Separating Oblivious and Adaptive Differential Privacy Under Continual Observation

Mark Bun
Boston University, MA, USA

Marco Gaboardi
Boston University, MA, USA

Connor Wagaman
Boston University, MA, USA

Abstract

We resolve an open question of Jain, Raskhodnikova, Sivakumar, and Smith (ICML 2023; [17]) by exhibiting a problem separating differential privacy under continual observation in the oblivious and adaptive settings. The *continual observation* (a.k.a. *continual release*) model formalizes privacy for streaming algorithms, where data is received over time and output is released at each time step. In the oblivious setting, privacy need only hold for data streams fixed in advance; in the adaptive setting, privacy is required even for streams that can be chosen adaptively based on the streaming algorithm's output.

We describe the first explicit separation between the oblivious and adaptive settings. The problem showing this separation is based on the correlated vector queries problem of Bun, Steinke, and Ullman (SODA 2017; [4]). Specifically, we present an $(\varepsilon, 0)$-DP algorithm for the oblivious setting that remains accurate for exponentially many time steps in the dimension of the input. On the other hand, we show that every (ε, δ)-DP adaptive algorithm fails to be accurate after releasing output for only a constant number of time steps.

2012 ACM Subject Classification Security and privacy → Privacy-preserving protocols; Theory of computation → Streaming models; Theory of computation → Adversary models; Theory of computation → Theory of database privacy and security

Keywords and phrases differential privacy, continual observation, continual release, streaming algorithms, adaptive algorithms

Digital Object Identifier 10.4230/LIPIcs.FORC.2026.22

Related Version *Full Version*: https://arxiv.org/abs/2603.11029

1 Introduction

Differential privacy (DP) [10] is the standard framework for guaranteeing individual privacy when releasing statistics about a sensitive dataset. DP was initially considered in the "batch model," where a trusted server holds a static dataset and privately answers queries posed by a data analyst. However, many real-world datasets change over time, and a well-developed line of work, beginning with [5, 11], has investigated privacy guarantees for evolving datasets. In the standard formulation of this setting, at each time step a new individual's data arrives and an output is released, with accuracy evaluated on the prefix of the dataset seen so far. Privacy requires indistinguishability of the entire output sequence for every pair of input datasets that differ by one individual.

Ensuring privacy in this streaming setting (known as *continual observation* or *continual release*) is challenging, since the arrival of one person's data may affect the algorithm's output at all subsequent time steps. Additionally, privacy for this setting is most meaningful when it accounts for adaptively selected inputs, where privacy must hold against an adversary who


© Mark Bun, Marco Gaboardi, and Connor Wagaman;
licensed under Creative Commons License CC-BY 4.0

can specify the input stream in response to previous outputs from the algorithm. (Initial formulations of the continual observation model focused on the oblivious setting, where the input stream is fixed in advance but revealed to the algorithm timestep-by-timestep.)

This adaptive setting was first investigated implicitly by [25] (later by [21]) and formally defined by [17], and it is especially relevant to machine learning. Stochastic gradient descent relies on accurately estimating the cumulative sum of individual gradients, where – as in the continual observation setting – one observation arrives at each time step (i.e., one gradient is computed at each iteration) and output is released at each time step (i.e., the model is updated). Moreover, privacy for this problem must hold against adaptive inputs even when the input dataset is static, since the points at which we compute gradients are chosen based on the gradients computed so far. These connections between adaptive continual observation and private learning [25, 19] have led to an explosion of work on the continual observation setting – see [22] for a survey of some results.

In their work defining the adaptive model, [17] present two problems showing strong separations between the batch model and the oblivious continual observation model (specifically, these problems require additive error that is larger by a factor of $\widetilde{\Omega}(T^{1/3})$, where T is the number of time steps at which output is produced). They note, "It is open to separate the two continual release models in the sense our work separates the batch model and the continual release model: by providing problems that require a large error blowup in the more demanding model." That is, they ask,

> *Is there a problem separating oblivious and adaptive differential privacy under continual observation?* [17]

We answer this question in the affirmative.

1.1 Our Results

We show a problem that separates oblivious and adaptive differential privacy under continual observation. Informally, we prove the following theorem.

