The power of synergy in differential privacy: Combining a small curator with local randomizers

Amos Beimel* Aleksandra Korolova† Kobbi Nissim‡ Or Sheffet§
Uri Stemmer¶

December 20, 2019

Abstract

Motivated by the desire to bridge the utility gap between local and trusted curator models of differential privacy for practical applications, we initiate the theoretical study of a hybrid model introduced by "Blender" [Avent et al., USENIX Security '17], in which differentially private protocols of n agents that work in the local-model are assisted by a differentially private curator that has access to the data of m additional users. We focus on the regime where $m \ll n$ and study the new capabilities of this (m,n)-*hybrid* model. We show that, despite the fact that the hybrid model adds no significant new capabilities for the basic task of simple hypothesis-testing, there are many other tasks (under a wide range of parameters) that can be solved in the hybrid model yet cannot be solved either by the curator or by the local-users separately. Moreover, we exhibit additional tasks where at least one round of *interaction* between the curator and the local-users is necessary – namely, no hybrid model protocol without such interaction can solve these tasks. Taken together, our results show that the combination of the local model with a small curator can become part of a promising toolkit for designing and implementing differential privacy.

*Dept. of Computer Science, Ben-Gurion University. amos.beimel@gmail.com
†Dept. of Computer Science, University of Southern California. korolova@usc.edu
‡Dept. of Computer Science, Georgetown University. kobbi.nissim@georgetown.edu
§Faculty of Engineering, Bar-Ilan University. or.sheffet@biu.ac.il
¶Dept. of Computer Science, Ben-Gurion University and Google Research. u@uri.co.il

1 Introduction

Data has become one of the main drivers of innovation in applications as varied as technology, medicine [33], and city planning [34, 53]. However, the collection and storage of personal data in the service of innovation by companies, researchers, and governments poses significant risks for privacy and personal freedom. Personal data collected by companies can be breached [21]; subpoenaed by law enforcement in broad requests [44]; mis-used by companies' employees [37, 20]; or used for purposes different from those announced at collection time [64]. These concerns alongside data-hungry company and government practices have propelled privacy to the frontier of individuals' concerns, societal and policy debates, and academic research.

The *local model of differential privacy* [66, 43] has recently emerged as one of the promising approaches for achieving the goals of enabling data-driven innovation while preserving a rigorous notion of privacy for individuals that also addresses the above challenges. The *differential privacy* aspect provides each participating individual (almost) with the same protection she would have had her information not been included [25], a guarantee that holds even with respect to all powerful adversaries with access to multiple analyses and rich auxiliary information. The *local* aspect of the model means that this guarantee will continue to hold even if the curator's data store is fully breached. From the utility perspective, the deployment of local differential privacy protocols by Google, Apple, and Microsoft demonstrate that the local differential privacy model is a viable approach in certain applications, without requiring trust in the companies or incurring risks from hackers and intelligence agencies [28, 55, 35, 1, 22].

The adoption of the local model by major industry players has motivated a line of research in the theory of local differential privacy (e.g., [9, 8, 16, 59, 40, 42, 41]). Alongside algorithmic improvements, this body of work highlighted the wide theoretical and practical gap between utility achievable in the more traditional trusted curator model of differential privacy (where the curator ensures the privacy of its output but can perform computations on raw individual data) and that achievable in the local differential privacy model. In particular, the number of data points necessary to achieve a particular level of accuracy in the local model is significantly larger than what is sufficient for the same accuracy in the curator model (see, e.g., [10, 43, 56, 9]). This has negative consequences. First, data analysis with local differential privacy becomes the privilege of the data-rich, handicapping smaller companies and helping to cement monopolies. Second, in an effort to maintain accuracy the entities deploying local differential privacy are tempted to use large privacy loss parameters [60], ultimately putting into question the privacy guarantee [36].

New models for differentially private computation have recently emerged to alleviate the (inevitable) low accuracy of the local model, of which we will discuss the *shuffle model* [38, 14, 19, 6, 5, 32] and the *hybrid model* [2].[1] In the shuffle model, it is assumed that the curator receives data in a manner disassociated from a user identifier (e.g., after the raw data has been stripped of identifiers and randomly shuffled). Recent work has proved that protocols employing shuffling can provide better accuracy than local protocols and sometimes match the accuracy of the curator model [19, 6, 5, 32].[2] Although the shuffle model is a promising approach for bridging the gap between the local and trusted curator model, it suffers from two weaknesses: it requires individuals' trust in the shuffler (which itself may be subject to breaches, subpoenaes, etc. and the infrastruc-

[1] Approaches which weaken differential privacy itself or justify the use of large privacy loss parameters are outside our scope and deserve a separate discussion.

[2] Furthermore, the shuffle model provides new "privacy amplification" tools that can be used in the design of differentially private algorithms [27, 6].

ture for which may not exist), and, as highlighted by [6], its privacy guarantees to an individual rely on the assumption that sufficiently many other individuals do not deviate from the protocol.

The focus of this work is on a generalization of the hybrid model introduced by [2], where a majority of individuals that participate in a local differential privacy computation is augmented with a small number of individuals who contribute their information via a trusted curator. From a practical point of view, this separation is aligned with current industry practices, and the small number of individuals willing to contribute via a curator can be employees, technology enthusiasts, or individuals recruited as alpha- or beta-testers of products in exchange for early access to its features or decreased cost [48, 49, 51].

Furthermore, unlike Blender [2], in an effort to explore a rich trust and communication model, and anticipate development of future technologies and practices, we do not assume that the curator trusted by the opt-in users and the eventual data recipient (whom we call the referee) are the same entity (see Figure 1). The detailed discussion of the benefits of this modeling assumption appears after the formal model definition in Section 2.2.

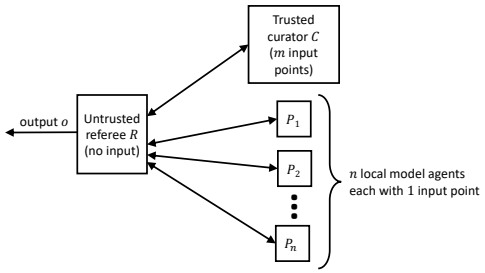

Figure 1: The hybrid model.

1.1 This work: The power of the hybrid model

We initiate a theoretical study of the extent to which the hybrid model improves on the utility of the local model by addition of a small curator (and vice versa, improves on the utility of a small curator by addition of parties that participate in a local model computation). We ask whether there are tasks that cannot be computed privately by a small curator, or in the local model (even with many participants), but are feasible in a model that combines both. We answer this question in the affirmative.

A concatenation of problems (direct-sum). Our first example is composed of two independent learning tasks, each task can be privately learned in one model, however, cannot be privately learned in the other model. Each data point is parsed as two points (x, y) that are labeled $(\mathsf{Par}_k(x), \mathsf{Thr}_t(y))$ where the former is a parity function $\mathsf{Par}_k(x) = \langle k, x \rangle$ and the latter is a threshold function $\mathsf{Thr}_t(y) = \mathbb{1}_{\{y \geq t\}}$ (see Section 2.4 for complete definitions). For a choice of parameters, known sample complexity bounds imply that the parity learning part of the task can be performed by the curator but cannot be privately computed in the local model with sub-exponential number of messages [43], and, conversely, that the threshold learning part cannot be performed by the curator [11, 30] but can be performed by the local model parties. It follows that for our choice of parameters the combined task cannot be performed neither by the curator nor by the local model with sub-exponential number of rounds (as the number of local agents is small, each agent in the local model must send many messages), but is feasible in the hybrid model without interaction (see Appendix A for a detailed analysis).

Select-then-estimate. In our second example, each input point x is sampled i.i.d. from an unknown distribution over $\{-1, 1\}^d$. Letting $\mu = E[x]$, the goal is to output a pair $(i, \hat{\mu}_i)$ where i approximately maximizes μ_i (i.e., $\mu_i \geq \max_j \mu_j - \alpha$) and $\hat{\mu}_i$ approximates μ_i (i.e., $|\hat{\mu}_i - \mu_i| \leq \alpha'$),

where $\alpha' < \alpha$. That is, once we found a good coordinate i, we want to approximate its quality with smaller error. A number of sample complexity bounds apply here (in particular, [61] and Appendix C), yielding a wide range of parameters where the task cannot be performed by the curator alone, nor by the local model parties alone. The hybrid model, however, allows for a solution where the curator first identifies i such that $\mu_i \geq \max_j \mu_j - \alpha$ and the estimation of μ_i within accuracy α' is then performed by the local model parties (see Appendix B for a detailed analysis).

The select-then-estimate problem is a sub-component of many statistical and data-mining applications [4, 58, 2]. It is also of core interest in genome-wide association studies (GWAS), where the inputs are genomes of patients affected by a particular disease and the goal is to (approximately) identify disease-related genome locations and estimate their significance [3, 67, 15]. Solving the problem while ensuring privacy is particularly challenging when the feature dimension is large compared to the number of available samples, which is the case for GWAS [13, 3, 39, 57]. As genomic data is particularly sensitive, the hybrid model of differential privacy appears appropriate for it from the privacy perspective – the majority of the data would be analyzed with the guarantee of local differential privacy, and only a small number of data points would be entrusted to a curator, whose analysis's output should also be differentially private [65]. As our example suggests, the hybrid model may be useful also from the utility perspective.

We next present and study tasks that require protocols with different interaction patterns involving both the curator and the local agents; that is, the referee needs to relay messages from the curator to the local model agents in one problem or vice versa in a second problem.

Learning parity XOR threshold. This task, which is a twist on the above concatenation of problems, combines two independent learning tasks. Rather than merely concatenating, in the learning parity XOR threshold problem points are labeled by $\mathsf{Par}_k(x) \oplus \mathsf{Thr}_t(y)$. A simple argument shows that (for specifically chosen parameters) the task cannot be performed by the curator alone or by the local model agents with sub-exponential number of rounds. The task is, however, feasible in the hybrid model without interaction. Observe that the local model agents can complete the task once the parity part is done, and that $\mathsf{Thr}_t(y) = 0$ for the lower half of the points y or $\mathsf{Thr}_t(y) = 1$ for the upper half of the points y (or both). These observations suggest a protocol where the curator first performs two parity learning tasks (splitting points according to y values), and the task is then completed by the local model agents. This requires communication (as the referee needs to relay a message from the curator to the local model agents), and it may seem that this interaction is necessary for the task. However, in Section 3 we show that this is not the case, by demonstrating a non-interactive protocol where all parties send a message to the referee at the same round.

1-out-of-2^d-parity. Our next task can be solved in the hybrid model (but neither by a small curator model nor in the local model with sub-exponential number of rounds). This task requires interaction, first with the local model agents and then with the curator. The task consists of a multitude of parity problems, only one of which – determined by the input – needs to be solved. The curator is capable of solving the parity task privately, however, the curator's limited number of samples does not allow solving all parity problems privately, nor does it allow distinguishing which problem to solve. The local model agents cannot solve parity (with sub-exponential number of rounds) [43] but can recover which problem needs to be solved (via a heavy hitters computa-

tion [16]). Thus, the referee needs to first interact with the local agents and then the curator. See Section 4.

Parity-chooses-secret. The third task in this part can be solved in the hybrid model (but neither with a small curator, nor in the local model with sub-exponential number of rounds). The task requires interaction in the reverse order from the previous task: first with the curator, then with the local model agents. The input to this task contains shares of a large collection of secrets and the goal is to recover one of these secrets. The secret that should be recovered is determined by the input as the solution to a parity problem. The curator can solve the parity problem but does not have enough information to recover any of the secrets. The local model agents receive enough information to recover all secrets, but doing so would breach privacy (as implied by [46]). They cannot solve the parity problem with sub-exponential number of rounds. In the hybrid model protocol, the curator first solves the parity problem privately, and relates the identity of the secret to be recovered through the referee to the local model agents who then can send enough information to recover the required secret. See Section 5.

The latter two tasks highlight information and private-computation gaps between the curator and the local model agents. The local model agents receive enough information to solve the task, but lack the ability to privately solve an essential sub-task (when they are not allowed to use exponentially many rounds). The curator does not receive enough information to solve the task (even non-privately), but can solve the hard sub-task.

When the hybrid model does not help much. Although most of the results in this work are on the positive side, demonstrating that cleverly utilizing a curator in synergy with local agents can allow for new capabilities, we also show that for one of the most basic tasks – namely, basic hypothesis testing – the hybrid model has no significant advantage over what can be done separately in either the local model with m agents or in the curator model with database of size n. We show that if for two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ there is a private protocol in the hybrid model that given i.i.d. samples from $\mathcal{D}_j$ correctly identifies j, then there is a private protocol with the same sample size either in the local model or in the curator model that correctly identifies j (with some small loss in the success probability). We then consider two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ over the the domain $\{0,1\}$, where in $\mathcal{D}_0$ the probability of 1 is strictly less than $1/2$ and in $\mathcal{D}_1$ the probability of 1 is strictly greater than $1/2$ and identify values of m and n such that in the hybrid model, where the curator has m samples and there are n agents (each holding one sample), no protocol exists that can differentiate whether the $m+n$ inputs were sampled i.i.d. from $\mathcal{D}_0$ or from $\mathcal{D}_1$. Since computing the sum of i.i.d. sampled bits from $\mathcal{D}_0$ or $\mathcal{D}_1$ can distinguish between these distributions, the above results imply that for computing the sum, the hybrid model is no better than each model separately. See Section 6.

A new lower bound for selection in the local model. As mentioned above, our analysis for the select-then-estimate task relies on lower bounds on the sample complexity of selection in the local model. Ullman [61] gives a (tight) lower bound of $\Omega(\frac{d \log d}{\alpha^2 \varepsilon^2})$ samples for the non-interactive case. In Appendix C we show that for *interactive* local model protocols, the number of messages in such protocol is $\Omega(d^{1/3})$. For example, if the curator interacts with the local model parties so that each party sends t messages, then the number of parties must be at least $\Omega(d^{1/3}/t)$. The proof is

by a generic reduction from any private PAC learning task to selection, which preserves sample complexity. The bound is obtained by applying the reduction from privately learning parity and applying bounds on the sample complexity of privately learning parity from [43].

1.2 Discussion and future work

Our results show that the combination of the local model with a small curator can become part of a promising toolkit for designing and implementing differential privacy. More work is needed to develop the theory of this model (and possibly introduce variants), and, in particular, characterize which tasks can benefit from it. From an algorithms design perspective, now that we know that the hybrid model can lead to significant improvements over both the curator and local models, an exciting open question is understanding what other non-trivial algorithms can be designed that take advantage of the synergy.

