
Differentially-Private Federated Linear Bandits

Abhimanyu Dubey and Alex Pentland
Media Lab and Institute for Data, Systems and Society
Massachusetts Institute of Technology
{dubeya, pentland}@mit.edu

Abstract

The rapid proliferation of decentralized learning systems mandates the need for differentially-private cooperative learning. In this paper, we study this in context of the contextual linear bandit: we consider a collection of agents cooperating to solve a common contextual bandit, while ensuring that their communication remains private. For this problem, we devise FEDUCB, a multiagent private algorithm for both centralized and decentralized (peer-to-peer) federated learning. We provide a rigorous technical analysis of its utility in terms of regret, improving several results in cooperative bandit learning, and provide rigorous privacy guarantees as well. Our algorithms provide competitive performance both in terms of pseudoregret bounds and empirical benchmark performance in various multi-agent settings.

1 Introduction

The multi-armed bandit is the classical sequential decision-making problem, involving an agent sequentially choosing actions to take in order to maximize a (stochastic) reward [44]. It embodies the central *exploration-exploitation* dilemma present in sequential decision-making. Practical applications of the multi-armed bandit range from recommender systems [52] and anomaly detection [11] to clinical trials [15] and finance [24]. Increasingly, however, such large-scale applications are becoming *decentralized*, as their data is often located with different entities, and involves cooperation between these entities to maximize performance [14, 22]. This paradigm is now known as *federated learning*¹.

The objective of the federated paradigm is to allow cooperative estimation with larger amounts of data (from multiple clients, devices, etc.) while keeping the data decentralized [27]. There has been a surge of interest in this problem from both academia and industry, owing to its overall applicability. Most research on provably private algorithms in the federated setting has been on distributed supervised learning [28] and optimization [20]. The contextual bandit problem, however, is a very interesting candidate for private methods, since the involved contexts and rewards both typically contain sensitive user information [38]. There is an increasing body of work on online learning and multi-armed bandits in cooperative settings [13, 31, 39], and private single-agent learning [41, 38], but methods for private *federated* bandit learning are still elusive, despite their immediate applicability.

Contributions. In this paper, we study the *federated* contextual bandit problem under constraints of differential privacy. We consider two popular paradigms: (a) *centralized* learning, where a central controller coordinates different clients [27, 49] and (b) *decentralized* peer-to-peer learning *with delays*, where agents communicate directly with each other, without a controller [39, 31, 30, 32]. We provide a rigorous formulation of (ϵ, δ) -differential privacy in the *federated* contextual bandit, and present two variants of FEDUCB, the first federated algorithm ensuring that each agent is private with respect to the data from all other agents, and provides a parameter to control communication.

¹Originally, federated learning referred to the algorithm proposed in [28] for supervised learning, however, now the term broadly refers to the distributed cooperative learning setting [27].

Next, we prove rigorous bounds on the cumulative group pseudoregret obtained by FEDUCB. In the centralized setting, we prove a high probability regret bound of $\tilde{O}(d^{3/4}\sqrt{MT/\varepsilon})$ which matches the non-private bound in terms of its dependence on T and M . In the decentralized case, we prove a corresponding regret bound of $\tilde{O}(d^{3/4}\sqrt{(\text{diameter}(\mathcal{G}))MT/\varepsilon})$, where \mathcal{G} is the communication network between agents. In addition to the regret bounds, we present a novel analysis of communication complexity, and its connections with the privacy budget and regret.

2 Related Work

Multi-Agent and Distributed Bandits. Bandit learning in multi-agent distributed settings has received attention from several academic communities. Channel selection in distributed radio networks consider the (context-free) multi-armed bandit with collisions [35, 37, 36] and cooperative estimation over a network with delays [31, 30, 32]. For the contextual case, recent work has considered non-private estimation in networks with delays [12, 13, 48, 29]. A closely-related problem is that of bandits with side information [8, 5], where the single learner obtains multiple observations every round, similar to the multi-agent communicative setting. Our work builds on the remarkable work of Abbasi-Yadkori *et al.*[1], which in turn improves the LinUCB algorithm introduced in [33].

Differential Privacy. Our work utilizes *differential privacy*, a cryptographically-secure privacy framework introduced by Dwork [16, 18] that requires the behavior of an algorithm to fluctuate only slightly (in probability) with any change in its inputs. A technique to maintain differential privacy for the continual release of statistics was introduced in [10, 17], known as the *tree-based* algorithm that privatizes the partial sums of n entries by adding at most $\log n$ noisy terms. This method has been used to preserve privacy across several online learning problems, including convex optimization [26, 25], online data analysis [23], collaborative filtering [6] and data aggregation [9]. In the single-agent bandit setting, differential privacy using tree-based algorithms have been explored in the multi-armed case [43, 40, 46] and the contextual case [41]. In particular, our work uses several elements from [41], extending their single-agent results to the federated multi-agent setting. For the multi-agent multi-armed (i.e., context-free) bandit problem, differentially private algorithms have been devised for the centralized [45] and decentralized [14] settings. Empirically, the advantages of privacy-preserving contextual bandits has been demonstrated in the work of Malekzadeh *et al.*[38], and Hannun *et al.*[21] consider a centralized multi-agent contextual bandit algorithm that use secure multi-party computations to provide privacy guarantees (both works do not have any regret guarantees).

To the best of our knowledge, this paper is the first to investigate differential privacy for contextual bandits in the federated learning setting, in both centralized and decentralized environments. Our work is most closely related to the important work of [1, 41], extending it to federated learning.

3 Background and Preliminaries

We use boldface to represent vectors, e.g., \mathbf{x} , and matrices, e.g., \mathbf{X} . For any matrix \mathbf{X} , its Gram matrix is given by $\mathbf{X}^\top \mathbf{X}$. Any symmetric matrix \mathbf{X} is positive semi-definite (PSD) if $\mathbf{y}^\top \mathbf{X} \mathbf{y} \geq 0$ for any vector \mathbf{y} , and we denote PSD by $\mathbf{X} \succcurlyeq \mathbf{0}$. For any PSD matrix \mathbf{X} we denote the ellipsoid \mathbf{X} -norm of vector \mathbf{y} as $\|\mathbf{y}\|_{\mathbf{X}} = \sqrt{\mathbf{y}^\top \mathbf{X} \mathbf{y}}$. For two matrices \mathbf{X} and \mathbf{Y} , $\mathbf{X} \succcurlyeq \mathbf{Y}$ implies that $\mathbf{X} - \mathbf{Y} \succcurlyeq \mathbf{0}$. For any PSD matrix \mathbf{X} and vectors \mathbf{u}, \mathbf{v} , we denote the \mathbf{X} -ellipsoid inner product as $\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{X}} = \mathbf{u}^\top \mathbf{X} \mathbf{v}$, and drop the subscript when $\mathbf{X} = \mathbf{I}$. We denote the set $\{a, a + 1, \dots, b - 1, b\}$ by the shorthand $[a, b]$, and if $a = 1$ we refer to it as $[b]$. We denote the γ^{th} power of graph \mathcal{G} as \mathcal{G}_γ (\mathcal{G}_γ has edge (i, j) if the shortest distance between i and j in \mathcal{G} is $\leq \gamma$). $N_\gamma(v)$ denotes the set of nodes in \mathcal{G} at a distance of at most γ (including itself) from node $v \in \mathcal{G}$.

Federated Contextual Bandit. This is an extension of the linear contextual bandit [33, 1] involving a set of M agents. At every trial $t \in [T]$, each agent $i \in [M]$ is presented with a *decision set* $\mathcal{D}_{i,t} \subset \mathbb{R}^d$ from which it selects an action $\mathbf{x}_{i,t} \in \mathbb{R}^d$. It then obtains a reward $y_{i,t} = \mathbf{x}_{i,t}^\top \boldsymbol{\theta}^* + \eta_{i,t}$ where $\boldsymbol{\theta}^* \in \mathbb{R}^d$ is an unknown (but fixed) parameter and $\eta_{i,t}$ is a noise parameter sampled i.i.d. every trial for every agent. The objective of the agents is to minimize the cumulative *group pseudoregret*:²

$$\mathcal{R}_M(T) = \sum_{i=1, t=1}^{M, T} \langle \mathbf{x}_{i,t}^* - \mathbf{x}_{i,t}, \boldsymbol{\theta}^* \rangle, \text{ where } \mathbf{x}_{i,t}^* = \arg \max_{\mathbf{x} \in \mathcal{D}_{i,t}} \langle \mathbf{x}, \boldsymbol{\theta}^* \rangle \text{ is optimal.}$$

²The *pseudoregret* is an expectation (over the randomness of $\eta_{i,t}$) of the stochastic quantity *regret*, and is more amenable to high-probability bounds. However, a bound over the pseudoregret can also bound the regret with high probability, e.g., by a Hoeffding concentration (see, e.g., [47]).

In the single-agent setting, the optimal regret is $\tilde{O}(\sqrt{dT})$, which has been matched by a variety of algorithms, most popular of them being the ones based on upper confidence bounds (UCB) [1, 33, 4]. In *non-private* distributed contextual bandits, a group pseudoregret of $\tilde{O}(\sqrt{dMT})$ has been achieved [48, 49], which is the order-optimality we seek in our algorithms as well.

Assumptions. We assume: (a) bounded action set: $\forall i, t, \|\mathbf{x}_{i,t}\| \leq L$, (b) bounded mean reward: $\forall \mathbf{x}, \langle \boldsymbol{\theta}^*, \mathbf{x} \rangle \leq 1$, (c) bounded target parameter: $\|\boldsymbol{\theta}^*\| \leq S$, (d) sub-Gaussian rewards: $\forall i, t, \eta_{i,t}$ is σ -sub-Gaussian, (e) bounded decision set: $\forall i, t \mathcal{D}_{i,t}$ is compact, (f) bounded reward: $\forall i, t, |y_{i,t}| \leq B$.³

Differential Privacy. The contextual bandit problem involves two sets of variables that any agent must private to the other participating agents – the available decision sets $(\mathcal{D}_{i,t})_{t \in [T]}$ and observed rewards $(y_{i,t})_{t \in [T]}$. The adversary model assumed here is to prevent any two colluding agents j and k to obtain non-private information about any specific element in agent i 's history. That is, we assume that each agent is a trusted entity that interacts with a new user at each instant t . Therefore, the context set $(\mathcal{D}_{i,t})$ and outcome $(y_{i,t})$ are sensitive variables that the user trusts only with the agent i . Hence, we wish to keep $(\mathcal{D}_{i,t})_{t \in [T]}$ private. However, the agent only stores the chosen actions $(\mathbf{x}_{i,t})_{t \in [T]}$ (and not all of $\mathcal{D}_{i,t}$), and hence making our technique differentially private with respect to $((\mathbf{x}_{i,t}, y_{i,t}))_{t \in [T]}$ will suffice. We first denote two sequences $S_i = ((\mathbf{x}_{i,t}, y_{i,t}))_{t \in [T]}$ and $S'_i = ((\mathbf{x}'_{i,t}, y'_{i,t}))_{t \in [T]}$ as t -neighbors if for each $t' \neq t, (\mathbf{x}_{i,t}, y_{i,t}) = (\mathbf{x}'_{i,t}, y'_{i,t})$. We can now provide the formal definition for federated differential privacy:

Definition 1 (Federated Differential Privacy). *In a federated learning setting with $M \geq 2$ agents, a randomized multiagent contextual bandit algorithm $A = (A_i)_{i=1}^M$ is (ε, δ, M) -federated differentially private under continual multi-agent observation if for any i, j s.t. $i \neq j$, any t and set of sequences $\mathbf{S}_i = (S_k)_{k=1}^M$ and $\mathbf{S}'_i = (S_k)_{k=1, k \neq i}^M \cup S'_i$ such that S_i and S'_i are t -neighboring, and any subset of actions $\mathcal{S}_j \subset \mathcal{D}_{j,1} \times \mathcal{D}_{j,2} \times \dots \times \mathcal{D}_{j,T}$ of actions, it holds that:*

$$\mathbb{P}(A_j(\mathbf{S}_i) \in \mathcal{S}_j) \leq e^\varepsilon \cdot \mathbb{P}(A_j(\mathbf{S}'_i) \in \mathcal{S}_j) + \delta.$$

Our notion of federated differential privacy is formalizing the standard intuition that “the action chosen by any agent must be sufficiently impervious (in probability) to any single (\mathbf{x}, y) pair from any other agent”. Here, we essentially lift the definition of joint differential privacy [41] from the individual (\mathbf{x}, y) level to the entire history $(\mathbf{x}_t, y_t)_t$ for each participating agent. Note that our definition in its current form does not require each algorithm to be private with respect to its own history, but only the histories belonging to other agents, i.e., each agent can be trusted with its own data. This setting can also be understood as requiring all *outgoing communication* from any agent to be *locally differentially private* [51] to the personal history $(\mathbf{x}_t, y_t)_t$. We can alternatively relax this assumption and assume that the agent cannot be trusted with its own history, in which case the notion of joint or local DP at the individual level (i.e., (\mathbf{x}_t, y_t)) must be considered, as done in [51, 41].