▶ **Theorem 1** (Theorems 7 and 8, informal). *There is a problem $\mathcal{P}^{d,T}$ parametrized by $d, T \in \mathbb{N}$ with the following properties.*

1. *For all $\varepsilon \in \left(0, \frac{3}{2}\right]$ there exists $T = 2^{\Omega(\varepsilon^4 d)}$ such that, for all sufficiently large $d \in \mathbb{N}$, there is an algorithm $\mathcal{A}$ that is $(\varepsilon, 0)$-DP under **oblivious** continual observation and accurately answers $\mathcal{P}^{d,T}$ for T time steps.*
2. *There exists some $T = O(1)$ such that, for all sufficiently large $d \in \mathbb{N}$, there is no algorithm $\mathcal{B}$ that is $(\frac{1}{5}, \frac{1}{20})$-DP under **adaptive** continual observation and accurately answers $\mathcal{P}^{d,T}$ for T time steps.*

The problem yielding this separation is inspired by the *correlated vector queries* problem of [4]. Rather than focusing on the setting where data arrives over time and the same query is asked at every time step, that paper focuses on a related but orthogonal setting where the dataset is fixed and a new query is asked at every time step. They show exponential separations between the offline, oblivious online, and adaptive online query settings. To show the latter separation, they use correlated vector queries, a search problem where, for a private dataset $x \in \{\pm 1\}^d$, each query specifies a set $V \subseteq \{\pm 1\}^d$ and asks for an output $y \in \{\pm 1\}^d$ that is α-correlated with x but nearly uncorrelated with all vectors in V (beyond correlation explained by x). In the oblivious setting, a single randomized response vector suffices to answer exponentially many such queries simultaneously. In contrast, in the adaptive setting, each new query can constrain the next answer to be nearly orthogonal to all previous answers,

forcing the mechanism to reveal fresh information about x at every step. After sufficiently many adaptive queries, this accumulated information enables reconstruction of a highly correlated estimate of x, contradicting differential privacy.

Our approach is similar, but does not follow in a black-box way from their results. In [4], each time step allows the analyst to specify an entirely new query – namely, an arbitrary collection $V \subseteq \{\pm 1\}^d$ of vectors that simultaneously constrain the next answer.

In contrast, our continual observation setting is structurally more restricted. The private dataset arrives in two stages. First, $b \in \{\pm 1\}^d$ arrives during a setup phase, during which no output is produced (or, alternatively, during which arbitrary output is produced and any answer is considered accurate). Then, during the arrival phase, a total of T vectors $v_1, \ldots, v_T \in \{\pm 1\}^d$ arrive, with only *one* new vector arriving per time step. At time t, the mechanism must output a vector $y^{(t)}$ that is α-correlated with b while remaining nearly orthogonal to the prefix $v_1, \ldots, v_t$. Thus, rather than answering a sequence of independently specified search queries (each of which may contain many new constraints), we repeatedly answer the *same* correlation task as the constraint set grows incrementally over time.

This distinction has important consequences for proving the lower bound. In [4], producing many nearly independent α-correlated vectors directly enables reconstruction of the dataset x, yielding a contradiction to differential privacy. In our setting, however, reconstructing an arbitrary estimate of b is not by itself sufficient to violate privacy. Instead, the reduction must be tailored to the continual observation model and show that the adaptive adversary can use the evolving constraint sequence to recover a specific *challenge bit* embedded in b that distinguishes the neighboring inputs. Establishing this stronger form of reconstruction also requires additional structure beyond a black-box use of the lower bound in [4].

In the oblivious setting, the entire sequence $v_1, \ldots, v_T$ is fixed in advance, so releasing b via randomized response satisfies all constraints simultaneously with high probability. In the adaptive setting, each v_{t+1} depends on previous outputs, forcing the mechanism to generate progressively less correlated and more independent views of b. We show that even under the one-constraint-per-round restriction, this adaptive process enables recovery of the challenge bit after only $O(1)$ steps, yielding the desired separation.

1.2 Related Work

As already described, our work relies heavily on the definition and open question of [17], and the work of [4]. Other separation results between various streaming settings also inspired the question we address, including the separation of (nonprivate) oblivious and adaptive streaming from [20] and the work of [7] on separating privacy under offline and oblivious continual observation, though we do not make use of their techniques in our results.

There is also a rich literature on privacy under continual observation that inspires and situates our work. The continual observation (a.k.a. continual release) model was first proposed by [11, 5]. The connections between (adaptive) continual observation and private learning shown by [25, 19] have led to a significant line of theoretical and empirical work in this space, with a focus on computing sums under continual observation – see [22] for a survey. Recent works have developed new adaptively private algorithms with improved error guarantees and computational costs [8, 6, 9], and there is a line of work on improving constant factors on the additive error for computing sums [3, 13, 15, 14]. Other work has investigated tasks related to counting distinct elements [16] and various generalizations thereof [23, 2, 1]. There has also been work on continual observation algorithms for graph data in the node-private [18] and edge-private [12, 24, 27] settings.

