{T}^0$, meaning $\boldsymbol{T}^k \subseteq \boldsymbol{A}^1 \cup \cdots \cup \boldsymbol{A}^k$. By construction, there is no overlap between $\boldsymbol{T}^0$ and $\boldsymbol{T}^k$, so $\mathrm{dist}(\boldsymbol{T}^0, \boldsymbol{T}^n) = m$ with probability 1.

Finally, we observe this process is fully symmetric in the following sense: If we remapped the entire domain X according to any permutation $\sigma : X \to X$, then all the probabilities remain the same:
$$\Pr[\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k = T^0, \ldots, T^k] = \Pr[\boldsymbol{T}^0, \ldots, \boldsymbol{T}^k = \sigma(T^0), \ldots, \sigma(T^k)],$$

where $\sigma(T) = \{\sigma(x) \mid x \in T\}$. Due to this symmetry, $\boldsymbol{T}^{i-1}$ and $\boldsymbol{T}^i$ are equally likely to be any two sets with distance d, and so are distributed according to $\mathcal{J}_d$. Similarly, $\boldsymbol{T}^0$ and $\boldsymbol{T}^k$ are equally likely to be any two sets with distance m, and as $\boldsymbol{T}^0, \boldsymbol{T}^k \sim \text{Unif}\left(\binom{X}{m}\right)$ conditioned on $\boldsymbol{T}^0 \cap \boldsymbol{T}^k = \emptyset$. The proof then follows by Proposition 6.9. □

7 Proof of Theorem 2: DP satisfies the axioms

First, we will need the following theorem to amplify (ε, δ)-differential privacy. This theorem is a small modification of Theorem 6.2 of [BGH+23] (mentioned in Section 4) and is proved in Section A.2. The result of [BGH+23] makes the error probability of the algorithm go from β to $O(\beta \log(\beta))$, whereas our modified Theorem achieves error probability $O(\beta)$, albeit with a slightly worse dependency on $\log(1/\beta)$ in the sample complexity.

Lemma 7.1. *There is a universal constant $0.1 \geq \alpha > 0$ such that the following holds. Let $\mathcal{M} : X^n \to Y$ be $(0.1, \alpha^2/n^3)$-differentially private. Then for every $\varepsilon, \delta > 0$; there exists an (ε, δ)-differentially private algorithm $\mathcal{M}' : X^m \to Y$ solving T using*

$$m = O\Big(\frac{\log(1/\beta)^2 + \log(1/\beta)\log(1/\delta)}{\varepsilon}\Big) \cdot n^2$$

samples and which is $(\beta, 5\beta)$-equivalent to $\mathcal{M}$.

We now define our stability measure $\mathcal{P}_{\text{DP}}$ as follows:

Definition 12. *For an algorithm $\mathcal{M}$ taking n samples let ε be the smallest value such that $\mathcal{M}$ is $(\varepsilon, \varepsilon^2/n^3)$-DP. Recalling that α is the constant of Lemma 7.1 we set*

$$\mathcal{P}_{\text{DP}}(\mathcal{M}) = \left(\frac{\varepsilon}{\alpha}\right)^{5/4}.$$

We now prove that $\mathcal{P}_{\text{DP}}$ fits all four of axioms. Before doing so we will need the following fact.

Fact 7.2. *Let $\mathcal{M}$ be an algorithm taking n samples. If $\mathcal{P}_{\text{DP}}(\mathcal{M}) \leq 1$, then $\mathcal{M}$ is $(0.1, \frac{\alpha^2}{n^3})$-DP.*

Lemma 7.3. *$\mathcal{P}_{\text{DP}}$ respects Axiom 1.*

Proof. We observe for any S, S' that differ in one coordinate and permutation $\pi : [n] \to [n]$, that $\pi(S)$ and $\pi(S')$ still differ in one coordinate. Similarly, for any $\sigma : X \to X$, $\sigma(S)$ and $\sigma(S')$ differ in (at most) one coordinate. Therefore, if $\mathcal{M}$ is (ε, δ)-DP, it is still (ε, δ)-DP after preprocessing. □

To prove that DP satisfies Axiom 2, we will need the following fact:

Fact 7.4. *Let $\mathcal{M} : X^n \to Y$, $\mathcal{A} : Y \to X^n$ be algorithms. Fix $S \in X^n, i \in [n]$ and $x \in X$. If $\mathcal{M}$ is (ε, δ)-differentially private then we have:*

$$\Pr[S_i \in \mathcal{A}(\mathcal{M}(S))] \leq e^\varepsilon \Pr[S_i \in \mathcal{A}(\mathcal{M}(S_{x \to i}))] + \delta,$$

where $S_{x \to i}$ denotes S with i-th element set to x.

Proof. S and $S_{x \to i}$ differ on at most 1 coordinate. The result thus follows from $\mathcal{M}$ being (ε, δ)-differentially private. □

Lemma 7.5. $\mathcal{P}_{\mathrm{DP}}$ *respects Axiom 2.*

Proof. Let $\mathcal{D}$ an arbitrary distribution over X with $\|\mathcal{D}\|_\infty \leq 1/100n^2$. Let $\mathcal{M} : X^n \to Y$ be an (ε, δ)-differentially private algorithm and let $\mathcal{A} : Y \to X^n$. We will show that:

$$\mathop{\mathbb{E}}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{S}' \leftarrow \mathcal{A}(\mathcal{M}(\boldsymbol{S}))}} \left[\sum_{x \in \boldsymbol{S}} \mathbb{1}[x \in \boldsymbol{S}'] \right] \leq \frac{e^\varepsilon}{100} + \delta n.$$

In particular, if $\mathcal{P}_{\mathrm{DP}}(\mathcal{M}) \leq 1$, which implies $\mathcal{M}$ is $(\varepsilon = 0.1, \delta = \alpha^2/n^3)$-DP, we have $\frac{e^\varepsilon}{100} + n\delta \leq 0.1$ proving the lemma. We want to bound:

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

Since $\mathcal{M}$ is (ε, δ)-DP, for any $i \in [n]$ and $x \in X$ we have, by Fact 7.4, that:

$$\mathop{\Pr}_{\boldsymbol{S} \sim \mathcal{D}^n} [\boldsymbol{S}_i \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}))] \leq e^\varepsilon \mathop{\Pr}_{\boldsymbol{S} \sim \mathcal{D}^n} [\boldsymbol{S}_i \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}_{x \to i}))] + \delta.$$

Where $S_{x \to i}$ is the set obtained by setting the i-th element of S to x. This implies that:

$$\mathop{\Pr}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{S}' \leftarrow \mathcal{A}(\mathcal{M}(\boldsymbol{S}))}} [\boldsymbol{S}_i \in \boldsymbol{S}'] = \mathop{\Pr}_{\boldsymbol{S} \sim \mathcal{D}^n} [\boldsymbol{S}_i \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}))]$$

$$\leq e^\varepsilon \mathop{\Pr}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{x} \sim \mathcal{D}}} [\boldsymbol{S}_i \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}_{\boldsymbol{x} \to i}))] + \delta.$$

$$= e^\varepsilon \mathop{\Pr}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{x} \sim \mathcal{D}}} [\boldsymbol{x} \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}))] + \delta.$$

Where the last line follows from the symmetry of $\boldsymbol{S}_i$ and $\boldsymbol{x}$. Hence, we have that:

$$\mathop{\mathbb{E}}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{S}' \leftarrow \mathcal{A}(\mathcal{M}(\boldsymbol{S}))}} \left[\sum_{x \in \boldsymbol{S}} \mathbb{1}[x \in \boldsymbol{S}'] \right] = \sum_{i=1}^n \left(e^\varepsilon \mathop{\Pr}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{x} \sim \mathcal{D}}} [\boldsymbol{x} \in \mathcal{A}(\mathcal{M}(\boldsymbol{S}))] + \delta \right)$$

$$\leq n \cdot e^\varepsilon \sup_{S \in X^n} \left(\mathrm{Pr}_{\boldsymbol{x} \sim \mathcal{D}}[\boldsymbol{x} \in S] \right) + \delta n.$$

Since $\|D\|_\infty \leq \frac{1}{100n^2}$, we have that for any $S \in X^n$, $\mathrm{Pr}_{\boldsymbol{x} \sim \mathcal{D}}[\boldsymbol{x} \in S] \leq \frac{n}{100n^2}$. From which we can conclude that Axiom 2 holds since:

$$\mathop{\mathbb{E}}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{S}' \leftarrow \mathcal{A}(\mathcal{M}(\boldsymbol{S}))}} \left[\sum_{x \in \boldsymbol{S}} \mathbb{1}[x \in \boldsymbol{S}'] \right] \leq \frac{e^\varepsilon}{100} + \delta n. \qquad \square$$

Before proving Axiom 3 holds, we recall advanced composition for (ε, δ)-differential privacy.

Theorem (DP-Strong-Composition, Theorem 3.20 in [DR14]).

For all $\varepsilon, \delta \geq 0$, if $\mathcal{M}^1, \ldots, \mathcal{M}^\ell$ are (ε, δ)-differentially private, then the composed algorithm $(\mathcal{M}^1, \ldots, \mathcal{M}^\ell)$ is (ε', δ')-differentially private where $\delta' := 2\ell\delta$ and

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

Lemma 7.6. $\mathcal{P}_{\mathrm{DP}}$ *respects Axiom 3 with composition constant $c = 5/8$.*

Proof. Let β be a suitably large constant, we define $\varepsilon' := \ell^c \cdot \varepsilon \cdot \left(\beta \log^2(n)\right)^{2c}$. Say we have algorithms $\mathcal{M}^1, \ldots, \mathcal{M}^\ell$ each taking n samples with $\varepsilon \geq \mathcal{P}_{\mathrm{DP}}(\mathcal{M}^i)$. For ε' to be less than 1 we need

$$\varepsilon \leq \ell^{-c \cdot} \left(\beta \log^2(n)\right)^{-2c}.$$

This means that each $\mathcal{M}^i$ is $(\mu, \mu^2/n^3)$-DP, where $\mu = \alpha\varepsilon^{4/5}$. By our choice of c we have

$$\mu = \frac{\alpha}{\beta \log^2(n)} \cdot \ell^{-1/2}.$$

We denote by $\mathcal{C}$ the composed algorithm $(\mathcal{M}^1, \ldots, \mathcal{M}^\ell)$, to prove the Axiom holds, it remains to show $\mathcal{P}_{\mathrm{DP}}(\mathcal{C}) \leq 1$. This is equivalent to showing $(0.1, \alpha^2/n^3)$-differentially private. First note that $\mu \leq 1$ and $\mu^2\ell \leq \alpha < 1$, which implies

$$\mu(e^\mu - 1) \leq 2\mu^2 \text{ and } \mu^2\ell \leq \mu\sqrt{\ell}.$$

Using Strong-Composition, we can conclude $\mathcal{C}$ is $(\varepsilon^\star, \delta^\star)$-DP where $\delta^\star = 2\ell\delta$. First note that

$$\ell\delta = \ell\mu^2/n^3 = \frac{1}{n^3}\left(\frac{\alpha}{\beta \log^2(n)}\right)^2. \tag{8}$$

From the above, we have $\delta^\star \leq \alpha^2/n^3$ by picking β to be suitably large.

$$\begin{aligned}
\varepsilon^\star &= \mu\sqrt{2\ell \ln(1/\ell\delta)} + \ell\mu(e^\mu - 1) \\
&\leq \mu\sqrt{2\ell \ln(1/\ell\delta)} + 2\ell\mu^2 \\
&\leq \mu\sqrt{2\ell \ln(1/\ell\delta)} + 2\mu\sqrt{\ell} \\
&\leq 3\mu\sqrt{2\ell \ln(1/\ell\delta)} \\
&\leq 3\sqrt{2}\alpha \cdot \frac{\ln(1/\ell\delta)}{\beta \log^2(n)} \\
&\leq \frac{\ln(1/\ell\delta)}{\beta \log^2(n)}
\end{aligned}$$

Where the last line follow by $\alpha \leq 0.1$. By our bound on $\ell\delta$ in Equation (8), we can choose β to be a suitably large constant to have $\varepsilon^\star \leq 0.1$. Hence, we have that $\mathcal{C}$ is $(0.1, \alpha^2/n^3)$-differentially private as needed. $\square$

Lemma 7.7. $\mathcal{P}_{\mathrm{DP}}$ *respects* Axiom 4 *with the following parameters:* $p(n, 1/\beta) = \Delta \cdot \frac{n^{10}}{\beta}$ *for some large enough constant* $\Delta \geq 1$. *The resulting algorithm* $\mathcal{M}'$ *is* $(\beta, 5\beta)$-*equivalent to* $\mathcal{M}$.