Selection bias. In this work we assume that the inputs for the curator and for the local model agents come from the same distribution. However, the recruitment of individuals for participating via a curator can create unintended differences between the input distributions seen by the curator and the entire population, and hence lead to biases, an issue which is outside the scope of the current work. Selection bias remains an important issue that needs to be addressed.

Approximate differential privacy. Our separations between the hybrid model and the curator and local models hold for pure differential privacy (i.e., ε-differential privacy). Specifically, we use lower bounds for ε-differential private learning of the threshold functions in the curator model [11, 30]; these lower bounds do not hold for (ε, δ)-differential private learning of the threshold functions [12, 17]. We also use lower bounds for ε-differential private learning of parity in the local model [43]; it is open if these lower bounds hold for fully interactive (ε, δ)-differential private learning protocols of parity. Possible separations for (ε, δ)-differential privacy are left for future research.

2 Preliminaries

2.1 Protocols for differentially private computations

Let X be a data domain. We consider a model where the inputs and the computation are distributed among parties $P_1, \ldots, P_n$. Each party is an interactive randomized functionality: it can receive messages from the other parties, perform a randomized computation, and send messages to the other parties. At the beginning of a computation, each party P_i receives its input $\mathbf{x}_i = (x_{i,1}, \ldots, x_{i,\ell_i}) \in X^{\ell_i}$. I.e., the input of party P_i consists of a sequence of $\ell_i \geq 0$ entries taken from the data domain X, and the entire joint input to the protocol is $(x_{1,1} \ldots, x_{1,\ell_1}, x_{2,1}, \ldots, x_{2,\ell_2}, \ldots, x_{n,1}, \ldots, x_{n,\ell_n})$. The parties engage in a randomized interactive protocol $\Pi = (\Pi_{P_1}, \ldots, \Pi_{P_n})$, where a message sent by a party P_i in some round is computed according to Π_{P_i} and depends on its input $\mathbf{x}_i$, its random coins, and the sequence of messages it has seen in previous rounds. When a party P_i halts, it writes its output to a local output register $\mathbf{o}_{P_i}$. The number of messages in a protocol is the number of rounds multiplied by the number of parties.

Definition 2.1. *We say that* $\mathbf{x} = (x_1, \ldots, x_\ell) \in X^\ell$ *and* $\mathbf{x}' = (x'_1, \ldots, x'_\ell) \in X^\ell$ *are* neighboring *if they differ on at most one entry, i.e., there exist* $i^* \in [\ell]$ *such that* $x_i = x'_i$ *for* $i \in [\ell] \setminus \{i^*\}$.

Definition 2.2. *We say that two probability distributions* $\mathcal{D}_0, \mathcal{D}_1 \in \Delta(\Omega)$ *are* (ε, δ)*-close and write* $\mathcal{D}_0 \approx_{\varepsilon,\delta} \mathcal{D}_1$ *if*

$$\Pr_{t \sim \mathcal{D}_b}[t \in T] \leq e^\varepsilon \cdot \Pr_{t \sim \mathcal{D}_{1-b}}[t \in T] + \delta$$

for all measurable events $T \subset \Omega$ *and* $b \in \{0, 1\}$.

We are now ready to define what it means for a protocol to be differentially private in a fully malicious setting, i.e., in a setting where an arbitrary adversary controls the behavior of all but one party. Intuitively, a protocol is differentially private in a fully malicious setting if there do not exist a party P_i and an adversary A controlling $P_1, \ldots, P_{i-1}, P_{i+1}, \ldots, P_n$ such that A can "trick" P_i to act non-privately. More formally, we model the adversary as an interactive randomized functionality. For a party P_i, define A_{P_i} to be a randomized functionality as follows.

1. An input of A_{P_i} consists of a sequence of ℓ_i entries taken from the data domain, $\mathbf{x} \in X^{\ell_i}$.
2. A_{P_i} simulates an interaction between party P_i with $\mathbf{x}$ as its input, and A. The simulated P_i interacts with A following the instructions of its part in the protocol, Π_{P_i}. The adversary A interacts with P_i sending messages for parties $P_1, \ldots, P_{i-1}, P_{i+1}, \ldots, P_n$. However, A does not necessarily adhere to the protocol Π.
3. The simulation continues until A halts with an output $\mathbf{o}_A$, at which time A_{P_i} halts and outputs $\mathbf{o}_A$.

Definition 2.3 (Multiparty differential privacy [25, 43, 10, 62])**.** *A protocol* Π *is* (ε, δ)*-differentially private if for all* $i \in [n]$*, for all interactive randomized functionalities* A*, and all neighboring* $\mathbf{x}, \mathbf{x}' \in X^{\ell_i}$*,* $A_{P_i}(\mathbf{x}) \approx_{\varepsilon,\delta} A_{P_i}(\mathbf{x}')$*. When* $\ell_1 = \ell_2 = \cdots = \ell_n = 1$ *we say that the protocol operates in the* local model*, and when* $n = 1$ *we say that the protocol (or the algorithm) operates in the* curator model*. We say that a protocol* Π *is* ε*-differentially private if it is* $(\varepsilon, 0)$*-differentially private.*

Comparison to previous definitions. In contrast to [10, 62], our definition applies also to a malicious adversary that can send arbitrary messages. The definition of [43] also applies to a malicious adversary, however it requires that each answer of an agent preserves ε-differential privacy (i.e., if there are d rounds, the protocol is $d\varepsilon$-differentially private). In contrast, the definitions of [10, 62] and our definition measures the privacy of the entire transcript of an agent.

Note that (i) Restricting the outcome $\mathbf{o}_A$ to binary results in a definition that is equivalent to Definition 2.3. (ii) It is possible to consider a relaxed version of Definition 2.3 where the adversary A is "semi-honest" by requiring A to follow the protocol Π. (iii) Definition 2.3 differs from definitions of security in the setting of secure multiparty computation as the latter also state correctness requirements with respect to the outcome of the protocol. The difference between the setting is that secure multiparty computation implements a specified functionality[3] whereas differential privacy limits the functionality to hide information specific to individuals, but does not specify it.

[3] Furthermore, secure multiparty computation is silent with respect to the chosen functionality, regardless whether it is "privacy preserving" or "secure".

2.2 The hybrid model

A computation in the (m,n)-hybrid model is defined as the execution of a randomized interactive protocol $\Pi = (\Pi_C, \Pi_{P_1}, \ldots, \Pi_{P_n}, \Pi_R)$ between three types of parties: a (single) curator C, n "local model" agents $P_1, \ldots, P_n$, and a referee R. The referee has no input, the curator C receives m input points $\mathbf{x} = (x_1, \ldots, x_m) \in X^m$ and the n "local model" agents $P_1, \ldots, P_n$ each receive a single input point $y_i \in X$. We use the notation $\mathbf{D}$ to denote the joint input to the computation, i.e., $\mathbf{D} = (x_1, \ldots, x_m, y_1, \ldots, y_n)$.

The communication in a hybrid model protocol is restricted to messages exchanged between the referee R and the other parties $C, P_1, \ldots, P_n$ (i.e., parties $C, P_1, \ldots, P_n$ cannot directly communicate among themselves). Parties $C, P_1, \ldots, P_n$ have no output, whereas when the execution halts the referee R writes to a special output register $\mathbf{o}$. See Figure 1. We require the protocol $\Pi = (\Pi_C, \Pi_{P_1}, \ldots, \Pi_{P_n}, \Pi_R)$ to satisfy differential privacy as in Definition 2.3.

The hybrid model is a natural extension of well-studied models in differential privacy. Setting $n = 0$ we get the trusted curator model (as C can perform any differentially private computation), and setting $m = 0$ we get the local model. In this work, we are interested in the case $0 < m \ll n$, because in this regime, the hybrid model is closest in nature to the local model. Furthermore, in many applications, once m is comparable to n it is possible to drop parties $P_1, \ldots, P_n$ from the protocol without a significant loss in utility.

Comparing with Blender [2], where the curator C and the referee R are the same party, we observe that the models are equivalent in their computation power – every differentially private computation in one model is possible in the other (however, the models may differ in the number of interaction rounds needed). Nevertheless, the separation between the curator and the referee has merits as we now discuss.

On the separation between the curator and referee. From a theory point of view, it is useful to separate these two parties as this allows to examine effects related to the order of interaction between the parties (e.g., whether the referee communicates first with the curator C or with the local model parties $P_1, \ldots, P_n$).

Moreover, by separating the curator and referee, the hybrid model encapsulates a richer class of trust models than [2], and, in particular, includes a trust model where data sent to the curator is not revealed to the referee. In an implementation this may correspond to a curator instantiated by a publicly trusted party, or by using technologies such as secure multiparty computation, or secure cryptographic hardware which protects data and computation from external processes [50].

The curator-referee separation also makes sense from a practical point of view within a company. It is reasonable that only a small fraction of a company's employees, with appropriate clearance and training, should be able to access the raw data of those who contribute their data to the trusted curator model, whereas the majority of employees should only see the privacy-preserving version of it [54].

Remark 2.4 (A note on public randomness)**.** Some of our protocols assume the existence of a shared random string. In an implementation, shared randomness can be either set up offline or be chosen and broadcast by the referee. We stress that the privacy of our protocols does not depend on the shared random string actually being random. Furthermore, all our lower bounds hold even when the local agents hold a shared (public) random string.

2.3 Parity and threshold functions

A *concept* $c : X \to \{0,1\}$ is a predicate that labels *examples* taken from the domain X by either 0 or 1. A *concept class* C over X is a set of concepts (predicates) mapping X to $\{0,1\}$. Let $b, c \in \mathbb{N}$ be parameters. The following two concept classes will appear throughout the paper:

- $\mathsf{Threshold}_b = \left\{\mathsf{Thr}_t : t \in \{0,1\}^b\right\}$ where $\mathsf{Thr}_t : \{0,1\}^b \to \{0,1\}$ is defined as $\mathsf{Thr}_t(x) = \mathbb{1}_{\{x \geq t\}}$, where we treat strings from $\{0,1\}^b$ as integers in $\left\{0, \ldots, 2^b - 1\right\}$.
- $\mathsf{Parity}_c = \{\mathsf{Par}_k : k \in \{0,1\}^c\}$ where $\mathsf{Par}_k : \{0,1\}^c \to \{0,1\}$ is defined as $\mathsf{Par}_k(x) = \langle k, x \rangle = \oplus_{j=1}^{c} k_j \cdot x_j$.

2.4 Preliminaries from learning theory and private learning

We now define the probably approximately correct (PAC) model of [63]. Given a collection of labeled examples, the goal of a learning algorithm (or protocol) is to *generalize* the given data into a concept (called a "hypothesis") that accurately predicts the labels of fresh examples from the underlying distribution. More formally:

Definition 2.5. *The* generalization error *of a hypothesis $h : X \to \{0,1\}$ w.r.t. a target concept c and a distribution $\mathcal{D}$ is defined as* $\mathrm{error}_{\mathcal{D}}(c,h) = \Pr_{x \sim \mathcal{D}}[h(x) \neq c(x)]$.

Definition 2.6 (PAC Learning [63])**.** *Let C be a concept class over a domain X, and let Π be a protocol in which the input of every party is a collection of (1 or more) labeled examples from X. The protocol Π is called an* (α, β)-PAC learner *for C if the following holds for all concepts $c \in C$ and all distributions $\mathcal{D}$ on X: If the inputs of the parties are sampled i.i.d. from $\mathcal{D}$ and labeled by c, then, with probability at least $1 - \beta$, the outcome of the protocol is a hypothesis $h : X \to \{0,1\}$ satisfying* $\mathrm{error}_{\mathcal{D}}(c,h) \leq \alpha$.

The sample complexity *of the protocol is the total number of labeled examples it operates on. That is, if there are n parties where party P_i gets as input ℓ_i labeled examples, then the sample complexity of the protocol is $\ell_1 + \cdots + \ell_n$.*

A PAC learning protocol that is restricted to only output hypotheses that are themselves in the class C is called a *proper learner*; otherwise, it is called an *improper learner*. A common technique for constructing PAC learners is to guarantee that the resulting hypothesis h has small *empirical error* (as defined below), and then to argue that such an h must also have small generalization error.

Definition 2.7. *The* empirical error *of a hypothesis $h : X \to \{0,1\}$ w.r.t. a labeled sample $S = (x_i, y_i)_{i=1}^m$ is defined as* $\mathrm{error}_S(h) = \frac{1}{m}|\{i : h(x_i) \neq y_i\}|$. *The* empirical error *of h w.r.t. an unlabeled sample $S = (x_i)_{i=1}^m$ and a concept c is defined as* $\mathrm{error}_S(h,c) = \frac{1}{m}|\{i : h(x_i) \neq c(x_i)\}|$.

Indeed, in some of our constructions we will use building blocks that aim to minimize the *empirical* error, as follows.

Definition 2.8. *Let $c : X \to \{0,1\}$ be a concept and let $\mathbf{D} \in (X \times \{0,1\})^n$ be a labeled database. We say that $\mathbf{D}$ is* consistent *with c if for every $(x,y) \in \mathbf{D}$ it holds that $y = c(x)$.*

Definition 2.9. *Let C be a concept class over a domain X, and let Π be a protocol in which the input of every party is a collection of (1 or more) labeled examples from X. The protocol Π is called an* (α, β)-empirical learner *for C if for every concept $c \in C$ and for every joint input to the protocol $\mathbf{D}$ that is consistent with c, with probability at least $1 - \beta$, the outcome of the protocol is a hypothesis $h : X \to \{0,1\}$ satisfying* $\mathrm{error}_{\mathbf{D}}(h) \leq \alpha$.

We will be interested in PAC-learning protocols that are also differentially private. Specifically,

Definition 2.10 ([43])**.** *A* private learner *is a protocol* Π *that satisfies both Definitions 2.3 and 2.6. Similarly, a* private empirical learner *is a protocol* Π *that satisfies both Definitions 2.3 and 2.9.*

Dwork et al. [24] and Bassily et al. [7] showed that if a hypothesis h is the result of a differentially private computation on a random sample, then the empirical error of h and its generalization error are guaranteed to be close. We will use the following multiplicative variant of their result [52], whose proof is a variant of the original proof of [7].

Theorem 2.11 ([24, 7, 52, 29])**.** *Let* $\mathcal{A} : X^n \to 2^X$ *be an* (ε, δ)*-differentially private algorithm that operates on a database of size* n *and outputs a predicate* $h : X \to \{0,1\}$*. Let* $\mathcal{D}$ *be a distribution over* X*, let* S *be a database containing* n *i.i.d. elements from* $\mathcal{D}$*, and let* $h \leftarrow \mathcal{A}(S)$*. Then,*

$$\Pr_{\substack{S\sim\mathcal{D}\\ h\leftarrow\mathcal{A}(S)}}\left[e^{-2\varepsilon}\cdot h(\mathcal{D}) - h(S) > \frac{10}{\varepsilon n}\log\left(\frac{1}{\varepsilon\delta n}\right)\right] < O(\varepsilon\delta n).$$

We next state known impossibility results for privately learning threshold and parity function.