1.3 Open Questions

A possible direction for future work is to ask, *Is there a "more natural" problem separating DP under oblivious and adaptive continual observation?*

Future work could also look for a further separation of the models: *Is there some problem $\mathcal{P}$ such that every algorithm $\mathcal{A}$ that is (ε, δ)-DP under oblivious continual observation and solves $\mathcal{P}$ accurately is blatantly nonprivate under adaptive continual observation?*

The work of [8] makes some progress on this question. In particular, they show that any algorithm that is $(\varepsilon, 0)$-DP under oblivious continual observation remains $(\varepsilon, 0)$-DP under adaptive continual observation (though, as we show in Theorems 7 and 8, accuracy may fail). Thus, any positive answer to this question must use a problem that requires an approximate-DP algorithm for accuracy. The work of [8] also answers a related but orthogonal question, showing that some algorithms can be private under oblivious continual observation yet blatantly nonprivate under adaptive continual observation. As a simplified example of their demonstration, consider an algorithm that holds a sensitive dataset, releases a random $r \in \{0,1\}^\lambda$ at time $t = 1$, and then outputs 0 thereafter unless the incoming point is equal to r, in which case it reveals the dataset. This is $\left(0, \frac{T}{2^\lambda}\right)$-DP under oblivious continual observation, but in the adaptive setting the adversary can set the second stream element to r and force blatantly nonprivate behavior at $t = 2$.

2 Preliminaries

We begin with some background on differential privacy, in both the non-streaming ("batch") model and the streaming ("continual observation") settings. Two datasets $x, x' \in \mathcal{X}^n$ are *neighbors* if they differ in the data of a single individual.

▶ **Definition 2** ((ε, δ)-indistinguishability)**.** *Let $\varepsilon > 0$ and $\delta \in [0, 1]$. Two random variables R_1, R_2 over outcome space $\mathcal{Y}$ are* (ε, δ)-indistinguishable *(denoted $R_1 \approx_{\varepsilon,\delta} R_2$) if, for all $Y \subseteq \mathcal{Y}$, we have*

$$\Pr[R_1 \in Y] \leq e^\varepsilon \Pr[R_2 \in Y] + \delta \quad \text{and} \quad \Pr[R_2 \in Y] \leq e^\varepsilon \Pr[R_1 \in Y] + \delta.$$

▶ **Definition 3** (Differential privacy (DP) [10])**.** *Let $\varepsilon > 0$ and $\delta \in [0, 1]$. A randomized algorithm $\mathcal{M} : \mathcal{U}^* \to \mathcal{Y}$ is* (ε, δ)-DP *if for all pairs of neighboring inputs $x, x' \in \mathcal{U}^*$, the distributions $\mathcal{M}(x)$ and $\mathcal{M}(x')$ are (ε, δ)-indistinguishable: $\mathcal{M}(x) \approx_{\varepsilon,\delta} \mathcal{M}(x')$.*

2.1 Differential Privacy under Continual Observation

Differential privacy was generalized to the streaming setting by [5, 11] and requires that the entire sequence of outputs remain indistinguishable up to a change in one individual's data in the input stream. The oblivious setting, where an input stream is fixed up front but is made visible to the algorithm timestep-by-timestep, was first defined by [11, 5]. The adaptive setting, where each input is specified immediately prior to being revealed to the algorithm (and can thus be chosen in response to output from the algorithm), was first considered implicitly by [25] and formally defined by [17].

▶ **Definition 4** (DP under **oblivious** continual observation [11, 5])**.** *Let $\mathcal{M}$ be an algorithm that, given a data stream $\mathbf{x} = (x_1, \ldots, x_T)$, produces an output stream $\mathbf{a} = (a_1, \ldots, a_T)$. The algorithm $\mathcal{M}$ is* (ε, δ)-DP under **oblivious** continual observation *if, for all neighboring streams $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^T$ (i.e., $x_t \neq x'_t$ for at most one $t \in [T]$), we have $\mathcal{M}(\mathbf{x}) \approx_{\varepsilon,\delta} \mathcal{M}(\mathbf{x}')$.*

▶ **Definition 5** (DP under **adaptive** continual observation [17]). *The* view *of $\mathcal{A}dv$ in the privacy game $\Pi_{\mathcal{M},\mathcal{A}dv}$ (Algorithm 1) consists of $\mathcal{A}dv$'s internal randomness and the transcript of messages it sends and receives. Let $V^{(\mathsf{side})}_{\mathcal{M},\mathcal{A}dv}$ denote $\mathcal{A}dv$'s view at the end of the game run with input* $\mathsf{side} \in \{L, R\}$.

An algorithm $\mathcal{M}$ is (ε, δ)-DP under **adaptive** continual observation *if, for all adversaries $\mathcal{A}dv$, we have $V^{(L)}_{\mathcal{M},\mathcal{A}dv} \approx_{\varepsilon,\delta} V^{(R)}_{\mathcal{M},\mathcal{A}dv}$.*

Algorithm 1 Privacy game $\Pi_{\mathcal{M},\mathcal{A}dv}$ for the adaptive continual release model.

