

Private and Byzantine-Proof Cooperative Decision-Making

Abhimanyu Dubey

Massachusetts Institute of Technology
dubeya@mit.edu

Alex Pentland

Massachusetts Institute of Technology
pentland@mit.edu

ABSTRACT

The cooperative bandit problem is a multi-agent decision problem involving a group of agents that interact simultaneously with a multi-armed bandit, while communicating over a network with delays. The central idea in this problem is to design algorithms that can efficiently leverage communication to obtain improvements over acting in isolation. In this paper, we investigate the stochastic bandit problem under two settings - (a) when the agents wish to make their communication private with respect to the action sequence, and (b) when the agents can be byzantine, i.e., they provide (stochastically) incorrect information. For both these problem settings, we provide upper-confidence bound algorithms that obtain optimal regret while being (a) differentially-private and (b) tolerant to byzantine agents. Our decentralized algorithms require no information about the network of connectivity between agents, making them scalable to large dynamic systems. We test our algorithms on a competitive benchmark of random graphs and demonstrate their superior performance with respect to existing robust algorithms. We hope that our work serves as an important step towards creating distributed decision-making systems that maintain privacy.

KEYWORDS

multi-armed bandits; sequential decision-making; differential privacy; robustness

ACM Reference Format:

Abhimanyu Dubey and Alex Pentland. 2020. Private and Byzantine-Proof Cooperative Decision-Making. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

Cooperative decision-making in multiagent systems has been a prominent area of scientific inquiry, with interest from a variety of disciplines from robotics and control systems to macroeconomics and game theory. Applications of cooperative decision-making range from robotics [9, 21] to sensor networks [8, 20].

However, as the availability of large-scale dense data (often collected through a network of sources) increases, the problem of cooperative learning among multiple agents becomes increasingly relevant, moving beyond systems that are “effectively” single-agent to much larger, real-time systems that are decentralized and have agents operating independently. Examples of such applications include seamless personalization across multiple devices in IoT networks [19, 33, 43], that constantly share data between agents.

In such a framework, it is imperative to ensure that information is securely and robustly shared between agents, and individual concerns around privacy are not violated in the pursuit of performance on inference problems. Most generally, the learning setting assumed in such problems and applications is sequential decision-making, where observations are provided in a sequence, and each agent takes actions with the knowledge of this stream of information.

The multi-armed bandit [38] is a classic framework to investigate sequential decision-making in. It provides an elegant mathematical framework to make precise statements about the *exploration-exploitation* dilemma that is central to sequential decision-making, and has thus become the environment of choice across a wide variety of application domains including, most prominently, online advertising [17]. An extension is the *cooperative* multiagent bandit, where a group of agents collectively interact with the same decision problem, and must *cooperate* to obtain optimal performance.

However, in a decentralized environment, cooperative decision-making comes with several caveats. Since it is impossible to control any individual agents’ behavior from a centralized server, enforcing important constraints such as private computation with respect to observations becomes a difficult technical challenge. Additionally, if any agent is malicious (i.e. provides incorrect or adversarial communication), achieving optimal performance becomes non-trivial.

In this paper, we consider two problems in the context of cooperative multiagent decision-making. In the multi-armed (context-free) bandit environment, we investigate (a) privacy in inter-agent communication and (b) tolerance to *byzantine* or malicious agents that deliberately provide stochastically corrupt communication. Our contributions can be listed as follows.

1. We provide a multi-agent UCB-like algorithm PRIVATE MULTIAGENT UCB for the cooperative bandit problem, that guarantees communication of the reward sequence of any agent to be private to any other receiving agent, and provides optimal group regret. Our algorithm is completely decentralized, i.e., its performance or operation does not require any knowledge of the agents’ communication network, and each agent chooses actions autonomously (i.e., without a central server). It also maintains privacy of the source agents’ messages regardless of the individual behavior of any receiving agent (i.e., it is robust to defection in individual behavior).

2. We provide a multi-agent UCB-like algorithm, dubbed BYZANTINE-PROOF MULTIAGENT UCB for the cooperative bandit problem with *byzantine* agents, that corrupt their messages following Huber’s ϵ -contamination model, and provide optimal (in terms of communication overhead) group regret bounds for its performance. Like our private algorithm, this too is completely decentralized, and operates without any knowledge of the communication network.

3. We validate the theoretical bounds on the group regret of our algorithm on a family of random graphs under a variety of different parameter settings, and demonstrate their superiority over existing single-agent algorithms and robustness to parameter changes.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

The paper is organized as follows. First, we provide an overview of the background and preliminaries essential to our problem setup. Next, we examine the private setting, followed by the byzantine setting. We then provide results on our experimental setup and survey related work in this problem domain before closing remarks.

2 BACKGROUND AND PRELIMINARIES

The Cooperative Stochastic Bandit. The multi-armed stochastic bandit problem is an online learning problem that proceeds in rounds. At every round $t \in \{1, 2, \dots, T\}$ (denoted in shorthand as $[T]$), the learner selects an arm $A_t \in [K]$ and the bandit draws a random reward X_t from a corresponding reward distribution ν_{A_t} with mean μ_{A_t} . We assume the rewards are drawn from probability densities with bounded support in $[0, 1]^1$, a typical assumption in the bandit literature. The objective is to minimize the *Regret* $R(T)$.

$$R(T) = T \cdot \max_{k \in [K]} \mu_k - \sum_{t \in [T]} \mu_{A_t}. \quad (1)$$

The cooperative multi-armed bandit problem is a distributed extension of the stochastic bandit problem. In this setting, a group of M agents that communicate over a (connected, undirected) network $\mathcal{G} = (V, E)$ face an identical stochastic K -armed bandit, and must cooperate to collectively minimize their group regret $R_{\mathcal{G}}(T)$.

$$R_{\mathcal{G}}(T) = TM \max_{k \in [K]} \mu_k - \sum_{m \in [M]} \sum_{t \in [T]} \mu_{A_{m,t}}. \quad (2)$$

Here, we denote the action taken by learner m at time T as $A_{m,t}$ and the corresponding reward received as $X_{m,t}$. $d(m, m')$ denotes the length of the shortest path between agents m and m' . For any T , $n_k(T)$ denotes the total number of times any arm k is pulled across all agents, and $n_k^m(T)$ denotes the times agent m has pulled the arm. In the cooperative bandit setting, we denote the policy of each agent $m \in [M]$ by $(\pi_{m,t})_{t \in [T]}$. The collective policy for all agents is denoted by $\Pi = (\pi_{m,t})_{m \in [M], t \in [T]}$. Since \mathcal{G} is not assumed to be complete (i.e. each agent can communicate with every other agent), it introduces a delay in how information flows between agents. However, to provide an understanding of the limits of cooperative decision-making, consider the scenario when \mathcal{G} is complete. In this setting, after every trial of the problem, agents can instantaneously communicate all required information, and therefore the problem can effectively be thought of as a centralized agent pulling M arms at every trial. Anantharam *et al.*[2] provide a lower bound on the group regret achievable in this setting, given as follows.