Fact 2.12 ([11, 30])**.** Let $b \in \mathbb{N}$. Any ε-differentially private (α,β)-PAC learner for $\mathsf{Threshold}_b$ requires $\Omega(\frac{b}{\varepsilon\alpha})$ many samples.

Fact 2.13 ([43])**.** Let $c \in \mathbb{N}$. In any ε-differentially private (α,β)-PAC learning protocol for Parity_c *in the local model* the number of messages is $\Omega(2^{c/3})$. This holds even when the underlying distribution is restricted to be the uniform distribution.

Fact 2.13 implies, for example, that when there are $\mathrm{poly}(c)$ agents the number of rounds is $2^{\Omega(c)}$. It is open if there exists an ε-private protocol (or an (ε,δ)-private protocol) for learning Parity_c in the local model with $\mathrm{poly}(c)$ agents and any number of rounds.

Remark 2.14. The proof of Fact 2.13 in [43] is stated in a weaker model, where in each round the referee sends an ε_i-differentially private local randomizer to an agent and the agent sends the output of this randomizer on its input to the referee, such that $\varepsilon_1 + \cdots + \varepsilon_\ell \le \varepsilon$. However, in their proof they only use the fact that $\varepsilon_i \le \varepsilon$ in every round, thus, their lower bound proof also applies to our model.

Our protocols use the private learner of [43] for parity functions, a protocol of [8] for answering all threshold queries, a protocol of [16] for heavy hitters, and a protocol of [31] for approximating a quantile. These are specified in the following theorems.

Theorem 2.15 ([8])**.** *Let* $\alpha,\beta,\varepsilon \le 1$*, and let* $b \in \mathbb{N}$*. There exists a non-interactive* ε*-differentially private protocol in the local model with* $n = O\left(\frac{b^3}{\alpha^2\varepsilon^2}\cdot\log\left(\frac{b}{\alpha\beta\varepsilon}\right)\right)$ *agents in which the input of every agent is a single element from* $\{0,1\}^b$ *and the outcome is a function* $q : \{0,1\}^b \to [0,1]$ *such that for every joint input to the protocol* $\mathbf{D} \in (\{0,1\}^b)^n$*, with probability at least* $1-\beta$*, the outcome* q *is such that* $\forall w \in \{0,1\}^b$ *we have* $q(w) \in |\{x \in \mathbf{D} : x \le w\}|/|\mathbf{D}| \pm \alpha$*.*

Remark 2.16. Theorem 2.15 does not appear explicitly in [8], but it is implicit in their analysis. In more details, Bassily et al. [8] presented a protocol, named `TreeHist`, for identifying *heavy*

hitters in the input database $\mathbf{D} \in (\{0,1\}^b)^n$. TreeHist works by privately estimating for each possible prefix $p \in \{0,1\}^{b'}$ (for $b' \leq b$) the number of input items that agree on the prefix p. Once these estimated counts are computed, [8] simply identified the input items $p \in \{0,1\}^b$ with large multiplicities in the data. Theorem 2.15 is obtained from the protocol TreeHist (with basically the same analysis) by observing that these estimated counts (for every possible prefix) in fact give estimations for the number of input items within any given *interval*. This observation has been used several times in the literature, see, e.g., [26].

Theorem 2.17 ([43])**.** *Let $\alpha, \beta, \varepsilon \leq 1$, and let $c \in \mathbb{N}$. There exists an ε-differentially private algorithm in the curator model that (α,β)-PAC learns and (α,β)-empirically learns* Parity_c *properly with sample complexity $O\left(\frac{c}{\alpha\varepsilon}\log(\frac{1}{\beta})\right)$.*

In fact, the algorithm of [43] privately produces a hypothesis with small error (w.h.p.) for every fixed input sample that is consistent with some parity function.

Theorem 2.18 (Heavy hitters protocol [16])**.** *There exist constants $\lambda_1, \lambda_2 > 0$ such that the following holds. Let $\beta, \varepsilon \leq 1$ and X be some finite domain. There exists a non-interactive ε-differentially private protocol in the local model with n agents in which the input of each agent is a single element from X and the outcome is a list* Est *of elements from X such that for every joint input to the protocol $\mathbf{D} \in X^n$, with probability at least $1-\beta$, every x that is an input of at least $\frac{\lambda_1}{\varepsilon}\sqrt{n\log\left(\frac{|X|}{\beta}\right)}$ agents appears in* Est*, and vice versa, every element x in* Est *is an input of at least $\frac{\lambda_2}{\varepsilon}\sqrt{n\log\left(\frac{|X|}{\beta}\right)}$ agents.*

Theorem 2.19 ([31, Theorem 17])**.** *Let $\mathcal{P}$ be any distribution on the real line. Fix any $p^* \in (0,1)$ and let $Q_{\min}, Q_{\max}, q^*$ be such that $\Pr_{x\sim\mathcal{P}}[x \leq q^*] = p^*$ and $q^* \in [Q_{\min}, Q_{max}]$. For any $\varepsilon > 0$ and for any $\lambda_{\text{quant}}, \tau_{\text{dist}}, \beta_{conf} \in (0, 1/2)$, there exists an interactive protocol in the local model with $T = \lceil \log_2(\frac{Q_{\max}-Q_{\min}}{\tau_{\text{dist}}})\rceil$ rounds that takes N i.i.d. draws from $\mathcal{P}$, where $N \geq \frac{8T}{\lambda_{\text{quant}}^2}\left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\right)^2\log(4T/\beta_{conf})$, and with probability at least $1-\beta_{\text{conf}}$ it returns $\tilde{q}$ such that either $|\tilde{q}-q^*| \leq \tau_{\text{dist}}$ or the probability mass $\mathcal{P}$ places in-between $\tilde{q}$ and q^* is at most λ_{quant}.*

3 Learning parity XOR threshold

In this section we present a learning task that cannot be solved privately in the curator model or in the local model, but can be solved in the hybrid model (without interaction). The task we consider in this section – parity XOR threshold – is similar to the simpler task of the direct product of parity and threshold discussed in Appendix A. In this section we design a non-interactive protocol in the hybrid model for the parity XOR threshold task, which is more involved than the trivial protocol for the parity and threshold task. This demonstrates that non-interactive protocols in the hybrid model may have more power than one might initially suspects.

Fix $b, c > 0$, and let $k \in \{0,1\}^c$ and $t \in \{0,1\}^b$ be parameters. Define the function $f_{b,c}^{k,t} : \{0,1\}^c \times \{0,1\}^b \to \{0,1\}$ as follows: $f_{b,c}^{k,t}(x,y) = \mathsf{Par}_k(x) \oplus \mathsf{Thr}_t(y) = \langle k, x\rangle \oplus \mathbb{1}_{\{y\geq t\}}$ (recall that we treat strings in $\{0,1\}^b$ as integers in $\left\{0,1,\ldots,2^b-1\right\}$). Define the concept class ParityThresh as follows:

$$\mathsf{ParityThresh}_{b,c} = \left\{f_{b,c}^{k,t} \;:\; k \in \{0,1\}^c \text{ and } t \in \{0,1\}^b\right\}.$$

We first show that every differentially private algorithm (even in the curator model) for learning ParityThresh must have sample complexity $\Omega(b)$.

Lemma 3.1. *Every ε-differentially private algorithm for $(\frac{1}{4},\frac{1}{4})$-PAC learning $\mathsf{ParityThresh}_{b,c}$ must have sample complexity $\Omega(b)$.*

Proof. Let $\mathcal{A}$ be an ε-differentially private algorithm for $(\frac{1}{4},\frac{1}{4})$-PAC learning $\mathsf{ParityThresh}_{b,c}$ with sample complexity m. We now use $\mathcal{A}$ to construct an ε-differentially private algorithm $\mathcal{B}$ for $(\frac{1}{4},\frac{1}{4})$-PAC learning $\mathsf{Threshold}_b$ with the same sample complexity. By Fact 2.12 this will show that $m=\Omega(b)$.

Algorithm $\mathcal{B}$ is simple: it takes an input database $S=\{(y_i,\sigma_i)\}_{i=1}^m \in \left(\{0,1\}^b\times\{0,1\}\right)^m$ and runs $\mathcal{A}$ on the database $\hat{S}=\left\{(\vec{0},y_i,\sigma_i)\right\}_{i=1}^m \in \left(\{0,1\}^c\times\{0,1\}^b\times\{0,1\}\right)^m$ to obtain a hypothesis $\hat{h}$. Then, algorithm $\mathcal{B}$ returns the hypothesis h defined as $h(y)=\hat{h}(\vec{0},y)$. Note that changing one element of S changes exactly one element of $\hat{S}$, and hence algorithm $\mathcal{B}$ is ε-differentially private.

We next show that algorithm $\mathcal{B}$ is a $(\frac{1}{4},\frac{1}{4})$-PAC learner for $\mathsf{Threshold}_b$. To that end, fix a target distribution $\mathcal{D}$ on $\{0,1\}^b$ and fix a target concept Thr_t (where $t\in\{0,1\}^b$). Suppose that S contains i.i.d. samples from $\mathcal{D}$ that are labeled by Thr_t, and consider the following distribution $\hat{\mathcal{D}}$: To sample an element from $\hat{\mathcal{D}}$ we sample $y\sim\mathcal{D}$ and return $\vec{0}\circ y\in\{0,1\}^{b+c}$. Now observe that $\hat{S}$ contains m i.i.d. samples from $\hat{\mathcal{D}}$ which are labeled by $f_{b,c}^{\vec{0},t}\in\mathsf{ParityThresh}_{b,c}$. Therefore, by the utility properties of $\mathcal{A}$, with probability at least 3/4 the hypothesis $\hat{h}$ satisfies $\mathrm{error}_{\hat{\mathcal{D}}}\left(\hat{h},f_{b,c}^{\vec{0},t}\right)\le\frac{1}{4}$. In that case,

$$\frac{1}{4}\le \mathrm{error}_{\hat{\mathcal{D}}}\left(\hat{h},f_{b,c}^{\vec{0},t}\right)=\Pr_{y\sim\mathcal{D}}\left[\hat{h}\left(\vec{0},y\right)\neq f_{b,c}^{\vec{0},t}\left(\vec{0},y\right)\right]=\Pr_{y\sim\mathcal{D}}\left[h(y)\neq\mathsf{Thr}_t(y)\right]=\mathrm{error}_{\mathcal{D}}(h,\mathsf{Thr}_t).$$

This shows that $\mathcal{B}$ is a $(\frac{1}{4},\frac{1}{4})$-PAC learner for $\mathsf{Threshold}_b$, as required. □

We next show that no protocol in the local model can learn ParityThresh, unless the number of exchanged messages is very large.

Lemma 3.2. *In every ε-differentially private protocol in the local model for $(\frac{1}{4},\frac{1}{4})$-PAC learning $\mathsf{ParityThresh}_{b,c}$ the number of messages is $\Omega(2^{c/3})$.*

The proof of Lemma 3.2 is analogous to the proof of Lemma 3.1 (using Fact 2.13 instead of Fact 2.12).

So, privately learning $\mathsf{ParityThresh}_{b,c}$ in the curator model requires $\Omega(b)$ labeled examples, and privately learning it in the local model requires $\Omega(2^{c/3})$ messages. We now show that $\mathsf{ParityThresh}_{b,c}$ can be learned privately by a non-interactive protocol in the hybrid model with roughly $O(c)$ examples for the curator and with roughly $O(b^3)$ local agents. We will focus on the case where $c\ll b$. Recall that a function $f_{b,c}^{k,t}(x,y)\in\mathsf{ParityThresh}_{b,c}$ is defined as $f_{b,c}^{k,t}(x,y)=\mathsf{Par}_k(x)\oplus\mathsf{Thr}_t(y)$. The difficulty in learning ParityThresh in the hybrid model is that we could only learn the threshold part of the target function using the local agents (since if $c\ll b$ then the curator does not have enough data to learn it), but the threshold label is "hidden" from the local agents (because it is "masked" by the parity bit that the local agents cannot learn). This false intuition might lead to the design of an *interactive* protocol, in which the referee first obtains some information from the curator and then passes this information to the local agents, which would allow them to learn the threshold part of the target function. We now show that such an interaction is not needed, and design a *non-interactive* protocol in which the local agents and the curator communicate with the referee only once, simultaneously.

Lemma 3.3. *There exists a non-interactive ε-differentially private protocol in the (m,n)-hybrid model for (α,β)-PAC learning* $\mathsf{ParityThresh}_{b,c}$ *where $m = O\left(\frac{c}{\alpha^4\varepsilon}\log(\frac{1}{\alpha\beta})\right)$ and $n = O\left(\frac{b^3}{\alpha^4\varepsilon^2}\cdot\log\left(\frac{b}{\alpha\beta\varepsilon}\right)\right)$.*

Proof. We begin by describing a non-interactive protocol Π. The (joint) input to the protocol is a database $\mathbf{D}$ where every point in $\mathbf{D}$ is of the form $(x_i,y_i,\sigma_i)\in\{0,1\}^c\times\{0,1\}^b\times\{0,1\}$. At a high level, the protocol works by using the local agents to obtain an approximation to the CDF of the (marginal) distribution on the y_i's (this approximation is given to the referee). In addition, the trusted curator solves $1/\alpha$ parity leaning problems. In more details, the trusted curator sorts its database according to the y_i's, divides its database into $1/\alpha$ chunks, and then applies a private learner for parity functions on each of the chunks. The trusted curator sends the referee the resulting $1/\alpha$ parity functions. The referee then defines the final hypothesis h that, given a point (x,y), first uses the approximation to the CDF (obtained fro the local agents) to match this input point to one of the chunks, and then uses the parity function obtained for that chunk from the trusted curator to predict the label of the point.

The key observation here is that the threshold part of the target function is *constant* on all but at most one of the chunks defined by the trusted curator. As we show, applying a learner for parity on such a "consistent chunk" results in a good predictor for the labels of elements of that chunk. Hence, provided that the approximation for the CDF of the y_i's is accurate enough, this results in an accurate learner for ParityThresh. We now formally present the protocol Π.