Input: time horizon $T \in \mathbb{N}$, $\mathsf{side} \in \{L, R\}$ (not known to $\mathcal{A}dv$).
for $t = 1$ to T **do**
  $\mathcal{A}dv$ outputs $\mathsf{type}_t \in \{\mathsf{challenge}, \mathsf{regular}\}$; challenge is chosen once during the game.
  if $\mathsf{type}_t = \mathsf{regular}$ **then**
    $\mathcal{A}dv$ outputs $x_t \in \mathcal{X}$ which is sent to $\mathcal{M}$.
  if $\mathsf{type}_t = \mathsf{challenge}$ **then**
    $t^* \leftarrow t$.
    $\mathcal{A}dv$ outputs $(x_t^{(L)}, x_t^{(R)}) \in \mathcal{X}^2$.
    $x_t^{(\mathsf{side})}$ is sent to $\mathcal{M}$.
  $\mathcal{M}$ outputs a_t which is given to $\mathcal{A}dv$.

[7] defines an offline version of Definition 4, in which an algorithm receives the entire input stream $\mathbf{x}$ at once, rather than item-by-item, and then outputs a full sequence $(a_1, \ldots, a_T)$. We say an algorithm for this setting is (ε, δ)-*DP under* ***offline*** *continual observation.*

3 Separating Oblivious and Adaptive Privacy

We show there is a problem that is "easier" to solve in the oblivious continual observation setting than in the adaptive continual observation setting – that is, for a specified privacy guarantee, we can give accurate answers at far more time steps in the former setting than in the latter setting. The problem yielding this separation is inspired by the problem of *correlated vector queries* described by [4] and defines accuracy with a very similar loss function. However, whereas we work with a data stream and ask the same query at every time step, [4] works with a static dataset and asks a different query at every time step (see Section 1.1 for a further discussion on differences between our settings).

In Definition 6 we introduce the problem and accuracy definition yielding our separation result. We next describe an algorithm for the oblivious setting that releases accurate answers at many time steps (Section 3.1), and then give an upper bound on the number of time steps at which accurate answers can be released in the adaptive setting (Section 3.2).

▶ **Definition 6** (Problem separating online and online adaptive models). *Let $\alpha \in (0, 1)$, $d \in \mathbb{N}$, and $T \in \mathbb{N}$ be parameters; we denote by $\mathcal{P}^{\alpha,d,T}$ an instance of the separation problem with these parameters. Sensitive data arrives in two phases: a* setup phase *and an* arrival phase. *Outputs are produced only during the arrival phase.*

- ***Setup phase:*** *d "one-bit" individuals $b_1, \ldots, b_d \in \{\pm 1\}$ arrive; no outputs are returned.*
- ***Arrival phase:*** *T "vector" individuals $v_1, \ldots, v_T \in \{\pm 1\}^d$ arrive. After each arrival $t \in [T]$, the algorithm outputs $y^{(t)} \in \{\pm 1\}^d$.*

Let $b = (b_1, \ldots, b_d)$ and $v_{[t]} = (v_1, \ldots, v_t)$. The (boolean-valued) loss at time t is

$$L_t(b, y^{(t)}) \iff \left|\langle y^{(t)} - \alpha b, b\rangle\right| \le \tfrac{\alpha^2 d}{100} \ \wedge\ \forall v \in v_{[t]},\ \left|\langle y^{(t)} - \alpha b, v\rangle\right| \le \tfrac{\alpha^2 d}{100}.$$

An algorithm $\mathcal{M}$ is β-accurate for $\mathcal{P}^{\alpha,d,T}$ *if, with probability at least $1 - \beta$, the entire sequence of outputs $y^{(1)}, \ldots, y^{(T)}$ satisfies the loss function.*

Informally, the loss function says that the output vector y should be close to the vector αb (e.g., for small α, the vector y must have a small but precise correlation with b), and have as little correlation as possible with each $v \in v_{[t]}$ (roughly, some $v \in v_{[t]}$ may themselves be close to αb, so correlation with these vectors is permitted, but the answer should otherwise be nearly orthogonal to all $v \in v_{[t]}$).

3.1 Algorithm for Oblivious Continual Observation

We show there is a DP algorithm for this problem in the oblivious setting that can run accurately for exponentially many time steps (in the dimension of the vector individuals).

The algorithm $\mathcal{M}$ for this setting is inspired by the algorithm in the proof of [4, Theorem 4.1]. At a high level, $\mathcal{M}$ runs a randomized response-like algorithm [26] once on each one-bit individual, stores the resulting vector y, and returns this same vector y at every time step. Privacy roughly follows from the privacy of randomized response, and accuracy follows from Hoeffding's inequality.