THEOREM 2.1 (LOWER BOUND ON CENTRALIZED ESTIMATION[2]). *For the stochastic bandit with K arms, each with mean μ_k , let $\Delta_k = \max_{k \in [K]} \mu_k - \mu_k$ and $D(\mu_*, \mu_k)$ denote the Kullback-Liebler divergence between the optimal arm $k^* = \arg \max_{k \in [K]} \mu_k$ and arm k . Then, for any algorithm that pulls M arms at every trial $t \in [T]$, the total regret incurred satisfies*

$$\liminf_{T \rightarrow \infty} R(T) \geq \left(\sum_{k: \Delta_k > 0} \frac{\Delta_k}{D(\mu_*, \mu_k)} \right) \ln T.$$

¹This can be extended to any bounded support by renormalization, and we choose this setup for notational convenience.

This implies that the per-agent regret in this case is $O(M^{-1} \ln T)$. While the above result holds for centralized estimation, we see that one cannot hope for (a) a dependence on T that grows sublogarithmically, and (b) a per-agent regret with a better dependence on M than $O(M^{-1})$, since delays will only limit the amount of information possessed by any agent. We will demonstrate that this bound can be matched (up to additive constants, in some algorithmic settings) in the two problems we consider, for any connected \mathcal{G} .

Communication Protocol. While the existing work on the cooperative bandit problem assumes communication via a gossip or consensus protocol [23, 24], this protocol requires complete knowledge of the graph \mathcal{G} , which is not feasible in a decentralized setting. In this paper, we therefore adopt a variant of the Local communication protocol that is utilized widely in distributed optimization [31], asynchronous online learning [36] and non-stochastic bandit settings [4, 10]. Under this protocol, at any time t , an agent $m \in [M]$ creates a message $\mathbf{q}_m(t)$ that it broadcasts to all its neighbors in \mathcal{G} . $\mathbf{q}_m(t)$ is a tuple of the following form.

$$\mathbf{q}_m(t) = \langle m, t, f_1(H_m(t)), \dots, f_L(H_m(t)) \rangle. \quad (3)$$

Here $H_m(t)$ denotes the history of action-reward pairs obtained by the agent until time t , i.e. $H_m(t) = \{X_{m,1}, A_{m,1}, \dots, X_{m,t}, A_{m,t}\}$, and f_1, \dots, f_L are real-valued functions for some message length $L > 0$. When a message is created, it forwarded from agent-to-agent γ times before it is discarded. Hence, at any time t , the oldest message being propagated in the network would be one created at time $t - \gamma$. As long as a message has not already been received by an agent previously, the agent will broadcast the messages it receives (including its own) to all its neighbors. Under this protocol, we see that two agents that are a distance γ apart can communicate messages to each other, and they will receive messages from each other at a delay of γ trials. An important concept we will use in the remainder of the paper, based on this protocol is the *power graph* of order γ of \mathcal{G} , denoted as \mathcal{G}_γ , which is a graph that contains an edge (i, j) if there is a path of length at most γ between i and j in \mathcal{G} .

In the single-agent setting, the policy of the learner $\pi = (\pi_t)_{t \in [T]}$ is a sequence of functions π_t that define a probability distribution over the space of actions, conditioned on the history of decision-making of the learner, i.e. $\pi_t = \pi_t(\cdot \in [K] | X_1, A_1, \dots, X_t, A_t)$. In the decentralized multi-agent setting, we have a set of policies $\Pi = (\pi_m)_{m \in [M]}$ for each agent $m \in [M]$, where $\pi_m = (\pi_{m,t})_{t \in [T]}$ defines a probability distribution over the space of actions conditioned on both the history of the agents and the messages it has received from other agents until time t . If for any pair of agents m, m' in \mathcal{G} , $d(m, m')$ denotes the distance between the two agents in \mathcal{G} , the policy for agent m at time t is of the following form.

$$\pi_{m,t} = \pi_{m,t} \left(\cdot \in [K] | H_m(t) \cup_{m' \in \mathcal{N}(m, \mathcal{G}_\gamma), t \in T} (q_{m'}(t - d(m, m'))) \right) \quad (4)$$

Here $\mathcal{N}_\gamma(m)$ denotes the γ -neighborhood of agent m in \mathcal{G} , and use the short notation $I_m(t) = \cup_{m' \in \mathcal{N}_\gamma(m), t \in T} (q_{m'}(t - d(m, m')))$ to denote the set of all messages possessed by the agent m at time t . We will now describe the first class of algorithms that are designed for privacy in the cooperative setting.

3 PRIVATE COOPERATIVE BANDITS

We begin by first providing an overview of ϵ -differential privacy and the techniques we will use in our algorithm.

3.1 Differential Privacy

Differential privacy, proposed by Dwork [13] is a formalisation of the amount of information that is leaked about the *input* of an algorithm to an adversary observing its *output*. Existing work in the single agent [30, 40] and the (centralized) multi-agent [40] bandit setting has assumed the algorithms only to have actions (outputs) that are differentially-private with respect to their rewards (inputs). However, since we are considering a decentralized setting, we assume here that each learner acts in isolation, while communicating with other agents by sharing messages. Its policy, therefore, is a function of its own decision history and the messages it receives from other agents. By making the *messages* differentially-private with respect to the reward sequence an agent observes, we can ensure that any policy using these messages is also private with respect to the rewards obtained by other agents. We use this to define differential privacy in the decentralized bandit context.

Definition 3.1 (ϵ -differentially private message protocol). A message $\mathbf{q}_m(t)$ composed of L functions $(f_i)_{i \in [L]}$ for agent m at time t is ϵ -differentially private with respect to the rewards it obtains if for all histories $H_m(t)$ and $H'_m(t)$ that differ in at most one reward sample, we have for all $i \in [L]$, $S \subseteq \mathbb{R}$:

$$e^{-\epsilon} \leq \frac{\Pr(f_i \in S | H_m(t))}{\Pr(f_i \in S | H'_m(t))} \leq e^\epsilon$$

There are several advantages of having such a mechanism to introduce privacy in the cooperative decision-making process. First, we see that unlike centralized or single-agent bandit algorithms that inject noise as a part of the central algorithm itself, our definition *requires* each message to *individually* be private, regardless of the recipient agents' algorithm. Such a protocol ensures that if any learner decides to reveal information about the reward sequence obtained by its neighbors, it will not be able to do so. In the later sections we will demonstrate that for any agent, guaranteeing the differential privacy of the messages it generates with respect to its own rewards suffices to ensure that our algorithm is differentially private with respect to the rewards of other agents. We now introduce concepts crucial to the development of our algorithm.