The same guarantee can be obtained if each agent A_i is (ε, δ) -differentially private (in the standard contextual bandit sense, see [41]) with respect to any other agent j 's observations, for all j . A composition argument [16]⁴ over all M agents would therefore provide us $(\sqrt{2M \log(1/\delta')}\varepsilon + M\varepsilon(e^\varepsilon - 1), M\delta + \delta')$ -differential privacy with respect to the overall sequence. To keep the notation simple, however, we adopt the (ε, δ, M) format.

4 Federated LinUCB with Differential Privacy

In this section, we introduce our algorithm for federated learning with differential privacy. For the remainder of this section, for exposition, we consider the single-agent setting and drop an additional index subscript we use in the actual algorithm (e.g., we refer to the action at time t as \mathbf{x}_t and not $\mathbf{x}_{i,t}$ for agent i). We build on the celebrated LinUCB algorithm, an application of the optimism heuristic to the linear bandit case [33, 1], designed for the single-agent problem. The central idea of the algorithm is, at every round t , to construct a *confidence set* \mathcal{E}_t that contains $\boldsymbol{\theta}^*$ with high probability, followed by computing an upper confidence bound on the reward of

³Assuming bounded rewards is required for the privacy mechanism, and is standard in the literature [3, 2].

⁴Under the stronger assumption that each agent interacts with a completely different set of individuals, we do not need to invoke the composition theorem (as $\mathbf{x}_{1,t}, \mathbf{x}_{2,t}, \dots, \mathbf{x}_{M,t}$ are independent for each t). However, in the case that one individual could potentially interact simultaneously with all agents, this is not true (e.g., when for some $t, \mathcal{D}_{i,t} = \mathcal{D}_{j,t} \forall i, j$) and we must invoke the k -fold composition Theorem[17] to ensure privacy.

Algorithm 1 CENTRALIZED FEDUCB($D, M, T, \rho_{\min}, \rho_{\max}$)

```
1: Initialization:  $\forall i$ , set  $\mathbf{S}_{i,1} \leftarrow M\rho_{\min}\mathbf{I}$ ,  $\mathbf{s}_{i,1} \leftarrow \mathbf{0}$ ,  $\widehat{\mathbf{Q}}_{i,0} \leftarrow \mathbf{0}$ ,  $\mathbf{U}_{i,1} \leftarrow \mathbf{0}$ ,  $\widehat{\mathbf{u}}_{i,1} \leftarrow \mathbf{0}$ .
2: for For each iteration  $t \in [T]$  do
3:   for For each agent  $i \in [M]$  do
4:     Set  $\mathbf{V}_{i,t} \leftarrow \mathbf{S}_{i,t} + \mathbf{U}_{i,t}$ ,  $\tilde{\mathbf{u}}_{i,t} \leftarrow \mathbf{s}_{i,t} + \widehat{\mathbf{u}}_{i,t}$ .
5:     Receive  $\mathcal{D}_{i,t}$  from environment.
6:     Compute regressor  $\widehat{\boldsymbol{\theta}}_{i,t} \leftarrow \mathbf{V}_{i,t}^{-1}\tilde{\mathbf{u}}_{i,t}$ .
7:     Compute  $\beta_{i,t}$  following Proposition 2.
8:     Select  $\mathbf{x}_{i,t} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{D}_{i,t}} \langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_{i,t} \rangle + \beta_{i,t} \|\mathbf{x}\|_{\mathbf{V}_{i,t}^{-1}}$ .
9:     Obtain  $y_{i,t}$  from environment.
10:    Update  $\mathbf{U}_{i,t+1} \leftarrow \mathbf{U}_{i,t} + \mathbf{x}_{i,t}\mathbf{x}_{i,t}^\top$ ,  $\mathbf{u}_{i,t+1} \leftarrow \mathbf{u}_{i,t} + \mathbf{x}_{i,t}y_{i,t}$ .
11:    Update  $\widehat{\mathbf{Q}}_{i,t} \leftarrow \widehat{\mathbf{Q}}_{i,t-1} + [\mathbf{x}_{i,t}^\top y_{i,t}]^\top [\mathbf{x}_{i,t}^\top y_{i,t}]$ 
12:    if  $\log \det(\mathbf{V}_{i,t} + \mathbf{x}_{i,t}\mathbf{x}_{i,t}^\top) + M(\rho_{\max} - \rho_{\min})\mathbf{I} - \log \det(\mathbf{S}_{i,t}) \geq D/\Delta t_i$  then
13:      SYNCHRONIZE  $\leftarrow$  TRUE.
14:    end if
15:    if SYNCHRONIZE then
16:       $[\forall \text{ AGENTS}]$  Agent sends  $\widehat{\mathbf{Q}}_{i,t} \rightarrow$  PRIVATIZER and gets  $\widehat{\mathbf{U}}_{i,t+1}, \widehat{\mathbf{u}}_{i,t+1} \leftarrow$  PRIVATIZER.
17:       $[\forall \text{ AGENTS}]$  Agent communicates  $\widehat{\mathbf{U}}_{i,t+1}, \widehat{\mathbf{u}}_{i,t+1}$  to controller.
18:       $[\text{CONTROLLER}]$  Compute  $\mathbf{S}_{t+1} \leftarrow \sum_{i=1}^M \widehat{\mathbf{U}}_{i,t+1}$ ,  $\mathbf{s}_{t+1} \leftarrow \sum_{i=1}^M \widehat{\mathbf{u}}_{i,t+1}$ .
19:       $[\text{CONTROLLER}]$  Communicate  $\mathbf{S}_{t+1}, \mathbf{s}_{t+1}$  back to agent.
20:       $[\forall \text{ AGENTS}]$   $\mathbf{S}_{i,t+1} \leftarrow \mathbf{S}_{t+1}$ ,  $\mathbf{s}_{i,t+1} \leftarrow \mathbf{s}_{t+1}$ .
21:       $[\forall \text{ AGENTS}]$   $\widehat{\mathbf{Q}}_{i,t+1} \leftarrow \mathbf{0}$ .
22:    else
23:       $\mathbf{S}_{i,t+1} \leftarrow \mathbf{S}_{i,t}$ ,  $\mathbf{s}_{i,t+1} \leftarrow \mathbf{s}_{i,t}$ ,  $\Delta t_i \leftarrow \Delta t_i + 1$ .
24:       $\Delta t_i \leftarrow 0$ ,  $\mathbf{U}_{i,t+1} \leftarrow \mathbf{0}$ ,  $\widehat{\mathbf{u}}_{i,t+1} \leftarrow \mathbf{0}$ .
25:    end if
26:  end for
27: end for
```

each action within the decision set \mathcal{D}_t , and finally selecting the action with the largest UCB, i.e., $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{D}_t} (\max_{\boldsymbol{\theta} \in \mathcal{E}_t} \langle \mathbf{x}, \boldsymbol{\theta} \rangle)$. The confidence set is an ellipsoid centered on the regularized linear regression estimate (for $\mathbf{X}_{<t} = [\mathbf{x}_1^\top \mathbf{x}_2^\top \dots \mathbf{x}_{t-1}^\top]^\top$ and $\mathbf{y}_{<t} = [y_1 \ y_2 \dots y_{t-1}]^\top$):

$$\mathcal{E}_t := \left\{ \boldsymbol{\theta} \in \mathbb{R}^d : \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}_t\|_{\mathbf{V}_t} \leq \beta_t \right\}, \text{ where } \widehat{\boldsymbol{\theta}}_t := \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} [\|\mathbf{X}_{<t}\boldsymbol{\theta} - \mathbf{y}_{<t}\|_2^2 + \|\boldsymbol{\theta}\|_{\mathbf{H}_t}^2].$$

The regression solution can be given by $\widehat{\boldsymbol{\theta}}_t := (\mathbf{G}_t + \mathbf{H}_t)^{-1}\mathbf{X}_{<t}^\top\mathbf{y}_{<t}$, where $\mathbf{G}_t = \mathbf{X}_{<t}^\top\mathbf{X}_{<t}$ is the Gram matrix of actions, \mathbf{H}_t is a (time-varying) regularizer, and β_t is an appropriately chosen exploration parameter. Typically in non-private settings, the regularizer is constant, i.e., $\mathbf{H}_t = \lambda\mathbf{I} \forall t$, $\lambda > 0$ [1, 33], however, in our case, we will carefully select \mathbf{H}_t to introduce privacy, using a strategy similar to [41]. Given $\mathbf{V}_t = \mathbf{G}_t + \mathbf{H}_t$, let $\text{UCB}_t(\mathbf{x}; \boldsymbol{\theta}) = \langle \boldsymbol{\theta}, \mathbf{x} \rangle + \beta_t \|\mathbf{x}\|_{\mathbf{V}_t^{-1}}$.

In the federated setting, since there are M learners that have distinct actions, the communication protocol is a key component of algorithm design: communication often creates *heterogeneity* between agents, e.g., for any two agents, their estimators ($\widehat{\boldsymbol{\theta}}_t$) at any instant are certainly distinct, and the algorithm must provide a control over this heterogeneity, to bound the group regret. We additionally require that communication between agents is (ε, δ) -private, making the problem more challenging.

4.1 Centralized Environment with a Controller

We first consider the centralized communication environment where there exists a controller that coordinates communication between different agents, as is typical in large-scale distributed learning. We consider a set of M agents that each are interacting with the contextual bandit, and periodically communicate with the controller, that synchronizes them with other agents. We present the algorithm CENTRALIZED FEDUCB in Algorithm 1 that details our approach.

Algorithm 1 works as follows. Consider an agent i , and assume that synchronization had last taken place at instant t' . At any instant $t > t'$, the agent has two sets of parameters - (A) the first being all observations up to instant t' for all M agents (including itself) and (B) the second being its own observations from instant t' to t . Since (A) includes samples from other agents, these

are privatized, and represented as the Gram matrix $\mathbf{S}_{t'+1} = \sum_{i \in [M]} \widehat{\mathbf{U}}_{i,t'+1}$ and reward vector $\mathbf{s}_{t'+1} = \sum_{i \in [M]} \widehat{\mathbf{u}}_{i,t'+1}$. Algorithm 1 privatizes its own observations as well (for simplicity in analysis) and hence \mathbf{S}, \mathbf{s} are identical for all agents at all times. Moreover, since the group parameters are noisy variants of the original parameters, i.e., $\widehat{\mathbf{U}}_{i,t} = \mathbf{G}_{i,t} + \mathbf{H}_{i,t}$ and $\widehat{\mathbf{u}}_{i,t} = \mathbf{u}_{i,t} + \mathbf{h}_{i,t}$ (where $\mathbf{H}_{i,t}$ and $\mathbf{h}_{i,t}$ are perturbations), we can rewrite $\mathbf{S}_t, \mathbf{s}_t$ as (for any instant $t > t'$),

$$\mathbf{S}_t = \sum_{i \in [M]} \left(\sum_{\tau=1}^{t'} \mathbf{x}_{i,\tau} \mathbf{x}_{i,\tau}^\top + \mathbf{H}_{i,t'} \right), \mathbf{s}_t = \sum_{i \in [M]} \left(\sum_{\tau=1}^{t'} y_{i,\tau} \mathbf{x}_{i,\tau} + \mathbf{h}_{i,t'} \right). \quad (1)$$

When we combine the group parameters with the local (unsynchronized) parameters, we obtain the final form of the parameters for any agent i as follows (for any instant $t > t'$):

$$\mathbf{V}_{i,t} = \sum_{\tau=t'}^{t-1} \mathbf{x}_{i,\tau} \mathbf{x}_{i,\tau}^\top + \mathbf{S}_t, \tilde{\mathbf{u}}_{i,t} = \sum_{\tau=t'}^{t-1} y_{i,\tau} \mathbf{x}_{i,\tau} + \mathbf{s}_t \quad (2)$$