▶ **Theorem 7** (Accuracy for oblivious continual observation)**.** *For every $0 < \alpha < 1/2$, there exists some $T = 2^{\Omega(\alpha^4 d)}$ such that, for every sufficiently large $d \in \mathbb{N}$, there is an algorithm $\mathcal{M}$ that is $(1/T)$-accurate for $\mathcal{P}^{\alpha,d,T}$ and is $(3\alpha, 0)$-DP under **oblivious** continual observation.*

Proof of Theorem 7. We first define the algorithm $\mathcal{M}$ for this setting. Let $y = (y_1, \cdots, y_d)$ be the random vector obtained by independently setting, for each $i \in [d]$,

$$y_i = \begin{cases} +b_i & \text{w.p. } \frac{1+\alpha}{2} \\ -b_i & \text{w.p. } \frac{1-\alpha}{2}. \end{cases}$$

Define $\mathcal{M}$ as the algorithm that releases this same vector y at every time step $t \in [T]$.

We now analyze the privacy of $\mathcal{M}$. Since the output of $\mathcal{M}$ depends only on $(b_1, \ldots, b_d)$, it suffices to consider two datasets b and b' that differ on some index i. Since each coordinate of y is generated independently, it suffices to bound the likelihood ratio of outputs for coordinate i as

$$\frac{\Pr[\mathcal{M}(b_i) = y_i]}{\Pr[\mathcal{M}(b'_i) = y_i]} \leq \frac{(1+\alpha)/2}{(1-\alpha)/2} = \frac{1+\alpha}{1-\alpha}.$$

We use the fact that, for $x \in [0, 1)$ we have $\ln(1+x) \leq x$ and $\ln(1-x) \geq -\frac{x}{1-x}$. This gives $\ln\left(\frac{1+\alpha}{1-\alpha}\right) \leq 3\alpha$. A symmetric argument shows that the natural log of the likelihood ratio is bounded below by -3α, so the algorithm is $(3\alpha, 0)$-DP.

We next analyze the algorithm's accuracy. This analysis follows almost verbatim from the accuracy analysis of [4, Theorem 4.1], which we repeat below for completeness.

To prove accuracy, observe that since the output y does not depend on the "vector" individuals, we can analyze $\mathcal{M}$ as if the "vector" individuals were fixed and given all at once. First, observe that $\mathbb{E}[y] = \alpha b$. Thus we have

$$\mathop{\mathbb{E}}_{y}[\langle y - \alpha b, b\rangle] = 0 \quad \text{and} \quad \forall t \in [T],\ \mathop{\mathbb{E}}_{y}[\langle y - \alpha b, v_t\rangle] = 0.$$

Since b and every vector $v_1, \ldots, v_T$ is fixed independently of y, and the coordinates of y are computed using independent randomness, both $\langle y, b\rangle$ and $\langle y, v_t\rangle$ for all $t \in [T]$ are the sum

of d independent $\{\pm 1\}$-valued random variables. Thus, we can apply Hoeffding's inequality[1] and a union bound to conclude there is some absolute constant $c > 0$ such that

$$\Pr_y\left[|\langle y - \alpha b, b\rangle| > \frac{\alpha^2 d}{100}\right] \leq 2\exp\left(-c\alpha^4 d\right), \text{ and}$$
$$\Pr_y\left[\exists t \in [T] \text{ s.t. } |\langle y - \alpha b, v_t\rangle| > \frac{\alpha^2 d}{100}\right] \leq 2T\exp\left(-c\alpha^4 d\right).$$

The theorem now follows by setting an appropriate choice of $T = 2^{\Omega(\alpha^4 d)}$ such that $2(T+1)\cdot\exp\left(-c\alpha^4 d\right) \leq 1/T$. Thus $\mathcal{M}$ is $(1/T)$-accurate for $\mathcal{P}^{\alpha,d,T}$, completing the proof. ◀

3.2 Error Lower Bound for Adaptive Continual Observation

We now show that no DP algorithm for the adaptive setting can run accurately for even some constant (in the dimension of the input) number of time steps.

▶ **Theorem 8** (Error required for adaptive continual observation). *For every $0 < \alpha < 1/2$, there exists some $T = O(1/\alpha^2)$ such that for all sufficiently large $d \in \mathbb{N}$, there is no algorithm $\mathcal{M}$ that is $(1/100)$-accurate for the problem $\mathcal{P}^{\alpha,d,T}$ and is $\left(\frac{1}{5}, \frac{1}{20}\right)$-DP under **adaptive** continual observation.*

Our proof is inspired by the structure of the attack used in proving [4, Theorem 4.2]. However, our result does not follow directly from [4], and moreover does not seem to follow directly from their attack.

At a high level, our attack uses the following strategy. Let $y^{(t)}$ be the output produced after the arrival of vector individual v_t. Set the next vector-valued individual $v_{t+1} = y^{(t)}$. Because $v_2, \ldots, v_t$ contains all outputs $y^{(1)}, \ldots, y^{(t-1)}$, every accurate algorithm will need to recompute a new output in response to seeing $v_{t+1} = y^{(t)}$ as the next vector individual (otherwise the output will be too close to v_{t+1} to satisfy the loss function). Intuitively, this recomputed query will leak information about the private dataset. We show that after T queries for some $T = O(1/\alpha^2)$, a "reconstruction lemma" [4, Lemma 4.3] means that the private dataset can be reconstructed with high probability, which violates privacy.