- The local agents on a (distributed) input $D=(x_i,y_i,\sigma_i)_{i=1}^n\in\left(\{0,1\}^c\times\{0,1\}^b\times\{0,1\}\right)^n$:
  Run the protocol from Theorem 2.15 on the (distributed) database $\hat{D}=(y_1,y_2,\ldots,y_n)$ with privacy parameter ε and utility parameters α^2,β. At the end of the execution, the referee obtains a function $q:\{0,1\}^b\to[0,1]$ that approximates all threshold queries w.r.t. $\hat{D}$.
- The curator on input $S=(x_i,y_i,\sigma_i)_{i=1}^m\in\left(\{0,1\}^c\times\{0,1\}^b\times\{0,1\}\right)^m$:
  - Sort S according to the y_i's in non-decreasing order.
  - Divide S into blocks of size αm: $S_1,S_2,\ldots,S_{1/\alpha}$. For $\ell\in[1/\alpha]$ we denote $S_\ell=(x_{\ell,i},y_{\ell,i},\sigma_{\ell,i})_{i=1}^{\alpha m}$.
  - For every $\ell\in[1/\alpha]$, apply an $\alpha\varepsilon$-differentially private $(\alpha^2,\alpha\beta)$-PAC learner for Parity on the database $\hat{S}_\ell=(x_{\ell,i}\circ 1,\sigma_{\ell,i})_{i=1}^{\alpha m}\in\left(\{0,1\}^{c+1}\times\{0,1\}\right)^{\alpha m}$ to obtain a vector $k_\ell\in\{0,1\}^{c+1}$ (using Theorem 2.17).
  - Send $k_1,\ldots,k_{1/\alpha}$ to the referee.
- The referee:
  - Obtain the function q and the vectors $k_1,\ldots,k_{1/\alpha}$.
  - Define a hypothesis $h:\{0,1\}^c\times\{0,1\}^b\to\{0,1\}$ as $h(x,y)=\langle x\circ 1,k_{I(y)}\rangle$, where $I(y)=\left\lceil\frac{q(y)}{\alpha}\right\rceil$.
  - Output h.

The privacy properties of the protocol Π are straightforward, as both the local agents and the curator apply ε-differentially private computations: The local agents apply the algorithm from Theorem 2.15, and the curator applies an $\alpha\varepsilon$-differentially private computation on each of the blocks $S_1,\ldots,S_{1/\alpha}$ (note that changing one element of S can change at most one element of each of these blocks).

We now proceed with the utility analysis. Fix a target function $f_{b,c}^{k^*,t^*} \in \mathsf{ParityThresh}_{b,c}$ and fix a target distribution $\mathcal{D}$ on $\{0,1\}^c \times \{0,1\}^b$. We use $\mathcal{D}_c$ and $\mathcal{D}_b$ to denote the marginal distributions on $\{0,1\}^c$ and $\{0,1\}^b$, respectively. We will make the simplifying assumption that $\mathcal{D}_b$ does not give too much weight on any single point in $\{0,1\}^b$, specifically, $\Pr_{w\sim\mathcal{D}_b}[w = y] \le \beta/m^2$ for every $y \in \{0,1\}^b$. This assumption can be enforced by padding every example with $O(\log(m/\beta))$ uniformly random bits.

Let S and D (the inputs to the curator and the local agents) be sampled i.i.d. from $\mathcal{D}$ and labeled by $f_{b,c}^{k^*,t^*}$. We next show that w.h.p. the resulting hypothesis h has low empirical error on S. By standard generalization arguments, such an h also has low generalization error.

First observe that there is at most one index $\ell^* \in [1/\alpha]$ such that $\mathsf{Thr}_{t^*}(y_{\ell^*,1}) \neq \mathsf{Thr}_{t^*}(y_{\ell^*,\alpha m})$. In all other blocks S_ℓ we have that $\mathsf{Thr}_{t^*}(\cdot)$ is constant on all the $y_{\ell,i}$'s of that block. We will show that w.h.p. the hypothesis h has small empirical error on every such block. Fix $\ell \neq \ell^*$, and let $\nu \in \{0,1\}$ be the value of $\mathsf{Thr}_{t^*}(\cdot)$ on the $y_{\ell,i}$'s of the ℓth block (that is, for every $i \in [\alpha m]$ we have $\mathsf{Thr}_{t^*}(y_{\ell,i}) = \nu$). Recall that since the elements of S are labeled by $f_{b,c}^{k^*,t^*}$, for every $i \in [\alpha m]$ we have that

$$\sigma_{\ell,i} = f_{b,c}^{k^*,t^*}(x_{\ell,i}, y_{\ell,i}) = \langle k^*, x_{\ell,i}\rangle \oplus \mathsf{Thr}_{t^*}(y_{\ell,i}) = \langle k^*, x_{\ell,i}\rangle \oplus \nu = \langle k^* \circ \nu,\ x_{\ell,i} \circ 1\rangle.$$

Hence, the elements of $\hat{S}_\ell$ are all labeled by the parity function defined by $k^* \circ \nu$. Therefore, as k_ℓ is the outcome of the learner form Theorem 2.17 on $\hat{S}_\ell$, for $m \ge O\left(\frac{c}{\alpha^2\varepsilon}\log(\frac{1}{\alpha\beta})\right)$, with probability at least $1-\alpha\beta$ we have that $\mathrm{error}_{\hat{S}_\ell}(\mathsf{Par}_{k_\ell}) \le \alpha^2$, that is, $\langle k_\ell,\ x \circ 1\rangle$ is a good predictor for the label of the elements in block S_ℓ.

Recall that the hypothesis h matches inputs (x,y) to the vectors $k_1,\ldots,k_{1/\alpha}$ using the function q obtained from the local agents, that is, on input (x,y), the hypothesis uses $k_{\lceil q(y)/\alpha\rceil}$. Therefore, to complete the proof we need to show that most of the elements from block S_ℓ are matched by the hypothesis h to the vector k_ℓ. To that end, let $\#_S(w) = |\{(x,y,\sigma) \in S : y \le w\}|$, and consider the following event:

$$\textbf{Event } \boldsymbol{E}_1: \qquad \forall w \in \{0,1\}^b \text{ it holds that } \left|q(w) - \frac{1}{m}\cdot \#_S(w)\right| \le 4\alpha^2.$$

We first conclude the proof assuming that Event E_1 occurs. Fix $\ell \neq \ell^*$, and recall that the elements of S (and in particular the elements of S_ℓ) are sorted in a non-decreasing order according to their y_i's. Now fix $8\alpha^2 m \le i \le \alpha m - 8\alpha^2 m$. By our simplifying assumption (that the distribution $\mathcal{D}_b$ does not put a lot of mass on any single point), we may assume that all the y_i's in S are distinct, which happens with probability at least $1-\beta$. In that case, we have that $\#_S(y_{\ell,i}) = \max\{0,\ell-1\}\cdot \alpha m + i$, and hence,

$$\max\{0,\ell-1\}\cdot\alpha + 8\alpha^2 \le \frac{1}{m}\#_S(y_{\ell,i}) \le \max\{0,\ell-1\}\cdot\alpha + \alpha - 8\alpha^2.$$

By Event E_1 we get that

$$\max\{0,\ell-1\}\cdot\alpha + 4\alpha^2 \le q(y_{\ell,i}) \le \max\{0,\ell-1\}\cdot\alpha + \alpha - 4\alpha^2,$$

and so, $\left\lceil \frac{q(y_{\ell,i})}{\alpha}\right\rceil = \ell$. That is, for all but at most $16\alpha^2 m$ elements of the block S_ℓ we get that $h(x_{\ell,i}, y_{\ell,i}) = \langle x_{\ell,i} \circ 1,\ k_\ell\rangle = \mathsf{Par}_{k_\ell(x_{\ell,i},y_{\ell,i})}$. Recall that Par_{k_ℓ} errs on at most $\alpha^2 m$ elements of S_ℓ, and so the hypothesis h errs on at most $17\alpha^2 m$ elements of the block S_ℓ. That is, h errs on at most

$17\alpha^2 m$ elements of every block S_ℓ for $\ell \neq \ell^*$, and might err on all of S_{ℓ^*} which is of size αm. So, h errs on at most $\frac{1}{\alpha} \cdot 17\alpha^2 m + \alpha m = 18\alpha m$ elements of S. Standard generalization bounds now state that, except with probability at most β, we get that $\text{error}_{\mathcal{D}}(h, f_{b,c}^{k^*,t^*}) \leq O(\alpha)$ (in particular, this follows from the generalization properties of differential privacy; see Section 2.4 for more details). Overall, with probability at least $1 - O(\beta)$ the resulting hypothesis has generalization error at most $O(\alpha)$.

It remains to show that Event E_1 occurs with high probability. First, for $n \geq O\left(\frac{b^3}{\alpha^4\varepsilon^2} \cdot \log\left(\frac{b}{\alpha\beta\varepsilon}\right)\right)$, Theorem 2.15 ensures that with probability at least $1 - \beta$ the function q is such that $\forall w \in \{0,1\}^b$ it holds that $\left|q(w) - \frac{1}{n}\#_{\hat{D}}(w)\right| \leq \alpha^2$, where $\#_{\hat{D}}(w) = |\{y \in \hat{D} : y \leq w\}|$. Second, by standard generalization arguments, assuming that n and m are big enough, we would also have that $\frac{1}{n}\#_{\hat{D}}(w)$ and $\frac{1}{m}\#_S(w)$ are both within α^2 from $\Pr_{y\sim\mathcal{D}_b}[y \leq w]$. Specifically, by the Dvoretzky-Kiefer-Wolfowitz inequality [23, 45], assuming that n and m are at least $\Omega\left(\frac{1}{\alpha^4}\log(\frac{1}{\beta})\right)$, this happens with probability at least $1-\beta$. Assuming that this is the case, by the triangle inequality we have that Event E_1 holds. This shows that Event E_1 happens with probability at least $1 - 3\beta$, and completes the proof. □

We remark that it is possible to design a more efficient learner for ParityThresh (in terms of sample complexity) by constructing a protocol in which there are multiple rounds of communication between the referee and the local agents (but this communication is still independent from the message that the curator sends to the referee). This will be illustrated in Appendix A. We summarize our possibility and impossibility results w.r.t. learning ParityThresh in the next theorem (which follows from Lemma 3.1 and Lemma 3.2 and from Lemma 3.3).

Theorem 3.4. *Let $c \in \mathbb{N}$ and $b = c^2$. Then there is a non-interactive $\frac{1}{4}$-differentially private $(\frac{1}{4}, \frac{1}{4})$-PAC learner for $\text{ParityThresh}_{b,c}$ in the (m,n)-hybrid model with $m = O(c)$ samples for the curator and $n = O(c^6 \log c)$ local agents. However, every such learner in the local model with $o(2^{(n/\log n)^{1/6}})$ local agents requires $2^{\Omega((n/\log n)^{1/6})}$ rounds, and every such learner in the curator model requires $\Omega(m^2)$ samples.*

4 The 1-out-of-2^d-parity task

In this section we describe a task that cannot be privately solved neither in the curator model nor in the local model with sub-exponential number of rounds. In the hybrid model, this task can only be solved with interaction, first with the local agents and then with the curator. In this task there are many instances of the parity problem and the referee needs to solve only one instance, which is determined by the inputs. The local agents can determine this instance (using a heavy hitters protocol) and the curator can now solve this instance. The curator cannot solve all instances since this will exceed its privacy budget, and by the definition of the task the curator will not have enough information to determine the instance; thus interaction with both the local agents and the curator is required.

Definition 4.1 (The 1-out-of-2^d-parity task)**.** *The inputs in the 1-out-of-2^d-parity task are generated as follows:*

1. **Input:** *2^d strings $(r_j)_{j\in\{0,1\}^d}$, where $r_j \in \{0,1\}^c$ for every $j \in \{0,1\}^d$, and $m+1$ elements $s_1, \ldots, s_{m+1} \in \{0,1\}^d$.*

2. *Set* $s = s_1 \oplus \cdots \oplus s_{m+1}$.[4]

3. *Each sample* $x_1, \ldots, x_m$ *and* $y_1, \ldots, y_n$ *is generated independently as follows:*

   - *with probability half choose* $x \in_R \{0,1\}^c$ *with uniform distribution and output* $(x, (\langle x, r_j \rangle)_{j \in \{0,1\}^d})$ *(that is, every point contains a string* x *of length* c *and* 2^d *bits which are the inner products of* x *and each of the* r_j*'s).*
   - *with probability half choose* $t \in_R [m+1]$ *with uniform distribution and output* (t, s_t) *(that is, every point contains a number* t *and the* t*-th string* s_t*).*

The goal of the referee in the 1-out-of-2^d*-parity task is for every* $(r_j)_{j \in \{0,1\}^d}$ *and* $s_1, \ldots, s_{m+1}$ *to recover* r_s *with probability at least* $1 - \beta$*, where the probability is over the generation of the inputs in Step 3 and the randomness of the parties in the protocol.*

We start by describing a protocol for this task.

Lemma 4.2. *Let* $\beta > 1/m$ *and assume that* $m = \Omega\left(\frac{c \log(1/\beta)}{\varepsilon}\right)$ *and* $n = \Omega\left(\frac{m^2}{\varepsilon^2} \log(\frac{m2^d}{\beta})\right)$*. The 1-out-of-*$2^d$*-parity task can be solved in the* (m, n)*-hybrid model by an* ε*-differentially private protocol with three rounds, where in the first round each local agent sends one message to the referee (without seeing any other messages), in the second round the referee sends one message to the curator, and in the third round the curator sends one message to the referee.*

Proof. The protocol is as follows: In the first round the local agents send messages according to the ε-differentially private heavy hitters protocol of Theorem 2.18 (from [16]) with the inputs (t, s_t) and $\beta/3$, that is, a protocol that returns with probability at least $1 - \beta/3$ all values that are inputs of "many" agents. If the input of a local agent is not (t, s_t) for some t, then it executes the protocol with some default input $\perp$. The referee reconstructs the $m+1$ strings $s_1, \ldots, s_{m+1}$ (with probability at least $1 - 2\beta/3$), computes $s = s_1 \oplus \cdots \oplus s_{m+1}$, and sends it to the curator. The curator privately solves the parity task with inputs $(x, \langle x, r_s \rangle)$ using the algorithm of Theorem 2.17 (from [43]) with $\alpha = 1/4$ and $\beta/3$. Since we use ε-differentially private algorithms, each operating on different inputs, the resulting protocol is ε-differentially private.