Then, with a suitable sequence $(\beta_{i,t})_t$, the agent selects the action following the linear UCB objective:

$$\mathbf{x}_{i,t} = \arg \max_{\mathbf{x} \in \mathcal{D}_{i,t}} \left(\langle \widehat{\boldsymbol{\theta}}_{i,t}, \mathbf{x} \rangle + \beta_{i,t} \|\mathbf{x}\|_{\mathbf{V}_{i,t}^{-1}} \right) \text{ where } \widehat{\boldsymbol{\theta}}_{i,t} = \mathbf{V}_{i,t}^{-1} \tilde{\mathbf{u}}_{i,t}. \quad (3)$$

The central idea of the algorithm is to therefore carefully perturb the Gram matrices $\mathbf{V}_{i,t}$ and the reward vector $\mathbf{u}_{i,t}$ with random noise $(\mathbf{H}_{i,t}, \mathbf{h}_{i,t})$ based on the sensitivity of these elements and the level of privacy required, as is typical in the literature [41]. First, each agent updates its local (unsynchronized) estimates. These are used to construct the UCB in a manner identical to the standard OFUL algorithm [1]. If, for any agent i , the log-determinant of the local Gram matrix exceeds the synchronized Gram matrix ($\mathbf{S}_{i,t}$) by an amount $D/\Delta t_i$ (where Δt_i is the time since the last synchronization), then it sends a signal to the controller, that synchronizes *all* agents with their latest action/reward pairs. The synchronization is done using a *privatized* version of the Gram matrix and rewards, carried out by the subroutine PRIVATIZER (Section 4.3, Alg. 2). This synchronization ensures that the *heterogeneity* between the agents is controlled, allowing us to control the overall regret and limit communication as well:

Proposition 1 (Communication Complexity). *If Algorithm 1 is run with threshold D , then total rounds of communication n is upper bounded by $2\sqrt{(dT/D)} \cdot \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}) + 4$.*

Now, the perturbations $\mathbf{H}_{i,t}, \mathbf{h}_{i,t}$ are designed keeping the privacy setting in mind. In our paper, we defer this to the subsequent section in a subroutine known as PRIVATIZER, and concentrate on the performance guarantees first. The PRIVATIZER subroutine provides suitable perturbations based on the privacy budget (ε and δ), and is inspired by the single-agent private algorithm of [41]. In this paper, we assume these budgets to be identical for all agents (however, the algorithm and analysis hold for unique privacy budgets as well, as long as a lower bound on the budgets is known). In turn, the quantities ε and δ affect the algorithm (and regret) via the quantities ρ_{\min}, ρ_{\max} , and κ which can be understood as spectral bounds on $\mathbf{H}_{i,t}, \mathbf{h}_{i,t}$.

Definition 2 (Sparsely-accurate ρ_{\min}, ρ_{\max} and κ). *Consider a subsequence $\bar{\sigma}$ of $[T] = 1, \dots, T$ of size n . The bounds $0 \leq \rho_{\min} \leq \rho_{\max}$ and κ are $(\alpha/2nM, \bar{\sigma})$ -accurate for $(\mathbf{H}_{i,t})_{i \in [M], t \in \bar{\sigma}}$ and $(\mathbf{h}_{i,t})_{i \in [M], t \in \bar{\sigma}}$, if, for each round $t \in \bar{\sigma}$ and agent i :*

$$\|\mathbf{H}_{i,t}\| \leq \rho_{\max}, \|\mathbf{H}_{i,t}^{-1}\| \leq 1/\rho_{\min}, \|\mathbf{h}_{i,t}\|_{\mathbf{H}_{i,t}^{-1}} \leq \kappa; \text{ with probability at least } (1 - \alpha/2nM).$$

The motivation for obtaining accurate bounds ρ_{\min}, ρ_{\max} and κ stems from the fact that in the non-private case, the quantities that determine regret are not stochastic conditioned on the obtained sequence $(\mathbf{x}_t, y_t)_{t \in [T]}$, whereas the addition of stochastic regularizers in the private case requires us to have control over their spectra to achieve any meaningful regret. To form the UCB, recall that we additionally require a suitable exploration sequence $\beta_{i,t}$ for each agent, which is defined as follows.

Definition 3 (Accurate $(\beta_{i,t})_{i \in [M], t \in [T]}$). *A sequence $(\beta_{i,t})_{i \in [M], t \in [T]}$ is (α, M, T) -accurate for $(\mathbf{H}_{i,t})_{i \in [M], t \in [T]}$ and $(\mathbf{h}_{i,t})_{i \in [M], t \in [T]}$, if it satisfies $\|\tilde{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}^*\|_{\mathbf{V}_{i,t}} \leq \beta_{i,t}$ with probability at least $1 - \alpha$ for all rounds $t = 1, \dots, T$ and agents $i = 1, \dots, M$ simultaneously.*

Proposition 2. Consider an instance of the problem where synchronization occurs exactly n times on instances $\bar{\sigma}$, up to and including T trials, and ρ_{\min}, ρ_{\max} and κ are $(\alpha/2nM)$ -accurate. Then, for Algorithm 1, the sequence $(\beta_{i,t})_{i \in [M], t \in [T]}$ is (α, M, T) -accurate where:

$$\beta_{i,t} := \sigma \sqrt{2 \log(2/\alpha) + d \log(\det(\mathbf{V}_{i,t})) - d \log(M\rho_{\min})} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M}.$$

The key point here is that for the desired levels of privacy (ε, δ) and synchronization rounds (n) , we can calculate appropriate ρ_{\min}, ρ_{\max} and κ , which in turn provide us a UCB algorithm with guarantees. Our basic techniques are similar to [1, 41], however, the multi-agent setting introduces several new aspects in the analysis, such as the control of heterogeneity.

Theorem 1 (Private Group Regret with Centralized Controller). Assuming Proposition 2 holds, and synchronization occurs in at least $n = \Omega(d \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}))$ rounds, Algorithm 1 obtains the following group pseudoregret with probability at least $1 - \alpha$:

$$\mathcal{R}_M(T) = O\left(\sigma\sqrt{MTd} \left(\log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}) + \sqrt{\log 2/\alpha} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M}\right)\right).$$

This regret bound is obtained by setting $D = 2Td(\log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}) + 1)^{-1}$, which therefore ensures $O(M \log T)$ total communication (by Proposition 1). However, the remarks next clarify a more sophisticated relationship between communication, privacy and regret:

Remark 1 (Communication Complexity and Regret). Theorem 1 assumes that communication is essentially $O(M \log T)$. This rate (and corresponding D) is chosen to provide a balance between privacy and utility, and in fact, can be altered depending on the application. In the Appendix, we demonstrate that when we allow $O(MT)$ communication (i.e., synchronize every round), the regret can be improved by a factor of $O(\sqrt{\log(T)})$ (in both the gap-dependent and independent bounds), to match the single-agent $\tilde{O}(\sqrt{dMT})$ rate. Similarly, following the reasoning in [48], we can show that with $O(M^{1.5}d^3)$ communication (i.e., independent of T), we incur regret of $O(\sqrt{dMT} \log^2(MT))$.

Theorem 1 demonstrates the relationship between communication complexity (i.e., number of synchronization rounds) and the regret bound for a fixed privacy budget, via the dependence on the bounds ρ_{\min}, ρ_{\max} and κ . We now present similar results (for a fixed privacy budget) on group regret for the decentralized setting, which is more involved, as the delays and lack of a centralized controller make it difficult to control the heterogeneity of information between agents. Subsequently, in the next section, we will present results on the privacy guarantees of our algorithms.

4.2 Decentralized Peer-to-Peer Environment

In this environment, we assume that the collection of agents communicate by directly sending messages to each other. The communication network is denoted by an undirected (connected) graph $\mathcal{G} = (V, E)$, where, edge $e_{ij} \in E$ if agents i and j can communicate directly. The protocol operates as follows: every trial, each agent interacts with their respective bandit, and obtains a reward. At any trial t , after receiving the reward, each agent v sends the message $\mathbf{m}_{v,t}$ to all its neighbors in \mathcal{G} . This message is forwarded from agent to agent γ times (taking one trial of the bandit problem each between forwards), after which it is dropped. This communication protocol, based on the *time-to-live* (delay) parameter $\gamma \leq \text{diam}(\mathcal{G})$ is a common technique to control communication complexity, known as the LOCAL protocol [19, 34, 42]. Each agent $v \in V$ therefore also receives messages $\mathbf{m}_{v',t-d(v,v')}$ from all the agents v' such that $d(v, v') \leq \gamma$, i.e., from all agents in $N_\gamma(v)$.

There are several differences in this setting compared to the previous one: first, since agents receive messages from different other agents based on their position in \mathcal{G} , they generally have heterogenous information throughout (no global sync). Second, information does not flow instantaneously through \mathcal{G} , and messages can take up to γ rounds to be communicated (delay). This requires (mainly technical) changes to the centralized algorithm in order to control the regret. We present the pseudocode our algorithm DECENTRALIZED FEDUCB in the Appendix, but highlight the major differences from the centralized version as follows. The first key algorithmic change from the earlier variant is the idea of subsampling, inspired by the work of Weinberger and Ordentlich [50]. Each agent i maintains γ total estimators $\bar{\theta}_{i,g}, g = 1, \dots, \gamma$, and uses each estimator in a round-robin fashion, i.e., at $t = 1$ each agent uses $\bar{\theta}_{i,1}$, and at $t = 2$, each agent uses $\bar{\theta}_{i,2}$, and so on. These estimators (and their associated

Algorithm 2 PRIVATIZER($\varepsilon, \delta, M, T$) for any agent i

- 1: **Initialization:**
 - 2: If communication rounds n are fixed *a priori*, set $m \leftarrow 1 + \lceil \log n \rceil$, else $m \leftarrow 1 + \lceil \log T \rceil$.
 - 3: Create binary tree \mathcal{T} of depth m .
 - 4: **for** node n in \mathcal{T} **do**
 - 5: Sample noise $\widehat{\mathbf{N}} \in \mathbb{R}^{(d+1) \times (d+1)}$, where $\widehat{\mathbf{N}}_{ij} \sim \mathcal{N}(0, 16m(L^2 + 1)^2 \log(2/\delta)^2 / \varepsilon^2)$.
 - 6: Store $\mathbf{N} = (\widehat{\mathbf{N}} + \widehat{\mathbf{N}}^\top) / \sqrt{2}$ at node n .
 - 7: **end for**
 - 8: **Runtime:**
 - 9: **for** each communication round $t \leq n$ **do**
 - 10: Receive $\widehat{\mathbf{Q}}_{i,t}$ from agent, and insert it into \mathcal{T} (see [26], Alg. 5).
 - 11: Compute $\widehat{\mathbf{M}}_{i,t+1}$ using the least nodes of \mathcal{T} (see [26], Alg. 5).
 - 12: Set $\widehat{\mathbf{U}}_{i,t+1} = \mathbf{U}_{i,t+1} + \mathbf{H}_{i,t}$ as top-left $d \times d$ submatrix of $\widehat{\mathbf{M}}_{i,t+1}$.
 - 13: Set $\widehat{\mathbf{u}}_{i,t+1} = \mathbf{u}_{i,t+1} + \mathbf{h}_{i,t}$ as first d entries of last column of $\widehat{\mathbf{M}}_{i,t+1}$.
 - 14: Return $\widehat{\mathbf{U}}_{i,t+1}, \widehat{\mathbf{u}}_{i,t+1}$ to agent.
 - 15: **end for**
-

$\mathbf{V}_{i,g}, \tilde{\mathbf{u}}_{i,g}$) are updated in a manner similar to Alg. 1: if the log det of the g^{th} Gram matrix exceeds a threshold $D/(\Delta_{i,g} + 1)$, then the agent broadcasts a request to synchronize. Each agent j within the γ -clique of i that receives this request broadcasts its own $\mathbf{V}_{j,g}, \tilde{\mathbf{u}}_{j,g}$. Therefore, each $\mathbf{V}_{i,g}, \tilde{\mathbf{u}}_{i,g}$ is updated with action/reward pairs from *only* the trials they were employed in (across all agents). This ensures that if a signal to synchronize the g^{th} set of parameters has been broadcast by an agent i , all agents within the γ -clique of i will have synchronized their g^{th} parameters by the next 2 rounds they will be used again (i.e., 2γ trials later). We defer most of the intermediary results to the Appendix for brevity, and present the primary regret bound (in terms of privacy parameters ρ_{\min}, ρ_{\max} and κ):