Our proof relies heavily on the following reconstruction lemma from [4].

▶ **Lemma 9** (Reconstruction Lemma 4.3 from [4]). *Fix parameters $p, q \in [0,1]$. Let $x \in \{\pm 1\}^d$ and $y^{(1)}, \cdots, y^{(k)} \in \{\pm 1\}^d$ be vectors such that*

$$\forall 1 \leq j \leq k, \quad \langle y^{(j)}, x\rangle \geq pd, \quad \text{and}$$
$$\forall 1 \leq j < j' \leq k, \quad |\langle y^{(j)}, y^{(j')}\rangle| \leq qd.$$

If we let $\tilde{x} = \text{sign}(\sum_{j=1}^k y^{(j)}) \in \{\pm 1\}^d$ be the coordinate-wise majority of $y^{(1)}, \ldots, y^{(k)}$, then

$$\langle \tilde{x}, x\rangle \geq \left(1 - \frac{2}{p^2 k} - \frac{2(q-p^2)}{p^2}\right) d.$$

We also use Fact 10 from the literature on differential privacy, which shows the following relationship between (ε, δ)-indistinguishability and total variation distance.

[1] We use the following statement of Hoeffding's inequality: if $Z_1, \ldots, Z_n$ are independent $\{\pm 1\}$-valued random variables, and $Z = \sum_{i=1}^n Z_i$, then $\Pr\left[|Z - \mathbb{E}[Z]| > C\sqrt{n}\right] \leq 2e^{-C^2/2}$.

▶ **Fact 10** (D_{TV} and (ε,δ)-indistinguishability). *Let $\varepsilon > 0$ and $\delta \in [0,1]$, and let P and Q be two probability distributions such that $P \approx_{\varepsilon,\delta} Q$. Then*

$$D_{TV}\big(P,Q\big) \le (e^{\varepsilon}-1)+\delta.$$

Moreover, if $\varepsilon \le 1$, then $D_{TV}\big(\mathcal{A}(x),\mathcal{A}(x')\big) \le 2\varepsilon+\delta$.

We now prove Theorem 8.

Proof of Theorem 8. We show there exists some $T = O(1/\alpha^2)$ such that the output $y^{(1)},\dots,y^{(T)}$ of any algorithm $\mathcal{M}$ that is (1/100)-accurate for $\mathcal{P}^{\alpha,d,T}$ can be used to identify the value of the challenge bit in the privacy game (Algorithm 1) with probability at least $\frac{3}{4}$. In other words, the total variation distance between output distributions on two inputs that differ on the challenge bit is at least $\frac{1}{2}$, which by Fact 10 contradicts the fact that the views of the adversary must be $\big(\frac{1}{5},\frac{1}{20}\big)$-indistinguishable for two inputs that differ in the challenge bit. In this proof, assume for contradiction that $\mathcal{M}$ is $\big(\frac{1}{5},\frac{1}{20}\big)$-DP under adaptive continual observation and (1/100)-accurate.

Define the following adversary $\mathcal{A}dv$. The adversary fixes a uniformly random bit string $b \in \{\pm 1\}^d$. The adversary chooses a uniformly random index $\mathsf{challenge} \in [d]$ and, for that index, chooses a uniformly random pair $\mathsf{pair} \in \{(-1,+1),(+1,-1)\}$. At each one-bit step $i \in [d]$, the adversary outputs type $\mathsf{regular}$ and sends b_i to the privacy game, with the exception of time step $\mathsf{challenge}$, at which the adversary presents $\mathsf{challenge}$ and sends pair. After completing the one-bit phase (at one-bit step d), the adversary sends a uniformly random bit string $v_1 \in \{\pm 1\}^d$ as the first vector individual, and receives output $y^{(1)}$. At all subsequent time steps $t \in \{2,\dots,T\}$, the adversary sends $v_t = y^{(t-1)}$.

We now show that no algorithm $\mathcal{M}$ guaranteeing (1/100)-accuracy for $\mathcal{P}^{\alpha,d,T}$ is $\big(\frac{1}{5},\frac{1}{20}\big)$-DP under adaptive continual observation. The subsequent steps follow from the proof of [4, Theorem 4.3]. Let $p = 99\alpha/100$ and $q = 51\alpha^2/50$. Since $\mathcal{M}$ is assumed to be accurate for T time steps, with probability 99/100, we obtain vectors $y^{(1)},\dots,y^{(T)} \in \{\pm 1\}^d$ such that

$$\begin{aligned}
\forall t \in [T] \quad \langle y^{(t)}, b\rangle &\ge \langle \alpha b, b\rangle - \Big|\langle y^{(t)} - \alpha b, b\rangle\Big| \\
&\ge \alpha d - \frac{\alpha^2 d}{100} \\
&\ge pd, \quad \text{and also}
\end{aligned}$$