Proof. let $\mathcal{M} : X^n \to Y$. Assume $\mathcal{P}_{\mathrm{DP}}(\mathcal{M}) \leq 1$ which implies $\mathcal{M}$ is $(0.1, \alpha^2/n^3)$-DP. Let $k \geq p(n, \frac{1}{\beta})$ and $m := kn$, we first use Lemma 7.1 with the following parameters:

$$\varepsilon = n^2 \log^3(m)/m \text{ and } \delta = \varepsilon^2/m^3$$

The resulting algorithm $\mathcal{M}'$ is $(\beta, 5\beta)$-equivalent to $\mathcal{M}$. Recall that $m \geq k \geq \Delta n^{10}/\beta$ and Δ is large enough. So $\mathcal{M}'$ uses:

$$O\Big(\frac{\log(1/\beta)^2 + \log(1/\beta)\log(1/\delta)}{\varepsilon}\Big) \cdot n^2 = O\left(\frac{\log^2(m)}{n^2 \log^3(m)/m}\right) \cdot n^2 \leq m \text{ samples.}$$

We also have that $\mathcal{M}'$ is $(\varepsilon, \varepsilon^2/m^3)$-DP. Using $m/k = n$ and $k \geq n^{10}$ we have

$$\varepsilon \leq n^2 \log^3(m)/m = \frac{m \log^3(m)}{k^2} = \frac{n \log^3(kn)}{k} \leq \frac{2 \log^3(k)}{k^{9/10}}.$$

Since $k \geq \Delta$ for some large enough constant Δ, we can conclude $\varepsilon \leq \frac{1}{k^{4/5}}$. And thus $\mathcal{P}_{\mathrm{DP}}(\mathcal{M}') = (\varepsilon/\alpha)^{5/4} \leq \alpha^{-5/4}/k = O(1/k)$. $\qquad\square$

8   Minimality of our axioms and proof of Theorem 3

As discussed in Section 3.3, removing any of our axioms would lead to ill-behaved notions of privacy. In this section, we formalize the claims made there. We will furthermore show that these ill-behaved notions of privacy all allow algorithms that solve one of the following two tasks:

Definition 13 (The task FINDELEMENT). *For any domain X, the task* FINDELEMENT *is defined as follows: An algorithm given i.i.d. samples* $\boldsymbol{S} \sim \mathcal{D}^n$ *from an unknown* $\mathcal{D}$ *should output some x for which* $\mathcal{D}(x) > 0$.

The next task is similar, but only defined on distributions with one heavy element, and the algorithm's task is to output a different element in the support.

Definition 14 (The task FINDLIGHTELEMENT). *For any domain X, the task* FINDLIGHTELEMENT *is defined only on distributions $\mathcal{D}$ where there is some x satisfying $0.7 \leq \mathcal{D}(x) \leq 0.9$. To solve it, an algorithm given i.i.d. samples* $\boldsymbol{S} \sim \mathcal{D}^n$ *should output some $y \neq x$ satisfying $\mathcal{D}(y) > 0$.*

We will show in Section 8.5 that neither FINDELEMENT nor FINDLIGHTELEMENT can be solved by a DP algorithm using less than $O(\sqrt{|X|})$ many samples. Hence, by showing that the removal of any one axiom allows for solving either FINDELEMENT or FINDLIGHTELEMENT with $O(1)$ samples, we complete the proof of Theorem 3.

8.1 Removing each requirement of Axiom 1

Since Axiom 1 has two requirements, we will show that removing each requirement on its own is enough to allow ill-behaved privacy measures. In both cases, the privacy measure we define will have the following special structure.

Definition 15 (Binary privacy measure). *A privacy measure $\mathcal{P}$ is* binary *if there some set of "good" algorithms $\mathcal{G}$ and*

$$\mathcal{P}(\mathcal{M}) = \begin{cases} 0 & \text{if } \mathcal{M} \in \mathcal{G} \\ 2 & \text{otherwise.} \end{cases}$$

The reason binary privacy measures are easy to analyze is because our axioms say very little about how a privacy measure should behave at privacy levels about 1 (i.e. none of Axioms 2 to 4 apply in that regime. Note the choice to not require our axioms to enforce much when the privacy level is high was not artificial: Even DP starts to behave different when $\varepsilon > 1$ (for example, DP only satisfies strong composition when $\varepsilon \leq 1$ and for larger ε, satisfies linear composition).

We now state how binary privacy measures interact with our axioms:

Claim 8.1 (Axiom 1 for binary privacy measures). *A binary privacy measure with set $\mathcal{G}$ satisfies Axiom 1 if and only if:*

1. $\mathcal{G}$ ***is closed under reordering inputs*** *For any algorithm $\mathcal{M} : X^n \to Y$ and permutation $\pi : [n] \to [n]$ $\mathcal{M} \in \mathcal{G} \iff \mathcal{M} \circ \pi \in \mathcal{G}$.*

2. ***Remapping the domain maintains*** $\mathcal{G}$***:*** *For any mapping $\sigma : X \to X$ and algorithm $\mathcal{M} : X^n \to Y$, $\mathcal{M} \in \mathcal{G} \implies \mathcal{M} \circ \sigma \in \mathcal{G}$.*

Proof. This is immediate from the definition of Axiom 1 and Definition 15. $\square$

Claim 8.2 (Axiom 2 for binary privacy measures). *A binary privacy measure with set $\mathcal{G}$ satisfies Axiom 2 if and only if every $\mathcal{M} \in \mathcal{G}$ is not blatantly non-private.*

Proof. This is immediate from the definition of Axiom 2 and Definition 15. $\square$

Claim 8.3 (Axiom 3 for binary privacy measures). *A binary privacy measure with set $\mathcal{G}$ satisfies Axiom 3 iff for any $\mathcal{M}_1, \ldots, \mathcal{M}_k : X^n \to Y$ in $\mathcal{G}$, the algorithm that takes as input $S \in X^n$ and outputs $(\mathcal{M}_1(S), \ldots, \mathcal{M}_k(S))$ is also in $\mathcal{G}$*

Proof. This is immediate from the definition of Axiom 3 and Definition 15. $\square$

Claim 8.4 (Axiom 4 for binary privacy measures). *A binary privacy measure with set $\mathcal{G}$ satisfies Axiom 4 iff there exists some polynomial $p : \mathbb{R}^2 \to \mathbb{R}$, so that, for any $\mathcal{M} : X^n \to Y$ in $\mathcal{G}$, failure probability $\beta > 0$, and $k \geq p(n, 1/\beta)$, there exists some $(\beta, \beta' = O(\beta))$-equivalent algorithm $\mathcal{M}'$ taking in $m := kn$ that is also in $\mathcal{G}$.*

Proof. This is immediate from the definition of Axiom 4 and Definition 15. $\square$

8.1.1 Removing the requirement that reordering the input maintains privacy

In this case, we show that a binary measure privacy where the good algorithms are those that only depend on the first half of their dataset satisfies the remaining axioms.

Definition 16 (First-half-only-privacy measure). *The* first-half-only-privacy measure $\mathcal{P}_{half}$ *is the binary privacy measure for set $\mathcal{G}$ defined as follows: An $\mathcal{M} : X^n \rightarrow Y$ is in $\mathcal{G}$ iff $n \geq 2$ and there exists some (possibly randomized) algorithm $f : X^{\lfloor n/2 \rfloor} \rightarrow Y$ for which $\mathcal{M}(S)$ and $f(S_{\leq \lfloor n/2 \rfloor})$ are equal in distribution for all S.*

Claim 8.5. *$\mathcal{P}_{half}$ satisfies Axioms 2 to 4 and also the "remapping the domain maintains privacy" part of Axiom 1.*

Proof. This is immediate from Claims 8.1 to 8.4. For Claim 8.4, given any $\mathcal{M} : X^n \rightarrow Y$, we take $\mathcal{M}' : X^m \rightarrow Y$ to be the algorithm that runs $\mathcal{M}$ on the first n points of its dataset. Since $\mathcal{M}$ depends on only the first half of its dataset, the same will be true of $\mathcal{M}'$ (indeed $\mathcal{M}'$ will depend on *less* than the first half of its dataset). $\square$

We also observe that this privacy measure allows an algorithm which solves FINDELEMENT.

Fact 8.6. *The algorithm $\mathcal{M} : X^n \rightarrow Y$ which, on input S, outputs S_1 both is $\mathcal{P}_{half}$-private and solves* FINDELEMENT *with failure probability 0.*

8.1.2 Removing the requirement that remapping the domain maintains privacy

We now will remove the second part of Axiom 1 and give a different ill-behaved definition of privacy. This definition will allow algorithms to behave arbitrarily when there is one very heavy element in their sample; however, if there is no such heavy element, the algorithm must output a single uninformative symbol $y^\star$.

Definition 17 (Heavy-elements-only-privacy). *We say a dataset $S \in X^n$ is "heavy" if there is a single $x \in X$ appearing at least $0.6n$ times in S. The* heavy-elements-only-privacy measure $\mathcal{P}_{heavy}$ *is the binary privacy measure with $\mathcal{M} : X^n \rightarrow Y \in \mathcal{G}$ iff[7] $n \geq 40$ and the following holds: There is some special output $y^\star \in Y$ s.t. for any $S \in X^n$ that is not heavy, $\mathcal{M}(S)$ always outputs $y^\star$.*

Claim 8.7. *$\mathcal{P}_{heavy}$ satisfies Axioms 2 to 4 and also the "reordering the input maintains privacy" part of Axiom 1.*

Remark 1 (Permuting vs remapping the domain). Claim 8.7 shows that if we fully remove the part of Axiom 1 that requires privacy is maintained whenever the domain is remapped according to all $\sigma : X \rightarrow X$, then $\mathcal{P}_{heavy}$ is a valid (and ill-behaved) measure of privacy. We actually observe something stronger: An alternative definition of Axiom 1 is to only require privacy is maintained when the domain is *permuted* according to some bijective function $\sigma : X \rightarrow X$. In this case, $\mathcal{P}_{heavy}$ would satisfy all the axioms. Therefore, $\mathcal{P}_{heavy}$ justifies why Axiom 1 requiring privacy is preserved for all remappings rather than permutations of the domain is essential.

Our proof of Claim 8.7 will use the following result.

---

[7]This requirement that $n \geq 40$ is only to make the proof of Claim 8.8 easier.

Claim 8.8. *There exists an absolute constant $c \leq 20$ such that for, $n \geq 40$, $m \geq 2n + 1$, and distribution $\mathcal{D}$,*

$$\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\boldsymbol{S} \text{ is heavy}] \geq \frac{1}{c} \cdot \Pr_{\boldsymbol{S}' \sim \mathcal{D}^m}[\boldsymbol{S}' \text{ is heavy}].$$

Our proof of Claim 8.8 will use two (nontrivial) bounds on the behavior of hypergeometric random variables.

