
Cooperative Multi-Agent Bandits with Heavy Tails

Abhimanyu Dubey¹ Alex Pentland¹

Abstract

We study the heavy-tailed stochastic bandit problem in the cooperative multi-agent setting, where a group of agents interact with a common bandit problem, while communicating on a network with delays. Existing algorithms for the stochastic bandit in this setting utilize confidence intervals arising from an averaging-based communication protocol known as *running consensus*, that does not lend itself to robust estimation for heavy-tailed settings. We propose MP-UCB, a decentralized multi-agent algorithm for the cooperative stochastic bandit that incorporates robust estimation with a message-passing protocol. We prove optimal regret bounds for MP-UCB for several problem settings, and also demonstrate its superiority to existing methods. Furthermore, we establish the first lower bounds for the cooperative bandit problem, in addition to providing efficient algorithms for robust bandit estimation of location.

1. Introduction

The multi-agent cooperative bandit is an increasingly relevant decision-making problem in the era of large-scale and distributed inference. In this setting, a group of agents, $v \in V$, are each faced with a decision-making problem (e.g., a multi-armed bandit), while simultaneously communicating with other agents over a network. The motivation for studying this problem stems from the rise in decentralized computing systems, and technical interest arises from designing algorithms to leverage cooperation and accelerate collective decision-making (Landgren et al., 2016a; 2018).

Consider the case when the agents interact with the classical multi-armed stochastic bandit. For each agent $v \in V$, the problem proceeds in a series of synchronous rounds $t = 1, 2, \dots, T$. In each round, each agent selects an action

$A_{v,t} \in \mathcal{A}$, where the space of actions \mathcal{A} is assumed to be finite and countable ($|\mathcal{A}| = K$). It then obtains an i.i.d. reward $x_{v,t}$, with mean $\mu_{A_{v,t}}$, $A_{v,t} \in \mathcal{A}$. The overall objective of the agents is to minimize the *group regret* $R_{\mathcal{G}}(T)$.

$$R_{\mathcal{G}}(T) = \sum_{v \in \mathcal{G}} R_v(T) = \mu_* T |V| - \sum_{v \in V} \sum_{t=1}^T \mu_{A_{v,t}}. \quad (1)$$

Here, $\mu_* = \max_{a \in \mathcal{A}} \mu_a$ denotes the mean of the optimal arm, and $\Delta_a = \mu_* - \mu_a$ is the suboptimality of arm a from the optimal arm $A_* = \arg \max_{a \in \mathcal{A}} \mu_a$. In the cooperative bandit, each agent faces the same problem (i.e., \mathcal{A} is the same for each agent). In addition to the exploration-exploitation tradeoff succinctly captured by the conventional bandit problem, each agent faces the additional complexity of leveraging feedback from its neighbors. This feedback is typically available to the agent after a (variable) *delay*, based on the communication network between the agents. This communication network is conveniently determined by a graph $\mathcal{G} = (V, E)$, where V denotes the set of agents, and E is the set of edges, where edge $(i, j) \in E$ if agents i and j are neighbors. Messages are sent between trials of the bandit problem, and take $d(v, v')$ trials to be communicated from agent v to v' (where $d(v, v')$ is the distance between the agents v and v' in \mathcal{G}).

Research on this problem has exclusively been on reward distributions that are sub-Gaussian (Landgren et al., 2016a;b; Martínez-Rubio et al., 2019). While this is certainly applicable in several domains, increasing evidence suggests that assumptions of sub-Gaussianity may not hold for numerous applications specific to *distributed* decision-making, in problems such as distributed load estimation of internet traffic (Crovella et al., 1998; Hernandez-Campos et al., 2004), multi-agent modeling of supply chain networks (Thadakamaila et al., 2004), modeling information cascades in economic multi-agent models (De Vany et al., 1999; Kononov, 2010) and, among others, numerous problems in distributed modeling for social science (Barabasi, 2005; Eom & Jo, 2015). Moreover, in sub-Gaussian environments, if the multi-agent communication is noisy with a (small) non-zero bias, this can lead to model misspecification equivalent to the communicated variable exhibiting heavy tails (McCulloch & Neuhaus, 2011). It is therefore prudent to devise methods that are robust to such *heavy-*

¹Media Lab and Institute for Data, Systems and Society, Massachusetts Institute of Technology. Correspondence to: Abhimanyu Dubey <dubeya@mit.edu>.

tailed effects, which is the central theme of this paper. We define heavy-tailed random variables as follows.