The Laplace Mechanism. One of the main techniques to introduce differential privacy for univariate problems is the Laplace mechanism, that operates on the notion of *sensitivity*. For any domain \mathcal{U} and function $f : \mathcal{U}^* \rightarrow \mathbb{R}$, we can define *sensitivity* of f for two datasets $D, D' \in \mathcal{U}^*$ that are neighbors (differ only in one entry) as follows.

$$s(f) = \max_{D, D' \in \mathcal{U}^*} |f(D) - f(D')| \quad (5)$$

The Laplace mechanism operates by adding zero-mean noise to the output of f with a scale that is governed by the level of privacy required and the sensitivity of f , as determined by the following.

LEMMA 3.2 (THEOREM 4 OF DWORK AND ROTH [14]). *For any real function f of the input data D , adding a Laplace noise sample X with scale parameter β ensures that $f(D) + X$ is $(\beta/s(f))$ -differentially private with respect to D , where $s(f)$ is the sensitivity of f .*

The Laplace mechanism is the underlying approach that we utilize in guaranteeing privacy in the multi-agent setting. We will now describe our message-passing protocol, which uses this mechanism to guarantee differential privacy between any pair of agents.

3.2 Private Message-Passing

The general approach to privacy in the (single-agent) bandit setting involves using a binary tree or hybrid mechanism to compute the running sum of rewards for any arm [30, 40]. This mechanism involves maintaining a binary tree of individual rewards, where each node in the tree stores a partial sum. This implies that updating the tree would be logarithmic in the number of elements present, and a similar (logarithmic) bound on the sensitivity can be derived, therefore, for any arm k that has been pulled m times, one requires introducing Laplace noise $\mathcal{L}(\frac{\ln m}{\epsilon})$ to achieve ϵ -differential privacy. Under this mechanism, [40] provide a bound on the regret.

THEOREM 3.3 (REGRET OF DP-UCB ALGORITHM [40]). *The single-agent DP-UCB algorithm, when run for T trials, obtains regret of $O\left(K \ln T \cdot \max\left\{\frac{4\sqrt{8}}{\epsilon} \ln(\ln T), \frac{8}{\Delta_{\min}}\right\}\right)$.*

While this approach is feasible in the single-agent case, in the distributed setting, we will have to maintain M separate binary trees (one for each agent), which would create an overhead of a factor of M in the regret (since the sensitivity is increased by a factor of M). To mitigate this overhead we can alternatively utilize the *interval* mechanism, introduced in [40] to the multi-agent setting. Under this mechanism, the mean of an arm is updated only (approximately) T/ϵ times, which makes it possible to add Laplacian noise that is of a lower scale, greatly improving the regret obtained by the algorithm: [40] demonstrate that this mechanism obtains the optimal *single-agent* regret with an *additive* increase due to differential privacy. We demonstrate that using a message-passing algorithm that is inspired by the interval mechanism, we obtain a group regret that has a much smaller than $O(M)$ dependence on the number of agents, and is additive in terms of the privacy overhead.

For the private setting, in the message-passing protocol described in the earlier sections, consider for any agent m , the following message created at time $t \in [T]$.

$$\mathbf{q}_m(t) = \langle m, t, \gamma, \hat{\mathbf{v}}_m(t), \mathbf{n}_m(t) \rangle \quad (6)$$

Here, $\hat{\mathbf{v}}_m(t) = (\hat{v}_k^m(t))_{k \in [K]}$ is a vector of arm-wise reward means that comprise the empirical mean of rewards with specific Laplace noise added, and $\mathbf{n}_m(t) = (n_k^m(t))_{k \in [K]}$ is a vector of the number of times an arm k has been pulled by the agent m . Since the interval mechanism proceeds by making the updates to the mean estimate infrequent, this strategy is adopted in our message-passing protocol as well, which is described in the Lemma below. We describe the complete message-passing protocol in Algorithm 1.

LEMMA 3.4 (PRIVATE RELEASE OF MEANS [40]). *The interval w update the mean for each arm follows a series such that for the given value of $\epsilon \in (0, 1)$, $v \in (1, 1.5)$, $w_n = W_{n+1} - W_n$ with $W_0 = 0$, and W_{n+1} such that,*

$$W_{n+1} = \inf_{x' \in \mathbb{N}} \left\{ x' \geq W_n + 1 : \sum_{W_{n+1}}^{x'} \frac{1}{\sqrt{i}^v} \geq \frac{1}{\epsilon y'^v} \right\}.$$

Protocol Description: The protocol followed is straightforward: at any time t , the agent updates the sum of its own rewards with the reward obtained from the bandit algorithm for the arm pulled. Then, for a set of intervals w_0, \dots, w_T as defined in Lemma 3.4, the broadcasted version of this personal mean is updated to the latest version, with an additional Laplace noise $\mathcal{L}(n_k^m(t)^{v/2-1})$, for some constant $v \in (1, 1.5)$. This mechanism ensures that the broadcast of the personal mean is differentially-private to the reward sequence observed, for each arm. After this update procedure has concluded for each arm, the agent broadcasts this (and all previously incoming messages from other agents) to its neighbors using the system routine `SENDMESSAGES`, which is assumed to be constant-time. It then updates its own group mean estimates of the mean for each arm ($\hat{\mu}_k^m(t)$) using the incoming messages from all other agents. This is done by using the latest message from each agent in $\mathcal{N}_Y(m)$. It then discards stale messages (with life parameter $l = 0$), and returns the updated group means for all arms for the bandit algorithm to use. We now describe the privacy guarantee for the messages, and then describe the `PRIVATE-MULTI-UCB` algorithm and its privacy and regret guarantees.

LEMMA 3.5 (PRIVACY OF OUTGOING MESSAGES). *For any agent $m \in [M]$, $t \in [T]$, $v \in (1, 1.5)$, each outgoing message $q_m(t)$ is $n_k^m(t)^{-v/2}$ -differentially private w.r.t. the sequence of arm $k \in [K]$.*

PROOF. This follows directly from the fact that the only element that is dependent on the reward sequence of arm k is its noisy mean $\hat{\mu}_k^m(t)$, which is $n_k^m(t)^{-v/2}$ -differentially private since we add Laplace noise of scale $n_k^m(t)^{v/2-1}$ and that the sensitivity of the empirical mean is $n_k^m(t)^{-1}$. \square

3.3 Private Multi-Agent UCB

The private multi-agent UCB algorithm is straightforward. During a trial, any agent $m \in [M]$ constructs the group means for each arm using the reward samples from its own pulls and the pulls obtained from its neighbors in $\mathcal{N}_Y(m)$. This mean is given by the estimate $\hat{\mu}_k^m(t)$ for any arm k , as obtained in Algorithm 1. Using these samples, it then follows the popular UCB strategy, where it constructs an upper confidence-bound for each arm, and pulls the arm with the largest upper confidence bound. The complete algorithm is described in Algorithm 2. First, we describe the privacy guarantee with respect to its neighboring agents in $\mathcal{N}_Y(m)$, and then prove the regret (performance) guarantee on the group regret.