Fact 8.9 (Median of hypergeometric is close to its mean, [Sie01, CE09])**.** *Let $\boldsymbol{x}$ be any hypergeometric random variable with mean μ. Then, the median of $\boldsymbol{x}$ is either $\lfloor \mu \rfloor$ or $\lceil \mu \rceil$.*

Fact 8.10 (Log-concave distributions are spread out, [MP25, Ara24])**.** *Let $\mathcal{D}$ be a distribution supported on $\mathbb{N}$ that is log-concave (which includes all hypergeometric distributions) and has variance σ^2. Then, for all $x \in \mathbb{N}$,*

$$\mathcal{D}(x) \leq \frac{1}{\sqrt{1 + \sigma^2}}.$$

Proof of Claim 8.8. One way to draw $\boldsymbol{S} \sim \mathcal{D}^n$ is to first draw $\boldsymbol{S}' \sim \mathcal{D}^m$ and then draw $\boldsymbol{S}$ uniformly without replacement from $\boldsymbol{S}'$. We will show that for any fixed choice of heavy S', if we draw $\boldsymbol{S}$ uniformly without replacement from S', then $\boldsymbol{S}$ is heavy with probability at least $1/c$. Then, the desired result easily follows from the following series of inequalities:

$$\Pr[\boldsymbol{S} \text{ is heavy}] \geq \Pr[\boldsymbol{S} \text{ is heavy} \mid \boldsymbol{S}' \text{ is heavy}] \geq \frac{1}{c} \cdot \Pr[\boldsymbol{S}' \text{ is heavy}].$$

Since S' is heavy, there exists a single element appearing k times where $k \coloneqq pm$ and $p \geq 0.6$. Let $\boldsymbol{x}$ be the random variable indicating the number of times this element appears in S. Then, $\boldsymbol{x}$ is drawn from a hypergeometric distribution with mean $\mu \coloneqq pn$ and variance

$$\sigma^2 \coloneqq np(1 - p)\frac{m - n}{m - 1}.$$

Since $m \geq 2n + 1$, we will have $\sigma^2 \geq np(1 - p)/2$. Then, by Fact 8.9,

$$\Pr[\boldsymbol{x} \geq \lfloor pn \rfloor] \geq 1/2.$$

In particular, if $\lfloor pn \rfloor \geq 0.6n$ we are done. Otherwise, we will have $pn - 1 \leq 0.6n$ in which case p must be between 0.6 and $0.6 + 1/n$. For $n \geq 40$, this, means $0.6 \leq p \leq 0.625$, in which case the variance is at least $\sigma^2 \geq 40 \cdot 0.625 \cdot 0.375/2 \geq 4$. Then, we can bound,

$$\Pr[\boldsymbol{x} \geq 0.6n] \geq \Pr[\boldsymbol{x} \geq \lfloor pn \rfloor] - \Pr[\boldsymbol{x} \geq \lfloor pn \rfloor \text{ and } \boldsymbol{x} < 0.6n].$$

Since $pn \geq 0.6$, there is at most one value for x satisfying $x \geq \lfloor pn \rfloor$ and $x < 0.6n$. Using Fact 8.10, the probability $\boldsymbol{x}$ takes on this one value is at most $\frac{1}{\sqrt{1 + \sigma^2}} \leq \frac{1}{\sqrt{5}}$. Therefore,

$$\Pr[\boldsymbol{x} \geq 0.6n] \geq \frac{1}{2} - \frac{1}{\sqrt{5}} \geq \frac{1}{20}. \qquad \square$$

We are now ready to prove that $\mathcal{P}_{\text{heavy}}$ satisfies all of our axioms except for the "reordering the domain maintains privacy" part of Axiom 1.

Proof of Claim 8.7. The "reordering the input maintains privacy" part of Axiom 1 is immediate from Claim 8.1. Similarly, Axiom 3 holds immediately from Claim 8.3. Axiom 2 holds by Claim 8.2 and the observation that for any $\mathcal{D}$ satisfying $\|\mathcal{D}\|_\infty \leq 1/(100n^2)$ it is unlikely for $n \geq 40$ that $\boldsymbol{S} \sim \mathcal{D}^n$ is heavy.

The proof that Axiom 4 holds requires a bit more care: Let $\mathcal{M} : X^n \to Y$ be any algorithm that is $\mathcal{P}_{\text{heavy}}$-private. Then, it has an input $y^\star$ it outputs whenever it does not have a heavy input. For any $m \geq 2n + 1$, let $\mathcal{M}' : X^m \to Y$ be the algorithm that on input $S' \in X^m$ does the following:

1. It draws a uniform size-n subsample $\boldsymbol{S}$ without replacement from S' and defines $\boldsymbol{y} = \mathcal{M}(S)$.

2. If S' is heavy, it outputs $\boldsymbol{y}$. Otherwise, it outputs $y^\star$.

Its clear that $\mathcal{P}_{\text{heavy}}(\mathcal{M}') = 0$. Therefore all that remains is to show that $\mathcal{M}'$ is $(\beta, \beta' = O(\beta))$-equivalent to $\mathcal{M}$ for all β. We will prove this with $\beta' = (c + 1) \cdot \beta$ where c is the constant in Claim 8.8. For this, we first observe that the definition of $\mathcal{M}'$ immediately implies that for any distribution $\mathcal{D}$ there exists a coupling of $\boldsymbol{y}' := \mathcal{M}'(\boldsymbol{S}')$ and $\boldsymbol{y} := \mathcal{M}(\boldsymbol{S})$ where $\boldsymbol{S} \sim \mathcal{D}^n$ and $\boldsymbol{S}' \sim \mathcal{D}^m$ satisfying that either $\boldsymbol{y}' = \boldsymbol{y}$ or $\boldsymbol{y}' = y^\star$. Furthermore, the latter case occurs with probability at most $\Pr[\boldsymbol{S}'$ is heavy$]$.

Now, let $\mathcal{T}$ be any statistical task that $\mathcal{M}$ solves with failure probability β. On any input distribution $\mathcal{D}$, if $y^\star$ is a valid output for the task $\mathcal{T}$ on input distribution $\mathcal{D}$, then the failure probability of $\mathcal{M}'$ is strictly less than the failure probability of $\mathcal{D}$. In this case, we are done.

Therefore, it remains to handle the case where $y^\star$ is not a valid output of $\mathcal{T}$ on the distribution $\mathcal{D}$. Then it must be the case that $\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\boldsymbol{S}$ is heavy$] \leq \beta$, as otherwise, $\mathcal{M}$ would have too high of a failure probability. In this case, the failure probability of $\mathcal{M}'$ is at most the failure probability of $\mathcal{M}$ plus the probability that $\boldsymbol{S}' \sim \mathcal{D}^m$ is heavy. The first quantity is at most β, and the second at most $c \cdot \beta$ by Claim 8.8, giving the desired bound. $\qquad\square$

Next, we observe that this notion of privacy allows for solving the FINDLIGHTELEMENT task:

Fact 8.11. *For any $y^\star \in Y$ and $n \geq 40$, let $\mathcal{M} : X^n \to Y$ be the algorithm that does the following on input S.*

1. *If S is not heavy, it outputs $y^\star$.*

2. *If S is heavy, an arbitrary element that appears $\leq n/2$ times in S if one exists. Otherwise, also output $y^\star$.*

Then, $\mathcal{M}$ is $\mathcal{P}_{heavy}$-private and solves FINDLIGHTELEMENT with failure probability $\exp(-\Omega(n))$.

8.2 Removing Axiom 2

Removing this axiom leads to the easiest analysis, because we can just make everything private.

Definition 18 (The privacy measure that allows all algorithms). *We define $\mathcal{P}_{\text{all}}$ to be the privacy measure for which $\mathcal{P}_{\text{all}}(\mathcal{M}) = 0$ for all algorithms $\mathcal{M}$.*

Claim 8.12. $\mathcal{P}_{\text{all}}$ *satisfies Axioms 1, 3 and 4*

Proof. This is immediate. $\qquad\square$

Furthermore, since $\mathcal{P}_{\text{all}}$ allows all algorithms, it trivially allows algorithms solving FINDELEMENT (such as the algorithm from Fact 8.6).

8.3 Replacing Axiom 3 with linear composition

Here, we will show that even if we don't fully remove Axiom 3 but weaken it to *linear composition* (i.e. $c = 1$ in Axiom 3), it still allows for an ill-behaved notion of privacy. This measure of privacy will use is a scaling of *junta*-size.

Definition 19 (Juntas). *For any $k \in [n]$, an algorithm $\mathcal{M} : X^n \to Y$ is a k-junta if there exists a size-k $I \subseteq [n]$ and (possibly randomized) function $f : X^k \to Y$ for which, on any input $S \in X^n$, the output distribution $\mathcal{M}(S)$ is equal to that of $f(S_I)$ where*

$$S_I := (S_{I_1}, \ldots, S_{I_k}).$$

For example, the algorithm $\mathcal{M} : X^n \to X^2$ which outputs the first and last element of its dataset is a 2-junta.

Definition 20 (Junta privacy). *We define the* junta privacy measure, $\mathcal{P}_{\text{junta}}$ *as follows: For any $\mathcal{M} : X^n \to Y$, we set $\mathcal{P}_{\text{junta}}(\mathcal{M}) = 2k/n$ where k is the minimum value for which $\mathcal{M}$ is a k-junta.*

Claim 8.13. $\mathcal{P}_{\text{junta}}$ *satisfies Axioms 1, 2 and 4 and also Axiom 3 with $c = 1$.*

Proof. Axiom 1 is immediate since the definition of junta size is maintained under reorderings of the sample and remappings of the domain. Axiom 2 holds because $\mathcal{P}_{\text{junta}}$-private algorithm can only depend on half of their dataset, so the adversary cannot guess 0.9-fraction of the points. For Axiom 4 given any $\mathcal{M} : X^n \to Y$ and any $m \geq n$, we construct $\mathcal{M}' : X^m \to Y$ to just run $\mathcal{M}$ on the first n points of its size-m sample. If $\mathcal{M}$ is a k-junta, then $\mathcal{M}'$ will still be a k-junta, so $\mathcal{P}_{\text{junta}}(\mathcal{M}') \leq \mathcal{P}_{\text{junta}}(\mathcal{M}) \cdot \frac{n}{m}$. Furthermore, the output distribution of $\mathcal{M}'$ given a draw from $\mathcal{D}^m$ is identical to that of $\mathcal{M}$ given a draw of $\mathcal{D}^n$, so $\mathcal{M}'$ is $(\beta, \beta' = \beta)$-equivalent to $\mathcal{M}$ for any choice of β.

Finally, Axiom 3 with $c = 1$ holds because the composition of ℓ many k-juntas is always an ℓk junta. $\qquad\square$

Next, we observe this definition of privacy allows for solving FINDELEMENT.

Fact 8.14. *The algorithm $\mathcal{M} : X^n \to Y$ which, on input S, outputs S_1 both is $\mathcal{P}_{junta}$-private for any $n \geq 2$ and solves FINDELEMENT with failure probability 0.*

8.4 Removing Axiom 4

If we remove Axiom 4, we can just use a rescaled version of Definition 20. This rescaling converts linear composition to strong composition, in exchange for losing linear scalability.

Definition 21 (Square Root Junta privacy). *We define the* square root junta privacy measure, $\mathcal{P}_{\sqrt{\text{junta}}}$ *as follows: For any $\mathcal{M} : X^n \to Y$, we set $\mathcal{P}_{\text{junta}}(\mathcal{M}) = \sqrt{2k/n}$ where k is the minimum value for which $\mathcal{M}$ is a k-junta.*

Claim 8.15. $\mathcal{P}_{\sqrt{\text{junta}}}$ *satisfies Axioms 1 and 2 and also Axiom 3 with $c = 1/2$.*

Proof. Axioms 1 and 2 hold exactly as they did in Claim 8.13. For Axiom 3, we once again use that the composition of any ℓ many k-juntas is a ℓk junta. In this case, if $\mathcal{M}'$ is the composed algorithm and $p = \sqrt{2k/n}$ is the original privacy level, then,

$$\mathcal{P}_{\sqrt{\text{junta}}}(\mathcal{M}') \leq \sqrt{2k\ell/n} = \sqrt{\ell} \cdot p. \qquad\square$$

Once again, we have that this notion of privacy allows for solving the FINDELEMENT task.

Fact 8.16. *The algorithm $\mathcal{M} : X^n \rightarrow Y$ which, on input S, outputs S_1 both is $\mathcal{P}_{\sqrt{junta}}$-private for any $n \geq 2$ and solves* FINDELEMENT *with failure probability 0.*

8.5 DP-hardness of FindElement and FindLightElement

Lastly, we show that neither FINDELEMENT nor FINDLIGHTELEMENT have DP algorithms using less than $\approx \sqrt{|X|}$ samples. This completes the proof of Theorem 3.