Theorem 2 (Decentralized Private Group Regret). *Assuming Proposition 2 holds, and synchronization occurs in at least $n = \Omega(d(\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma)(1 + L^2)^{-1} \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}))$ rounds, decentralized FEDUCB obtains the following group pseudoregret with probability at least $1 - \alpha$:*

$$\mathcal{R}_M(T) = O\left(\sigma \sqrt{M(\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma)} T d \left(\log\left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{\gamma d \rho_{\min}}\right) + \sqrt{\log \frac{2}{\alpha}} + S \sqrt{M \rho_{\max}} + \kappa \sqrt{M} \right)\right).$$

Remark 2 (Decentralized Group Regret). *Decentralized FEDUCB obtains an identical dependence on the privacy bounds ρ_{\min}, ρ_{\max} and κ and horizon T as Algorithm 1, since the underlying bandit subroutines are identical for both. The key difference is in additional the leading factor of $\sqrt{\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma}$, which arises from the delayed spread of information: if \mathcal{G} is dense, e.g., complete, then $\gamma = 1$ and $\bar{\chi}(\mathcal{G}_\gamma) = 1$, since there is only one clique of \mathcal{G} . In the worst case, if \mathcal{G} is a line graph, then $\bar{\chi}(\mathcal{G}_\gamma) = M/\gamma$, giving an additional factor of M (i.e., it is as good as each agent acting individually). In more practical scenarios, we expect \mathcal{G} to be hierarchical, and expect a delay overhead of $o(1)$. Note that since each agent only communicates within its clique, the factor $\bar{\chi}(\mathcal{G}_\gamma)$ is optimal, see [12].*

4.3 Privacy Guarantees

We now discuss the privacy guarantees for both algorithms. Here we present results for the centralized algorithm, but our results hold (almost identically) for the decentralized case as well (see appendix). Note that each agent interacts with data from other agents only via the cumulative parameters \mathbf{S}_t and \mathbf{s}_t . These, in turn, depend on $\mathbf{Z}_{i,t} = \mathbf{U}_{i,t} + \mathbf{H}_{i,t}$ and $\mathbf{z}_{i,t} = \tilde{\mathbf{u}}_{i,t} + \mathbf{h}_{i,t}$ for each agent i , on the instances t that synchronization occurs.

Proposition 3 (see [16, 41]). *Consider $n \leq T$ synchronization rounds occurring on trials $\bar{\sigma} \subseteq [T]$. If the sequence $(\mathbf{Z}_{i,t}, \mathbf{z}_{i,t})_{t \in \bar{\sigma}}$ is (ε, δ) -differentially private with respect to $(\mathbf{x}_{i,t}, y_{i,t})_{t \in [T]}$, for each agent $i \in [M]$, then all agents are (ε, δ, M) -federated differentially private.*

Tree-Based Mechanism. Let x_1, x_2, \dots, x_T be a (matrix-valued) sequence of length T , and $s_i = \sum_{t=1}^i x_t$ be the incremental sum of the sequence that must be released privately. The tree-based mechanism [17] for differential privacy involves a trusted entity maintaining a binary tree \mathcal{T} (of depth $m = 1 + \lceil \log_2 T \rceil$), where the leaf nodes contain the sequence items x_i , and each parent node maintains the (matrix) sum of its children. Let n_i be the value stored at any node in the tree. The

mechanism achieves privacy by adding a noise h_i to each node, and releasing $n_i + h_i$ whenever a node is queried. Now, to calculate s_t for some $t \in [T]$, the procedure is to traverse the tree \mathcal{T} up to the leaf node corresponding to x_t , and summing up the values at each node on the traversal path. The advantage is that we only access at most m nodes, and add $m = O(\log T)$ noise (instead of $O(T)$).

The implementation of the private release of $(\mathbf{U}_{i,t}, \bar{\mathbf{u}}_{i,t})$ is done by the ubiquitous tree-based mechanism for partial sums. Following [41], we also aggregate both into a single matrix $\mathbf{M}_{i,t} \in \mathbb{R}^{(d+1) \times (d+1)}$, by first concatenating: $\mathbf{A}_{i,t} := [\mathbf{X}_{i,1:t}, \mathbf{y}_{i,1:t}] \in \mathbb{R}^{t \times (d+1)}$, and then computing $\mathbf{M}_{i,t} = \mathbf{A}_{i,t}^\top \mathbf{A}_{i,t}$. Furthermore, the update is straightforward: $\mathbf{M}_{i,t+1} = \mathbf{M}_{i,t} + [\mathbf{x}_{i,t}^\top \mathbf{y}_{i,t}]^\top [\mathbf{x}_{i,t}^\top \mathbf{y}_{i,t}]$. Recall that in our implementation, we only communicate in synchronization rounds (and not every round). Assume that two successive rounds of synchronization occur at time t' and t . Then, at instant t , each agent i inserts $\sum_{\tau=t'}^t [\mathbf{x}_{i,\tau}^\top \mathbf{y}_{i,\tau}]^\top [\mathbf{x}_{i,\tau}^\top \mathbf{y}_{i,\tau}]$ into \mathcal{T} , and computes $(\mathbf{U}_{i,t}, \bar{\mathbf{u}}_{i,t})$ by summing up the entire path up to instant t via the tree mechanism. Therefore, the tree mechanism accesses at most $m = 1 + \lceil \log_2 n \rceil$ nodes (where n total rounds of communication occur until instant T), and hence noise that ensures each node guarantees $(\varepsilon/\sqrt{8m \ln(2/\delta)}, \delta/2m)$ -privacy is sufficient to make the outgoing sequence (ε, δ) -private. This is different from the setting in the joint DP single-agent bandit [41], where observations are inserted *every* round, and not only for synchronization rounds.

To make the partial sums private, a noise matrix is also added to each node in \mathcal{T} . We utilize additive Gaussian noise: at each node, we sample $\widehat{\mathbf{N}} \in \mathbb{R}^{(d+1) \times (d+1)}$, where each $\widehat{\mathbf{N}}_{i,j} \sim \mathcal{N}(0, \sigma_N^2)$, and $\sigma_N^2 = 16m(L^2 + 1)^2 \log(2/\delta)^2/\varepsilon^2$, and symmetrize it (see step 6 of Alg. 2). The total noise $\mathbf{H}_{i,t}$ is the sum of at most m such terms, hence the variance of each element in $\mathbf{H}_{i,t}$ is $\leq m\sigma_N^2$. We can bound the operator norm of the top-left $(d \times d)$ -submatrix of each noise term. Therefore, to guarantee (ε, δ, M) -federated DP, we require that with probability at least $1 - \alpha/nM$:

$$\|\mathbf{H}_{i,t}\|_{\text{op}} \leq \Lambda = \sqrt{32}m(L^2 + 1) \log(4/\delta)(4\sqrt{d} + 2 \ln(2nM/\alpha))/\varepsilon.$$

Remark 3 (Privacy Guarantee). *The procedure outlined above guarantees that each of the n outgoing messages $(\mathbf{U}_{i,t}, \bar{\mathbf{u}}_{i,t})$ (where t is a synchronization round) for any agent i is (ε, δ) -differentially private. This analysis considers the L_2 -sensitivity with respect to a single differing observation, i.e., (x, y) and not the entire message itself, i.e., the complete sequence $(\mathbf{x}_{i,\tau}, \mathbf{y}_{i,\tau})_{\tau=t'}^t$ (where t' and t are successive synchronization rounds), which may potentially have $O(t)$ sensitivity and warrants a detailed analysis. While our analysis is sufficient for the user-level adversary model, there may be settings where privacy is required at the message-level as well, which we leave as future work.*

However, as noted by [41], this $\mathbf{H}_{i,t}$ would not always be PSD. To ensure that it is always PSD, we can shift each $\mathbf{H}_{i,t}$ by $2\Lambda \mathbf{I}$, giving a bound on ρ_{\max} . Similarly, we can obtain bounds on ρ_{\min} and κ :

Proposition 4. *Fix $\alpha > 0$. If each agent i samples noise parameters $\mathbf{H}_{i,t}$ and $\mathbf{h}_{i,t}$ using the tree-based Gaussian mechanism mentioned above for all n trials of $\bar{\sigma}$ in which communication occurs, then the following ρ_{\min}, ρ_{\max} and κ are $(\alpha/2nM, \bar{\sigma})$ -accurate bounds:*

$$\rho_{\min} = \Lambda, \rho_{\max} = 3\Lambda, \kappa \leq \sqrt{m(L^2 + 1) \left(\sqrt{d} + 2 \log(2nM/\alpha) \right) / (\sqrt{2}\varepsilon)}.$$

Remark 4 (Strategyproof Analysis). *The mechanism presented above assumes the worst-case communication, i.e., synchronization occurs every round, therefore at most T partial sums will be released, and $m = 1 + \lceil \log T \rceil$. This is not true for general settings, where infrequent communication would typically require $m = O(\log \log T)$. However, if any agent is byzantine and requests a synchronization every trial, m must be $1 + \lceil \log T \rceil$ to ensure privacy. In case the protocol is fixed in advance (i.e., synchronization occurs on a pre-determined set $\bar{\sigma}$ of n rounds), then we can set $m = 1 + \lceil \log n \rceil$ to achieve the maximum utility at the desired privacy budget.*

Remark 5 (Decentralized Protocol). *Decentralized FEDUCB obtains similar bounds for ρ_{\min}, ρ_{\max} and κ , with $m = 1 + \lceil \log(T/\gamma) \rceil$. An additional term of $\log(\gamma)$ appears in Λ and κ , since we need to now maintain γ partial sums with at most T/γ elements in the worst-case (see Appendix). Unsurprisingly, there is no dependence on the network \mathcal{G} , as privatization is done at the source itself.*

Corollary 1 ((ε, δ) -dependent Regret). *FEDUCB with the PRIVATIZER subroutine in Alg. 2, obtains $\tilde{O}(d^{3/4} \sqrt{MT/\varepsilon})$ centralized regret and $\tilde{O}(d^{3/4} \sqrt{(\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma)MT/\varepsilon})$ decentralized regret.*




Figure 1: A comparison of centralized FEDUCB on 3 different axes. Fig. (A) describes the variation in asymptotic per-agent regret for varying privacy budget ϵ (where $\delta = 0.1$); (B) describes the effect of n in private (solid) vs. non-private (dashed) settings; (C) describes the effect of d in per-agent regret in the private setting ($n = O(M \log T)$, $\epsilon = 1$, $\delta = 0.1$). Experiments averaged over 100 runs.

4.4 Experiments

We provide an abridged summary of our experiments (ref. Appendix for all experiments). Here, we focus on the centralized environment, and on the variation of the regret with communication complexity and privacy budget. For all experiments, we assume $L = S = 1$. For any d , we randomly fix $\theta^* \in \mathcal{B}_d(1)$. Each $\mathcal{D}_{i,t}$ is generated as follows: we randomly sample $K \leq d^2$ actions \mathbf{x} , such that for $K - 1$ actions $0.5 \leq \langle \mathbf{x}, \theta^* \rangle \leq 0.6$ and for the optimal \mathbf{x}^* , $0.7 \leq \langle \mathbf{x}^*, \theta^* \rangle \leq 0.8$ such that $\Delta \geq 0.1$ always. $y_{i,t}$ is sampled from $\text{Ber}(\langle \mathbf{x}_{i,t}, \theta^* \rangle)$ such that $\mathbb{E}[y_{i,t}] = \langle \mathbf{x}_{i,t}, \theta^* \rangle$ and $|y_{i,t}| \leq 1$. Results are in Fig. 1, and experiments are averaged on 100 trials.

Experiment 1: Privacy Budget. In this setting, we set $n = O(M \log T)$, $d = 10$ (to balance communication and performance), and plot the average per-agent regret after $T = 10^7$ trials for varying M and ϵ , while keeping $\delta = 0.1$. Figure 1A describes the results, competitive even at large privacy budget.