We next argue that with probability at least $1 - \beta$, the referee reconstructs r_s. Note that for a fixed $t \in [m+1]$, the expected number of times that (t, s_t) is an input of agents $P_1, \ldots, P_n$ is $n/2(m+1)$. By a simple Chernoff bound, with probability $1 - \beta/3$ for all t, the value (t, s_t) is an input of at least $n/4(m+1)$ parties. By Theorem 2.18, with probability at least $1 - \beta/3$, each value that is an input of at least $O\left(1/\varepsilon \sqrt{n \log(\frac{m2^d}{\beta/3})}\right)$ agents will appear in the list computed by the referee. By the assertion of the lemma, $O\left(1/\varepsilon \sqrt{n \log(\frac{m2^d}{\beta/3})}\right) < \frac{n}{4(m+1)}$. Thus, with probability at least $1 - 2\beta/3$, the referee reconstructs all s_t's and reconstructs the correct value s. Furthermore, as $m = \Omega(\frac{c \log(3/\beta)}{\varepsilon})$, the algorithm of Theorem 2.17 returns, with probability at least $1 - \beta/3$, a string r such that $\Pr[\langle x, r \rangle \neq \langle x, r_s \rangle] \leq 1/4$ under the uniform distribution on $x \in \{0,1\}^c$. Since for $r \neq r_s$ this probability is exactly $1/2$, we get that $r = r_s$ with probability at least $1 - \beta$. □

[4] The strings $s_1, \ldots, s_{m+1}$ are an $m+1$-out-of-$m+1$ secret sharing of s, that is, together they determine s, but every subset of them gives no information on s.

We next prove that, unless the database is big, the 1-out-of-2^d-parity task requires interaction. To prove this result, we first convert a protocol for the 1-out-of-2^d-parity task to a private algorithm in the trusted curator model that recovers all strings $(r_j)_{j\in\{0,1\}^d}$. We then prove, using a simple packing argument, that, unless the database is "big", such algorithm cannot exist. For our proof, we define the all-2^d-parity task as the task in which all inputs are of the form $(x,(\langle x,r_j\rangle)_{j\in\{0,1\}^d})$ and the goal of the referee is to reconstruct all strings $(r_j)_{j\in\{0,1\}^d}$.

Claim 4.3. *Let $m<n$. If there is an ε-differentially private protocol for the 1-out-of-2^d-parity problem in the (m,n)-hybrid model in which the curator and the referee can exchange many messages and then the referee simultaneously sends one message to each local agent and gets one answer from each agent, then there is an ε-differentially private algorithm in the trusted curator model for the all-2^d-parity problem for a database of size $O(nd)$.*

Proof. Let Π be an ε-differentially private protocol with the above interaction pattern for the 1-out-of-2^d-parity task in the (m,n)-hybrid model in which the referee reconstructs r_s with probability at least $1-\beta$. We construct, in three steps, an algorithm $\mathcal{A}$ for the all-2^d-parity task in the trusted curator.

First, we construct from Π a protocol Π' in the $(O(md),O(dn))$-hybrid model that reconstructs r_s with error probability at most $\beta/2^d$ (e.g., execute Π with disjoint inputs $O(d)$ times and take the value r_s that is returned in the majority of the executions).

Next, we construct from Π' a protocol Π'' for the the all-2^d-parity task in the $(O(md),O(nd))$-hybrid model (with error probability at most β). In Π'', the parties holding inputs of the all-2^d-parity problem simulate Π' on inputs for the 1-out-of-2^d-parity task as follows:

- The curator on input $(x_i,(\langle x_i,r_j\rangle)_{j\in\{0,1\}^d})_{i=1}^m$:
  - Chooses random $s_1,\ldots,s_{m+1}\in_R\{0,1\}^d$.
  - For each $i\in[m]$, with probability 1/2 replaces its i-th input by (t_i,s_{t_i}) for a uniformly distributed $t_i\in_R[m+1]$.
  - Exchanges messages with the referee as specified by Π' on the new input. In addition it sends to the referee $s_1,\ldots,s_{m+1}$ and an index ℓ such that (ℓ,s_ℓ) does not appear in its new input.
- The referee after getting the message from the curator:
  - Chooses a set $A\subseteq[n]$ with uniform distribution.
  - For every $i\notin A$, sends its message in Π' to the i-th agent and gets an answer M_i from the agent.
  - For every $i\in A$, chooses a random $q_i\in_R[m+1]$. Let $B=\{i\in A: q_i=\ell\}$.
  - For every $i\in A\setminus B$, computes (without any interaction) its message in Π' to agent P_i and the answer M_i of agent P_i with input (q_i,s_{q_i}).
  - For every $i\in B$ and $s\in\{0,1\}^d$, computes (without any interaction) its message in Π' to agent P_i and the answer $M_{i,s}$ of agent P_i with input $(\ell,s\oplus\bigoplus_{k\neq\ell}s_k)$.
  - For every $s\in\{0,1\}^d$, reconstructs r_s from the messages of the curator in Π', $(M_i)_{i\notin B}$, and $(M_{i,s})_{i\in B}$.

As the curator holds m samples and there are $m+1$ values $s_1,\ldots,s_{m+1}$, there exists an index ℓ such that (ℓ, s_ℓ) does not appear in the new input of the curator. Thus, the referee for every $s \in \{0,1\}^d$ can choose a value s'_ℓ such that it is consistent with the messages of the curator in Π' and $s = s'_\ell \oplus \bigoplus_{k\neq\ell} s_k$. Furthermore, each of $x_1,\ldots,x_m,y_1,\ldots,y_n$ is replaced with probability half with a value (t,s_t) for a uniformly distributed t, thus, these inputs are distributed as required for the 1-out-of-2^d-parity task. This implies that for every $s \in \{0,1\}^d$ the referee reconstructs r_s from the messages of the curator in Π', $(M_i)_{i\notin B}$, and $(M_{i,s})_{i\in B}$ with probability at least $1-\beta/2^d$. By the union bound, the referee correctly reconstructs all $(r_j)_{j\in\{0,1\}^d}$ with probability at least $1-\beta$.

Finally, we construct the desired algorithm $\mathcal{A}$ from Π''. The trusted curator simply simulates the referee, the curator, and the agent in Π'', that is, it take its database with $O((m+n)d)$ samples and partitions it to $(x_1,\ldots,x_{O(md)})$ (the input of the curator) and $y_1,\ldots,y_{O(nd)}$, computes without any interaction a random transcript of Π'' on these inputs, and reconstructs the output $(r_j)_{j\in\{0,1\}^d}$. Since the transcript preserves ε-differential privacy and computing the output is post-processing, algorithm $\mathcal{A}$ is ε-differential private. □

Claim 4.4. *If there is exists an ε-differentially private algorithm in the trusted curator model for the all-2^d-parity task with strings of length c, then $n = \Omega\left(\frac{c2^d+\ln(1-\beta)}{\varepsilon}\right)$.*

Proof. The proof is by a simple packing argument. For every strings $(r_j)_{j\in\{0,1\}^d}$, with probability at least $1-\beta$, the algorithm returns $(r_j)_{j\in\{0,1\}^d}$ when the samples are generated with $(r_j)_{j\in\{0,1\}^d}$. By the group privacy of ε-differential privacy, with probability at least $e^{-n\varepsilon}(1-\beta)$ the algorithm returns $(r_j)_{j\in\{0,1\}^d}$ when the samples are generated with $(0^c)_{j\in\{0,1\}^d}$. As there are 2^{c2^d} options for $(r_j)_{j\in\{0,1\}^d}$ and the above events are disjoint, $2^{c2^d}e^{-n\varepsilon}(1-\beta) \leq 1$, i.e., $n = \Omega\left(\frac{c2^d+\ln(1-\beta)}{\varepsilon}\right)$. □

Lemma 4.5. *Let $m < n$. If there is an ε-differentially private protocol for the 1-out-of-2^d-parity task in the (m,n)-hybrid model with $\beta = 1/4$ in which the curator and the referee can exchange many messages and then the referee simultaneously sends one message to each local agent and gets one answer from each agent, then $n = \Omega(c2^d/d\varepsilon)$.*

Proof. By Claim 4.3, if there exists an ε-differentially protocol in the (m,n)-hybrid model for the 1-out-of-2^d-parity task, then there exists an ε-differentially private algorithm in trusted curator model for the all-2^d-parity task with database of size $O(dn)$. Thus, by Claim 4.4 with $\beta = 1/4$, $dn = \Omega(\frac{c2^d}{\varepsilon})$. □

Lemma 4.5 is valid also if the local agents are allowed to hold a shared (public) random string as this string can be sent by the referee to each agent as part of its message (without adding extra rounds of communication).

We summarize the possibility and impossibility results for the 1-out-of-2^d-parity task in the following theorem, where, for convenience, we choose specific parameters that highlight these results.

Theorem 4.6. *Let $\varepsilon = 1/4$, $\beta = 1/4$. For every integer c, there are $d = \Theta(\log c)$, $m = \Theta(c)$, and $n = \Theta(c^2 \log c)$ such that*

1. *There exists an ε-differentially private protocol for the 1-out-of-2^d-parity task with strings of length c in the (m,n)-hybrid model where first each local agent sends one message to the referee and then the referee exchanges one message with the curator.*

2. *There does not exist an ε-differentially private protocol for this task in the (m,n)-hybrid model in which the referee first exchanges messages with the curator and then simultaneously exchanges one message with the local agents.*

3. *In any ε-differentially private protocol for this task in the local model with n agents the number of rounds is $2^{\Omega(c)} = 2^{\Omega(\sqrt{n/\log n})}$.*

4. *There is no algorithm in the trusted curator model that solves this task with m examples.*

Proof. Item 1 follows directly from Lemma 4.2. For Item 2, by Lemma 4.5, with $\varepsilon = 1/4$, $\beta = 1/4$, and $d = 2\log c + \log\log c$

$$n = \Omega\left(\frac{c2^d}{d\varepsilon}\right) = \Omega(c^3),$$

contradicting the choice of $n = \Theta(c^2 \log c)$.

For the impossibility result in Item 3, recall that by Fact 2.13 the number of messages sent to the referee in an ε-differentially private learning protocol in the local model for parity of strings of length c with respect to the uniform distribution is $2^{\Omega(c)}$. By simple simulation, an ε-differentially private protocol in the local model for the 1-out-of-2^d-parity task implies an ε-differentially private protocol in the local model for learning parity with respect to the uniform distribution (with the same number of messages). Specifically, since the number of agents is $n = O(c^2 \log c)$, the number of rounds is $2^{\Omega(c)}/(c^2 \log c) = 2^{\Omega(\sqrt{n/\log n})}$.

For Item 4, observe that a curator receiving m input points obtains less than $m+1$ shares of s and hence obtains no information about r_s. Hence, such a curator cannot solve the 1-out-of-2^m-parity task alone, even without privacy constraints. □

5 The parity-chooses-secret task

We now present another task that cannot be privately solved neither in the curator model nor in the local model with sub-exponential number of rounds. This task can be solved in the hybrid model; however, it requires interaction, this time first with the curator and then with the local agents. This task (as well as 1-out-of-2^d-parity task) highlights both the information and private-computation gaps between the curator and the local model agents. The local model agents receive enough information to solve the task, but lack the ability to privately solve an essential sub-task. The curator does not receive enough information to solve the task (even non-privately), however the curator can be used to privately solve the hard sub-task. Once the hard sub-task is solved, this information is forwarded to the local agents, which now can solve the task.

Definition 5.1 (The parity-chooses-secret task)**.** *The inputs in the parity-chooses-secret task are generated as follows:*

1. **Input:** *A string $r \in \{0,1\}^c$ and 2^c vectors of $m+1$ bits: a vector $(s_{j,1},\ldots,s_{j,m+1}) \in \{0,1\}^{m+1}$ for every $j \in \{0,1\}^c$.*

2. *Set $s_j = s_{j,1} \oplus \cdots \oplus s_{j,m+1}$ for every for $j \in \{0,1\}^c$, i.e., s_j is a random bit shared via an $m+1$-out-of-$m+1$ secret-sharing scheme, with the shares being $s_{j,1},\ldots,s_{j,m+1}$.*

3. *Each sample $x_1,\ldots,x_m$ and $y_1,\ldots,y_n$ is generated independently as follows:*

- *Choose $x \in_R \{0,1\}^c$ and $t \in_R [m+1]$ and output $(x, \langle x, r\rangle, t, (s_{j,t})_{j\in\{0,1\}^c})$ (that is, every point contains a string of length c, its inner product with r, an integer t, and the t-th share of each s_j).*

The goal of the referee in the parity-chooses-secret task is for every r and every $\left((s_{j,1},\ldots,s_{j,m+1})\right)_{j\in\{0,1\}^c}$ to recover s_r with probability at least $1-\beta$, where the probability is over the generation of the inputs in Step 3 and the randomness of the parties.

We start by describing a protocol for the parity-chooses-secret task.

Lemma 5.2. *Let $\beta > 1/m$ and assume that $m = \Omega\left(\frac{c\log(1/\beta)}{\varepsilon}\right)$ and $n = \Omega\left(\frac{m^2}{\varepsilon^2}\log(\frac{n}{\varepsilon})\right)$. The parity-chooses-secret task can be solved in the (m,n)-hybrid model by an ε-differentially private protocol with three rounds, where in the first round the curator sends one message to the referee, in the second round the referee sends one message to the local agents, and in the third round each local agent sends one message to the referee.*

Proof. The protocol is as follows: First, the curator learns r by executing the ε-differentially private algorithm for learning parity of [43] (see Theorem 2.17) with $\alpha = 1/4$, $\beta/3$, and the m inputs $(x, \langle x, r\rangle)$. The curator sends r to the referee, which forwards it to the local agents. Next each local agent sends a messages according to the ε-differentially private heavy hitters protocol of [16] (see Theorem 2.18) with the input $(t, s_{r,t})$ and $\beta/3$. The referee recovers $(1, s_{r,1}),\ldots,(m+1, s_{r,m+1})$ from the messages of the heavy-hitters protocol, and outputs $ys_r = s_{r,1} \oplus \cdots \oplus s_{r,m+1}$. Since we use ε-differentially private algorithms, each operating on different inputs, the resulting protocol is ε-differentially private.