Claim 8.17 (DP-hardness of FINDELEMENT). *For any domain X and $n \leq \sqrt{|X|}/100$, there is no $(\varepsilon = 1, \delta = 1/(10n))$-algorithm $\mathcal{M} : X^n \rightarrow X$ which solves* FINDELEMENT *with failure probability at most $1/2$.*

Proof. For any algorithm $\mathcal{M} : X^n \rightarrow X$, let us define

$$p(\mathcal{M}) \coloneqq \Pr_{\boldsymbol{S} \sim X^n}[\mathcal{M}(\boldsymbol{S}) \in S].$$

We first argue that any $\mathcal{M}$ that solves FINDELEMENT with failure probability at most $1/2$ satisfies $p(\mathcal{M}) \geq 1/3$.

For this, draw $\boldsymbol{X}'$ to a uniform size $|X|/10$ subset of $|X|$. Then, let us consider the average failure probability of $\mathcal{M}$ over a $\boldsymbol{\mathcal{D}} \coloneqq \mathrm{Unif}(\boldsymbol{X}')$,

$$\mathop{\mathbb{E}}_{\boldsymbol{X}' \sim \binom{X}{|X|/10}} \left[\mathop{\mathbb{E}}_{\boldsymbol{S} \sim \mathrm{Unif}(\boldsymbol{X}')^n} [\Pr[\mathcal{M}(\boldsymbol{S}) \notin \boldsymbol{X}']] \right] \leq 1/2,$$

where the $1/2$ upper bound is because, if $\mathcal{M}$ has a failure probability of at most $1/2$ on any single distribution, it also has a failure probability of at most $1/2$ on average over any set of distributions. Next, will switch the order of the above expectation:

$$\mathop{\mathbb{E}}_{\boldsymbol{X}' \sim \binom{X}{|X|/10}} \left[\mathop{\mathbb{E}}_{\boldsymbol{S} \sim \mathrm{Unif}(\boldsymbol{X}')^n} [\Pr[\mathcal{M}(\boldsymbol{S}) \notin \boldsymbol{X}']] \right] = \mathop{\mathbb{E}}_{\boldsymbol{S} \sim \mathrm{Unif}(X)^n} \left[\mathop{\mathbb{E}}_{\boldsymbol{X}' | \boldsymbol{S}} \left[\Pr[\mathcal{M}(\boldsymbol{S}) \notin \boldsymbol{X}'] \right] \right].$$

Observe that conditioning on $\boldsymbol{S} = S$ conditions on all the values in S being part of $\boldsymbol{X}'$. In particular, for any $x \notin S$, if we draw $\boldsymbol{X}' \mid \boldsymbol{S} = S$ the probability that $x \in \boldsymbol{X}'$ is at most $1/10$. Therefore,

$$\mathop{\mathbb{E}}_{\boldsymbol{S} \sim \mathrm{Unif}(X)^n} \left[\mathop{\mathbb{E}}_{\boldsymbol{X}' | \boldsymbol{S}} \left[\Pr[\mathcal{M}(\boldsymbol{S}) \notin \boldsymbol{X}'] \right] \right] \geq (1 - p(\mathcal{M})) \cdot 0.9.$$

Hence, we have the desired condition that $p(\mathcal{M}) \geq 1/3$.

Finally, we'll show that no algorithm with $p(\mathcal{M}) \geq 1/3$ can be (ε, δ)-DP. For this, we will argue that for any such $\mathcal{M}$, there exists some neighboring datasets S, S' and element x s.t.

$$\Pr[\mathcal{M}(S) = x] \leq \frac{1}{100n} \qquad \text{and} \qquad \Pr[\mathcal{M}(S') = x] \geq \frac{1}{5n}. \tag{9}$$

We show the existence of such an S, S', x via the probabilistic method: Let us draw $\boldsymbol{S} \sim \mathrm{Unif}(X)^n$, $\boldsymbol{x} \sim \mathrm{Unif}(X)$, and $\boldsymbol{i} \sim \mathrm{Unif}([n])$ independently. Then, we set $\boldsymbol{S}'$ to be identical to $\boldsymbol{S}$ except with the $\boldsymbol{i}^{\text{th}}$ element replaced with $\boldsymbol{x}$. We will show that,

$$\Pr_{\boldsymbol{S}, \boldsymbol{x}} \left[\Pr_{\text{randomness of } \mathcal{M}} [\mathcal{M}(\boldsymbol{S}) = \boldsymbol{x}] \geq 1/100n \right] \leq 1/100n \tag{10}$$

For this, we use that $\boldsymbol{S}$ and $\boldsymbol{x}$ are independent, meaning

$$\mathop{\mathbb{E}}_{\boldsymbol{S},\boldsymbol{x}}\left[\Pr_{\text{randomness of }\mathcal{M}}[\mathcal{M}(\boldsymbol{S}) = \boldsymbol{x}]\right] = \Pr[\mathcal{M}(\boldsymbol{S}) = \boldsymbol{x}] \leq \frac{1}{|X|}.$$

As long as $|X| \geq (100n)^2$, we obtain Equation (10) by Markov's inequality.

Next, we will show that,

$$\Pr_{\boldsymbol{S}',\boldsymbol{x}}\left[\Pr_{\text{randomness of }\mathcal{M}}[\mathcal{M}(\boldsymbol{S}') = \boldsymbol{x}] \geq 1/5n\right] \geq 1/10n. \tag{11}$$

For this, observe the distribution of $\boldsymbol{S}'$ is simply $\text{Unif}(X)^n$ and that $\boldsymbol{i}$ is independent of $\boldsymbol{S}'$. Then, by the definition of $p(\mathcal{M})$,

$$\Pr_{\boldsymbol{S}'\sim\text{Unif}(X)^n,\boldsymbol{i}\sim\text{Unif}([n])}\left[\mathcal{M}(\boldsymbol{S}') = \boldsymbol{S}'_{\boldsymbol{i}}\right] \geq \frac{p(\mathcal{M})}{n} \geq \frac{1}{3n}.$$

Using the fact that $\boldsymbol{S}'_{\boldsymbol{i}} = \boldsymbol{x}$, we can rewrite this as

$$\mathop{\mathbb{E}}_{\boldsymbol{S}',\boldsymbol{x}}\left[\Pr_{\text{randomness of }\mathcal{M}}[\mathcal{M}(\boldsymbol{S}') = \boldsymbol{x}]\right] = \Pr[\mathcal{M}(\boldsymbol{S}') = \boldsymbol{x}] \geq \frac{1}{3n}.$$

The above inequality implies Equation (11) by reverse Markov. Combining Equation (11) and Equation (10), we have that with nonzero probability over $\boldsymbol{S}, \boldsymbol{S}', \boldsymbol{x}$, that Equation (9) holds. Since $\boldsymbol{S}$ and $\boldsymbol{S}'$ are guaranteed to be neighbors, this implies $\mathcal{M}$ is not (ε, δ)-DP. $\qquad\square$

We will prove a similar bound for FINDLIGHTELEMENT. This proof will use a reduction to Claim 8.17.

Claim 8.18 (DP-hardness of FINDLIGHTELEMENT). *For any domain X and $n \leq \frac{\sqrt{|X|-1}}{100}$, there is no $(\varepsilon = 1, \delta = 1/(10n))$-algorithm $\mathcal{M} : X^n \to X$ which solves* FINDLIGHTELEMENT *with failure probability at most $1/2$.*

Proof. Suppose we had an algorithm $\mathcal{M} : X^n \to X$ that was (ε, δ)-DP and solved FINDLIGHTELEMENT with failure probability at most $1/2$ over domain X. Let X' be created by removing any one element, $x^\star$ from X (i.e. $X' \cup \{x^\star\}$). We will show that the algorithm $\mathcal{M}' : (X')^n \to X'$ in Figure 3 is (ε, δ)-DP and solves FINDELEMENT with failure probability at most $1/2$. The desired result then follows from Claim 8.17.

We first show that $\mathcal{M}'$ solves FINDELEMENT with failure probability at most $1/2$. Observe that for any distribution $\mathcal{D}'$ over X', if we draw $\boldsymbol{S}' \sim (\mathcal{D}')^n$ and then create the sample $\boldsymbol{S}$ as in Figure 3, then the distribution of $\boldsymbol{S}$ is $\mathcal{D}^n$ where $\mathcal{D}$ is the distribution that puts 0.7 mass on $x^\star$ and the remaining 0.3 mass is distributed according to $\mathcal{D}'$. Hence, since $\mathcal{M}$ solves FINDLIGHTELEMENT with failure probability at most $1/2$, the output $\boldsymbol{x} = \mathcal{M}(\boldsymbol{S})$ must be some element satisfying $\mathcal{D}'(\boldsymbol{x}) > 0$ with probability at least $1/2$. Therefore $\mathcal{M}'$ solves FINDELEMENT with failure probability at most $1/2$

Lastly, we show that $\mathcal{M}'$ is (ε, δ)-DP whenever $\mathcal{M}$ is (ε, δ)-DP. For any $Y' \subseteq X'$ and $(S^{(1)})', (S^{(2)})' \in (X')^n$ differing in one coordinate, we wish to show that

$$\Pr[\mathcal{M}'((S^{(1)})') \in Y'] \leq e^\varepsilon \cdot \Pr[\mathcal{M}'((S^{(2)})') \in Y'] + \delta.$$

35

Input: An (ε, δ)-DP algorithm $\mathcal{M} : (X)^n \to (X)$, solving FINDLIGHTELEMENT with failure probability at most β, single element $x^\star \in X$, and sample $S' \in (X')^n$ where $X' := X \setminus \{x^\star\}$.

Output: The output of an algorithm $\mathcal{M}'(S')$ where $\mathcal{M}' : (X')^n \to (X')$ is (ε, δ)-DP and solves FINDELEMENT with failure probability at most β.

1. Draw $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \sim \text{Ber}(0.7)$.
2. Construct the sample $\boldsymbol{S} \in X^n$, by setting, for all $i \in [n]$,

$$\boldsymbol{S}_i := \begin{cases} x^\star & \text{if } \boldsymbol{b}_i = 1 \\ S'_i & \text{otherwise.} \end{cases}$$

3. Let $\boldsymbol{x} = \mathcal{M}(\boldsymbol{S})$. If $\boldsymbol{x} \neq x^\star$, output $\boldsymbol{x}$. Otherwise, output a uniform element of X'.

Figure 3: Reduction from FINDELEMENT to FINDLIGHTELEMENT

Since the choices of $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n$ in Figure 3 are independent of the input sample, it suffices to show that for any fixed choice of $b := b_1, \ldots, b_n$,

$$\Pr[\mathcal{M}'((S^{(1)})') \in Y' \mid \boldsymbol{b} = b] \leq e^\varepsilon \cdot \Pr[\mathcal{M}'((S^{(2)})') \in Y' \mid \boldsymbol{b} = b] + \delta. \tag{12}$$

Fix any choice of b. Then, let $S^{(1)}$ and $S^{(2)}$ respectively be the samples S constructed in Figure 3 when given inputs $(S^{(1)})'$ and $(S^{(2)})'$ respectively, and when $\boldsymbol{b} = b$. Then, we observe that $S^{(1)}$ and $S^{(2)}$ also differ in at most one coordinate. Therefore, since $\mathcal{M}$ is (ε, δ)-DP, for any $Y \subseteq X$,

$$\Pr[\mathcal{M}((S^{(1)})) \in Y] \leq e^\varepsilon \cdot \Pr[\mathcal{M}((S^{(2)})) \in Y] + \delta.$$

Since the output of $\mathcal{M}'(S')$ is formed by postprocessing the output of $\mathcal{M}(S)$, the above implies Equation (12) holds for all fixed values of b. This means that $\mathcal{M}'$ is indeed (ε, δ)-DP, and the desired result follows from the impossibility result of Claim 8.17. □

9   Acknowledgments

References

[AACM+22] 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. 1

[ABS24]   Maryam Aliakbarpour, Mark Bun, and Adam Smith. Optimal hypothesis selection in (almost) linear time. *Advances in Neural Information Processing Systems*, 37:141490–141527, 2024. 6.3