LEMMA 3.6 (PRIVACY GUARANTEE). *After t trials, any agent $m \in [M]$ is (ϵ'_m, δ'_m) -differentially private with respect to the reward sequence of any arm observed by any other agent $m' \in \mathcal{N}_Y(m)$ communicating via Algorithm 1 with privacy parameter $v \in (1, 1.5)$, where, for $\epsilon \in (0, 1]$, $\delta'_m \in (0, 1]$, ϵ'_m satisfies*

$$\epsilon'_m \leq \min \left(\epsilon \frac{(t - d(m, m'))^{1-v/2}}{1 - v/2}, 2\epsilon\zeta(v) + \sqrt{2\epsilon\zeta(v) \ln(1/\delta')} \right).$$

PROOF. We can see that for any arm k , the estimate $\hat{\mu}_k^m(t)$ is composed of the sum of rewards from all other agents in $\mathcal{N}_Y(m)$. However, with respect to the reward sequence of any single agent $m' \in \mathcal{N}_Y(m)$, this term only depends on the differential privacy of the outgoing messages $q_{m'}$ obtained by agent m until time $t -$

Algorithm 1 PRIVATE MESSAGE-PASSING PROTOCOL

```

1: Input: Agent  $m \in [M]$ , Iteration  $t \in [T]$ , series  $W = (w_0, w_1, \dots, w_T)$ , such that  $W(i) = w_i$ , series counter  $i_k^m$  for each  $k \in [K]$ ;  $i_k^m = w_0$  if  $t = 0 \forall k$ , set of existing messages  $Q_m(t-1)$ ,  $Q_m(t) = \phi$  if  $t = 0$ , privacy constant  $v \in (1, 1.5)$ , existing reward sums  $s_{m,k}(t) \forall k \in [K]$  such that  $s_{m,k}(0) = 0 \forall k$ .
2: Obtain  $X_{m,t}, A_{m,t}$  from bandit algorithm.
3: Set  $s_{m,k}(t) = s_{m,k}(t-1) + X_{m,t}$  for  $k = A_{m,t}$ .
4: Set  $n_k^m(t) = n_k^m(t-1) + 1$  for  $k = A_{m,t}$ .
5: for Arm  $k$  in  $[K]$  do
6:   if  $n_k^m(t) = W(i_k^m)$  then
7:     Set  $\hat{v}_m^k(t) = s_{m,k}(t)/n_k^m(t) + \mathcal{L}(n_k^m(t)^{v/2-1})$ .
8:     Set  $i_k^m = i_k^m + 1$ .
9:   else
10:    Set  $\hat{v}_m^k(t) = v_m^k(t-1)$ .
11:   end if
12: end for
13: Set  $q_m(t) = \langle m, t, \hat{v}_m(t), \mathbf{n}_m(t) \rangle$ .
14: Set  $Q_m(t) = Q_m(t-1) \cup \{q_m(t)\}$ .
15: for Each neighbor  $m'$  in  $\mathcal{N}_1(m)$  do
16:   SENDMESSAGES( $m, m', Q_m(t)$ ).
17: end for
18: for Each neighbor  $m'$  in  $\mathcal{N}_1(m)$  do
19:    $Q' = \text{RECEIVEMESSAGES}(m', m)$ 
20:    $Q_m(t) = Q_m(t) \cup Q'$ .
21: end for
22: Set  $N_k^m(t) = n_k^m(t)$ ,  $\hat{\mu}_k^m(t) = s_{m,k}(t) \forall k \in [K]$ .
23: for  $q' = \langle m', t', x'_1, \dots, x'_K, a'_1, \dots, a'_K \rangle \in Q_m(t)$  do
24:   if ISLATESTMESSAGE( $q'$ ) then
25:     for Arm  $k \in [K]$  do
26:       Set  $N_k^m(t) = N_k^m(t) + a'_k$ .
27:       Set  $\hat{\mu}_k^m(t) = \hat{\mu}_k^m(t) + a'_k \cdot x'_k$ .
28:     end for
29:   end if
30: end for
31: for Arm  $k \in [K]$  do
32:    $\hat{\mu}_k^m(t) = \hat{\mu}_k^m(t)/N_k^m(t)$ .
33: end for

```

$d(m, m')$ (since it takes at least $d(m, m')$ trials for a message from m' to reach m). However, since for each arm, a new mean is released only t/f times (where f is the interval in Algorithm 1), a k -fold composition theorem [15] identical to [40] can be applied with $t' = t - d(m, m')$. Using Theorem 3.4 and then Corollary 3.1 of [40] subsequently provides us the result. \square

REMARK 1 (ROBUSTNESS TO STRATEGY). *It is important to note that if all agents $m \in [M]$ follow the protocol in Algorithm 1, then the privacy guarantee is sufficient, regardless of the algorithm any agent individually uses to make decisions. This is true since for any agent, the complete sequence of messages it receives from any other agent is differentially private with respect to the origin's reward sequence for any arm and at any instant of the problem. Hence, if the agent chooses to ignore other agents and follow another arbitrary decision-making algorithm, the privacy guarantee only gets stronger,*

Algorithm 2 PRIVATE MULTI-AGENT UCB

```

1: Input: Agent  $m \in [M]$ , trial  $t \in [T]$ , arms  $k \in [K]$ , mean
    $\hat{\mu}_k^m(t)$  and count  $n_k^m(t)$  estimates for each arm  $k \in [K]$ , from
   Algorithm 1.
2: if  $t \leq K \lceil 1/\epsilon \rceil$  then
3:    $A_{m,t} = t \bmod K$ .
4: else
5:   for Arm  $k \in [K]$  do
6:      $\text{UCB}_k^m(t) = \sqrt{\frac{2 \ln t}{N_k^m(t)}}$ .
7:   end for
8:    $A_{m,t} = \arg \max_{k \in [K]} \left\{ \hat{\mu}_k^m(t) + \text{UCB}_k^{(m)}(t) \right\}$ .
9: end if
10:  $X_{m,t} = \text{PULL}(A_{m,t})$ .
11: return  $A_{m,t}, X_{m,t}$ .

```

never worse (since it cannot gain any additional information from the messages). Therefore, this communication protocol is robust to any agent’s individual decision-making policy.

REMARK 2 (ROBUSTNESS TO DEFECTION). *The communication protocol crucially involves broadcasting rewards with privacy alterations at the source, and not at the destination. Consider an alternate protocol where the agents broadcasted their true personal average rewards (without noise), and the receiving agent added an appropriately scaled Laplacian noise on receiving a message. While this protocol would yield slightly better guarantees (both on privacy and regret, since the noise will be one Laplace sample, and not the sum of several samples in the current case), it is not robust to defection i.e. if the agent decides to learn the individual reward sequence, it can. Since our mechanism adds noise at the source, it maintains robustness under defection.*

THEOREM 3.7. *If all agents $m \in [M]$ each follow the PRIVATE MULTI-AGENT UCB algorithm with the messaging protocol described in Algorithm 1, then the group regret incurred after T trials obeys the following upper bound.*

$$R_{\mathcal{G}}(T) \leq \sum_{k:\Delta_k > 0} \bar{\chi}(\mathcal{G}_\gamma) \left(\frac{8 \ln T}{\Delta_k} \right) + \left(\sum_{k:\Delta_k > 0} \Delta_k \right) \left(\chi(\mathcal{G}_\gamma) (M\gamma + 2) + M \left(\frac{1}{\epsilon} + \zeta(1.5) \right) \right).$$

Here, $\bar{\chi}$ is the clique covering number, and ζ is the Riemann zeta function.