We next argue that with probability at least $1-\beta$, the referee reconstructs s_r. As $m = \Omega(\frac{c\log(1/\beta)}{\varepsilon})$, the algorithm of [43] (see Theorem 2.17) returns, with probability at least $1-\beta/3$, a string r' such that $\Pr[\langle x, r'\rangle \neq \langle x, r\rangle] \leq 1/4$ under the uniform distribution on $x \in \{0,1\}^c$. Since for $r \neq r'$ this probability is exactly $1/2$, we get that $r = r'$ with probability at least $1-\beta/3$. We complete the proof by showing that, once the local agents and the referee know r, the referee reconstructs with probability at least $2\beta/3$ all values $(1, s_{r,1}),\ldots,(m+1, s_{r,m+1})$. Note that for a fixed $t \in [m+1]$, the expected number of times that $(t, s_{r,t})$ is an input of agents $P_1,\ldots,P_n$ is $n/(m+1)$. By a simple Chernoff bound, with probability $1-\beta/3$, for all t the value $(t, s_{r,t})$ is an input of at least $n/2(m+1)$ parties. The protocol of [16] (see Theorem 2.18) guarantees that, with probability at least $1-\beta/3$, each value that is an input of at least $O\left(1/\varepsilon\sqrt{n\log(\frac{n}{\beta/3})}\right)$ agents will appear in the list computed by the referee. By the assertion of the lemma, $O\left(1/\varepsilon\sqrt{n\log(\frac{n}{\beta/3})}\right) < \frac{n}{2(m+1)}$. Thus, with probability at least $1-\beta$, the referee reconstructs $s_{r,1},\ldots,s_{r,m+1}$ and reconstructs the correct value s_r. □

We next prove that, unless the database is big, the parity-chooses-secret task requires interaction. Furthermore, we rule-out protocols in which first the referee simultaneously sends one message to each local agent, then receives an answer from each local agent, and finally exchanges (possibly many) messages with the curator. In particular, we rule-out the communication pattern used in Lemma 4.2 for the 1-out-of-2^d-parity task. To prove this result, we first convert a protocol Π for the parity-chooses-secret task with the above communication pattern to a protocol Π' in the hybrid model with the same communication pattern for a similar task (which we call the parity-chooses-secret' task, defined below). We then convert the protocol Π' to a non-interactive

protocol Π'' in the local model for another related task, and complete the proof by showing an impossibility result for the related task.

We define the parity-chooses-secret' task as the task in which the input of the curator is generated as in the parity-chooses-secret task and the input of each local agent only contains shares, that is, it is of the form $(t,(s_{j,t})_{j\in\{0,1\}^c})$. The goal of the referee in remains the same – to recover s_r.

Claim 5.3. *Assume that $m=\Omega\left(\frac{c\log(1/\beta)}{\varepsilon}\right)$. If there is an ε-differentially private protocol for the parity-chooses-secret task in the (m,n)-hybrid model with error at most β in which in the first round the referee simultaneously sends one message to each local agent, in the second round gets an answer from each agent, and then the referee and the curator exchange (possibly many) messages, then there is a 2ε-differentially private protocol for the parity-chooses-secret' task in the (m,n)-hybrid model with error at most 2β with the same communication pattern.*

Proof. Let Π be an ε-differentially private protocol for the parity-chooses-secret task in the (m,n)-hybrid model with the communication pattern as in the claim in which the referee reconstructs s_r with probability at least $1-\beta$. We construct from Π a 2ε-differentially private protocol Π' with the same communication pattern for the the parity-chooses-secret' task in which the referee reconstructs s_r with probability at least $1-2\beta$. In Π', each agent P_i, holding an input $(t,(s_{j,t})_{j\in\{0,1\}^c})$, chooses with uniform distribution a string $x_i\in_{\mathrm{R}}\{0,1\}^c$ and sends two messages of Π, one message, denoted $M_{i,0}$, for the input $(x_i,0,t,(s_{j,t})_{j\in\{0,1\}^c})$ and one message, denoted $M_{i,1}$, for the input $(x_i,1,t,(s_{j,t})_{j\in\{0,1\}^c})$. In addition, the agent sends x_i to the referee. The referee sends the messages that it gets from the local agents (i.e., $(x_i,M_{i,0},M_{i,1})_{i\in[n]}$) and its random string to the curator. The curator does as follows:

- Privately learns r by executing the ε-differentially private algorithm of [43] (see Theorem 2.17) for learning parity with $\alpha=1/4$, β, and the m inputs $(x,\langle x,r\rangle)$.
- For each agent P_i, computes $b_i=\langle x_i,r\rangle$ and $M_i=M_{i,b_i}$ (that is, the curator chooses the correct message from the two messages the agent sends).
- Simulates the communication between the curator and the referee in Π assuming that the curator gets the messages $(M_i)_{i\in[n]}$ in the first round and reconstructs s_r as the referee reconstructs it in Π.
- Sends s_r to the referee.

As each party executes two ε-differentially private algorithm on its input, the resulting protocol is 2ε-differentially private. Each agent P_i chooses x_i with uniform distribution (as in the parity-chooses-secret task). Furthermore, as m is big enough, with probability at least $1-\beta$, the curator computes the correct value r. Thus, $(x_i,b_i,t,(s_{j,t})_{j\in\{0,1\}^c})$ is an input distributed as required in the parity-chooses-secret task, and the curator reconstructs s_r with probability at most $1-2\beta$. □

We recall a result of [46] showing that the mutual information between an the input and output of a differential private algorithm is low. Recall that the *entropy* $H(X)$ of a random variable X is defined as

$$H(X)\triangleq-\sum_{x,\Pr[X=x]>0}\Pr[X=x]\log\Pr[X=x].$$

It can be proved that $0 \leq H(X) \leq \log(|\text{support}(X)|)$, where $|\text{support}(X)|$ is the size of the support of X (the number of values with probability greater than zero). The upper bound $|\text{support}(X)|$ is obtained if and only if the distribution of X is uniform and the lower bound is obtained if and only if X is deterministic. Given two random variables X and Y (possibly dependent), the *conditioned entropy* of X given Y is defined as $H(X|Y) \triangleq H(XY) - H(Y)$. From the definition of the conditional entropy, the following properties can be proved:

$$H(XY) \leq H(X) + H(Y),$$

and for 3 random variables X, Y, Z

$$H(X|YZ) \leq H(X|Y).$$

The *mutual information* between X and Y is defined as

$$I(X;Y) \triangleq H(X) - H(X|Y) = H(X) + H(Y) - H(XY).$$

Theorem 5.4 (Differential privacy implies low mutual information [46])**.** *Let $A : X^n \to Y$ be an ε-differentially private mechanism. Then for every random variable V distributed on X^n, we have $I(V; A(V)) \leq 1.5\varepsilon n$.*

Lemma 5.5. *Assume that $m = \Omega\left(\frac{c\log(1/\beta)}{\varepsilon}\right)$. If there is an ε-differentially private protocol for the parity-chooses-secret task in the (m,n)-hybrid model with error at most β in which in the first round the referee simultaneously sends one message to each local agent, in the second round gets an answer from each agent, and then the referee and the curator exchange (possibly many) messages, then $n \geq \frac{(1-2\beta)2^c}{3\varepsilon}$.*

Proof. We convert the protocol for the parity-chooses-secret task to a protocol Π'' in the local model with n agents, where the input of each agent contains 2^c bits $(s_j)_{j\in\{0,1\}^c}$. If the inputs of all agents are equal, then for every $r \in \{0,1\}^c$ the referee should output the bit s_r with probability at least $1-\beta$. We will show at the end of the proof that such protocol can exist only if n is big.

By Claim 5.3, under the assumption of the lemma there is a 2ε-differentially private protocol Π' for the parity-chooses-secret' task in the (m,n)-hybrid model with error at most 2β and communication pattern is as in the lemma. We construct the following protocol Π'' in the local model with n agents:

- **Input of each agent P_i:** $(s_j)_{j\in\{0,1\}^c}$.
- The referee chooses with uniform distribution 2^c vectors of $m+1$ bits: a vector $(s_{j,1},\ldots,s_{j,m+1}) \in_R \{0,1\}^{m+1}$ for every $j \in \{0,1\}^c$.
- The referee chooses with uniform distribution m indices $t_1,\ldots,t_m \in_R [m+1]^m$. Let ℓ be an index that does not appear in this list.
- The referee chooses with uniform distribution m strings $(x_1,\ldots,x_m) \in_R (\{0,1\}^c)^m$.
- The referee sends $((s_{j,1},\ldots,s_{j,m+1}))_{j\in\{0,1\}^c}$ and ℓ to each agent.
- Each agent P_i chooses with uniform distribution an index $t \in [m+1]$. If $t \neq \ell$, it sends to the referee its message in Π' on input $t, (s_{j,t})_{j\in\{0,1\}^c}$. If $t = \ell$, it sends to the referee its message in protocol Π' on input $\ell, \left(s_j \oplus \bigoplus_{k\neq\ell} s_{j,k}\right)_{j\in\{0,1\}^c}$. Denote the message of P_i by M_i.

- For every $r \in \{0,1\}^c$, the referee does the following:
  - Computes (without interaction) the communication exchanged in Π' between the referee and the curator with input

    $$(x_1, \langle x_1, r\rangle, t_1, (s_{j,t_1})_{j\in\{0,1\}^c}), \ldots, (x_m, \langle x_m, r\rangle, t_m, (s_{j,t_m})_{j\in\{0,1\}^c}).$$

    Denote this communication by $M_{C,r}$.
  - The referee reconstructs s_r from the messages $M_{C,r}, M_1, \ldots, M_n$ using the reconstruction function of Π'.

Protocol Π'' is 2ε-differentially private, since Π' is 2ε-differentially private. Furthermore, if all the inputs of the local agents are equal, then $M_{C,r}, M_1, \ldots, M_n$ are distributed as in Π', thus, for every $r \in \{0,1\}^c$, the referee reconstructs s_r with probability at least $1-\beta$.

We complete the proof by showing that n must be large enough in Π'' (hence, also in Π). Assume we execute protocol Π'' when $(s_j)_{j\in\{0,1\}^c}$ is choosen with uniform distribution and denote its input by $(s_j)_{j\in\{0,1\}^c}$ and its output by $(s'_j)_{j\in\{0,1\}^c}$. As the output in Π'' is computed from the transcript of Π', the post-processing property of differential privacy implies that the algorithm that first executes protocol Π' and then computes the output from the transcript is 2ε-differentially private. By Theorem 5.4,

$$I\left((s_j)_{j\in\{0,1\}^c}; (s'_j)_{j\in\{0,1\}^c}\right) \leq 3\varepsilon n. \tag{1}$$

On the other hand, $\Pr[s_{j_0} = s'_{j_0}] \geq 1-2\beta$ for a given $j_0 \in \{0,1\}^c$, thus

$$H(s_{j_0}|(s'_j)_{j\in\{0,1\}^c}) \leq H(s_{j_0}|s'_{j_0}) \leq 2\beta,$$

and

$$H\left((s_j)_{j\in\{0,1\}^c}|(s'_j)_{j\in\{0,1\}^c}\right) \leq \sum_{j_0\in\{0,1\}^c} H\left(s_{j_0}|(s'_j)_{j\in\{0,1\}^c}\right) \leq 2\beta 2^c.$$

Thus,

$$I\left((s_j)_{j\in\{0,1\}^c}; (s'_j)_{j\in\{0,1\}^c}\right) = H\left((s_j)_{j\in\{0,1\}^c}\right) - H\left((s_j)_{j\in\{0,1\}^c}|(s'_j)_{j\in\{0,1\}^c}\right) \geq 2^c - 2\beta 2^c = (1-2\beta)2^c. \tag{2}$$

Inequalities (1) and (2) imply that $(1-2\beta)2^c \leq 3\varepsilon n$. □

We summarize the possibility and impossibility results for the parity-chooses-secret task in the next theorem.

Theorem 5.6. *Let $\varepsilon = 1/4$, $\beta = 1/4$. For every integer c, there are $m = \Theta(c)$ and $n = \Theta(c^2 \log c)$ such that*

1. *There exists an ε-differentially private protocol for the parity-chooses-secret task with strings of length c in the (m,n)-hybrid model where first the curator sends one message to the referee and then the referee simultaneously exchanges one message with each local agent.*
2. *There does not exist an ε-differentially private protocol for this task in the (m,n)-hybrid model in which the referee first simultaneously exchanges one message with the local agents and then exchanges messages with the curator.*

3. *In any ε-differentially private protocol for this task in the local model with n agents the number of rounds is $2^{\Omega(c)}$.*

4. *There is no algorithm in the trusted curator model that solves this task with m examples.*

Proof. Item 1 follows directly from Lemma 5.2. Item 2 is implied by Lemma 5.5, since $n \ll 2^c$. Item 3 follows from Fact 2.13 as in the proof of Theorem 4.6.

For Item 4, observe that a curator receiving m input points obtains less than $m+1$ shares of $(s_1,\ldots,s_j)$ and hence obtains no information about s_r. That is, such a curator cannot solve the parity-chooses-secret task alone, even without privacy constraints. □

6 A Negative Result: Basic hypothesis testing

Here, we show that for one of the most basic tasks, differentiating between two discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$, the hybrid model gives no significant added power.

Definition 6.1 (The simple hypothesis-testing task)**.** *Let $0 < \beta < 1$ be a parameter, X be a finite domain, and $\mathcal{D}_0$ and $\mathcal{D}_1$ be two distributions over X. The input of the hypothesis-testing task is composed of i.i.d. samples from $\mathcal{D}_j$ for some $j \in \{0,1\}$ and the goal of the referee is to output $\hat{j}$ s.t. $\Pr[\hat{j} = j] \geq 1-\beta$.*

Theorem 6.2. *If there exists an ε-differentially private protocol in the (m,n)-hybrid model for testing between distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ with success probability $1/2+\gamma$, then either there exists an ε-differentially private protocol for this task in the curator model that uses m samples and succeeds with probability at least $1/2+\gamma/4$ or there exists an ε-differentially private protocol for this task in the local model with n agents that succeeds with probability at least $1/2+\gamma/4$.*

Proof. Let Π be a protocol guaranteed by the lemma, that is, when the inputs of the curator and the local agents are drawn from $\mathcal{D}_j$, the referee in Π returns j with probability at least $1/2+\gamma$. Consider an execution of the protocol when the inputs of the curator are drawn from $\mathcal{D}_0$ and the inputs of the local agents are drawn from $\mathcal{D}_1$ and let p be the probability that the referee in Π returns 1 in this case.

We first assume that $p \geq 1/2$ and show that there exists an ε-differentially private protocol Π^{local} for this task in the local model with n agents that succeeds with probability at least $1/2+\gamma/4$. The referee in protocol Π^{local} with probability $\gamma/2$ outputs 1 and with probability $1-\gamma/2$ draws m samples from $\mathcal{D}_0$, executes protocol Π, where the referee simulates the messages of the curator using the m samples, and returns the output of Π.

We next analyze this protocol. If the inputs of the local agents are drawn from $\mathcal{D}_1$, then the probability that the referee in protocol Π returns 1 is at least $1/2$ and the probability that the referee in Π^{local} returns 1 is at least $\gamma/2 + (1-\gamma/2)\cdot 1/2 = 1/2+\gamma/4$. If the inputs of the local agents are drawn from $\mathcal{D}_0$, then the probability that the referee in Π^{local} returns 0 is at least $(1-\gamma/2)\cdot(1/2+\gamma) \geq 1/2+\gamma/4$ (since $\gamma \leq 1/2$).