Experiment 2: Communication. To test the communication-privacy trade-off, we plot the regret curves for $M = 100$ agents on $n = O(M^{1.5})$, $O(M \log T)$, $O(MT)$ communication for both private ($\epsilon = 1$) and non-private settings. We observe a tradeoff as highlighted in Remark 1.

Experiment 3: Dependence on d . Finally, as an ablation, we provide the average per-agent regret curves for $M = 100$ by varying the dimensionality d . We observe essentially a quadratic dependence.

5 Discussion and Conclusion

The relentless permeance of intelligent decision-making on sensitive user data necessitates the development of technically sound and robust machine learning systems; it is difficult to stress enough the importance of protecting user data in modern times. There is a significant effort [7, 22] in creating decentralized decision-making systems that benefit from pooling data from multiple sources while maintaining user-level privacy. This research provides an instance of such a system, along with guarantees on both its utility and privacy. From a technical perspective, we make improvements along several fronts – while there has been prior work on multi-agent private linear bandits [21, 38] our work is the first to provide rigorous guarantees on private linear bandits in the multi-agent setting. Our work additionally provides the first algorithm with regret guarantees for contextual bandits in decentralized networks, extending the work of many on multi-armed bandits [39, 31, 30, 14, 12].

There are several unresolved questions in this line of work, as highlighted by our algorithm itself. In the decentralized case, our algorithm obtains a communication overhead of $O(\sqrt{\chi(\mathcal{G}_\gamma)} \cdot \gamma)$, which we comment is an artefact of our proof technique. We conjecture that the true dependence scales as $O(\sqrt{\alpha(\mathcal{G}_\gamma)} + \gamma)$ (where $\alpha(\cdot)$ is the independence number of the graph), that would be obtained by an algorithm that bypasses subsampling, with a more sophisticated synchronization technique. Additionally, the lower bound obtained for private contextual bandit estimation [41] has an $O(\epsilon^{-1})$ dependence on the privacy budget. In the federated setting, given the intricate relationship between communication and privacy, an interesting line of inquiry would be to understand the communication-privacy tradeoff in more detail, with appropriate lower bounds for fixed communication budgets.

Broader Impact

We envision a world that runs on data-dependent inference. Given the ease of availability of sensitive data, large-scale centralized data sources are not feasible to uphold the best interests and safety of the average user, and decentralized inference mechanisms are a feasible alternative balancing utility with important normative values as privacy, transparency and accountability. This research is a part of such a broader research goal, to create decentralized systems that maintain the security of the end-user in mind. We believe that methods like ours are a starting point for better private algorithms that can be readily deployed in industry, with appropriate guarantees as well. While there are several subsequent steps that can be improved both theoretically and practically, the design for this research from the ground-up has been to ensure robustness and security (sometimes) in lieu of efficiency.

Funding Disclosures

This project was funded by the MIT Trust:Data Consortium.

A Full Proofs

Proposition 5 (Communication Complexity, Proposition 1 of main paper). *If Algorithm 1 is run with threshold D , then total rounds of communication n is upper bounded by $2\sqrt{(dT/D) \cdot \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min})} + 4$.*

Proof. We denote the number of (common) bandit trials between two rounds of communication as an epoch. Let $n' = \sqrt{\frac{DT}{d \log(\rho_{\max}/\rho_{\min} + TL^2/(d\rho_{\min}))}} + 1$. There can be at most $\lceil T/n' \rceil$ rounds of communication such that they occur after an epoch of length n' . On the other hand, if there is any round of communication succeeding an epoch (that begins, say at time t) of length $< n'$, then for that epoch, $\log \frac{\det(\mathbf{S}_{i,t+n'})}{\det(\mathbf{S}_{i,t})} > D/n'$. Let the communication occur at a set of rounds t'_1, \dots, t'_n . Now, since:

$$\sum_{i=1}^{n-1} \log \frac{\det(\mathbf{S}_{i,t'_{i+1}})}{\det(\mathbf{S}_{i,t'_i})} = \log \frac{\det(\mathbf{S}_{i,T})}{\det(\mathbf{S}_{i,0})} \leq d \log(\rho_{\max}/\rho_{\min} + TL^2/(d\rho_{\min})), \quad (4)$$

We have that the total number of communication rounds succeeding epochs of length less than n' is upper bounded by $\log \frac{\det(\mathbf{S}_{i,T})}{\det(\mathbf{S}_{i,0})} \leq d \log(\rho_{\max}/\rho_{\min} + TL^2/(d\rho_{\min})) \cdot (n'/D)$. Combining both the results together, we have the total rounds of communication as:

$$n \leq \lceil T/n' \rceil + \lceil d \log(\rho_{\max}/\rho_{\min} + TL^2/(d\rho_{\min})) \cdot (n'/D) \rceil \quad (5)$$

$$\leq T/n' + d \log(\rho_{\max}/\rho_{\min} + TL^2/(d\rho_{\min})) \cdot (n'/D) + 2 \quad (6)$$

Replacing n' from earlier gives us the final result. \square

Proposition 6 ($\beta_{i,t}$, Proposition 2 of main paper). *Consider an instance of the problem where synchronization occurs exactly n times on instances $\bar{\sigma}$, up to and including T trials, and ρ_{\min}, ρ_{\max} and κ are $(\alpha/2nM)$ -accurate. Then, for Algorithm 1, the sequence $(\beta_{i,t})_{i \in [M], t \in [T]}$ is (α, M, T) -accurate where:*

$$\beta_{i,t} := \sigma \sqrt{2 \log(2/\alpha) + d \log(\det(\mathbf{V}_{i,t})) - d \log(M\rho_{\min})} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M}.$$

Proof. Let $\mathbf{X}_{i,<t}$ denote the set of all observations available to the agent (including private communication from other agents). Furthermore, let the noise-free Gram matrix of all observations as $\mathbf{G}_{i,t} = \sum_{\tau=1}^t \mathbf{x}_{i,\tau} \mathbf{x}_{i,\tau}^\top + \sum_{\tau=1}^{t_s} \sum_{j=1, j \neq i}^M \mathbf{x}_{j,\tau} \mathbf{x}_{j,\tau}^\top$ (where t_s is the last synchronization iteration). We also have that $\mathbf{V}_{i,t} = \mathbf{G}_{i,t} + \sum_{j=1}^M \mathbf{H}_{j,t}$, and for any t between synchronization rounds t_s and t_{s+1} , $\mathbf{H}_{j,t} = \mathbf{H}_{j,t_s}$ for all j . By definition, $\tilde{\boldsymbol{\theta}}_{i,t} = \mathbf{V}_{i,t}^{-1} \tilde{\mathbf{u}}_{i,t}$, $\tilde{\mathbf{u}}_{i,t} = \mathbf{u}_{i,t} + \sum_{j \in [M]} \mathbf{h}_{j,t}$ and $\mathbf{u}_{i,t} = \mathbf{X}_{i,<t}^\top \mathbf{y}_{<t}$. Therefore, we have that,

$$\begin{aligned} \boldsymbol{\theta}^* - \tilde{\boldsymbol{\theta}}_{i,t} &= \boldsymbol{\theta}^* - \mathbf{V}_{i,t}^{-1} \left(\mathbf{X}_{i,<t}^\top \mathbf{y}_{<t} + \sum_{j \in [M]} \mathbf{h}_{j,t} \right) \\ &= \boldsymbol{\theta}^* - \mathbf{V}_{i,t}^{-1} \left(\mathbf{X}_{i,<t}^\top \mathbf{X}_{i,<t} \boldsymbol{\theta}^* + \mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t} + \sum_{j \in [M]} \mathbf{h}_{j,t} \right) \\ &= \boldsymbol{\theta}^* - \mathbf{V}_{i,t}^{-1} \left(\mathbf{V}_{i,t} \boldsymbol{\theta}^* - \sum_{j \in [M]} \mathbf{H}_{j,t} \boldsymbol{\theta}^* + \mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t} + \sum_{j \in [M]} \mathbf{h}_{j,t} \right) \\ &= \mathbf{V}_{i,t}^{-1} \left(\sum_{j \in [M]} \mathbf{H}_{j,t} \boldsymbol{\theta}^* - \mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t} - \sum_{j \in [M]} \mathbf{h}_{j,t} \right). \end{aligned}$$

Multiplying both sides by $\mathbf{V}_{i,t}^{1/2}$ gives

$$\begin{aligned}
\mathbf{V}_{i,t}^{1/2} (\boldsymbol{\theta}^* - \tilde{\boldsymbol{\theta}}_{i,t}) &= \mathbf{V}_{i,t}^{-1/2} \left(\sum_{j \in [M]} \mathbf{H}_{j,t} \boldsymbol{\theta}^* - \mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t} - \sum_{j \in [M]} \mathbf{h}_{j,t} \right) \\
\Rightarrow \|\boldsymbol{\theta}^* - \tilde{\boldsymbol{\theta}}_{i,t}\|_{\mathbf{V}_{i,t}} &= \left\| \sum_{j \in [M]} \mathbf{H}_{j,t} \boldsymbol{\theta}^* - \mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t} - \sum_{j \in [M]} \mathbf{h}_{j,t} \right\|_{\mathbf{V}_{i,t}^{-1}} \quad (\text{Applying } \|\cdot\|) \\
&\leq \left\| \sum_{j \in [M]} \mathbf{H}_{j,t} \boldsymbol{\theta}^* \right\|_{\mathbf{V}_{i,t}^{-1}} + \|\mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t}\|_{\mathbf{V}_{i,t}^{-1}} + \left\| \sum_{j \in [M]} \mathbf{h}_{j,t} \right\|_{\mathbf{V}_{i,t}^{-1}} \\
&\hspace{15em} (\text{Triangle inequality})
\end{aligned}$$

Making the substitution $\mathbf{H}_t = \sum_{j \in [M]} \mathbf{H}_{j,t}$,

$$\begin{aligned}
&\leq \|\mathbf{H}_t \boldsymbol{\theta}^*\|_{\mathbf{H}_t^{-1}} + \|\mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t}\|_{\mathbf{V}_{i,t}^{-1}} + \|\mathbf{h}_{j,t}\|_{\mathbf{H}_t^{-1}} \\
&\hspace{15em} (\text{Since } \mathbf{V}_{i,t} \succcurlyeq \sum_{j \in [M]} \mathbf{H}_{j,t}) \\
&\leq \|\boldsymbol{\theta}^*\|_{\mathbf{H}_t} + \|\mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t}\|_{(\mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I})^{-1}} + \sum_{j \in [M]} \|\mathbf{h}_{j,t}\|_{\mathbf{H}_t^{-1}} \\
&\hspace{15em} (\text{Since } \forall i \in [M], \mathbf{V}_{i,t} \succcurlyeq \mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I})
\end{aligned}$$

Now, note that since we only require at most Mn different noise matrices, we only need the noise sequences \mathbf{H} by a union bound over all TM rounds (T rounds per agent), we can say that simultaneously for all $i \in [M], t \in [T]$, with probability at least $1 - \alpha/2$, $\|\boldsymbol{\theta}^*\|_{\mathbf{H}_t} \leq \sqrt{\|\mathbf{H}_t\|} \|\boldsymbol{\theta}^*\| \leq S\sqrt{M\rho_{\max}}$ and $\|\sum_{j \in [M]} \mathbf{h}_{j,t}\|_{\mathbf{H}_t^{-1}} \leq \kappa\sqrt{M}$. Next, by Theorem 1 of Abbasi-Yadkori *et al.* [1], we have that, with probability at least $1 - \alpha/2$ simultaneously,

$$\begin{aligned}
\|\mathbf{X}_{i,<t}^\top \boldsymbol{\eta}_{<t}\|_{(\mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I})^{-1}} &\leq \sigma \sqrt{2 \log \frac{2}{\alpha} + \log \frac{\det(\mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I})}{\det(M\rho_{\min} \mathbf{I})}} \\
&\leq \sigma \sqrt{2 \log \frac{2}{\alpha} + d \log \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{tL^2}{d\rho_{\min}} \right)}.
\end{aligned}$$

The last step follows from (a) noting that $\forall i \in [M], \mathbf{G}_{i,t} + M\rho_{\max} \mathbf{I} \succcurlyeq \mathbf{V}_{i,t} \succcurlyeq \mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I} \Rightarrow \det(\mathbf{G}_{i,t} + M\rho_{\max} \mathbf{I}) \geq \det(\mathbf{V}_{i,t}) \geq \det(\mathbf{G}_{i,t} + M\rho_{\min} \mathbf{I})$, and (b) the trace-determinant inequality. Putting it all together, we have that all $\beta_{i,t}$ are bounded by $\bar{\beta}_t$, given by,

$$\bar{\beta}_t := \sigma \sqrt{2 \log \frac{2}{\alpha} + d \log \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{tL^2}{d\rho_{\min}} \right)} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M}.$$