PROOF. (Sketch) The proof proceeds by decomposing \mathcal{G}_γ into its minimum clique partitions, and analysing each clique individually. We then derive a UCB-like decomposition for the group regret of each clique. We first bound the delay in information propagation through this clique, and then, using concentration results for bounded and Laplace random variables, we bound the number of pulls of a suboptimal arm within the clique, and subsequently bound the total group regret. A full proof is presented in the Appendix. \square

REMARK 3 (DEVIATION FROM OPTIMAL RATES). *Upon examining the bound, we see that there are two $O(1)$ terms - the first, $O(\bar{\chi}(\mathcal{G}_\gamma)M\gamma)$ is the loss from delays in communication, since information flows*

between two agents with delays, and path lengths between agents are (on average) larger than 1 (unless \mathcal{G} is connected). When \mathcal{G} is connected, this reduces to a constant $M + 2$, that matches the single-agent UCB case for MT total pulls. The second term, $M \left(\frac{1}{\epsilon} + \zeta(1.5) \right)$, arises from the level of differential privacy required. When $\epsilon = 1$, i.e. no privacy, this term also reduces to a constant of M , matching the single-agent non-private case with MT pulls. Finally, consider the overhead in the logarithmic term, $\chi(\mathcal{G}_\gamma)$. This too, arises from the delays in the network. When \mathcal{G} is complete, this reduces to 1, matching the optimal lower bound.

REMARK 4 (DEPENDENCE ON γ AND ASYMPTOTIC OPTIMALITY). *In addition to the dependencies discussed earlier, a major factor in controlling the asymptotic optimality is via the life parameter γ , that controls “how far” information flows from any single agent. This dependence is present within the bound as well, as we can tune the performance of the algorithm greatly by tuning γ . If we set $\gamma = \text{diameter}(\mathcal{G})$, then $\chi(\mathcal{G}_\gamma) = 1$ (since \mathcal{G}_γ will be fully connected), which implies that the group regret incurred asymptotically ($\lim T \rightarrow \infty$) matches the lower bound. Alternatively, if we were to run the algorithm with $\gamma = 0$, i.e. no communication, then we obtain a regret of $O(M \ln T)$, that is equivalent to the group regret obtained by M agents operating in isolation. Interestingly, there is no dependence of \mathcal{G} and γ on the privacy loss. This is intuitive, since our mechanism introduces suboptimality at the source, and not at the destinations.*

We defer experimental evaluations to the end of the paper. We now describe the algorithm and message-passing protocol for the setup with byzantine agents.

4 BYZANTINE-PROOF BANDITS

In this section, we examine the cooperative multi-agent bandit problem where agents cannot be trusted. In this setting, there exist *byzantine* agents, that, at any trial, instead of reporting their true means, give a random sample from an alternate (but fixed) distribution, with some probability ϵ . Once a message is created, however, we assume that it is received correctly by all subsequent agents, and any corruption occurs only at the source.

One might consider the setting where in addition to messages being corrupted at the source, it is possible for a message to be corrupted by any intermediary byzantine agent with the same probability ϵ . From a technical perspective, this setting is not different than the first setting, since the probability of any incoming message being corrupted becomes at most $\gamma\epsilon$ (by the union bound and that $d(m, m') \leq \gamma$ for any pair of agents m, m' that can communicate), and the remainder of the analysis is identical henceforth. Therefore, we simply consider the first setting. We now formalize the problem and provide some technical tools for this setting.

Problem Setting and Messaging Protocol. For the byzantine-proof setting, in the message-passing protocol described in the earlier sections, consider for any agent m , the following message created at time $t \in [T]$.

$$\mathbf{q}_m(t) = \langle m, t, A_{m,t}, \hat{X}_{m,t} \rangle \quad (7)$$

Where, for byzantine agents, $\hat{X}_{m,t} = X_{m,t}$ with probability $(1 - \epsilon)$, and a random sample from an unknown (but fixed) distribution Q with probability ϵ . For non-byzantine agents, $\hat{X}_{m,t} = X_{m,t}$ with

probability 1. Contrasted to the protocol used in the previous problem setting, there are several key differences. First, instead of communicating reward means from all arms, we simply communicate just the individual rewards obtained along with the action taken. This is essential owing to the nature of our algorithm, as will be made precise later on (cf. Remark 6).

4.1 Robust Estimation

The general problem of estimating statistics of a distribution P from samples of a mixture distribution of the form $(1 - \epsilon)P + \epsilon Q$ for $\epsilon < 1/2$ is a classic contamination model in robust statistics, known as Huber’s ϵ -Contamination model [18]. In the univariate mean estimation setting, a popular approach is to utilize a *trimmed* estimator of the mean that is robust to outlying samples, that works as long as ϵ is small. We utilize the trimmed estimator to design our algorithm as well, defined as follows.

Definition 4.1 (Robust Trimmed Estimator). Consider a set of $2N$ samples $X_1, Y_2, \dots, X_N, Y_N$ of a mixture distribution $(1 - \epsilon)P + \epsilon Q$ for some $\epsilon \in [0, 1/2)$. Let $X_1^* \leq X_2^* \leq \dots \leq X_N^*$ denote a non-decreasing arrangement of X_1, \dots, X_N . For some confidence level $\delta \in (0, 1)$, let $\alpha = \max(\epsilon, \frac{\ln(1/\delta)}{N})$, and let \mathcal{Z} be the shortest interval containing $N \left(1 - 2\alpha - \sqrt{2\alpha \frac{\ln(4/\delta)}{N} - \frac{\ln(4/\delta)}{N}}\right)$ points of $\{Y_1^*, \dots, Y_N^*\}$.

Then, the univariate trimmed estimator $\hat{\mu}_R(\{X_1, Y_2, \dots, X_N, Y_N\}, \delta)$ can be given by

$$\hat{\mu}_R = \frac{1}{\sum_i^N \mathbb{1}\{X_i \in \mathcal{Z}\}} \sum_{i=1}^N X_i \mathbb{1}\{X_i \in \mathcal{Z}\}. \quad (8)$$

The fundamental idea behind the trimmed mean is to use half of the input points to identify quartiles, and then use the remaining half to estimate the “trimmed” mean around those quartiles. A natural question arises regarding the optimality of this estimator and its rate of concentration around the true mean of P . Recent work from [32] provides such a result.