[ACG⁺16]   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. 1

[App]   Differential privacy: Technical overview. Technical report, Apple Inc. White paper; documents Apple's local DP deployment and budgets. Accessed Aug 19, 2025. 1

[Ara24]   Heshan Aravinda. Entropy-variance inequalities for discrete log-concave random variables via degree of freedom. *Discrete Mathematics*, 347(1):113683, 2024. 8.10

[BBG18]   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. 3.2, 4

[BDRS18]   Mark Bun, Cynthia Dwork, Guy N Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 74–86, 2018. 6, 8, B.3

[BF16]   Raef Bassily and Yoav Freund. Typical stability. *arXiv preprint arXiv:1604.03336*, 2016. (document), 1

[BGH⁺23]   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, page 520–527, New York, NY, USA, 2023. Association for Computing Machinery. 2, 3.1, 3.2, 4, 6, 7, 22, A, A, A.1, A.2, A.2

[BNS16]   Mark Bun, Kobbi Nissim, and Uri Stemmer. Simultaneous private learning of multiple concepts. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, page 369–380, New York, NY, USA, 2016. Association for Computing Machinery. 6

[BS16]   Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of cryptography conference*, pages 635–658. Springer, 2016. 4

[Can20]   Clément L. Canonne. A short note on learning discrete distributions, 2020. 6.2

[CE09]   Joshua N Cooper and Robert B Ellis. Linearly bounded liars, adaptive covering codes, and deterministic random walks. *arXiv preprint arXiv:0909.0029*, 2009. 8.9

[CLN⁺16]   Rachel Cummings, Katrina Ligett, Kobbi Nissim, Aaron Roth, and Zhiwei Steven Wu. Adaptive learning with robust generalization guarantees. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*,

volume 49 of *Proceedings of Machine Learning Research*, pages 772–814, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. 22

[DKM⁺06]   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. 1

[DKY17]   Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017. 1

[DL01]   Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2001. 6.3

[DMNS06]   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. 1

[DR14]   Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, August 2014. 1, 1, 3.2, 5.2, 7, 5, A.1

[DR16]   Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016. 4

[DRS22]   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. 4

[DRV10]   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. 1, A.3

[Fel17]   Vitaly Feldman. A general characterization of the statistical query complexity. In *Conference on learning theory*, pages 785–830. PMLR, 2017. 2

[HWR13]   Robert Hall, Larry Wasserman, and Alessandro Rinaldo. Random differential privacy. *Journal of Privacy and Confidentiality*, 4(2), 2013. (document), 1, 4

[ILPS22]   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. 4, 23

[KKMN09]   Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing search queries and clicks privately. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, page 171–180, New York, NY, USA, 2009. Association for Computing Machinery. 6

[KKMV23]   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. 6, 4

[KL10]   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. 4

[KOV15]   Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR, 2015. 1

[LWX23]   Ao Liu, Yu-Xiang Wang, and Lirong Xia. Smoothed differential privacy. *Transactions on Machine Learning Research*, 2023. (document), 1, 4

[Mir17]   Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017. 4, 24, B.1

[MP25]   Arnaud Marsiglietti and Puja Pandey. A note on statistical distances for discrete log-concave measures. *Statistics & Probability Letters*, 216:110257, 2025. 8.10

[NH+21]   Joseph P Near, Xi He, et al. Differential privacy for databases. *Foundations and Trends® in Databases*, 11(2):109–225, 2021. 1

[Sie01]   Alan Siegel. Median bounds and their application. *Journal of Algorithms*, 38(1):184–236, 2001. 8.9

[Ste22]   Thomas Steinke. Composition of differential privacy & privacy amplification by subsampling. *arXiv preprint arXiv:2210.00597*, 2022. 4

[Su24]   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. 4

[Vad17]   Salil Vadhan. The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, pages 347–450. Springer, 2017. 1

[WBK19]   Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled renyi differential privacy and analytical moments accountant. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1226–1235. PMLR, 16–18 Apr 2019. B.2

[WLF16]   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. (document), 1, 4

[XZA+23]   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. 1

A From TV-Stability to DP

In this section, we give the proof of Lemma 6.1 which we recall below.

Lemma 6.1. *There is a universal constant $1 > \rho^\star > 0$ such that if $\mathcal{M} : X^n \to Y$ is an $\rho^\star$-TV-Stable algorithm, then there exists a $(\beta, 5\beta)$-equivalent algorithm $\mathcal{M}' : X^m \to Y$ which is (ε, δ)-differentially private using*

$$m = n \cdot O\left(\log(1/\beta) \cdot \frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon} \right) \; samples.$$

We firs recall the following notions of sample-perfect generalization and replicability:

Definition 22 (Sample perfectly generalization, [CLN+16, BGH+23]). *An algorithm $\mathcal{M} : X^m \to Y$ is said to be $(\beta, \varepsilon, \delta)$-sample perfectly generalizing, if for every distribution D over X , with probability at least $1 - \beta$ over the draw of $\boldsymbol{S}, \boldsymbol{S}' \sim D^n$ we have that for every $Y' \subseteq Y$:*

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

Definition 23 (Replicable algorithms [ILPS22]). *Let $\mathcal{M} : X^n \to Y$, be an algorithm taking a set of $S \in X^n$ and using internal randomness r. We say that coin tosses r are τ-good for A on distribution $\mathcal{D}$ if there exists a "canonical output" s_r such that*

$$\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}, r) = s_r] \geq 1 - \tau$$

We say that $\mathcal{M}$ is (ρ, τ)-replicable if, for every distribution $\mathcal{D}$, with probability at least $1 - \rho$, the coin toss $\boldsymbol{r}$ drawn by $\mathcal{M}$ is τ-good on distribution $\mathcal{D}$.

In [BGH+23], the authors show how to transform a $(\beta, \varepsilon, \delta)$-sample-perfectly generalizing (for small enough $\beta, \varepsilon, \delta$) algorithm into a $(0.1, 0.1)$-replicable one. Then then give a transformation from replicability to (ε, δ)-DP. We do the same: Lemma A.1 follows almost immediately from their result, though for completeness, we will spell out the steps. Lemma A.2 is follows a similar proof strategy as their result, but we obtain $(\beta, O(\beta))$-equivalence rather than $(\beta, O(\beta \log(1/\beta)))$-equivalence.

Lemma A.1. *There exists a constant $1 > \rho^* > 0$ such that if $\mathcal{M} : X^n \to Y$ is ρ^*-TV stable then there exists an algorithm $\mathcal{M}' : X^n \to Y$ which is $(0.1, 0.1)$-replicable. Furthermore, for any set S the distribution $\mathcal{M}'(S)$ is the same as that $\mathcal{M}(S)$.*

Lemma A.2. *Let $1 \geq \beta, \varepsilon, \delta \geq 0$ and $\mathcal{M} : X^n \to Y$ be $(0.1, 0.1)$-replicable. There exists a $(\beta, 5\beta)$-equivalent algorithm $\mathcal{M}' : X^m \to Y$, which is (ε, δ)-differentially private using*

$$m = n \cdot O\left(\log(1/\beta) \cdot \frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon} \right) \; samples.$$

Using the two above, the proof of Lemma 6.1 is straightforward.

Proof of Lemma 6.1. Let $\mathcal{M} : X^n \to Y$ be ρ-TV stable, where ρ^* is the constant of Lemma A.1. Then there exists a (β, β)-equivalent algorithm $\mathcal{M}' : X^n \to Y$ which is $(0.1, 0.1)$-replicable. The lemma then follows by applying Lemma A.2 to $\mathcal{M}'$. □

We now prove Lemma A.1, while we prove Lemma A.2 in the next subsection.

Proof of Lemma A.1. Let $\mathcal{D}$ by an arbitrary distribution over X. Since $\mathcal{M}$ is ρ-TV stable we have by Markov's inequality that:

$$\Pr_{\boldsymbol{S},\boldsymbol{S'}\sim\mathcal{D}^n}\left[d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{S}),\mathcal{M}(\boldsymbol{S'}))\geq\sqrt{\rho}\right]\leq\sqrt{\rho}.$$

However, if $d_{\mathrm{TV}}(\mathcal{M}(\boldsymbol{S}),\mathcal{M}(\boldsymbol{S'}))\leq\sqrt{\rho}$, we clearly have that for every $Y'\subseteq Y$:

$$\Pr[\mathcal{M}(\boldsymbol{S'})\in Y']=\Pr[\mathcal{M}(S)\in Y]\leq\Pr[\mathcal{M}(\boldsymbol{S'})\in Y']\pm\sqrt{\rho}.$$

As such, it's easy to see that $\mathcal{M}$ is $(\sqrt{\rho},0,\sqrt{\rho})$-sample perfectly generalizing. From this point, the proof follows from the work of [BGH$^+$23]. We can first show $\mathcal{M}$ is $(\rho^{1/4},0,\sqrt{\rho}+\rho^{1/4})$-perfectly generalizing (See Lemma 3.6 of [BGH$^+$23]). This in turns implies that there exists an algorithm $\mathcal{M'}:X^n\to Y$ which is $(10\cdot20(2\rho^{1/4}+\rho^{1/2}),\frac{1}{10})$-replicable (See Theorem 3.17 and Claim 2.21 of [BGH$^+$23]). $\square$

A.1 The proof of Lemma A.2

Our proof will follow the strategy of Theorem 3.1 of [BGH$^+$23]: the key idea is that if $\mathcal{M}$ is replicable, one can draw many random strings $\boldsymbol{r}_1,\ldots,\boldsymbol{r}_k$. For each $\boldsymbol{r}_i$, we can draw multiple fresh sets $\boldsymbol{S}_{i,1},\ldots,\boldsymbol{S}_{i,\ell}\sim D^n$ and run $\mathcal{M}(\cdot,\boldsymbol{r}_i)$ on each of these to get answers $\boldsymbol{y}_{i,j}$. In [BGH$^+$23], the authors then run DP-selection on $\mathcal{Y}=[\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{1,\ell},\ldots\boldsymbol{y}_{k,1},\ldots,\boldsymbol{y}_{k,\ell}]$ to output a frequent element from that set. Note that if a $\boldsymbol{r}_i$ was 0.1 good, then most of these runs of $\mathcal{M}(\circ,\boldsymbol{r}_i)$ will output its canonical answer $s_{\boldsymbol{r}_i}$, so there will be "frequent answer" in the dataset $\mathcal{Y}$. Correctness then follows by bounding how many of the elements in $\mathcal{Y}$ can "bad" answers to the statistical task we're trying to solve.

We proceed slightly differently, for each sets $[\boldsymbol{y}_{1,1},\ldots,\boldsymbol{y}_{1,\ell}],\ldots,[\boldsymbol{y}_{k,1},\ldots,\boldsymbol{y}_{k,\ell}]$, we run a slightly different version of DP-selection on each of these sets to get replies $\boldsymbol{h}_1,\ldots,\boldsymbol{h}_\ell$. In particular, our selection algorithm is such that it outputs $\perp$ with high probability when no element appeared many times in the dataset. Indeed, there should be few random strings r which are 0.5 replicable but with canonical answer s_r being a "bad" answer. Indeed, we can show that if $\boldsymbol{h}_i\neq\perp$, it's very likely that $\boldsymbol{h}_i$ is a "good" answer to the statistical task. As such, it suffices to return an arbitrary $\boldsymbol{h}_i$ which isn't $\perp$.

Before diving into the proof, we recall the following results:

Lemma A.3 ([DRV10])**.** *Let $\varepsilon,\delta>0$ if $\mathcal{M}_1,\ldots,\mathcal{M}_k$ are (ε,δ)-approximate DP. Then their composition $(\mathcal{M}_1,\ldots,\mathcal{M}_k)$ is $(k\varepsilon,k\delta)$-differentially private.*

Theorem 5 (Theorem 3.6 of [DR14])**.** *Let $f:X^n\to\mathbb{R}$ be a function with sensitivity 1 [8]. Then for any $O\subseteq\mathbb{R}$ and $S,S'\in X^n$ differing on single entry:*

$$\Pr_{\boldsymbol{\eta}\sim Lap(1/\varepsilon)}[f(S)+\boldsymbol{\eta}\in O]\leq e^\varepsilon\Pr_{\boldsymbol{\eta}\sim Lap(1/\varepsilon)}[f(S')+\boldsymbol{\eta}\in O].$$

Theorem 6 (DP-Selection [KKMN09, BNS16, BDRS18])**.** *There exists an (ε,δ)-differentially private algorithm DP-Select such that on input $S\in X^n$, the algorithm has the following property: If $n\geq O(\log(1/\delta)/\varepsilon)$ and the most common of the input dataset appears at least $0.2n$ more times than any other element, then DP-Select outputs this element with probability 1.*

---
[8] This means, that for any $S,S'\in X^n$ differ in a single entry $|f(S)-f(S')|\leq1$

With the above results, in mind we present our slightly modified DP-selection algorithm. Which as eluded above returns $\perp$ with high probability when no element appears 60% of the time.