□

Lemma 1. *The instantaneous pseudoregret $r_{i,t}$ obtained by any agent i at any instant t obeys the following:*

$$r_{i,t} \leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}}$$

Proof. At every round, each agent i selects an ‘‘optimistic’’ action $\mathbf{x}_{i,t}$ such that,

$$(\mathbf{x}_{i,t}, \bar{\boldsymbol{\theta}}_{i,t}) = \arg \max_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}_{i,t} \times \mathcal{E}_{i,t}} \langle \mathbf{x}, \boldsymbol{\theta} \rangle. \quad (7)$$

Let $\mathbf{x}_{i,t}^*$ be the optimal action at time t for agent i , i.e., $\mathbf{x}_{i,t}^* = \arg \max_{\mathbf{x} \in \mathcal{D}_{i,t}} \langle \mathbf{x}, \boldsymbol{\theta}^* \rangle$. We can then decompose the immediate pseudoregret $r_{i,t}$ for agent i as the following.

$$\begin{aligned}
r_{i,t} &= \langle \mathbf{x}_{i,t}^*, \boldsymbol{\theta}^* \rangle - \langle \mathbf{x}_{i,t}, \boldsymbol{\theta}^* \rangle \\
&\leq \langle \mathbf{x}_{i,t}, \bar{\boldsymbol{\theta}}_{i,t} \rangle - \langle \mathbf{x}_{i,t}, \boldsymbol{\theta}^* \rangle && \text{(Since } (\mathbf{x}_{i,t}, \bar{\boldsymbol{\theta}}_{i,t}) \text{ is optimistic)} \\
&= \langle \mathbf{x}_{i,t}, \bar{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}^* \rangle \\
&= \left\langle \mathbf{V}_{i,t}^{-1/2} \mathbf{x}_{i,t}, \mathbf{V}_{i,t}^{1/2} (\bar{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}^*) \right\rangle && (\mathbf{V}_{i,t} \succcurlyeq \mathbf{0}) \\
&\leq \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}} \|\bar{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}^*\|_{\mathbf{V}_{i,t}} && \text{(Cauchy-Schwarz)} \\
&\leq \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}} \left(\|\bar{\boldsymbol{\theta}}_{i,t} - \tilde{\boldsymbol{\theta}}_{i,t}\|_{\mathbf{V}_{i,t}} + \|\tilde{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}^*\|_{\mathbf{V}_{i,t}} \right) && \text{(Triangle inequality)} \\
&\leq 2\beta_{t,i} \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}} && \text{(Since } \bar{\boldsymbol{\theta}}_{i,t}, \boldsymbol{\theta}^* \in \mathcal{E}_{i,t}) \\
&\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}}. && \text{(By Proposition 2)}
\end{aligned}$$

□

Lemma 2 (Elliptical Potential, Lemma 3 of [1], Lemma 22 of [41]). *Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be vectors such that $\|\mathbf{x}\|_2 \leq L$. Then, for any positive definite matrix $\mathbf{U}_0 \in \mathbb{R}^{d \times d}$, define $\mathbf{U}_t := \mathbf{U}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$ for all t . Then, for any $\nu > 1$,*

$$\sum_{t=1}^n \|\mathbf{x}_t\|_{\mathbf{U}_{t-1}^{-1}}^2 \leq 2d \log_\nu \left(\frac{\text{tr}(\mathbf{U}_0) + nL^2}{d \det^{1/d}(\mathbf{U}_0)} \right).$$

Proof. We urge the readers to refer to [41] for a proof of this statement. □

Theorem 3 (Private Group Regret with Centralized Controller, Theorem 1 of the main paper). *Assuming Proposition 2 holds, and synchronization occurs in at least $n = \Omega(d \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}))$ rounds, Algorithm 1 obtains the following group pseudoregret with probability at least $1 - \alpha$:*

$$\mathcal{R}_M(T) = O\left(\sigma\sqrt{MTd} \left(\log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}) + \sqrt{\log 2/\alpha} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M}\right)\right).$$

Proof. Consider a hypothetical agent that takes the following MT actions $\mathbf{x}_{1,1}, \mathbf{x}_{2,1}, \dots, \mathbf{x}_{1,2}, \mathbf{x}_{2,2}, \dots, \mathbf{x}_{M-1,T}, \mathbf{x}_{M,T}$ sequentially. Let $\mathbf{W}_{i,t} = M\rho_{\min}\mathbf{I} + \sum_{j=1}^M \sum_{u=1}^{t-1} \mathbf{x}_{j,u} \mathbf{x}_{j,u}^\top + \sum_{j=1}^{i-1} \mathbf{x}_{j,t} \mathbf{x}_{j,t}^\top$ be the Gram matrix formed until the hypothetical agent reaches $\mathbf{x}_{i,t}$. By Lemma 2, we have that

$$\sum_{t=1}^T \sum_{i=1}^M \|\mathbf{x}_{i,t}\|_{\mathbf{W}_{i,t}^{-1}}^2 \leq 2d \log \left(1 + \frac{TL^2}{d\rho_{\min}} \right). \quad (8)$$

Now, in the original setting, let T_1, T_2, \dots, T_{p-1} be the trials at which synchronization occurs. After any round T_k of synchronization, consider the cumulative Gram matrices of all observations obtained until that round as $\mathbf{V}_k, k = 1, \dots, p-1$, regularized by $M\rho_{\min}\mathbf{I}$, i.e., $\mathbf{V}_k = \sum_{i \in [M]} \sum_{t=1}^{T_k} \mathbf{x}_{i,t} \mathbf{x}_{i,t}^\top + M\rho_{\min}\mathbf{I}$. Finally, let \mathbf{V}_p denote the (regularized) Gram matrix with all trials at time T , and $\mathbf{V}_0 = M\rho_{\min}\mathbf{I}$. Therefore, we have that $\det(\mathbf{V}_0) = (M\rho_{\min})^d$, and that $\det(\mathbf{V}_p) \leq \left(\frac{\text{tr}(\mathbf{V}_p)}{d}\right)^d \leq (M\rho_{\max} + MTL^2/d)^d$. Therefore, for any $\nu > 1$,

$$\log_\nu \left(\frac{\det(\mathbf{V}_p)}{\det(\mathbf{V}_0)} \right) \leq d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right).$$

Let $R = \left\lceil d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) \right\rceil$. It follows that in all but R periods between synchronization,

$$1 \leq \frac{\det(\mathbf{V}_k)}{\det(\mathbf{V}_{k-1})} \leq \nu. \quad (9)$$

We consider the event E to be the period k when Equation 9 holds. Now, for any $T_{k-1} \leq t \leq T_k$, consider the immediate pseudoregret for any agent i . By Lemma 1, we have

$$\begin{aligned}
r_{i,t} &\leq 2\bar{\beta}_T \|x_{i,t}\|_{\mathbf{V}_{i,t}^{-1}} \\
&\leq 2\bar{\beta}_T \|x_{i,t}\|_{(\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I})^{-1}} && (\mathbf{V}_{i,t} \succcurlyeq \mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I}) \\
&\leq 2\bar{\beta}_T \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}} \cdot \sqrt{\frac{\det(\mathbf{W}_{i,t})}{\det(\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I})}} \\
&\leq 2\bar{\beta}_T \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}} \cdot \sqrt{\frac{\det(\mathbf{V}_k)}{\det(\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I})}} && (\mathbf{V}_k \succcurlyeq \mathbf{W}_{i,t}) \\
&\leq 2\bar{\beta}_T \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}} \cdot \sqrt{\frac{\det(\mathbf{V}_k)}{\det(\mathbf{V}_{k-1})}} && (\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I} \succcurlyeq \mathbf{V}_{k-1}) \\
&\leq 2\nu\bar{\beta}_T \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}}. && (\text{Event } E \text{ holds})
\end{aligned}$$

Now, we can sum up the immediate pseudoregret over all such periods where E holds to obtain the total regret for these periods. With probability at least $1 - \alpha$,

$$\begin{aligned}
\text{Regret}(T, E) &= \sum_{i=1}^M \sum_{t \in [T]: E \text{ is true}} r_{i,t} \\
&\leq \sqrt{MT \left(\sum_{i=1}^M \sum_{t \in [T]: E \text{ is true}} r_{i,t}^2 \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{MT \left(\sum_{i=1}^M \sum_{t \in [T]: E \text{ is true}} \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}} \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{MT \left(\sum_{i=1}^M \sum_{t \in [T]} \|x_{i,t}\|_{\mathbf{W}_{i,t}^{-1}} \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{2MTd \log_\nu \left(1 + \frac{TL^2}{d\rho_{\min}} \right)}.
\end{aligned}$$

Now let us consider the periods in which E does not hold. In any such period between synchronization of length $t_k = T_k - T_{k-1}$, we have, for any agent i , the regret accumulated given by:

$$\begin{aligned}
\text{Regret}([T_{k-1}, T_k]) &= \sum_{t=T_{k-1}}^{T_k} \sum_{i=1}^M r_{i,t} \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i=1}^M \sqrt{t_k \sum_{t=T_{k-1}}^{T_k} \|x_{i,t}\|_{\mathbf{V}_{i,t}^{-1}}^2} \right) \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i=1}^M \sqrt{t_k \log_\nu \left(\frac{\det(\mathbf{V}_{i,t+t_k})}{\det(\mathbf{V}_{i,t})} \right)} \right) \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i=1}^M \sqrt{t_k \log_\nu \left(\frac{\det(\mathbf{G}_{i,t+t_k} + M\rho_{\max}\mathbf{I})}{\det(\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I})} \right)} \right)
\end{aligned}$$

By Algorithm 1, we know that for all agents, $t_k \log_\nu \left(\frac{\det(\mathbf{G}_{i,t+t_k} + M\rho_{\max}\mathbf{I})}{\det(\mathbf{G}_{i,t} + M\rho_{\min}\mathbf{I})} \right) \leq D$ (since there would be a synchronization round otherwise), therefore

$$\leq 2\nu\bar{\beta}_T M\sqrt{D}.$$

Now, note that of the total p periods between synchronizations, only at most R periods will not have event E be true. Therefore, the total regret over all these periods can be bound as,

$$\begin{aligned} \text{Regret}(T, \bar{E}) &\leq R \cdot 2\nu\bar{\beta}_T M\sqrt{D} \\ &\leq 2\nu\bar{\beta}_T M\sqrt{D} \left(d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right). \end{aligned}$$

Adding it all up together gives us,

$$\begin{aligned} \text{Regret}(T) &= \text{Regret}(T, E) + \text{Regret}(T, \bar{E}) \\ &\leq 2\nu\bar{\beta}_T \left[\sqrt{2MTd \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right)} + M\sqrt{D} \left(d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right) \right] \end{aligned}$$

Setting $D = 2Td \left(\log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right)^{-1}$, we have:

$$\begin{aligned} \text{Regret}(T) &\leq 4\nu\bar{\beta}_T \sqrt{2MTd \left(\log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right)} \\ &\leq 4\nu \left(\sigma \sqrt{2 \log \frac{2}{\alpha}} + d \log \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M} \right) \sqrt{2MTd \left(\log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right)} \end{aligned}$$

Optimizing over ν gives us the final result. Asymptotically, we have (setting $\nu = e$), with probability at least $1 - \alpha$,

$$\text{Regret}(T) = O \left(\sigma \sqrt{MTd} \left(\log \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + \sqrt{\log \frac{2}{\alpha}} \right) \right). \quad (10)$$

□

Theorem 4 (Decentralized Private Group Regret, Theorem 2 of the main paper). *Assuming Proposition 2 holds, and synchronization occurs in at least $n = \Omega(d(\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma)(1 + L^2)^{-1} \log(\rho_{\max}/\rho_{\min} + TL^2/d\rho_{\min}))$ rounds, decentralized FEDUCB obtains the following group pseudoregret with probability at least $1 - \alpha$:*

$$\mathcal{R}_M(T) = O \left(\sigma \sqrt{M(\bar{\chi}(\mathcal{G}_\gamma) \cdot \gamma)} Td \left(\log \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{\gamma d\rho_{\min}} \right) + \sqrt{\log \frac{2}{\alpha}} + S\sqrt{M\rho_{\max}} + \kappa\sqrt{M} \right) \right).$$