THEOREM 4.2 (CONFIDENCE BOUND FOR THE TRIMMED MEAN [32]). *Let $\delta \in (0, 1)$. Then for any distribution P with mean μ and finite variance σ^2 , we have, with probability at least $1 - \delta$,*

$$|\hat{\mu}_R - \mu| \leq \sigma\sqrt{\epsilon} + \sqrt{\frac{\sigma \ln(1/\delta)}{N}}.$$

Using these results, we can design an algorithm for the byzantine cooperative bandit problem, as described next.

4.2 Byzantine-Proof Multi-Agent UCB

The byzantine-proof algorithm works for an agent m by making a conservative assumption that all agents (including itself) are byzantine, and constructing a robust mean estimator under this assumption. It then uses this estimator and an associated upper confidence bound (UCB) to choose an arm, similar to the single-agent UCB algorithm. The assumption of all agents being byzantine primarily is made to aid in the analysis of the algorithm: from the perspective of any agent, if all M agents are byzantine, then it can expect to obtain approximately $O(TM\epsilon)$ corrupted messages.

Algorithm 3 BYZANTINE-PROOF MULTI-AGENT UCB

```

1: Input: Agent  $m$ , arms  $k \in [K]$ , mean estimator  $\hat{\mu}_R(n, \delta)$ 
2: Set  $S_k^m = \phi \forall k \in [K]$ ,  $Q_m(t) = \phi$ .
3: for For  $t \in [T]$  do
4:   if  $t \leq K$  then
5:      $A_{m,t} = t$ .
6:   else
7:     for Arm  $k \in [K]$  do
8:        $\hat{\mu}_k^{(m)} = \hat{\mu}_R(S_k^m, 1/t^2)$ .
9:        $\text{UCB}_k^{(m)}(t) = \sigma\sqrt{\epsilon} + \sqrt{\frac{\sigma \ln(1/\delta)}{|S_k^m(t)|}}$ .
10:    end for
11:     $A_{m,t} = \arg \max_{k \in [K]} \{\hat{\mu}_k^{(m)}(t) + \text{UCB}_k^{(m)}(t)\}$ .
12:  end if
13:   $X_{m,t} = \text{PULL}(A_{m,t})$ .
14:   $S_{A_{m,t}}^m = S_{A_{m,t}}^m \cup \{X_{m,t}\}$ 
15:   $Q_m(t) = \text{PRUNEDeADMESSAGES}(Q_m(t))$ .
16:   $Q_m(t) = Q_m(t) \cup \{\langle m, t, \gamma, A_{m,t}, X_{m,t} \rangle\}$ .
17:  Set  $l = l - 1 \forall \langle m', t', a', x' \rangle$  in  $Q_m(t)$ .
18:  for Each neighbor  $m'$  in  $\mathcal{N}_1(m)$  do
19:     $\text{SENDMESSAGES}(m, m', Q_m(t))$ .
20:  end for
21:   $Q_m(t+1) = \phi$ .
22:  for Each neighbor  $m'$  in  $\mathcal{N}_1(m)$  do
23:     $Q' = \text{RECEIVEMESSAGES}(m', m)$ 
24:     $Q_m(t+1) = Q_m(t+1) \cup Q'$ .
25:  end for
26:  for  $\langle m', t', a', x' \rangle \in Q_m(t+1)$  do
27:     $S_{a'}^m = S_{a'}^m \cup \{x'\}$ .
28:  end for
29: end for

```

If however, only a fraction $f < 1$ of agents are byzantine, then it can expect to obtain approximately $O(TMf\epsilon)$ corrupted messages, which would imply that (in expectation), all M agents are byzantine with $\epsilon' = f\epsilon$. This information can be incorporated at runtime as well, and hence we proceed with the conservative assumption. Moreover, if we overestimate ϵ during initialization, the performance of the algorithm will remain unchanged, since the robust mean estimator for any value ϵ is also applicable to any other mixture proportion $\epsilon' < \epsilon$ without any degradation in performance.

In this setting, for the our description of the algorithm, we present both the algorithm and the message-passing protocol in the same setup, unlike the previous case since there is no explicit distinction to be made between the decision-making algorithm and the message-passing algorithm, as they both operate together. The complete algorithm for all agents is described in Algorithm 3.

Algorithm Description. We will consider an agent $m \in [M]$. The agent maintains a set of rewards S_k^m for each arm k , which it updates with the reward samples it obtains from messages from other agents, and its own rewards. At every trial, for each arm, it uses the robust mean estimator $\hat{\mu}_R$ to compute the trimmed mean of all the reward samples present in S_k^m , and uses the number of samples to estimate an upper confidence bound as well. It then

chooses the arm with the largest sum of the (robust) upper confidence bound. After choosing an arm, it updates its set S_k^m and sends messages to all its neighbors, similar to the previous algorithm.

THEOREM 4.3. *If all agents $m \in [M]$ run algorithm 3 with mean estimator $\hat{\mu}_R$, then for $\epsilon < \Delta_{\min}/2\sigma$, the group regret after T iterations obeys the following upper bound.*

$$R_{\mathcal{G}}(T) \leq \bar{\chi}(\mathcal{G}_\gamma) \left(\sum_{k:\Delta_k > 0} \frac{4\sigma^2}{(\Delta_k - 2\sigma\sqrt{\epsilon})^2} \right) \ln T + (3M + \gamma\chi(\mathcal{G}_\gamma)(M-1)) \left(\sum_{k:\Delta_k > 0} \Delta_k \right).$$

Here, $\bar{\chi}(\cdot)$ refers to the minimum clique number.

PROOF. (Sketch). We proceed in a manner similar to the previous regret bound, by decomposing the group regret into a sum of clique-wise group regrets. We then construct similar UCB-like “exploration” vs “exploitation” events, and bound the probability of pulling an incorrect arm under these events using the confidence bounds of the robust mean estimator coupled with the worst-case delay within the network. The full proof is in the Appendix. \square

REMARK 5 (DEVIATION FROM OPTIMAL RATES). *When comparing the group regret bound to the optimal rate achievable in the single-agent case, there is an additive constant that arises from the delay in the network. Identical to the differentially private algorithm, this additive constant reduces to the constant corresponding to MT individual trials of the UCB algorithm when information flows instantly throughout the network, i.e. \mathcal{G} is connected. Additionally, we observe identical dependencies on other graph parameters in the leading term as well, except for the modified denominator $(\Delta_k - 2\sigma\sqrt{\epsilon})^2$. This term arises from the inescapable bias that the mixture distribution Q introduces into estimation, and [32] show that this bias is optimal and unimprovable in general. This can intuitively be explained by the fact that after a certain ϵ , there will be so many corrupted samples that it would be impossible to distinguish the best arm from the next best arm, regardless of the number of samples (since the observed, noisy distributions for each arms will become indistinguishable). We must also note that the dependence on γ on the asymptotic optimality remains identical to the previous algorithm, and when all agents are truthful ($\epsilon = 0$) and $\gamma = \text{diameter}(\mathcal{G})$, this algorithm obtains optimal multi-agent performance, matching the lower bound.*

REMARK 6 (SIMULTANEOUS PRIVACY AND ROBUSTNESS). *A natural question to ask is whether both privacy with respect to other agents’ rewards and robustness with respect to byzantine agents can be simultaneously achieved. This is a non-trivial question, since the fundamental approaches to both problems are conflicting. First, we must notice that to achieve differential privacy without damaging regret, one must convey summary statistics (such as the mean or sum of rewards) directly, with appropriately added noise. However, for robust estimation in byzantine-tolerance, communicating individual rewards is essential, to allow for an appropriate rate of concentration of the robust mean. These two approaches seem contradictory at a first inquiry, and hence we leave this line of research for future work.*

In the next section, we describe the experimental comparisons of our proposed algorithms with existing benchmark algorithms.