$$\begin{aligned}
\forall 1 \le t < t' \le T \quad \Big|\langle y^{(t)}, y^{(t')}\rangle\Big| &\le \Big|\langle \alpha b, y^{(t)}\rangle\Big| + \Big|\langle y^{(t')} - \alpha b, y^{(t)}\rangle\Big| \\
&\le \alpha\Big|\langle y^{(t)}, b\rangle\Big| + \frac{\alpha^2 d}{100} \\
&\le \alpha\left(|\langle \alpha b, b\rangle| + \Big|\langle y^{(t)} - \alpha b, b\rangle\Big|\right) + \frac{\alpha^2 d}{100} \\
&\le \alpha^2 d + \frac{\alpha^3 d}{100} + \frac{\alpha^2 d}{100} \\
&\le \frac{51}{50}\alpha^2 d = qd.
\end{aligned}$$

Thus, by Lemma 9, for $T \ge 1 + 100/\alpha^2$, $k \ge \max\{T,d\}$, and $\tilde{b} = \mathrm{sign}(\sum_{j=1}^{k} y^{(j)})$, we have

$$\langle \tilde{b}, b\rangle \ge \left(1 - \frac{2}{p^2 k} - \frac{2(q-p^2)}{p^2}\right) d$$

$$\geq \left(1 - \frac{2}{(99\alpha/100)^2 d} - \frac{2(51\alpha^2/50 - (99\alpha/100)^2)}{(99\alpha/100)^2}\right) d$$
$$= \left(1 - \frac{2(100/99)^2}{100} - 2\left(\frac{(51/50) - (99/100)^2}{(99/100)^2}\right)\right) d$$
$$\geq 0.89d.$$

Recall that the challenge bit's location and value were chosen by the adversary uniformly at random. Therefore, $\mathcal{M}$ cannot distinguish the challenge location from non-challenge locations, so because at least a 0.89 fraction of the one-bit individuals' values were reconstructed correctly, this means the probability that the challenge value was correctly reconstructed by the adversary is at least 0.89. This implies that the TV distance between adversary views is greater than $\frac{1}{2}$, which by Fact 10 contradicts the fact that the views of the adversary must be $\left(\frac{1}{5}, \frac{1}{20}\right)$-indistinguishable. ◀

References

1 Anders Aamand, Justin Y. Chen, and Sandeep Silwal. Skirting additive error barriers for private turnstile streams, 2026. doi:10.48550/arXiv.2602.10360.

2 Joel Daniel Andersson, Palak Jain, and Satchit Sivakumar. Improved accuracy for private continual cardinality estimation in fully dynamic streams via matrix factorization. *CoRR*, abs/2601.02257, 2026. doi:10.48550/arXiv.2601.02257.

3 Joel Daniel Andersson and Rasmus Pagh. A smooth binary mechanism for efficient private continual observation. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *NeurIPS*, 2023. URL: http://papers.nips.cc/paper_files/paper/2023/hash/99c41fb9fd53abfdd4a0259560ef1c9d-Abstract-Conference.html.

4 Mark Bun, Thomas Steinke, and Jonathan R. Ullman. Make up your mind: The price of online queries in differential privacy. *J. Priv. Confidentiality*, 9(1), 2019. doi:10.29012/JPC.655.

5 T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Trans. Inf. Syst. Secur.*, 14(3):26:1–26:24, 2011. doi:10.1145/2043621.2043626.

6 Christopher A. Choquette-Choo, Krishnamurthy Dj Dvijotham, Krishna Pillutla, Arun Ganesh, Thomas Steinke, and Abhradeep Guha Thakurta. Correlated noise provably beats independent noise for differentially private learning. In *ICLR*. OpenReview.net, 2024. URL: https://openreview.net/forum?id=xHmCdSArUC.

7 Edith Cohen, Xin Lyu, Jelani Nelson, Tamás Sarlós, and Uri Stemmer. Lower bounds for differential privacy under continual observation and online threshold queries. In Shipra Agrawal and Aaron Roth, editors, *COLT*, volume 247 of *Proceedings of Machine Learning Research*, pages 1200–1222. PMLR, 2024. URL: https://proceedings.mlr.press/v247/cohen24b.html.

8 Sergey Denisov, H. Brendan McMahan, John Rush, Adam D. Smith, and Abhradeep Guha Thakurta. Improved differential privacy for SGD via optimal private linear operators on adaptive streams. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *NeurIPS*, 2022. URL: http://papers.nips.cc/paper_files/paper/2022/hash/271ec4d1a9ff5e6b81a6e21d38b1ba96-Abstract-Conference.html.