Input: A set of S of n elements from X and parameters $\varepsilon, \delta > 0$.
1. Let max-freq(S) be the maximum frequency of any element in S.
2. Let
$$\mathbf{m} = \mathsf{max\text{-}freq}(S) + \boldsymbol{\eta} \text{ and } \boldsymbol{h} = \mathsf{DP\text{-}Select}(S, \varepsilon, \delta)$$
where $\boldsymbol{\eta} \sim Lap(1/\varepsilon)$.
3. Return $\boldsymbol{h}$ if $\mathbf{m} > 0.7n$ and $\perp$ otherwise.

Figure 4: Pick-Heavy : A differentially private algorithm to find an very frequent element of a dataset.

Lemma A.4. *The algorithm* Pick-Heavy *(See Figure 4) is $(2\varepsilon, 2\delta)$-DP. Furthermore if*

$$n \geq O\left(\frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon}\right)$$

the we have the following:

1. *If an element appears at least $0.6n$ times in S, the algorithm outputs $\perp$ or h at least $1 - \beta$.*

2. *If an element h appears at least $0.8n$ times in S, the algorithm outputs h with probability at least $1 - \beta$.*

3. *If no element appears more than $0.6n$ times in h, the algorithm outputs $\perp$ with probability at least $1 - \beta$.*

Proof. First, note that since, counting the maximal frequency of an element in S is a sensitivity 1 query, we have by Theorem 5 that line one line 2 computing $\boldsymbol{m}$ is done in $(\varepsilon, 0)$-differentially private way. Furthermore, since DP-Select is an (ε, δ)-differentially private algorithm, we have that line 2, by Lemma A.3, corresponds to running a $(2\varepsilon, 2\delta)$-DP algorithm to get the pair $(\boldsymbol{m}, \boldsymbol{h})$. Finally line 3 is just post-processing this pair, so we can conclude that the algorithm is $(2\varepsilon, 2\delta)$.

We now turn to the correctness of the algorithm. It is well known (See Fact 3.7 in [DR14]) that:

$$\Pr_{\boldsymbol{\eta} \sim Lap(1/\varepsilon)}[|\boldsymbol{\eta}| \geq t/\varepsilon] = e^{-t}.$$

Let $h^\star$ be the most frequent element in S. If $h^\star$ appears less than $0.6n$ times, the algorithm will return $\perp$ as long as $\boldsymbol{\eta} < 0.1n$. The proof of the Item (3) immediately follows from our choice of $n \geq O(\log(1/\beta)/\varepsilon)$.

Now assume $h^\star$ appears at least $0.6n$ times. Then, $h^\star$ appears at least $0.2n$ more times than every other element in S. Whenever this happens, the algorithm either returns $\perp$ or $h^\star$. From this Item (1) immediately follows.

Finally if $h^\star$ appears at least $0.8n$ times. By our choice of $n \geq O(\log(1/\beta)/\varepsilon)$, we have $\boldsymbol{\eta} > -0.1n$, with probability $1 - \beta$. As such the algorithm will return $h^\star$ with probability at least $1 - \beta$, proving Item (2). $\qquad\square$

> **Input:** A replicable algorithm $\mathcal{M} : X^n \to Y$. Sample to an unknown distribution $\mathcal{D}$ over X.
> And parameters $\varepsilon, \delta, \beta > 0$.
> **Initialization:** Let $k = O(\log(1/\beta))$ and $\ell = O(\frac{\log(\beta^{-1} \cdot \delta^{-1})}{\varepsilon})$
>
> 1. Draw $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_k$ random strings for $\mathcal{M}$.
> 2. For each $i \in [k]$:
>   - Draw sets $\boldsymbol{S}_{i,1}, \ldots, \boldsymbol{S}_{i,\ell}$ each made of n from $\mathcal{D}$.
>   - Let $\boldsymbol{T}^i = [\boldsymbol{y}_{i,1}, \ldots, \boldsymbol{y}_{i,\ell}]$ where $\boldsymbol{y}_{i,j} = \mathcal{M}(\boldsymbol{S}_{i,j}, \boldsymbol{r}_i)$.
>   - Run Pick-Heavy$(\boldsymbol{T}^i, \varepsilon/2, \delta/2)$ to get $\boldsymbol{h}^i$.
> 3. Output $\boldsymbol{h}^\star$ as a uniformly random $\boldsymbol{h}^i$ which isn't $\perp$. If no such $\boldsymbol{h}^i$ exists, output a uniformly random element of Y.

Figure 5: Replicability to DP reduction

Using Pick-Heavy, we can now prove Lemma A.2.

Proof of Lemma A.2. Given $\mathcal{M}$, we consider the algorithm $\mathcal{M}'$ obtained by using the transformation of figure Figure 5. We first prove (ε, δ)-differential privacy. For a given draw of $\boldsymbol{r}_1 \ldots, \boldsymbol{r}_k$, two neighboring input dataset S, S' can differs in at most one of the $\boldsymbol{y}_{i,j}$. Which means there is a unique $i^\star$ where S, S' can differ on $\boldsymbol{T}^{i^\star}$. By Lemma A.4, we have that each call to Pick-Heavy is $(2\varepsilon, 2\delta)$-differentially private. So by our choice of parameters for any $O \subseteq X \cup \{\perp\}$ we have:

$$\Pr[\boldsymbol{h}^{i^\star} \in O \text{ for } S] \le e^\varepsilon \Pr[\boldsymbol{h}^{i^\star} \in O \text{ for } S'] + \delta.$$

On the other hand, all other $\boldsymbol{h}^j$, $j \neq i^\star$ are identically distributed. So clearly on line 3 for any $O \subseteq X$ we have

$$\Pr[\boldsymbol{h}^\star \in O \text{ for } S] \le e^\varepsilon \Pr[\boldsymbol{h}^* \in O \text{ for } S'] + \delta.$$

We now turn the correctness of the algorithm. We assume $\beta \le 0.2$, as otherwise the bound of 5β on the correctness probability is trivial. Fix a distribution $\mathcal{D}$ over X and let Y^* be the set of "good answers" to the task $\mathcal{T}$ under the distribution $\mathcal{D}$.

We let C be the set of possible coin draws for $\mathcal{M}$ and $\mathcal{C}$ denote the distribution of the coin draws. Since $\mathcal{M}$ is replicable, we have if we draw a random string $\boldsymbol{r} \sim \mathcal{C}$, with probability at least 0.9, it is 0.9 good. Hence let G be the set of 0.9 good strings in C. We denote by B the set of strings r such that $\Pr_{S \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}, \boldsymbol{r}) \notin Y^*]$. Since $\mathcal{M}$'s output is correct with probability at least $1 - \beta$, we have

$$\Pr_{\boldsymbol{r} \sim \mathcal{C}} \left[\boldsymbol{r} \in B \right] \le 2\beta.$$

Finally, we let $P = C \setminus B$, we have $\Pr_{\boldsymbol{r} \sim \mathcal{C}}[\boldsymbol{r} \in P] \ge 0.9 - 2\beta \ge 0.5$. We will consider the distribution of $\boldsymbol{h}^i$ depending on wether $\boldsymbol{r}_i \in B, P$ or $R \setminus P, B$.

- First, assume we have $r \in P$. This implies that r is 0.1-replicable and the canonical answer s_r is in Y^*. Given ℓ independently drawn $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_\ell \sim \mathcal{D}^n$, by a standard Chernoff bound (and making the hidden constant in ℓ large enough), we have that with probability at least $1 - \beta/40$

43

less than 0.8ℓ of the elements in $\boldsymbol{T} = [\mathcal{M}(\boldsymbol{S}^1, r), \ldots, \mathcal{M}(\boldsymbol{S}_\ell, r)]$ are s_r. Assuming this holds, by making the hidden constant in ℓ large enough, we can conclude using Lemma A.4 (Item (2)) that Pick-Heavy$(\boldsymbol{T}, \varepsilon/2, \delta/2)$ returns s_r with probability at least $1 - e^{-0.1\ell\varepsilon} \geq 1 - \beta/40$. This implies that if $\boldsymbol{r}_i \in P$, line 2 of Figure 5 returns $\boldsymbol{h}^i \in Y^*$ with probability at least $1 - \beta/20 \geq 0.99$.

- If $r \in B$, given ℓ independently drawn $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_\ell \sim \mathcal{D}^n$, in the worst case we have that Pick-Heavy$[\mathcal{M}(\boldsymbol{S}^1, r) \ldots, \mathcal{M}(\boldsymbol{S}_\ell, r)]$ always returns $\boldsymbol{h} \in Y \setminus Y^*$.

- If $r \in R \setminus (B \cup G)$. Then, we must have $\Pr_{\boldsymbol{S} \sim \mathcal{D}^n}[\mathcal{M}(\boldsymbol{S}, r) \notin Y^*] \leq 0.5$. Hence, given ℓ independently drawn $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_\ell \sim \mathcal{D}^n$, it follows by a standard Chernoff bound, that the probability there exists an element $y \in Y \setminus Y^*$ which appears at least $0.6n$ times in $\boldsymbol{T} = [\mathcal{M}(\boldsymbol{S}^1, r), \ldots, \mathcal{M}(\boldsymbol{S}_\ell, r)]$ is at most $\beta/10$. Assuming no bad answer appears more than $0.6n$ times, we consider two cases:

  - If some element h appears in $\boldsymbol{T}$ $0.6n$ times, this element must be in Y^*. In this case Pick-Heavy can only return h or $\perp$. (See Item (1) of Lemma A.4).

  - If no element appears in $\boldsymbol{T}$ more than $0.6n$ times, Pick-Heavy$(\boldsymbol{T})$ returns $\perp$ with failure probability $e^{-0.1\ell\varepsilon} \leq \beta/10$.(See Item (3) of Lemma A.4)

  This implies that if $\boldsymbol{r}_i \in R \setminus (B \cup P)$, line 2 of Figure 5 returns $\boldsymbol{h}^i \in Y \setminus Y^*$ with probability at most $\beta/5$.

From the above discussion, we have that for each $i \in [k]$ on line 3 of Figure 5:

$$\Pr[\boldsymbol{h}^i \in Y^*] \geq \Pr[\boldsymbol{r}_i \in P](1 - \beta/20) \geq 0.49$$

and

$$\Pr[\boldsymbol{h}^i \in Y \setminus Y^*] \leq \Pr[\boldsymbol{r}_i \in R \setminus (B \cup P)] \cdot \beta/5 + \Pr[\boldsymbol{r}_i \in B] \leq \beta/5 + 2\beta \leq 2.2\beta.$$

As such the probability the algorithm outputs a random value in Y, meaning every $\boldsymbol{h}^i = \perp$ is at most $(1 - 0.6)^{O(k)} \leq \beta/10$. Otherwise, the algorithm returns some $\boldsymbol{h}^i \neq \perp$ in which case the probability we return a bad answer is at most $2.2\beta/0.49 \leq 4.9\beta$. So the algorithm of Figure 5 has a failure probability of $4.9\beta + \beta/10 \leq 5\beta$. $\qquad\qquad\square$

A.2 The proof of Lemma 7.1

Finally, we sketch the proof of Lemma 7.1 (recalled below for convenience). As we already we mentioned, the result is quite similar to that Theorem 6.2 of [BGH$^+$23]. The proof of Lemma 7.1 is identical, except for the last step where instead of calling the "Replicability to DP" theorem used by [BGH$^+$23] (Theorem 3.1 in [BGH$^+$23]) we instead use Lemma A.2 presented earlier in this section.