For the case that $p < 1/2$, it can be shown, using an analogous construction, that there exists an ε-differentially private protocol Π^{curator} for this task in the curator model with m samples that succeeds with probability at least $1/2+\gamma/4$. □

Notation. The *total variation distance* (also known as the statistical distance) of two discrete distributions $\mathcal{D}_0, \mathcal{D}_1$ over a domain X is $d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1) = \sup_{T \subset X} |\mathcal{D}_1(T) - \mathcal{D}_0(T)| = \frac{1}{2} \sum_{x \in X} |\mathcal{D}_1(x) - \mathcal{D}_0(x)|$. The *squared Hellinger distance* between two distributions $\mathcal{D}_0, \mathcal{D}_1$ over a domain X is defined as $d_{H^2}(\mathcal{D}_0, \mathcal{D}_1) = \frac{1}{2} \sum_{x \in X} \left(\sqrt{\mathcal{D}_0(x)} - \sqrt{\mathcal{D}_1(x)} \right)^2$.

For the rest of the discussion in this section, fix the domain $X = \{0, 1\}$, and some $\alpha > 0$. We define two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ where under $\mathcal{D}_1$ we have $\Pr_{x \sim \mathcal{D}_1}[x = 1] = \frac{1}{2}(1 + \alpha)$ and under $\mathcal{D}_0$ we have $\Pr_{x \sim \mathcal{D}_0}[x = 1] = \frac{1}{2}(1 - \alpha)$. It is a fairly simple calculation to see that $d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1) = \alpha$ and $\frac{\alpha^2}{2} \leq d_{H^2}(\mathcal{D}_0, \mathcal{D}_1) \leq \alpha^2$. We prove that for some setting of the parameters n and m, the hypothesis-testing task between $\mathcal{D}_0$ and $\mathcal{D}_1$ is impossible in the (m, n)-hybrid model.

Next, we cite two known results regarding differentially private simple hypothesis testing. The work of Joseph et al. [40] discusses sample complexity bounds for simple hypothesis testing the in local-model.

Fact 6.3 ([40, Theorem 5.3])**.** Let Π be an ε-differentially private protocol in the local model for distinguishing between a setting where the input of the n local agents is drawn i.i.d. from $\mathcal{D}_0$ vs. the input drawn i.i.d. from $\mathcal{D}_1$. Let Π^j denote the view of the protocol under n input point drawn i.i.d. from $\mathcal{D}_j$ for $j \in \{0, 1\}$. Then $d_{\mathrm{TV}}(\Pi^1, \Pi^0) \leq 50 n \varepsilon^2 d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1) = 50 n \varepsilon^2 \alpha^2$.

The work of Cannone et al. [18] gives tight sample complexity bounds for private hypothesis testing in the curator model. Although their result for general distributions is rather technical to state, doing so for the result for distributions supported on precisely two elements is fairly simple.

Fact 6.4 ([18, Theorem 3.5] reworded)**.** There exists a constant $C > 0$ such that any ε-DP algorithm for distinguishing w.p. ≥ 0.55 between a setting where the input of m datapoints is drawn i.i.d. from $\mathcal{D}_0$ vs. the input drawn i.i.d. from $\mathcal{D}_1$ requires

$$m \geq C \left(\frac{1}{d_{H^2}(\mathcal{D}_0, \mathcal{D}_1)} + \frac{1}{\varepsilon \cdot d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1)} \right) = \Omega\left(\frac{1}{\alpha^2} + \frac{1}{\varepsilon \cdot \alpha} \right)$$

Combining the above two facts with Theorem 6.2 we obtain the following as an immediate corollary.

Theorem 6.5. *There exists two constants c_0, c_1 such that in the (m, n)-hybrid model, with $m \leq c_0 \left(\frac{1}{\alpha^2} + \frac{1}{\varepsilon \alpha} \right)$ and $n \leq c_1 \cdot \frac{1}{\varepsilon^2 \alpha^2}$, there does not exists an ε-differentially private protocol that succeeds in determining whether all $m + n$ input points are drawn from $\mathcal{D}_0$ vs. $\mathcal{D}_1$ w.p. ≥ 0.75.*

Proof. Assume towards a contradiction that such a protocol Π exists. By Theorem 6.2, there is an ε-differentially private protocol that succeeds in determining whether all input points are drawn from $\mathcal{D}_0$ vs. $\mathcal{D}_1$ w.p. ≥ 0.5625 either in the trusted curator model with sample complexity m or in the local model with n agents. The former yields an immediate contradiction with Fact 6.4 and the latter yields an immediate contradiction with Fact 6.3 for a sufficiently small c_1. □

7 Acknowledgements

We thank Adam Smith for suggesting the select-then-estimate task discussed in the introduction. The work was done while the authors were at the "Data Privacy: Foundations and Applications" program held in spring 2019 at the Simons Institute for the Theory of Computing, UC Berkeley.

Work of A. B. and K. N. was supported by NSF grant No. 1565387 TWC: Large: Collaborative: Computing Over Distributed Sensitive Data. Work of A. B. was also supported by ISF grant no. 152/17, a grant from the Cyber Security Research Center at Ben-Gurion University, and ERC grant 742754 (project NTSC).

Work of A. K. was supported by NSF grant No. 1755992 CRII: SaTC: Democratizing Differential Privacy via Algorithms for Hybrid Models, a VMWare fellowship, and a gift from Google.

Work of O. S. was supported by grant #2017–06701 of the Natural Sciences and Engineering Research Council of Canada (NSERC). The bulk of this work was done when O. S. was affiliated with the University of Alberta, Canada

Work of U. S. was supported in part by the Israel Science Foundation (grant No. 1871/19).

References

[1] Apple. Learning with privacy at scale. *Apple Machine Learning Journal*, 1, 2017. https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html.

[2] Brendan Avent, Aleksandra Korolova, David Zeber, Torgeir Hovden, and Benjamin Livshits. BLENDER: Enabling local search with a hybrid differential privacy model. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 747–764. USENIX Association, 2017.

[3] C.-A. Azencott. Machine learning and genomics: precision medicine versus patient privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128):20170350, 2018.

[4] Mitali Bafna and Jonathan Ullman. The price of selection in differential privacy. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 151–168. PMLR, 2017.

[5] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. Differentially private summation with multi-message shuffling. *CoRR*, abs/1906.09116, 2019.

[6] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. The privacy blanket of the shuffle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Proceedings, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 638–667. Springer, 2019.

[7] Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, pages 1046–1059, 2016.

[8] Raef Bassily, Kobbi Nissim, Uri Stemmer, and Abhradeep Guha Thakurta. Practical locally private heavy hitters. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 2288–2296, 2017.

[9] Raef Bassily and Adam D. Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015*, pages 127–135, 2015.

[10] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference*, pages 451–468, 2008.

[11] Amos Beimel, Kobbi Nissim, and Uri Stemmer. Characterizing the sample complexity of private learners. In *ITCS*, pages 97–110. ACM, 2013.

[12] Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. *Theory of Computing*, 12(1):1–61, 2016.

[13] Bonnie Berger and Hyunghoon Cho. Emerging technologies towards enhancing privacy in genomic data sharing. *Genome Biology*, 20, 2019. https://genomebiology.biomedcentral.com/articles/10.1186/s13059-019-1741-0.

[14] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*, SOSP '17, pages 441–459. ACM, 2017.

[15] Charlotte Bonte, Eleftheria Makri, Amin Ardeshirdavani, Jaak Simm, Yves Moreau, and Frederik Vercauteren. Towards practical privacy-preserving genome-wide association study. *BMC bioinformatics*, 19(1):537, 2018.

[16] Mark Bun, Jelani Nelson, and Uri Stemmer. Heavy hitters and the structure of local privacy. In Jan Van den Bussche and Marcelo Arenas, editors, *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 435–447. ACM, 2018.

[17] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil P. Vadhan. Differentially private release and learning of threshold functions. In *FOCS*, pages 634–649, 2015.

[18] Clément L. Canonne, Gautam Kamath, Audra McMillan, Adam D. Smith, and Jonathan Ullman. The structure of optimal private tests for simple hypotheses. In *STOC*, pages 310–321, 2019.

[19] Albert Cheu, Adam D. Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 375–403. Springer, 2019.

[20] Catalin Cimpanu. Capital One hacker took data from more than 30 companies, new court docs reveal. *ZDNet*, Aug 14, 2019. https://www.zdnet.com/article/capital-one-hacker-took-data-from-more-than-30-companies-new-court-docs-reveal/.

[21] Privacy Rights Clearinghouse. Data Breaches, 2019. https://www.privacyrights.org/data-breaches.

[22] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3571–3580. Curran Associates, Inc., 2017.

[23] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, 27(3):642–669, 1956.

[24] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis. In *ACM Symposium on the Theory of Computing (STOC)*. ACM, June 2015.

[25] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[26] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *Symposium on Theory of Computing (STOC)*, pages 715–724. ACM, 2010.

[27] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 2468–2479. Society for Industrial and Applied Mathematics, 2019.

[28] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 1054–1067, New York, NY, USA, 2014. ACM.

[29] Vitaly Feldman and Thomas Steinke. Generalization for adaptively-chosen estimators via stable median. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017*, pages 728–757, 2017.

[30] Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. *SIAM J. Comput.*, 44(6):1740–1764, 2015.

[31] Marco Gaboardi, Ryan Rogers, and Or Sheffet. Locally private mean estimation: Z-test and tight confidence intervals. *CoRR*, abs/1810.08054, 2018.

[32] Badih Ghazi, Rasmus Pagh, and Ameya Velingker. Scalable and differentially private distributed aggregation in the shuffled model. *CoRR*, abs/1906.08320, 2019.

[33] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457(7232):1012–1014, 2009.

[34] Google Research Blog. Tackling urban mobility with technology. https://europe.googleblog.com/2015/11/tackling-urban-mobility-with-technology.html, Nov 18, 2015.

[35] Andy Greenberg. Apple's differential privacy is about collecting your data – but not your data. In *Wired*, June 13, 2016.

[36] Andy Greenberg. How one of Apple's key privacy safeguards falls short. *WIRED*, Sep 15, 2017. https://www.wired.com/story/apple-differential-privacy-shortcomings/.

[37] Alex Hern. Uber employees 'spied on ex-partners, politicians and Beyoncé'. *The Guardian*, Dec 13, 2016. https://www.theguardian.com/technology/2016/dec/13/uber-employees-spying-ex-partners-politicians-beyonce.

[38] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, pages 239–248. IEEE Computer Society, 2006.

[39] Aaron Johnson and Vitaly Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1079–1087. ACM, 2013.

[40] Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth. The role of interactivity in local differential privacy. *CoRR*, abs/1904.03564, 2019.

[41] Matthew Joseph, Jieming Mao, and Aaron Roth. Exponential separations in local differential privacy through communication complexity. *CoRR*, abs/1907.00813, 2019.

[42] Matthew Joseph, Aaron Roth, Jonathan Ullman, and Bo Waggoner. Local differential privacy for evolving data. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, pages 2381–2390, 2018.

[43] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, 2011.

[44] Aaron Mak. Do Tech Companies Really Need to Snoop Into Private Conversations to Improve Their A.I.? *Slate*, Feb 19, 2019. https://slate.com/technology/2019/02/reverse-location-search-warrants-google-police.html.

[45] Pascal Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, pages 1269–1283, 1990.

[46] Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil P. Vadhan. The limits of two-party differential privacy. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 81–90. IEEE Computer Society, 2010.

[47] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103. IEEE, Oct 20–23 2007.

[48] Chris Merriman. Microsoft reminds privacy-concerned Windows 10 beta testers that they're volunteers. In *The Inquirer, http://www.theinquirer.net/2374302*, Oct 7, 2014.

[49] Microsoft. Windows Insider Program Agreement. https://insider.windows.com/en-us/program-agreement, Sep 15, 2017.

[50] Ilya Mironov. Beyond software: Hardware-based security. https://simons.berkeley.edu/talks/beyond-software-hardware-based-security, May 8, 2019.

[51] Mozilla. Firefox Privacy Notice. https://www.mozilla.org/en-US/privacy/firefox/#pre-release, June 4, 2019.

[52] Kobbi Nissim and Uri Stemmer. Personal communication, 2017.

[53] NYC Department of Transportation. New York City Mobility Report. http://www.nyc.gov/html/dot/downloads/pdf/mobility-report-2016-print.pdf, 2016.

[54] Phillip Rogaway. The moral character of cryptographic work. *IACR Cryptology ePrint Archive*, 2015:1162, 2015.

[55] Stephen Shankland. How Google tricks itself to protect Chrome user privacy. In *CNET*, Oct 31, 2014.

[56] Elaine Shi, T.-H. Hubert Chan, Eleanor G. Rieffel, and Dawn Song. Distributed private data analysis: Lower bounds and practical constructions. *ACM Trans. Algorithms*, 13(4):50:1–50:38, 2017.

[57] Sean Simmons and Bonnie Berger. Realizing privacy preserving genome-wide association studies. *Bioinformatics*, 32(9):1293–1300, 2016.

[58] Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 552–563. IEEE, 2017.

[59] Uri Stemmer. Locally private k-means clustering. *CoRR*, abs/1907.02513, 2019.

[60] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and XiaoFeng Wang. Privacy loss in Apple's implementation of differential privacy on MacOS 10.12. *arXiv preprint arXiv:1709.02753*, 2017. https://arxiv.org/abs/1709.02753.

[61] Jonathan Ullman. Tight lower bounds for locally differentially private selection. Technical Report abs/1802.02638, arXiv, 2018.

[62] Salil Vadhan. *The Complexity of Differential Privacy*. Unpublished manuscript, 2016. https://privacytools.seas.harvard.edu/publications/complexity-differential-privacy.

[63] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.

[64] Giridhari Venkatadri, Elena Lucherini, Piotr Sapiezyński, and Alan Mislove. Investigating sources of PII used in Facebook's targeted advertising. In *Proceedings of the Privacy Enhancing Technologies Symposium (PETS'19)*, Jul 2019.

[65] Rui Wang, Yong Fuga Li, XiaoFeng Wang, Haixu Tang, and Xiaoyong Zhou. Learning your identity and disease from research papers: information leaks in genome wide association study. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pages 534–544. ACM, 2009.

[66] Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[67] Fei Yu and Zhanglong Ji. Scalable privacy-preserving data sharing methodology for genome-wide association studies: an application to idash healthcare privacy protection challenge. *BMC medical informatics and decision making*, 14(1):S3, 2014.