Proof. The proof for this setting is similar to the centralized variant. We can first partition the power graph \mathcal{G}_γ of the communication network \mathcal{G} into a clique cover $\mathcal{C} = \cup_i C_i$. The overall regret can then be decomposed as the following.

$$\begin{aligned} \text{Regret}(T) &= \sum_{i=1}^M \sum_{t=1}^T r_{i,t} \\ &= \sum_{C \in \mathcal{C}} \sum_{i \in C} \sum_{t=1}^T r_{i,t} \\ &= \sum_{C \in \mathcal{C}} \text{Regret}_C(T). \end{aligned}$$

Here, $\text{Regret}_C(T)$ denotes the cumulative pseudoregret of all agents within the clique C . It is clear that since there is no communication between agents in different cliques, their behavior is independent

Algorithm 3 DECENTRALIZED FEDUCB($D, M, T, \rho_{\min}, \rho_{\max}, \mathcal{G}, \gamma$)

```

1: Initialization:  $\forall i, \forall g \in [\gamma]$ , set  $\mathbf{S}_{i,1}^{(g)} \leftarrow \mathbf{0}, \mathbf{s}_{i,1}^{(g)} \leftarrow \mathbf{0}, \mathbf{H}_{i,0}^{(g)} \leftarrow \mathbf{0}, \mathbf{h}_{i,0}^{(g)} \leftarrow \mathbf{0}, \mathbf{U}_{i,1}^{(g)} \leftarrow \mathbf{0}, \bar{\mathbf{u}}_{i,1}^{(g)} \leftarrow \mathbf{0}$ .
2: for each iteration  $t \in [T]$  do
3:   for each agent  $i \in [M]$  do
4:     Set subsampling index  $g \leftarrow t \bmod \gamma$ .
5:     Run lines 4-13 of Algorithm 1 with  $\mathbf{S}_{i,t}^{(g)}, \mathbf{s}_{i,t}^{(g)}, \mathbf{U}_{i,t}^{(g)}, \bar{\mathbf{u}}_{i,t}^{(g)}, \mathbf{V}_{i,t}^{(g)}, \tilde{\mathbf{u}}_{i,t}^{(g)}, \mathbf{H}_{i,t}^{(g)}, \mathbf{h}_{i,t}^{(g)}$ 
6:     if  $\log \left( \frac{\det(\mathbf{V}_{i,t}^{(g)} + \mathbf{x}_{i,t}^{(g)}(\mathbf{x}_{i,t}^{(g)})^\top + M(\rho_{\max} - \rho_{\min})\mathbf{I})}{\det(\mathbf{S}_{i,t}^{(g)})} \right) \geq \frac{D}{(\Delta t_{i,g} + 1)(1 + L^2)}$  then
7:       REQUEST TO SYNCHRONIZE( $i, g$ )  $\leftarrow$  TRUE.
8:     end if
9:     if REQUEST TO SYNCHRONIZE( $i, g$ ) then
10:      BROADCAST Message SYNCHRONIZE( $i, g, t$ ) from agent  $i$  at time  $t$  to all neighbors.
11:    end if
12:    for message  $m$  received at time  $t$  by agent  $i$  do
13:      if  $m = \text{SYNCHRONIZE}(i', g', t')$  then
14:        if  $i'$  belongs to the same clique as  $i$  and  $t' \geq t - \gamma$  then
15:          Agent sends  $\hat{\mathbf{Q}}_{i,t}^{(g')} \rightarrow \text{PRIVATIZER}$  and gets  $\hat{\mathbf{U}}_{i,t+1}^{(g')}, \hat{\mathbf{u}}_{i,t+1}^{(g')} \leftarrow \text{PRIVATIZER}$ .
16:          Agent broadcasts  $\hat{\mathbf{U}}_{i,t+1}^{(g')}, \hat{\mathbf{u}}_{i,t+1}^{(g')}$  to all neighbors.
17:        end if
18:        if  $m = \hat{\mathbf{U}}_{i',t'+1}^{(g')}, \hat{\mathbf{u}}_{i',t'+1}^{(g')}$  then
19:          if  $i'$  belongs to the same clique as  $i$  and  $t' \geq t - \gamma$  then
20:            Agent updates  $\mathbf{S}_{i,t+1}^{(g')}, \mathbf{s}_{i,t+1}^{(g')}$  with  $\hat{\mathbf{U}}_{i',t'+1}^{(g')}, \hat{\mathbf{u}}_{i',t'+1}^{(g')}$ .
21:          end if
22:        end if
23:      end if
24:    end for
25:  end for
26: end for

```

and we can analyse each clique separately. Now, consider κ sequences given by s_1, \dots, s_κ , where $s_i = (i, i + \kappa, i + 2\kappa, \dots, i + (\lceil T/\kappa \rceil - 1)\kappa)$. For any clique C we can furthermore decompose the regret as follows.

$$\begin{aligned}
\text{Regret}_C(T) &= \sum_{i \in C} \sum_{t=1}^T r_{i,t} \\
&= \sum_{i \in C} \sum_{j=1}^{\kappa} \sum_{t \in s_j} r_{i,t} \\
&= \sum_{j=1}^{\gamma} \text{Regret}_{C,j}(T).
\end{aligned}$$

Here $\text{Regret}_{C,j}(T)$ denotes the cumulative pseudoregret of the j^{th} subsequence. We will now bound each of these regret terms individually, with an identical argument as the centralized case. This can be done since the behavior of the algorithm in each of these subsequences is independent: each sequence s_j corresponds to a different Gram matrix and least-squares estimate, and is equivalent to each agent running γ parallel bandit algorithms. We now bound each of the regret terms $\text{Regret}_{C,j}(T)$ via an identical argument as the centralized case. Let $\bar{\sigma}_j$ be the subsequence of $[T]$ containing every j^{th} index, i.e., $\bar{\sigma}_j = j, j + \gamma, j + 2\gamma, \dots, j + \lceil T/\gamma \rceil \gamma$. For any clique C , and index j we compare the pulls of each agent within C to the pulls taken by an agent pulling arms in a round-robin manner $(\mathbf{x}_{i,j})_{i \in C, j \in \bar{\sigma}_j}$. This corresponds to a total of $|C|T/\gamma$ pulls.

Now, it is crucial to see that according to Algorithm 3, if a signal to synchronize has been sent by any agent $i \in C$ at any time t (belonging to subsequence j), then the j^{th} parameter set $\mathbf{V}_{i,t}^{(j)}$ is used first at time t (by each agent in C), then at time $t + \gamma$ (at which time each agent in C broadcasts their parameters, since by this time each other agent will have received the signal to synchronize, as they are at most distance γ apart), and next at time $t + 2\gamma$, upon which they will be fully synchronized.

Now, if we denote the rounds $n_{C,j} \subset \bar{\sigma}_j$ as the rounds in which each agent broadcasted their j^{th} parameter sets, then for any $\tau \in n_{C,j}$, all agents in C have identical j^{th} parameter sets at instant $\tau + 1$, and for each instant $\tau - 1$, all agents in C will obey the synchronization threshold for the j^{th} set of parameters, (log-det condition).

Now, we denote by $\mathbf{W}_{i,t}^{(j,C)}$ the Gram matrix obtained by the hypothetical round-robin agent for subsequence j in clique C . By Lemma 2, we have that

$$\sum_{t \in \bar{\sigma}_j} \sum_{i \in C} \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}}^2 \leq 2d \log \left(1 + \frac{TL^2}{\gamma d \rho_{\min}} \right). \quad (11)$$

After any round T_k of synchronization, consider the cumulative Gram matrices of all observations obtained until that round as $\mathbf{V}_k^{(j,C)}$, $k = 1, \dots, n_{j,C} - 1$, regularized by $|C|\rho_{\min}\mathbf{I}$, i.e., $\mathbf{V}_k^{(j,C)} = \sum_{i \in C} \sum_{t \in \bar{\sigma}_j; t < T_k} \mathbf{x}_{i,t} \mathbf{x}_{i,t}^\top + |C|\rho_{\min}\mathbf{I}$. Finally, let \mathbf{V}_p denote the (regularized) j^{th} Gram matrix with all trials from agents within C at time T , and $\mathbf{V}_0^{(j,C)} = |C|\rho_{\min}\mathbf{I}$. Therefore, we have that $\det(\mathbf{V}_0^{(j,C)}) = (|C|\rho_{\min})^d$, and that $\det(\mathbf{V}_{n_{j,C}}) \leq \left(\frac{\text{tr}(\mathbf{V}_{n_{j,C}}^{(j,C)})}{d} \right)^d \leq (|C|\rho_{\max} + |C|TL^2/(\gamma d))^d$. Therefore, for any $\nu > 1$,

$$\log_\nu \left(\frac{\det(\mathbf{V}_{n_{j,C}}^{(j,C)})}{\det(\mathbf{V}_0^{(j,C)})} \right) \leq d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{\gamma d \rho_{\min}} \right).$$

Let $R = \left\lceil d \log_\nu \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{\gamma d \rho_{\min}} \right) \right\rceil$. It follows that in all but R periods between synchronization,

$$1 \leq \frac{\det(\mathbf{V}_k^{(j,C)})}{\det(\mathbf{V}_{k-1}^{(j,C)})} \leq \nu. \quad (12)$$

We consider the event E to be the period k when Equation 12 holds. Now, for any $T_{k-1} \leq t \leq T_k$, consider the immediate pseudogregret for any agent i . By Lemma 1, we have for any agent $i \in C$ and $t \in \bar{\sigma}_j$:

$$\begin{aligned} r_{i,t} &\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{V}_{i,t}^{(j,C)})^{-1}} \\ &\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{G}_{i,t}^{(j,C)} + |C|\rho_{\min}\mathbf{I})^{-1}} && (\mathbf{V}_{i,t}^{(j,C)} \succcurlyeq \mathbf{G}_{i,t}^{(j,C)} + |C|\rho_{\min}\mathbf{I}) \\ &\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} \cdot \sqrt{\frac{\det(\mathbf{W}_{i,t}^{(j,C)})}{\det(\mathbf{G}_{i,t}^{(j,C)} + |C|\rho_{\min}\mathbf{I})}} \\ &\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} \cdot \sqrt{\frac{\det(\mathbf{V}_k^{(j,C)})}{\det(\mathbf{G}_{i,t}^{(j,C)} + |C|\rho_{\min}\mathbf{I})}} && (\mathbf{V}_k^{(j,C)} \succcurlyeq \mathbf{W}_{i,t}^{(j,C)}) \\ &\leq 2\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} \cdot \sqrt{\frac{\det(\mathbf{V}_k^{(j,C)})}{\det(\mathbf{V}_{k-1}^{(j,C)})}} && (\mathbf{G}_{i,t}^{(j,C)} + |C|\rho_{\min}\mathbf{I} \succcurlyeq \mathbf{V}_{k-1}^{(j,C)}) \\ &\leq 2\nu\bar{\beta}_T \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} && (\text{Event } E \text{ holds}) \end{aligned}$$

Now, we can sum up the immediate pseudoregret over all such periods where E holds to obtain the total regret for these periods. With probability at least $1 - \alpha$,

$$\begin{aligned}
\text{Regret}_{C,j}(T, E) &= \sum_{i \in C} \sum_{t \in \bar{\sigma}_j: E \text{ is true}} r_{i,t} \\
&\leq \sqrt{\frac{|C|T}{\gamma} \left(\sum_{i \in C} \sum_{t \in \bar{\sigma}_j: E \text{ is true}} r_{i,t}^2 \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{\frac{|C|T}{\gamma} \left(\sum_{i \in C} \sum_{t \in \bar{\sigma}_j: E \text{ is true}} \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{\frac{|C|T}{\gamma} \left(\sum_{i \in C} \sum_{t \in \bar{\sigma}_j} \|\mathbf{x}_{i,t}\|_{(\mathbf{W}_{i,t}^{(j,C)})^{-1}} \right)} \\
&\leq 2\nu\bar{\beta}_T \sqrt{2 \frac{|C|T}{\gamma} d \log_{\nu} \left(1 + \frac{TL^2}{\gamma d \rho_{\min}} \right)}.
\end{aligned}$$