Lemma 7.1. *There is a universal constant $0.1 \geq \alpha > 0$ such that the following holds. Let $\mathcal{M} : X^n \to Y$ be $(0.1, \alpha^2/n^3)$-differentially private. Then for every $\varepsilon, \delta > 0$; there exists an (ε, δ)-differentially private algorithm $\mathcal{M}' : X^m \to Y$ solving T using*

$$m = O\Big(\frac{\log(1/\beta)^2 + \log(1/\beta)\log(1/\delta)}{\varepsilon}\Big) \cdot n^2$$

samples and which is $(\beta, 5\beta)$-equivalent to $\mathcal{M}$.

Sketch of the proof of Lemma 7.1. Let α be a enough small constant and $\mathcal{M} : X^n \to Y$ be a $(0.1, \alpha^2/n^3)$ differentially private algorithm. Following the first two steps of the proof of Theorem 6.2 of [BGH$^+$23], we can get an algorithm $\mathcal{M}' : X^m \to Y$ which is (β, β)-equivalent to $\mathcal{M}$ and using $m := O(n^2)$ samples. Furthermore, $\mathcal{M}'$ is $(0.1, 0.1)$-replicable. We can now apply Lemma A.2 to get a $(\beta, 5\beta)$-equivalent algorithm $\mathcal{M}^\star = X^{m'} \to Y$ which is (ε, δ)-differentially private using

$$m' = m \cdot O\left(\log(1/\beta) \cdot \frac{\log(1/\beta) + \log(1/\delta)}{\varepsilon}\right) \text{ samples.} \qquad \square$$

B   Rényi DP

In this section, we show that Rényi DP fits our axioms, that our axioms imply Rényi DP (up to a $\log \log |Y|$ factor), and that this additional factor is necessary.

Definition 24 (Rényi differential privacy, [Mir17])**.** *For any $\alpha > 1, \varepsilon > 0$, an algorithm $\mathcal{M} : X^n \to Y$ is (α, ε)-RDP if, for every $S, S' \in X^n$ differing in only one of the n coordinates,*

$$\frac{1}{\alpha - 1} \cdot \ln\left(\mathop{\mathbb{E}}_{\boldsymbol{y} \sim \mathcal{M}(S')}\left[\left(\frac{\Pr[\mathcal{M}(S) = \boldsymbol{y}]}{\Pr[\mathcal{M}(S') = \boldsymbol{y}]}\right)^\alpha\right]\right) \leq \varepsilon.$$

We recall some facts about Rényi-DP:

Fact B.1 ([Mir17])**.** *For any $\alpha > 0$, Réyni-differential privacy has the following properties:*

- *Let $\alpha > 1$. If $\mathcal{M} : X^n \to Y$ is (α, ε)-RDP, then for any $S, S' \in X^n$ differing in exactly one coordinate and $Y' \subseteq Y$ we have:*

$$\Pr[\mathcal{M}(S) \in Y] \leq \left(e^\varepsilon \Pr[\mathcal{M}(S') \in Y']\right)^{\frac{\alpha-1}{\alpha}}$$

- *(Post-Processing) If $\mathcal{M} : X^n \to Y$ is (α, ε)-RDP, then for any $\mathcal{M}' : Y \to Z$, the algorithm $\mathcal{M}' \circ \mathcal{M}$ is (α, ε)-RDP.*

- *(Non-Adaptive Composition) If $\mathcal{M}, \mathcal{M}'$ are (α, ε)-RDP algorithm then their non-adaptive composition $(\mathcal{M}, \mathcal{M}')$ is $(\alpha, 2\varepsilon)$-RDP.*

- *If M is $(\varepsilon, 0)$-DP, then it is (α, ε)-RDP.*

B.1   Axioms imply Rényi DP

We prove our axioms imply Rényi DP, albeit with now a (small) dependency on $|Y|$. In particular we prove the following:

Theorem 7. *Let $\mathcal{P}$ be any privacy measure satisfying Axioms 1 to 4 and $\mathcal{M} : X^n \to Y$ be any $\mathcal{P}$-private algorithm. For any $\varepsilon, \delta, \beta > 0$ and, c the constant in Axiom 3 and p the polynomial in Axiom 4, define*

$$m' := \tilde{O}\left(\frac{r^2 \cdot n^2 \cdot \log(\log(|Y|))}{\beta^2 \cdot \varepsilon}\right) \text{ where } r = \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}}\right).$$

there is an $(2, \varepsilon)$-Rényi differentially private algorithm $\mathcal{M}'$ using m' samples that is $(\beta, \beta' := O(\beta))$-equivalent to $\mathcal{M}$.

The proof will rely on the following lemma.

Lemma B.2. *Given a $(0.1, 0.1)$-replicable algorithm $\mathcal{M} : X^n \to Y$ solving a statistical task $\mathcal{T}$ with failure probability β, there exists is $(2, \varepsilon)$-Rényi differentially private algorithm $\mathcal{M}' : X^m \to Y$ solving task $\mathcal{T}$ with failure probability 5β and using*

$$m = n \cdot \tilde{O}\left(\frac{\log^2(\beta^{-1})\log(\log(|Y|))}{\varepsilon} \right) \; samples.$$

We delay the proof of the above, from which Theorem 7 follows easily.

Proof of Theorem 7. Let $\rho^\star$ be the constant of Lemma A.1. We first apply Theorem 4 to $\mathcal{M}$ to get a (β, β')-equivalent, $\beta' = O(\beta)$, and $\rho^\star$-TV-Stable algorithm $\mathcal{M}' : X^m \to Y$. We have that $\mathcal{M}'$ uses

$$m = \tilde{O}(r^2 \cdot n^2 / \beta^2) \text{ where } r = \max\left(n \cdot p(n, 1/\beta), n^{\frac{1}{1-c}} \right) \; samples.$$

The proof then follows by applying Lemma B.2 to $\mathcal{M}'$. $\qquad\square$

Theorem 8 (Stable selection from [BDRS18]). *For every $0 < \rho < 1$, and $\omega \geq 1/\sqrt{\rho}$, there exists an algorithm $\mathcal{M} : X^n \to X$ for the such that for every $1 < \alpha < \omega$, $\mathcal{M}$ is $(\alpha, \rho\alpha)$-Rényi-DP. The algorithm uses*

$$n = O\left(\omega \log\left(1 + \frac{\log(|Y|)}{\rho\omega} \right) \right) \; samples,$$

and if the most common element appears at least $0.2n$ more times than any other element, the algorithm outputs this element with probability at least $2/3$.

We will need the following corollary:

Corollary B.3. *Let $0 < \varepsilon < 1$, there exists an $(2, \varepsilon)$-Rényi differentially private algorithm RDP-Select such that on input $S \in X^n$, the algorithm has the following property: If*

$$n \geq \tilde{O}\left(\frac{\log(1/\beta)}{\varepsilon} \log(\log(|Y|)) \right),$$

and the most common of the input dataset appears at least $0.2n$ more times than any other element, the algorithm RDP-Select outputs that element with probability at least $1 - \beta$.

Proof. Let $k = O(\log(1/\beta)), \rho = \varepsilon/2k, \omega = O(1/\rho)$ and

$$n = O\left(\omega \log\left(1 + \frac{\log(|Y|)}{\rho\omega} \right) \right) = \tilde{O}\left(\frac{\log(1/\beta)}{\varepsilon} \log(\log(|Y|)) \right).$$

Let $\mathcal{M} : X^n \to X$ be the algorithm obtained from Corollary B.3 using ρ, ω as above. We have that $\mathcal{M}$ is $(2, \varepsilon/k)$-Rényi DP. We will consider the following algorithm $\mathcal{M}' : X^n \to X$, which on input $S \in X^n$ runs k copies of $\mathcal{M}(S)$ in parallel to obtain $(\boldsymbol{h}_1, \dots, \boldsymbol{h}_k)$. And finally outputs the most frequent element among $(\boldsymbol{h}_1, \dots, \boldsymbol{h}_k)$ (breaking ties arbitrarily).

First, by Non-adaptive Composition and Postprocessing (See Fact B.1), it's easy to see that $\mathcal{M}'$ is $(2, \varepsilon)$-Rényi DP. Now assume, that there exists some element $h^\star$ which appears at least $0.2n$ more times than any other elements in S. It follows by a standard Chernoff bound (and setting the hidden constant in k to be large enough) that with probability at least $1 - \beta$ strictly more than half of the k copies of $\mathcal{M}$ returned $h^\star$, in which case $\mathcal{M}'$ returns $h^\star$. $\qquad\square$

With the above algorithm, the proof of Lemma B.2 follows the same pattern as the proof of Lemma A.2. We thus only sketch the proof.

Proof sketch of Lemma B.2. If we replace DP-Selection in Pick-Heavy by RDP-Selection, using Fact B.1 it's easy to see that the resulting algorithm would be $(2, 2\varepsilon)$-RDP. Furthermore, the resulting algorithm the same correctness guarantees as the one we gave in Lemma A.4 (except the error probability is now $1 - 2\beta$). Finally, given a $(0.1, 0.1)$-replicable algorithm $\mathcal{M} : X^n \to Y$ we can use the same reduction as the one of Figure 5, but with the Rényi-DP version of RDP-Selection instead, and setting $\ell = \tilde{O}\left(\frac{\log(1/\beta)}{\varepsilon} \log(\log(|Y|))\right)$. By appropriately setting the hidden constants in ℓ and k the resulting algorithm would be $(\beta, 5\beta)$-equivalent to $\mathcal{M}$, and it would use

$$n \cdot \tilde{O}\left(\frac{\log(1/\beta)^2}{\varepsilon} \log(\log(|Y|))\right) \text{ samples.} \qquad \square$$

B.2 Rényi DP satisfies the axioms

We now define our stability measure $\mathcal{P}_{\mathrm{RDP}}$ as follows:

Definition 25. *For an algorithm $\mathcal{M}$ taking n samples let ε be the smallest value such that $\mathcal{M}$ is $(2, \varepsilon)$-RDP. We set*

$$\mathcal{P}_{\mathrm{RDP}}(\mathcal{M}) = \sqrt{\varepsilon}.$$

Lemma B.4. $\mathcal{P}_{\mathrm{RDP}}$ *respects Axiom 1.*

Proof. We observe for any S, S' that differ in one coordinate and permutation $\pi : [n] \to [n]$, that $\pi(S)$ and $\pi(S')$ still differ in one coordinate. Similarly, for any $\sigma : X \to X$, $\sigma(S)$ and $\sigma(S')$ differ in (at most) one coordinate. Therefore, if $\mathcal{M}$ is $(2, \varepsilon)$-RDP, it is still $(2, \varepsilon)$-RDP after preprocessing. $\square$

To prove RDP respects Axiom 2 we will need the following fact:

Fact B.5. *Let $\mathcal{M} : X^n \to Y$, $\mathcal{A} : Y \to X^n$ be algorithms. Fix $S \in X^n, i \in [n]$ and $x \in X$. If $\mathcal{M}$ is (α, ε)-Rényi differentially private then we have:*

$$\Pr[S_i \in \mathcal{A}(\mathcal{M}(S))] \leq \left(e^\varepsilon \Pr[S_i \in \mathcal{A}(\mathcal{M}(S_{x \to i}))]\right)^{\frac{\alpha-1}{\alpha}}$$

where $S_{x \to i}$ denotes S with i-th element set to x.