A Learning parity and threshold

We now define a problem which, due to known results regarding sample complexity lower bounds, cannot be privately learned neither in the curator model nor in the local model with sub-exponential number of rounds and yet can be learned in the hybrid model. Specifically, we make use of existing impossibility results for learning threshold functions in the curator model (see Fact 2.12) and for learning parity functions in the local model (see Fact 2.13). The learning problem we define is fairly natural – it is the concatenation of the two problems, resulting in a two-tuple label. All that is left is to set parameters so that either the curator or the local model agents fail to learn the suitable part of the problem.

For integers b, c, we define a concept class

$$\mathcal{C}_{b,c} = \left\{ c_{k,t} : \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^2 : k \in \{0,1\}^c, t \in \{0,1\}^b \right\},$$

where $c_{k,t}(x,y) = (\mathsf{Par}_k(x), \mathsf{Thr}_t(y))$.

Theorem A.1. *Let b be an integer and $0 < \varepsilon, \alpha, \beta < 1$ s.t. $8 \cdot 10^5 b \log(b/\beta\alpha)/(\alpha^3\varepsilon^2) < 2^{\sqrt{b}/4}$ and let $c = \sqrt{b}$, $m = 1000\sqrt{b}\log(1/\beta)/\varepsilon\alpha$, and $n = 8 \cdot 10^5 b \log(b/\beta\alpha)/(\alpha^3\varepsilon^2)$. The task of PAC-learning $\mathcal{C}_{b,c}$ with ε-differentially privacy can be solved in the (m,n)-hybrid model yet cannot be learned neither in the curator model with m samples nor in the local agents model with n agents and at most $2^{\Omega(\sqrt{b})}/n = \frac{\alpha^3\varepsilon^2 2^{\Omega(\sqrt{b})}}{\log 1/\beta\alpha}$ rounds.*

Proof. We begin with the latter part of the theorem. Assume towards contradiction that there exists an ε-differentialy private PAC-learner $\mathcal{A}$ for $\mathcal{C}_{b,c}$ in the curator model with a sample size of at most m. As we explain below, this implies that there exists an ε-differentialy private algorithm $\mathcal{A}'$ in the curator model for PAC-learning a Threshold problem that has access to at most m examples, contradicting Fact 2.12. The construction of $\mathcal{A}'$ from $\mathcal{A}$ is as follows: Algorithm $\mathcal{A}'$, given m inputs to the Threshold problem, picks $k \in \{0,1\}^c$ arbitrarily and pads each example with a $x \in \{0,1\}^c$ chosen with uniform distribution and the label $\langle x, k\rangle$ and then feeds the concatenated inputs to $\mathcal{A}$. By definition, the m input points are legal inputs to $\mathcal{A}$ and thus w.p. $\geq 1-\beta$ the algorithm produces a good h. Projecting h onto its second coordinate yields an hypothesis h' whose error in comparison to Thr_t is at most α. The proof that no differentialy private protocol in the local model with at most $2^{c/3}/n$ rounds can learn such h on its own is symmetric and follows from Fact 2.13.

We now show that there exists a protocol in the (m,n)-hybrid model that is capable of privately learning $\mathcal{C}_{b,c}$. The protocol itself is very straight-forward: the curator $(\frac{\alpha}{2}, \frac{\beta}{2})$-learns the parity portion of h (its second coordinate) and the local agents $(\frac{\alpha}{2}, \frac{\beta}{2})$-learn the threshold portion (h's first coordinate). In this protocol, the interaction of the referee with the curator is independent from the interaction of the referee with the local agents.

We use the Parity learning algorithm of Theorem 2.17 (from [43]), and, not surprisingly, m is set s.t. it sufficiently large to learn a good hypothesis w.p. $\geq 1-\frac{\beta}{2}$. The differentialy-private protocol in the local model for learning Threshold is folklore, so for completion we detail it here. Basically, it is composed of two steps: first finding all quantiles of $\{\frac{\alpha}{4}, \frac{2\alpha}{4}, \frac{3\alpha}{4}, \ldots, 1-\frac{\alpha}{4}\}$ of the distribution, and then figuring out in which of the quantiles there is a flip of the labels from 0 to 1.

In more details, here is the differentialy-private protocol in the local model for learning the threshold function. We partition the n agents into two sets. The first set is then further equipartitioned into $\lceil \frac{4}{\alpha} \rceil - 1$ sets, where in set i we find the $i \cdot \frac{\alpha}{4}$-quantile of the distribution over the y's, up

to an error of at most $\frac{\alpha}{10}$ using, say, the ε-differentialy-private binary-search protocol for quantiles in the local model of Theorem 2.19 (from [31]). In our setting, with a discrete distribution taking values between 0 and 2^b, we set $Q_{\max} = 2^b$ and $Q_{\min} = 0$, $\tau_{\text{dist}} = 1/4$, $\beta_{\text{conf}} = \frac{\beta\alpha}{16}$, and $\lambda_{\text{quant}} = \frac{\alpha}{10}$ as discussed above; hence a sample of size $10^5 \frac{b}{\varepsilon^2\alpha^2} \log(8b/\beta_{\text{conf}})$ is sufficient for a single quantile approximation w.p. $\geq 1 - \beta_{\text{conf}}$. We have set n such that a sample complexity of $(n/2)/(\lceil 4/\alpha \rceil - 1)$ suffices for finding the suitable quantile with an error probability of at most $\frac{\beta\alpha}{16}$. Once all quantiles of the form $i \cdot \frac{\alpha}{4}$ have been published, we turn to the latter half of the n local agents. We apply the heavy-hitters algorithm of Theorem 2.18 (from [16]) for the agents with 1-label, that is, the input of an agent with 1-label to the heavy-hitters protocol is its quantile and the input of an agent with 1-label is 0. Note again that $n/2$ is sufficiently large s.t. w.p. $\geq 1 - \frac{\beta}{4}$ we can assess the fraction of 1-labels in all bins. Furthermore, by definition of the threshold problem, any bin that contains only values $\geq t$ must be all 1-labeled. The agents then find the first quantile from which all agents have 1-label and output it as the suggested threshold. By definition, w.p. $\geq 1 - \frac{\beta}{2}$ we have that the error of this threshold is the worst-case width of any bin, i.e. at most $\frac{\alpha}{4} + 2 \cdot \frac{\alpha}{10} < \frac{\alpha}{2}$. Thus we have a ε-differentialy private protocol in the local model for $(\frac{\alpha}{2}, \frac{\beta}{2})$-learning thresholds with n agents. □

We remark that it is possible to design a *non-interactive* protocol in the hybrid model for PAC learning $\mathcal{C}_{b,c}$, at the expense of increasing the sample complexity of the protocols. See Section 3.

B The select-then-estimate task

Select-then-estimate is an example of a task that cannot be privately solved in the curator model or in the local model but can be solved (with interaction) in the hybrid model. We do not know whether the interaction is essential.

Definition B.1 (The select-then-estimate task)**.** *Let $X = \{-1,1\}^d$, and let $P \in \Delta(X)$ be an unknown distribution over X. Define $\mu = E_{x\sim P}[x]$. Note that $\mu = (\mu_1, \ldots, \mu_d) \in [-1,1]^d$. The inputs in the select-then-estimate task are $x_1, \ldots, x_n$ sampled i.i.d. from P.*

For parameters $\alpha < \alpha'$, the goal of the referee in the select-then-estimate task is to output $(i, \hat{\mu}_i)$ such that $\mu_i \geq max_j(\mu_j) - \alpha$ and $\left|\hat{\mu}_i - \mu_i\right| \leq \alpha'$.

We focus on the case when $\alpha' < \alpha$, that is, the selection i is done with the lesser accuracy α, and the estimation $\hat{\mu}_i$ is with the better accuracy α'. In the following discussion, we omit the dependency on the allowed failure probability for the task, β.

A number of sample complexity bounds are relevant for this problem for non-interactive protocols. Selecting a coordinate i such that $\mu_i \geq \max_j \mu_j - \alpha$ requires $\Theta(\frac{\log d}{\alpha^2} + \frac{\log d}{\alpha\varepsilon})$ samples in the trusted curator model under pure differential privacy [47, 4, 58], and $\Theta(\frac{d \log d}{\alpha^2\varepsilon^2})$ samples in the local model [61].[5] Once i is selected, estimating μ_i up to an error α' (i.e., computing $\hat{\mu}_i$ so that $|\hat{\mu}_i - \mu_i| \leq \alpha'$) requires $\Omega(\frac{1}{\alpha'^2})$ samples regardless of privacy, and $O(\frac{1}{\alpha'^2\varepsilon^2})$ samples suffice in the local model (using the randomized respond protocol of Warner [66]).

To exemplify the many settings where the hybrid model provides a solution for the task but neither curator alone nor the local model parties alone can solve it, consider the case where $\varepsilon = 0.1$,

[5] The bound in [61] is for non-interactive local model protocols. In Appendix C we show that for interactive local model protocols $n = \Omega(\frac{d}{\alpha\varepsilon})$.

$\alpha < 0.1$, and $\alpha/\sqrt{d} < \alpha' < \alpha/\log d$. With this choice of parameters, $m = O(\log d/\alpha^2)$ samples suffice for the curator to identify i, but not to perform the estimate because $m = o(1/\alpha'^2)$ as $1/\alpha'^2 = \Omega(\log^2 d/\alpha^2)$. Likewise, a choice of $n = O(1/\alpha'^2)$ allows for performing the estimate in the local model, but not for performing the selection, as $1/\alpha'^2 = O(d/\alpha^2)$.

The (m, n)-hybrid model with the above parameter choices, however, allows for a solution where the curator first identifies i such that $\mu_i \geq \max_j \mu_j - \alpha$ and then the estimation of μ_i within accuracy α' is performed by the local model parties.

By our analysis above, for the case $\varepsilon = 0.1, \alpha < 0.1$ we get that the greater the ratio n/m is, the better the improvement in accuracy as

$$\frac{\alpha}{\alpha'} \sim \frac{\sqrt{\log d/m}}{\sqrt{1/n}} = \sqrt{\frac{n \log d}{m}}.$$

Another way to look at it is to get a sense of how small the size of the curator contributors m can be in comparison to n in order to be helpful. It depends on the ratio of desired accuracies $\frac{\alpha'}{\alpha}$, as $\frac{m}{n} \sim (\frac{\alpha'}{\alpha})^2 \log d$, i.e., $\frac{\log d}{d} < \frac{m}{n} < \frac{1}{\log d}$.

The analysis can be repeated for other values of ε and α as well. For example, the (m, n) hybrid model can privately solve the select-then-estimate task when $\varepsilon = 1, \alpha = \frac{1}{d^{0.25}}, \alpha' = \frac{1}{d^{0.75-0.5t}}, m = \sqrt{d} \log d, n = d^{1.5-t} \log d$, for any choice of $0.5 < t < 1$, and neither the local model nor the curator can solve it to the same accuracies separately.

In the high-dimensional problems that are of interest for the select-then-estimate task, when often the dimension of the data d exceeds the number of data points available n, the parameter values for when the model is helpful illustrate its considerable potential.

C A lower bound for the selection function

We next present a simple lower bound of the number of messages sent in an ε-differentially private protocol for the selection problem.

Definition C.1 (The selection problem)**.** *The input for the selection problem is n i.i.d. samples $Y = (y_1, \ldots, y_n)$ from an unknown distribution $\mathcal{D}$ over $\{0,1\}^d$ with mean $\mu = (\mu_1, \ldots, \mu_d)$. The goal is to output a coordinate j such that*

$$E_{Y,j}[\mu_j] \geq \max_k \mu_k - \alpha.$$

Ullman [61] proved that in any ε-differentially private non-interactive protocol in the local model for the selection problem, the number of agents is $\Omega\left(\frac{d \log d}{\alpha^2 \varepsilon^2}\right)$. We prove a lower bound of $\Omega(d^{1/3})$ on the number messages sent to the referee in any 1-differentially private protocol for this task, even if interaction is allowed. For example, if each agent sends one message, then the number of agents is $\Omega(d^{1/3})$ (even if the protocol is interactive, e.g., the referee sends a message to P_1, gets a message from P_1, based on this message computes a message and sends it to P_2, and so on). As another example, if there are only $O(d^{1/6})$ agents, then the number of rounds in any 1-differentially private protocol for the selection problem is $\Omega(d^{1/6})$. To summarize, our lower bound applied to a larger family of protocols than the result of [61]; however, our bound is weaker.

The idea of our proof is simple: We show a reduction from private learning in the local model to privately solving the selection problem and use a known lower bounds for the parity problem [43] to obtain the lower bound.

Claim C.2. *Let $C = \{c_1, \ldots, c_d\}$ be a class of functions and n be such that $n \geq \frac{64}{\alpha}(\mathrm{VC}(C)\ln(\frac{128}{\alpha})+\ln(8))$. If there is an ε-differentially private protocol M in the local model for the selection problem with n agents, then there is an ε-differentially private proper $(4\alpha, 1/4)$-PAC learning protocol in the local model for the class C with n agents and the same number of messages.*

Proof. We convert a labeled example for C to an input for the selection problem: Given a labeled example (x, b), construct the input $\mathrm{convert}(x, b) = (y[1], \ldots, y[d])$, where if $c_i(x) = b$ then $y[i] = 1$, else $y[i] = 0$. Given $((x_1, b_1), \ldots, (x_n, b_n))$ – examples labeled by the some concept c_i – we construct $Y = (\mathrm{convert}(x_1, b_1), \ldots, \mathrm{convert}(x_n, b_n))$, execute M on Y, and return c_j, where j is the output of M.

Notice that if the samples are labeled by c_i, then the i-th coordinate in all points in Y is 1. Thus, M returns a coordinate j such that $E_{Y,j}[\mu_j] \geq 1 - \alpha$; in particular, the probability that μ_j is less than $1 - 4\alpha$ is at most 1/4. Thus, with probability at least 3/4, the learning algorithm has error at most 4α on the sample. By standard VC arguments, with probability at least 1/2 the concept c_j has error at most 8α on the distribution $\mathcal{D}$. □

Theorem C.3. *Suppose there is a 1-differentially private local protocol M in the local model for the selection problem with $\alpha = 1/10$. Then, the number of messages sent by the agents to the referee is $\Omega\left(d^{1/3}\right)$.*

Proof. Let r be the number of messages in the protocol M on inputs of size d. We consider the class of parity functions of strings of length c. We apply Claim C.2 to this class, where the examples are taken with uniform distribution. The length of the inputs for the selection problem is $d = 2^c$, the number of parity functions. This results in a 1-differentially private local protocol for learning the class of parity functions under the uniform distribution with r messages. By Fact 2.13 (from [43]), the number of messages in such protocol is $\Omega(2^{c/3}) = \Omega(d^{1/3})$. □