5 EXPERIMENTS

For our experimental setup, we consider rewards drawn from randomly initialized Bernoulli distributions, with $K = 5$ as our default operating environment. We initialize the connectivity graph \mathcal{G} as a sample from a random Erdos-Renyi [6] family with edge probability $p = 0.1$, and set the communication parameter $\gamma = \text{diameter}(\mathcal{G})/2$. We describe the experiments and associated benchmark algorithms individually.

Testing Private Multiagent UCB. In this setting, we compare the group regret $R_{\mathcal{G}}$ of $M = 200$ agents over 100 randomly initialized trials of the above setup. The benchmark algorithms we compare with are (a) single-agent UCB1 [3] running individually on each agent, and (b) DP-UCB-INT of [40] running individually on each agent, under the exact setup for ϵ and δ as described in [40]. We choose this benchmark as it is the state-of-the-art in the single-agent stochastic bandit case. The results of this experiment are plotted in Figure 1(a). We observe that the performance of our algorithm is significantly better than both private benchmarks, however, it incurs higher regret (as expected) than the non-private version.

Testing Byzantine-Proof Multiagent UCB. To test the performance of our byzantine-proof algorithm, we use the same setup as the previous case, with $M = 200$ agents, repeated over 100 randomly initialized trials, except we set the reward probability of each arm to be constrained within [0.3, 0.7], and set the contamination probability as $\epsilon = 10^{-3}$. The benchmark algorithms in this case are (a) the single-agent UCB1 [3] running individually on each agent, and (b) the byzantine-proof multiagent UCB with $\epsilon = 0$ (since no other work explicitly studies the Huber ϵ -contamination model as used in our setting). The results of this experiment are summarized in Figure 1(b). Similar to the previous case, we observe a large improvement in the group regret compared to the single-agent UCB1, and with no contamination, the performance is better.

Testing the effect of γ . To understand the effect of γ on the obtained regret, we repeated the same experiment ($M = 200$, trials of randomly generated Erdos-Renyi graphs with $p = 0.1$) with our two algorithms and compared their obtained group regret at $T = 1000$ trials. We observe a sharp decline as γ increases from 1 to $\text{diameter}(\mathcal{G})$, and matches the optimal group regret at $\gamma = \text{diameter}(\mathcal{G})$, as hypothesized by our regret bound. The results of this experiment are summarized in Figure 1(c).

6 RELATED WORK

Our work relates to a vast body of work across several fields in the sequential decision-making literature.

Multi-Agent Bandit Algorithms. Multi-agent bandit problems have been studied widely in the context of distributed control [5, 26–28]. In this line of work, the problem is competitive rather than cooperative, where a group of agents compete for a set of arms, and collisions occur when two or more agents select the same arms. Contextual bandit algorithms have also been explored the context of peer-to-peer networks and collaborative filtering over multiple agent recommendations [11, 16, 17, 22], where “agents” are clustered or averaged by their similarity in contexts. While some of these problem domains assume an underlying network structure between agents, none assume “delay” in the observed feedback between agents, and communication is instantaneous with




Figure 1: Experimental comparisons on random graphs. Each figure is constructed by averaging over 100 trials.

the knowledge of the graph, making these algorithms effectively single-agent. For the cooperative bandit problem, however, general settings have been investigated in both the stochastic case using a *consensus* protocol [23–25] and in the non-stochastic case using a message-passing protocol similar to ours [4, 10]. To the best of our knowledge, our work is the first to consider the cooperative bandit problem with differential privacy or byzantine agents.

Private Sequential Decision-Making. Our work uses the immensely influential *differential privacy* framework of Dwork [13], that has been the foundation of almost all work in private sequential decision-making. For the bandit setting, there has been significant research interest in the *single-agent* case, including work on private stochastic bandit algorithms [30, 37, 40], linear bandit algorithms [35] and in the nonstochastic case [41]. In the competitive multi-agent stochastic bandit case, Tossou and Dimitrakakis [39] provide a UCB-based algorithm based on Time-Division Fair Sharing (TDFS).

Robust Sequential Decision-Making. Robust estimation has a rich history in the bandit literature. Robustness to heavy-tailed reward distributions has been extensively explored in the stochastic multi-armed setting, from the initial work of Bubeck *et al.*[7] that proposed a UCB algorithm based on robust mean estimators, to the subsequent work of [12, 42, 44] on both Bayesian and frequentist algorithms for the same. Contamination in bandits has also been explored in the stochastic *single-agent* case, such as the work of [1] in best-armed identification under contamination and algorithms that are jointly optimal for the stochastic and nonstochastic case [34], that uses a modification of the popular EXP3 algorithm. Our work builds on the work in robust statistics of Huber [18] and Lugosi and Mendelson [29].

To the best of our knowledge, our work is the first to consider privacy and contamination in the context of the multi-agent cooperative bandit problem, with the additional benefit of being completely decentralized.

7 CONCLUSION

In this paper, we discussed the cooperative multi-armed stochastic bandit problem under two important practical settings – private communication and the existence of byzantine agents that follow an ϵ -contamination model. We provided two algorithms that are completely decentralized, and provide optimal group regret guarantees when run with certain parameter settings. Our work is the first to investigate real-world scenarios of cooperative decision-making, however, it does leave many open questions for future work.

First, we realise that achieving both differential privacy and byzantine-tolerance simultaneously is non-trivial, even in the stochastic case: differential privacy requires the computation of summary statistics (such as the sum or average of rewards) in order to provide feasible regret, however, for byzantine-tolerance we require the use of robust estimators, that explicitly require individual reward samples to compute, making their combination a difficult problem, which can be investigated in future work. Next, we analysed a very specific setting for byzantine agents, which can be generalized to adversarial corruptions, as done in the single-agent case. Finally, extensions of the stochastic multi-armed setting to contextual settings is also an important direction of research that our work opens up in multi-agent sequential decision-making.