9 Krishnamurthy Dj Dvijotham, H. Brendan McMahan, Krishna Pillutla, Thomas Steinke, and Abhradeep Thakurta. Efficient and near-optimal noise generation for streaming differential privacy. In *FOCS*, pages 2306–2317. IEEE, 2024. doi:10.1109/FOCS61266.2024.00135.

10 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016. doi:10.29012/jpc.v7i3.405.

11 Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In Leonard J. Schulman, editor, *STOC*, pages 715–724. ACM, 2010. doi:10.1145/1806689.1806787.

12 Alessandro Epasto, Quanquan C. Liu, Tamalika Mukherjee, and Felix Zhou. Sublinear space graph algorithms in the continual release model. In Alina Ene and Eshan Chattopadhyay, editors, *APPROX/RANDOM*, volume 353 of *LIPIcs*, pages 40:1–40:27. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.APPROX/RANDOM.2025.40.

13 Hendrik Fichtenberger, Monika Henzinger, and Jalaj Upadhyay. Constant matters: Fine-grained error bound on differentially private continual observation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 10072–10092. PMLR, 2023. URL: https://proceedings.mlr.press/v202/fichtenberger23a.html.

14 Monika Henzinger, Nikita P. Kalinin, and Jalaj Upadhyay. Normalized square root: Sharper matrix factorization bounds for differentially private continual counting. *CoRR*, abs/2509.14334, 2025. doi:10.48550/arXiv.2509.14334.

15 Monika Henzinger and Jalaj Upadhyay. Improved differentially private continual observation using group algebra. In Yossi Azar and Debmalya Panigrahi, editors, *SODA*, pages 2951–2970. SIAM, 2025. doi:10.1137/1.9781611978322.95.

16 Palak Jain, Iden Kalemaj, Sofya Raskhodnikova, Satchit Sivakumar, and Adam D. Smith. Counting distinct elements in the turnstile model with differential privacy under continual observation. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *NeurIPS*, 2023. URL: http://papers.nips.cc/paper_files/paper/2023/hash/0ef1afa0daa888d695dcd5e9513bafa3-Abstract-Conference.html.

17 Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam D. Smith. The price of differential privacy under continual observation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 14654–14678. PMLR, 2023. URL: https://proceedings.mlr.press/v202/jain23b.html.

18 Palak Jain, Adam Smith, and Connor Wagaman. Time-aware projections: Truly node-private graph statistics under continual observation. In *Oakland IEEE S&P*, pages 127–145. IEEE, 2024. doi:10.1109/SP54263.2024.00196.

19 Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling. In Marina Meila and Tong Zhang, editors, *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 5213–5225. PMLR, 2021. URL: http://proceedings.mlr.press/v139/kairouz21b.html.

20 Haim Kaplan, Yishay Mansour, Kobbi Nissim, and Uri Stemmer. Separating adaptive streaming from oblivious streaming using the bounded storage model. In Tal Malkin and Chris Peikert, editors, *CRYPTO*, volume 12827 of *Lecture Notes in Computer Science*, pages 94–121. Springer, 2021. doi:10.1007/978-3-030-84252-9_4.

21 Mathias Lécuyer, Riley Spahn, Kiran Vodrahalli, Roxana Geambasu, and Daniel Hsu. Privacy accounting and quality control in the sage differentially private ML platform. In Tim Brecht and Carey Williamson, editors, *SOSP*, pages 181–195. ACM, 2019. doi:10.1145/3341301.3359639.

22 Krishna Pillutla, Jalaj Upadhyay, Christopher A. Choquette-Choo, Krishnamurthy Dvijotham, Arun Ganesh, Monika Henzinger, Jonathan Katz, Ryan McKenna, H. Brendan McMahan, Keith Rush, Thomas Steinke, and Abhradeep Thakurta. Correlated noise mechanisms for differentially private learning. *CoRR*, abs/2506.08201, 2025. doi:10.48550/arXiv.2506.08201.

23 Yuan Qiu and Ke Yi. Differential privacy on fully dynamic streams. In *NeurIPS*, 2025. URL: https://openreview.net/forum?id=piM21sPyVL.

24 Sofya Raskhodnikova and Teresa Anna Steiner. Fully dynamic algorithms for graph databases with edge differential privacy. *Proc. ACM Manag. Data*, 3(2):99:1–99:28, 2025. doi:10.1145/3725236.

25 Abhradeep Guha Thakurta and Adam D. Smith. (Nearly) optimal algorithms for private online learning in full-information and bandit settings. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *NeurIPS*, pages 2733–2741, 2013. URL: https://proceedings.neurips.cc/paper/2013/hash/c850371fda6892fbfd1c5a5b457e5777-Abstract.html.

26 Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965. URL: http://www.jstor.org/stable/2283137.

27 Felix Zhou. Continual release of densest subgraphs: Privacy amplification & sublinear space via subsampling. In *SOSA*, 2026. doi:10.48550/arXiv.2510.11640.