Now let us consider the periods in which E does not hold for any subsample j . In any such period between synchronization of length $t_k = T_k - T_{k-1}$, we have, for any agent i , the regret accumulated given by:

$$\begin{aligned}
\text{Regret}_{C,j}([T_{k-1}, T_k]) &= \sum_{t=T_{k-1}}^{T_k} \sum_{i \in C} r_{i,t} \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i \in C} \sqrt{t_k \sum_{t=T_{k-1}}^{T_k} \|\mathbf{x}_{i,t}\|_{(\mathbf{V}_{i,t}^{(j)})^{-1}}^2} \right) \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i \in C} \sqrt{t_k \log_{\nu} \left(\frac{\det(\mathbf{V}_{i,t+k}^j)}{\det(\mathbf{V}_{i,t}^j)} \right)} \right) \\
&\leq 2\nu\bar{\beta}_T \left(\sum_{i \in C} \sqrt{t_k \log_{\nu} \left(\frac{\det(\mathbf{G}_{i,t+k}^j + |C|\rho_{\max}\mathbf{I})}{\det(\mathbf{G}_{i,t}^j + |C|\rho_{\min}\mathbf{I})} \right)} \right)
\end{aligned}$$

By Algorithm 1, we know that for all agents, $t_k \log_{\nu} \left(\frac{\det(\mathbf{G}_{i,t+k}^j + M\rho_{\max}\mathbf{I})}{\det(\mathbf{G}_{i,t}^j + M\rho_{\min}\mathbf{I})} \right) \leq D$ (since there would be a synchronization round otherwise), therefore

$$\leq 2\nu\bar{\beta}_T |C| \sqrt{D}.$$

Now, note that of the total p periods between synchronizations, only at most R periods will not have event E be true. Therefore, the total regret over all these periods can be bound as,

$$\begin{aligned}
\text{Regret}_{C,j}(T, \bar{E}) &\leq R \cdot 2\nu\bar{\beta}_T |C| \sqrt{D} \\
&\leq 2\nu\bar{\beta}_T |C| \sqrt{D} \left(d \log_{\nu} \left(\frac{\rho_{\max}}{\rho_{\min}} + \frac{TL^2}{d\rho_{\min}} \right) + 1 \right).
\end{aligned}$$

In a manner identical to the centralized case, we can obtain the total pseudoregret within a clique C for subsampling index j by choosing an appropriate value of D . Note that since the broadcast between agents happens at an additional delay of 1 trial, the value D must be scaled by a factor of $1 + \sup_{\mathbf{x} \in \mathcal{D}_{i,t}} \|\mathbf{x}\|_2^2$ to bound the additional term in the determinant (by the matrix-determinant lemma), giving us the extra $1 + L^2$ term from the proof. Finally, summing up the above over all $C \in \mathcal{C}$ and $j \in [\gamma]$ and noting that $|C| \leq M\forall C \in \mathcal{C}$, and that $|C| = \bar{\chi}(\mathcal{G}_{\gamma})$ gives us the final result. \square




Figure 2: An experimental comparison of centralized FEDUCB on varying the minimum gap between arms Δ , for various values of the privacy budget ρ_{\min} .

Proposition 7. Fix $\alpha > 0$. If each agent i samples noise parameters $\mathbf{H}_{i,t}$ and $\mathbf{h}_{i,t}$ using the tree-based Gaussian mechanism mentioned above for all n trials of $\bar{\sigma}$ in which communication occurs, then the following ρ_{\min}, ρ_{\max} and κ are $(\alpha/2nM, \bar{\sigma})$ -accurate bounds:

$$\rho_{\min} = \Lambda, \rho_{\max} = 3\Lambda, \kappa \leq \sqrt{m(L^2 + 1) \left(\sqrt{d} + 2 \log(2nM/\alpha) \right) / (\sqrt{2}\epsilon)}.$$

Proof. The proof follows directly from Shariff and Sheffet [41], Proposition 11. Instead of $\alpha/2n$, we replace the parameter as $\alpha/2nM$. \square

Remark 6 (Decentralized Protocol). *Decentralized FEDUCB obtains similar bounds for ρ_{\min}, ρ_{\max} and κ , with $m = 1 + \lceil \log(T/\gamma) \rceil$. An additional term of $\log(\gamma)$ appears in Λ and κ , since we need to now maintain γ partial sums with at most T/γ elements in the worst-case (see Appendix). Unsurprisingly, there is no dependence on the network \mathcal{G} , as privatization is done at the source itself.*

Discussion. In the decentralized version of the algorithm, each agent maintains γ sets of parameters that are used in a round-robin manner. Note that this implies that each parameter set is used at most T/γ times per agent. This implies that in the worst case, each agent will communicate at most T/γ partial sums related to this parameter set, hence needing at most $1 + \lceil \log(T/\gamma) \rceil$ separate nodes of the tree-based mechanism for the particular set of parameters, which leads to the additional $\log \gamma$ in the bounds as well. Next, when determining κ , each agent will not require M -accurate bounds, since it communicates with only $|C|$ other agents. However, in the worst case, a centrally-positioned node belongs to a small clique but can still communicate with all other $M - 1$ nodes (i.e., they can still obtain the partial sums broadcasted), and hence we maintain the factor M to ensure privacy.

Additional Experiments

In addition to the experiments in the main paper, we conduct ablations with variations in Δ . Figure 2 summarizes the results when run on $M = 10$ agents for different privacy budgets and arm gaps. As expected, the overall regret decreases as the gap increases, and the algorithm becomes less sensitive to privacy budget altogether.

References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [2] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1, 2012.
- [3] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson sampling. In *Artificial Intelligence and Statistics*, pages 99–107, 2013.
- [4] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [5] Swapna Buccapatnam, Atilla Eryilmaz, and Ness B Shroff. Stochastic bandits with side observations on networks. In *The 2014 ACM international conference on Measurement and modeling of computer systems*, pages 289–300, 2014.
- [6] Joseph A Calandrino, Ann Kilzer, Arvind Narayanan, Edward W Felten, and Vitaly Shmatikov. "you might also like:" privacy risks of collaborative filtering. In *2011 IEEE symposium on security and privacy*, pages 231–246. IEEE, 2011.
- [7] Igor Calzada and Esteve Almirall. Data ecosystems for protecting european citizens' digital rights. *Transforming Government: People, Process and Policy*, 2020.
- [8] Nicolo Cesa-Bianchi, Claudio Gentile, and Giovanni Zappella. A gang of bandits. In *Advances in Neural Information Processing Systems*, pages 737–745, 2013.
- [9] T-H Hubert Chan, Elaine Shi, and Dawn Song. Privacy-preserving stream aggregation with fault tolerance. In *International Conference on Financial Cryptography and Data Security*, pages 200–214. Springer, 2012.
- [10] TH Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. In *International Colloquium on Automata, Languages, and Programming*, pages 405–417. Springer, 2010.
- [11] Kaize Ding, Jundong Li, and Huan Liu. Interactive anomaly detection on attributed networks. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 357–365, 2019.
- [12] Abhimanyu Dubey and Alex Pentland. Cooperative multi-agent bandits with heavy tails. *arXiv preprint arXiv:2008.06244*, 2020.
- [13] Abhimanyu Dubey and Alex Pentland. Kernel methods for cooperative multi-agent contextual bandits. *arXiv preprint arXiv:2008.06220*, 2020.
- [14] Abhimanyu Dubey and Alex Pentland. Private and byzantine-proof cooperative decision-making. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 357–365, 2020.
- [15] Audrey Durand, Charis Achilleos, Demetris Iacovides, Katerina Strati, Georgios D Mitsis, and Joelle Pineau. Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis. In *Machine Learning for Healthcare Conference*, pages 67–82, 2018.
- [16] Cynthia Dwork. Differential privacy. *Encyclopedia of Cryptography and Security*, pages 338–340, 2011.
- [17] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724, 2010.
- [18] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [19] Pierre Fraigniaud. Locality in distributed graph algorithms, 2016.
- [20] Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.
- [21] Awni Hannun, Brian Knott, Shubho Sengupta, and Laurens van der Maaten. Privacy-preserving contextual bandits. *arXiv preprint arXiv:1910.05299*, 2019.

- [22] Thomas Hardjono and Alex Pentland. Data cooperatives: Towards a foundation for decentralized personal data management. *arXiv preprint arXiv:1905.08819*, 2019.
- [23] Moritz Hardt and Guy N Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 61–70. IEEE, 2010.
- [24] Xiaoguang Huo and Feng Fu. Risk-aware multi-armed bandit problem with application to portfolio selection. *Royal Society open science*, 4(11):171377, 2017.
- [25] Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 299–316. IEEE, 2019.
- [26] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1, 2012.
- [27] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- [28] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [29] Nathan Korda, Balázs Szörényi, and Li Shuai. Distributed clustering of linear bandits in peer to peer networks. In *Journal of machine learning research workshop and conference proceedings*, volume 48, pages 1301–1309. International Machine Learning Society, 2016.
- [30] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. Distributed cooperative decision-making in multiarmed bandits: Frequentist and bayesian algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 167–172. IEEE, 2016.
- [31] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. On distributed cooperative decision-making in multiarmed bandits. In *2016 European Control Conference (ECC)*, pages 243–248. IEEE, 2016.
- [32] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. Social imitation in cooperative multiarmed bandits: Partition-based algorithms with strictly local information. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 5239–5244. IEEE, 2018.
- [33] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- [34] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on computing*, 21(1):193–201, 1992.
- [35] Keqin Liu and Qing Zhao. Decentralized multi-armed bandit with multiple distributed players. In *2010 Information Theory and Applications Workshop (ITA)*, pages 1–10. IEEE, 2010.
- [36] Keqin Liu and Qing Zhao. Distributed learning in cognitive radio networks: Multi-armed bandit with distributed multiple players. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3010–3013. IEEE, 2010.
- [37] Keqin Liu and Qing Zhao. Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681, 2010.
- [38] Mohammad Malekzadeh, Dimitrios Athanasakis, Hamed Haddadi, and Ben Livshits. Privacy-preserving bandits. *arXiv preprint arXiv:1909.04421*, 2019.
- [39] David Martínez-Rubio, Varun Kanade, and Patrick Rebeschini. Decentralized cooperative stochastic multi-armed bandits. In *Advances in Neural Information Processing Systems*, 2019.
- [40] Nikita Mishra and Abhradeep Thakurta. (nearly) optimal differentially private stochastic multi-arm bandits. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 592–601. AUAI Press, 2015.
- [41] Roshan Shariff and Or Sheffet. Differentially private contextual linear bandits. In *Advances in Neural Information Processing Systems*, pages 4296–4306, 2018.

- [42] Jukka Suomela. Survey of local algorithms. *ACM Computing Surveys (CSUR)*, 45(2):24, 2013.
- [43] Abhradeep Guha Thakurta and Adam Smith. (nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems*, pages 2733–2741, 2013.
- [44] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [45] Aristide CY Tossou and Christos Dimitrakakis. Differentially private, multi-agent multi-armed bandits. In *European Workshop on Reinforcement Learning (EWRL)*, 2015.
- [46] Aristide CY Tossou and Christos Dimitrakakis. Algorithms for differentially private multi-armed bandits. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [47] Michal Valko, Nathaniel Korda, Rémi Munos, Ilias Flaounas, and Nelo Cristianini. Finite-time analysis of kernelised contextual bandits. *arXiv preprint arXiv:1309.6869*, 2013.
- [48] Yuanhao Wang, Jiachen Hu, Xiaoyu Chen, and Liwei Wang. Distributed bandit learning: How much communication is needed to achieve (near) optimal regret. *arXiv preprint arXiv:1904.06309*, 2019.
- [49] Yuanhao Wang, Jiachen Hu, Xiaoyu Chen, and Liwei Wang. Distributed bandit learning: Near-optimal regret with efficient communication. In *International Conference on Learning Representations*, 2020.
- [50] Marcelo J Weinberger and Erik Ordentlich. On delayed prediction of individual sequences. *IEEE Transactions on Information Theory*, 48(7):1959–1976, 2002.
- [51] Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, and Kwok-Yan Lam. Local differential privacy and its applications: A comprehensive survey. *arXiv preprint arXiv:2008.03686*, 2020.
- [52] Chunqiu Zeng, Qing Wang, Shekoofeh Mokhtari, and Tao Li. Online context-aware recommendation with time varying multi-armed bandit. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 2025–2034, 2016.