Proof. S and $S_{x \to i}$ differ on at most 1 coordinate. By postprocessing (second item of Fact B.1) meaning $\mathcal{A} \circ \mathcal{M}$ is also (α, ε)-RDP. The result then follows from the first item of Fact B.1. $\square$

Lemma B.6. $\mathcal{P}_{\mathrm{RDP}}$ *respects Axiom 2.*

Proof. Let $\mathcal{D}$ an arbitrary distribution over X with $||\mathcal{D}||_\infty \leq 1/100n^2$. Let $\mathcal{M} : X^n \to Y$ be an $(2, \varepsilon)$-Rényi differentially private algorithm and let $\mathcal{A} : Y \to X^n$. We will show that:

$$\mathop{\mathbb{E}}_{\substack{\boldsymbol{S} \sim \mathcal{D}^n \\ \boldsymbol{S'} \leftarrow \mathcal{A}(\mathcal{M}(\boldsymbol{S}))}} \left[\sum_{x \in \boldsymbol{S}} \mathbb{1}[x \in \boldsymbol{S'}]\right] \leq \frac{e^{\varepsilon/2}\sqrt{n}}{10}.$$

In particular, if $\mathcal{P}_{\mathrm{RDP}}(\mathcal{M}) \leq 1$, which implies $\mathcal{M}$ is $(2,1)$-RDP, we have $\frac{e^{\varepsilon/2}\sqrt{n}}{10} \leq 0.2n$ proving the lemma. We want to bound

$$\mathbb{E}_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}(\boldsymbol{S}))}}\left[\sum_{x\in\boldsymbol{S}}\mathbb{1}[x\in\boldsymbol{S}']\right] = \sum_{i=1}^{n}\Pr_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}(\boldsymbol{S}))}}[\boldsymbol{S}_i\in\boldsymbol{S}']$$

Since $\mathcal{M}$ is $(2,\varepsilon)$-RDP, for any $i\in[n]$ and $x\in X$ we have, by Fact B.5, that:

$$\Pr_{\boldsymbol{S}\sim\mathcal{D}^n}[\boldsymbol{S}_i\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}))] \leq e^{\varepsilon/2}\left(\Pr_{\boldsymbol{S}\sim\mathcal{D}^n}[\boldsymbol{S}_i\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}_{x\to i}))]\right)^{1/2}.$$

Where $S_{x\to i}$ is the set obtained by setting the i-th element of S to x. This implies that:

$$\Pr_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}(\boldsymbol{S}))}}[\boldsymbol{S}_i\in\boldsymbol{S}'] = \Pr_{\boldsymbol{S}\sim\mathcal{D}^n}[\boldsymbol{S}_i\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}))]$$

$$\leq e^{\varepsilon/2}\left(\Pr_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{x}\sim\mathcal{D}}}[\boldsymbol{S}_i\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}_{\boldsymbol{x}\to i}))]\right)^{1/2}.$$

$$= e^{\varepsilon/2}\left(\Pr_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{x}\sim\mathcal{D}}}[\boldsymbol{x}\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}))]\right)^{1/2}.$$

Where the last line follows from the symmetry of $\boldsymbol{S}_i$ and $\boldsymbol{x}$. Hence, we have that:

$$\mathbb{E}_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}(\boldsymbol{S}))}}\left[\sum_{x\in\boldsymbol{S}}\mathbb{1}[x\in\boldsymbol{S}']\right] = \sum_{i=1}^{n}e^{\varepsilon/2}\left(\Pr_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{x}\sim\mathcal{D}}}[\boldsymbol{x}\in\mathcal{A}(\mathcal{M}(\boldsymbol{S}))]\right)^{1/2}$$

$$\leq n\cdot e^{\varepsilon/2}\left(\sup_{S\in X^n}\left(\Pr_{\boldsymbol{x}\sim\mathcal{D}}[\boldsymbol{x}\in S]\right)\right)^{1/2}.$$

Since $\|D\|_{\infty} \leq \frac{1}{100n^2}$, we have that for any $S\in X^n$, $\Pr_{\boldsymbol{x}\sim\mathcal{D}}[\boldsymbol{x}\in S] \leq \frac{n}{100n^2}$. From which we can conclude that for all $n\in\mathbb{N}$, Axiom 2 holds since:

$$\mathbb{E}_{\substack{\boldsymbol{S}\sim\mathcal{D}^n \\ \boldsymbol{S}'\leftarrow\mathcal{A}(\mathcal{M}(\boldsymbol{S}))}}\left[\sum_{x\in\boldsymbol{S}}\mathbb{1}[x\in\boldsymbol{S}']\right] \leq \frac{e^{\varepsilon/2}\sqrt{n}}{10} \leq 0.2n. \qquad \square$$

Lemma B.7. $\mathcal{P}_{\mathrm{RDP}}$ *respects Axiom 3 with composition constant $c = 1/2$.*

Proof. We define $\varepsilon' := \ell^{1/2}\varepsilon$. Say we have algorithms $\mathcal{M}^1,\ldots,\mathcal{M}^\ell$ each taking n samples and with $\ell^{-1/2} \geq \mathcal{P}_{\mathrm{RDP}}(\mathcal{M}^i)$. By definition, this implies each $\mathcal{M}^i$ is $(2,\frac{1}{\ell})$-RDP. Then, by Fact B.1 (Non-adaptive composition), we have that the algorithm $\mathcal{M}' = (\mathcal{M}^1,\ldots,\mathcal{M}^\ell)$ is $(2,1)$-Renyi differentially private. Thus, $\mathcal{P}_{\mathrm{RDP}}(\mathcal{M}') \leq 1$, meaning $\mathcal{M}'$ is $\mathcal{P}_{\mathrm{RDP}}$-private. $\qquad \square$

Before proving $\mathcal{P}_{\mathrm{RDP}}$ fits Axiom 4, we will need the following result about amplification for Rényi-differential privacy:

Theorem (Rényi Amplification by subsampling, [WBK19]) *Given an algorithm $\mathcal{M} : X^n \to Y$. Let $m \geq n$ and consider the randomized algorithm $\mathcal{M}' : X^m \to Y$ which on input $S \in X^m$ subsamples without replacement n elements from S and runs $\mathcal{M}$ the subsampled dataset. If $\mathcal{M}$ is $(2, \varepsilon)$-RDP, then $\mathcal{M}'$ is $(2, \varepsilon')$-RDP where*

$$\varepsilon' \leq \log\left(1 + 4\left(\frac{n}{m}\right)^2 \cdot (e^\varepsilon - 1)\right).$$

We can now prove the following:

Lemma B.8. $\mathcal{P}_{\mathrm{RDP}}$ *respects Axiom 4 with $p(n, \frac{1}{\beta}) = 1$. For any choice of β, the resulting algorithm $\mathcal{M}'$ is (β, β)-equivalent to $\mathcal{M}$.*

Proof. Let $\mathcal{M} : X^n \to Y$ be $\mathcal{P}_{\mathrm{RDP}}$-private. We have that $\mathcal{M}$ is $(2, 1)$-RDP. Let such that Let $k \geq n$, by applying Rényi amplification by subsampling with $m = nk$, we get an algorithm $\mathcal{M}' : X^m \to Y$ which is $(2, \varepsilon)$-RDP where

$$\varepsilon \leq \log\left(1 + 4\left(\frac{n}{m}\right)^2 \cdot (e - 1)\right) \leq \log\left(1 + 8\left(\frac{n}{m}\right)^2\right) \leq 16(n/m)^2 = \frac{16}{k^2}.$$

As such we have that $\mathcal{P}_{\mathrm{RDP}}(\mathcal{M}') \leq \sqrt{\frac{16}{k^2}} = \frac{4}{k}$.

Furthermore, $\mathcal{M}'$ simply runs $\mathcal{M}$ on a subset of its input which is subsampled without replacement. So we clearly have that the distribution $\mathcal{M}'(\boldsymbol{S}')$ where $\boldsymbol{S}' \sim \mathcal{D}^m$ and $\mathcal{M}(\boldsymbol{S})$ $\boldsymbol{S} \sim \mathcal{D}^n$ are identical for any distribution $\mathcal{D}$ over X. So $\mathcal{M}'$ and $\mathcal{M}$ are (β, β) equivalent for every choice of β. $\qquad\square$

B.3 Tightness of Theorem 7

Theorem 7 converts an algorithm $\mathcal{M} : X^n \to Y$ that is $\mathcal{P}$-private into a new an equivalent RDP algorithm with a new sample size m' depends on $\log\log|Y|$. Here, we show that $\log\log|Y|$ term is necessary.

Lemma B.9 (Separation between (ε, δ)-DP and RDP). *For every output domain Y, there is a statistical task $\mathcal{T}$ satisfying*

1. *For every ε, δ and $n := \mathrm{poly}(1/\varepsilon, \log(1/\delta), \log(1/\beta))$, there is an (ε, δ)-DP algorithm $\mathcal{M} : X^n \to Y$ solving $\mathcal{T}$ with failure probability 0.*

2. *Any $(\alpha = 2, \varepsilon = 1)$-RDP algorithm $\mathcal{M}' : X^m \to Y$ that solves $\mathcal{T}$ with failure probability $1/2$ must use $m = \Omega(\log\log|Y|)$ samples.*

Since (ε, δ)-DP fits our axioms (Theorem 2), Lemma B.9 implies that the $\log\log|Y|$ term cannot be removed from Theorem 7. This separation is based on stable selection and is similar to the lower bound from [BDRS18], though we specialize to the $\alpha = 2$ case. Lemma B.9 follows from the following claim.

Claim B.10 (Weak group privacy for RDP). *Let $\mathcal{M} : X^m \to Y$ be any $(\alpha = 2, \varepsilon)$-RDP algorithm. Then, for any S, S' differing in k coordinates and $y \in Y$,*

$$\Pr[\mathcal{M}(S') = y] \geq (\Pr[\mathcal{M}(S) = y])^{2^k} \cdot e^{-(2^k - 1)\varepsilon}.$$

Proof. We prove this by induction on k. If $k = 0$, we have the $\Pr[\mathcal{M}(S') = y] \geq \Pr[\mathcal{M}(S) = y]$ which clearly holds because S and S' must be the same.

For $k \geq 1$, let $S^{(\text{mid})}$ be a sample that differs from S' in 1 coordinate and from S in $k - 1$ coordinates. Then, by rearranging the first item of Fact B.1,

$$\Pr[\mathcal{M}(S') = y] \geq e^{-\varepsilon} \cdot \Pr[\mathcal{M}(S^{(mid)}) = y]^2.$$

We apply the inductive hypothesis and obtain

$$\begin{aligned}
\Pr[\mathcal{M}(S') = y] &\geq e^{-\varepsilon} \cdot \Pr[\mathcal{M}(S^{(mid)}) = y]^2 \\
&\geq e^{-\varepsilon} \cdot \left((\Pr[\mathcal{M}(S) = y])^{2^{k-1}} \cdot e^{-(2^{k-1} - 1)\varepsilon} \right)^2 \\
&= (\Pr[\mathcal{M}(S) = y])^{2^k} \cdot e^{-(2^k - 1)\varepsilon}.
\end{aligned}$$

$\square$

Proof of Lemma B.9. Fix any output domain Y and let $X = Y$. Consider the statistical task $\mathcal{T}$ that is defined for any distribution $\mathcal{D}$ where there is some $x^\star$ with all the mass (i.e. $\mathcal{D}(x) = 1$). For such $\mathcal{D}$, the only valid response of the task $\mathcal{D}$ is this $x^\star$. Put differently, for an algorithm $\mathcal{M} : X^n \to Y$ to solve $\mathcal{T}$ with failure probability β is equivalent to $\mathcal{M}(\{x, x, \ldots, x\}) = x$ with probability at least $1 - \beta$ for all choices of x.

The first requirement of Lemma B.9 is satisfied by the selection algorithm of Theorem 6. Therefore, all that remains is to prove a lower bound against RDP algorithms. Fix any $(\alpha = 2, \varepsilon = 1)$-RDP algorithm $\mathcal{M}' : X^m \to Y$ that solves the task $\mathcal{T}$. Then, for each $y \in Y$, we have that $\mathcal{M}'(\{y, y, \ldots, y\}) = y$ with probability at least $1/2$. Then, since any set $S \in X^m$ differs from the set $\{y, y, \ldots, y\}$ in at most m coordinates, we can apply Claim B.10 to obtain for any set S and $y \in Y$,

$$\Pr[\mathcal{M}(S) = y] \geq (e^{-1}/2)^{2^m}.$$

Since it must be the case that $\sum_{y \in Y} \Pr[\mathcal{M}(S) = y] = 1$, we have that,

$$|Y| \cdot (e^{-1}/2)^{2^m} \leq 1.$$

This implies that $m \geq \Omega(\log \log |Y|)$, as desired. $\square$