REFERENCES

- [1] Jason Altschuler, Victor-Emmanuel Brunel, and Alan Malek. 2019. Best Arm Identification for Contaminated Bandits. *Journal of Machine Learning Research* 20, 91 (2019), 1–39.
- [2] Venkatesh Anantharam, Pravin Varaiya, and Jean Walrand. 1987. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-Part I: IID rewards. *IEEE Trans. Automat. Control* 32, 11 (1987), 968–976.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [4] Yogev Bar-On and Yishay Mansour. 2019. Individual Regret in Cooperative Nonstochastic Multi-Armed Bandits. *arXiv preprint arXiv:1907.03346* (2019).
- [5] Ilai Bistriz and Amir Leshem. 2018. Distributed multi-player bandits—a game of thrones approach. In *Advances in Neural Information Processing Systems*. 7222–7232.
- [6] Béla Bollobás and Oliver M Riordan. 2003. Mathematical results on scale-free random graphs. *Handbook of graphs and networks: from the genome to the internet* (2003), 1–34.
- [7] Sébastien Bubeck, Nicolo Cesa-Bianchi, and Gábor Lugosi. 2013. Bandits with heavy tail. *IEEE Transactions on Information Theory* 59, 11 (2013), 7711–7717.
- [8] John Byers and Gabriel Nasser. 2000. Utility-based decision-making in wireless sensor networks. In *Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing*. IEEE Press, 143–144.
- [9] Y Uny Cao, Alex S Fukunaga, and Andrew Kahng. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous robots* 4, 1 (1997), 7–27.
- [10] Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2019. Delay and cooperation in nonstochastic bandits. *The Journal of Machine Learning Research* 20, 1 (2019), 613–650.
- [11] Nicolo Cesa-Bianchi, Claudio Gentile, and Giovanni Zappella. 2013. A gang of bandits. In *Advances in Neural Information Processing Systems*. 737–745.
- [12] Abhimanyu Dubey and Alex ‘Sandy’ Pentland. 2019. Thompson Sampling on Symmetric Alpha-Stable Bandits. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 5715–5721. <https://doi.org/10.24963/ijcai.2019/792>
- [13] Cynthia Dwork. 2011. Differential privacy. *Encyclopedia of Cryptography and Security* (2011), 338–340.
- [14] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
- [15] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. 2010. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer*

- Science. IEEE, 51–60.
- [16] Claudio Gentile, Shuai Li, Purushottam Kar, Alexandros Karatzoglou, Giovanni Zappella, and Evans Etrud. 2017. On context-dependent clustering of bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1253–1262.
- [17] Claudio Gentile, Shuai Li, and Giovanni Zappella. 2014. Online clustering of bandits. In *International Conference on Machine Learning*. 757–765.
- [18] Peter J Huber. 2011. *Robust statistics*. Springer.
- [19] Harald R Kisch and Claudia Lage Rebello da Motta. 2017. Real World Examples of Agent based Decision Support Systems for Deep Learning based on Complex Feed Forward Neural Networks. In *COMPLEXIS*. 94–101.
- [20] Lawrence A Klein. 2004. *Sensor and data fusion: a tool for information assessment and decision making*. Vol. 324. SPIE press Bellingham eWA WA.
- [21] Jelle R Kok, Matthijs TJ Spaan, Nikos Vlassis, et al. 2003. Multi-robot decision making using coordination graphs. In *Proceedings of the 11th International Conference on Advanced Robotics, ICAR*, Vol. 3. 1124–1129.
- [22] Nathan Korda, Balázs Szörényi, and Li Shuai. 2016. Distributed clustering of linear bandits in peer to peer networks. In *Journal of machine learning research workshop and conference proceedings*, Vol. 48. International Machine Learning Societ, 1301–1309.
- [23] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. 2016. Distributed cooperative decision-making in multiarmed bandits: Frequentist and Bayesian algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 167–172.
- [24] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. 2016. On distributed cooperative decision-making in multiarmed bandits. In *2016 European Control Conference (ECC)*. IEEE, 243–248.
- [25] Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. 2018. Social Imitation in Cooperative Multiarmed Bandits: Partition-Based Algorithms with Strictly Local Information. In *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 5239–5244.
- [26] Keqin Liu and Qing Zhao. 2010. Decentralized multi-armed bandit with multiple distributed players. In *2010 Information Theory and Applications Workshop (ITA)*. IEEE, 1–10.
- [27] Keqin Liu and Qing Zhao. 2010. Distributed learning in cognitive radio networks: Multi-armed bandit with distributed multiple players. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 3010–3013.
- [28] Keqin Liu and Qing Zhao. 2010. Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing* 58, 11 (2010), 5667–5681.
- [29] Gabor Lugosi and Shahar Mendelson. 2019. Robust multivariate mean estimation: the optimality of trimmed mean. *arXiv preprint arXiv:1907.11391* (2019).
- [30] Nikita Mishra and Abhradeep Thakurta. 2015. (Nearly) optimal differentially private stochastic multi-arm bandits. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 592–601.
- [31] Ciamac Cyrus Moallemi. 2007. *A message-passing paradigm for optimization*. Vol. 68.
- [32] Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. 2019. A unified approach to robust mean estimation. *arXiv preprint arXiv:1907.00927* (2019).
- [33] Rodrigo Roman, Jianying Zhou, and Javier Lopez. 2013. On the features and challenges of security and privacy in distributed internet of things. *Computer Networks* 57, 10 (2013), 2266–2279.
- [34] Yevgeny Seldin and Aleksandrs Slivkins. 2014. One Practical Algorithm for Both Stochastic and Adversarial Bandits. In *ICML*. 1287–1295.
- [35] Roshan Shariff and Or Sheffet. 2018. Differentially private contextual linear bandits. In *Advances in Neural Information Processing Systems*. 4296–4306.
- [36] Jukka Suomela. 2013. Survey of local algorithms. 45, 2 (2013), 24.
- [37] Abhradeep Guha Thakurta and Adam Smith. 2013. (Nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems*. 2733–2741.
- [38] William R Thompson. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3/4 (1933), 285–294.
- [39] Aristide CY Tossou and Christos Dimitrakakis. [n. d.]. Differentially private, multi-agent multi-armed bandits.
- [40] Aristide CY Tossou and Christos Dimitrakakis. 2016. Algorithms for differentially private multi-armed bandits. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [41] Aristide Charles Yedia Tossou and Christos Dimitrakakis. 2017. Achieving privacy in the adversarial multi-armed bandit. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [42] Sattar Vakili, Keqin Liu, and Qing Zhao. 2013. Deterministic sequencing of exploration and exploitation for multi-armed bandit problems. *IEEE Journal of Selected Topics in Signal Processing* 7, 5 (2013), 759–767.
- [43] Lizhe Wang and Rajiv Ranjan. 2015. Processing distributed internet of things data in clouds. *IEEE Cloud Computing* 2, 1 (2015), 76–80.
- [44] Xiaotian Yu, Han Shao, Michael R Lyu, and Irwin King. 2018. Pure exploration of multi-armed bandits with heavy-tailed payoffs. In *Conference on Uncertainty in Artificial Intelligence*. 937–946.