Definition 1 (Heavy-Tailed Random Variables). *A random variable X is heavy-tailed if it does not admit a finite moment generating function, i.e., there is no $u_0 > 0$ such that,*

$$\forall |u| \leq u_0, M_X(u) \triangleq \mathbb{E}[\exp(uX)] < \infty.$$

X is $(1 + \varepsilon)$ -heavy tailed if $\mathbb{E}[|X|^t] = \infty$ for all $t > 1 + \varepsilon$.

Almost all aforementioned research on this problem proposes algorithms that average opinions via a *consensus* protocol (Landgren et al., 2016a), and ensure a group regret of $O(\ln T)$ in sub-Gaussian settings. However, these algorithms, as a consequence, employ the empirical mean, which, as we demonstrate in this paper, leads to $O(T^{2/3})$ regret under 2-heavy tails (i.e., finite variance, see Section 5). Moreover, the group regret achieved has a $O(\log |V|)$ dependence on $|V|$ (i.e., the number of agents), which we demonstrate to be suboptimal. To our knowledge, no prior work addresses $|V|$ -optimality for the stochastic problem¹.

Contributions. Our first contribution is the first (problem-dependent) asymptotic lower bound of $\Omega(K\Delta^{-1/\varepsilon} \ln T)$ on the *group regret* for the cooperative multi-armed stochastic bandit. This result holds for any connected graph \mathcal{G} and arbitrary communication protocol (following mild conditions). Note that this implies that agents on average can achieve a maximum of $O(\frac{1}{\sqrt{|V|}})$ reduction in regret when cooperating, compared to the earlier benchmark of $O(\frac{\ln |V|}{|V|})$ (Landgren et al., 2016a; Martínez-Rubio et al., 2019). These results generalize the lower bound for the stochastic bandit with multiple pulls, obtained by Anantharam et al. (1987), to the case when rewards are obtained after delays. In the heavy-tailed case, this lower bound matches the $\Omega(\Delta^{-1/\varepsilon} \ln T)$ rate obtained by Bubeck et al. (2013), and in the sub-Gaussian case, matches the $\Omega(\Delta^{-1} \ln T)$ problem-dependent rates (Agrawal & Goyal, 2012).

Next, we present an algorithm MP-UCB for the cooperative multi-agent stochastic bandit under heavy-tailed densities. The key concept utilized in the development of the algorithm is to provide an alternate technique to control the variance of the arm estimators across the network \mathcal{G} when the *consensus* protocol provides suboptimal guarantees. This is done by utilizing an alternate communication protocol titled LOCAL (Linial, 1992) to share information between agents, and then incorporating robust mean estimators to achieve optimal rates. In this process, we also outline a subroutine to efficiently compute the univariate robust mean for the bandit problem (i.e., when confidence δ changes with time).

¹There is recent research on the *adversarial* cooperative case (Cesa-Bianchi et al., 2019) with an $\Omega(\sqrt{|V|T})$ lower bound. For the stochastic case, we propose algorithms that achieve lower *problem-dependent* rates, matching our bound of $\Omega(K \ln T)$.

We demonstrate that MP-UCB achieves $O(\bar{\chi}(\mathcal{G}_\gamma)K \ln T)$ group regret when run in a completely decentralized manner (i.e., agents select actions independently), and $O(\alpha(\mathcal{G}_\gamma)K \ln T)$ when run in a centralized manner (i.e., some agents mimic others), similar to the regret bounds obtained by Cesa-Bianchi et al. (2019) for the adversarial case. Here, $\gamma \leq \text{diameter}(\mathcal{G})$ is a parameter controlling the density of communication, and \mathcal{G}_γ is the γ graph power of \mathcal{G} . $\bar{\chi}$ denotes the clique covering number, and $\alpha \leq \bar{\chi}$ denotes the independence number of a graph. These results are optimal, in the sense that when our algorithm is run with $\gamma = \text{diameter}(\mathcal{G})$, both variants obtain a group regret of $O(K \ln T)$, matching the lower bound. This $O(\ln |V|)$ improvement is achieved since the LOCAL protocol allows us to partition the power graph \mathcal{G}_γ in a manner that induces a constant regret overhead from the communication delay, in contrast to the *consensus* protocol that diffuses information slowly for sparse \mathcal{G} . Furthermore, when we allow $O(K)$ -sized messages per round, we demonstrate that MP-UCB obtains $O(\alpha(\mathcal{G}_\gamma)K \ln T)$ regret without knowledge of \mathcal{G} .

We evaluate our algorithms on a benchmark of real-world and random graphs. While we consider heavy-tailed densities, it can be easily seen that MP-UCB can be applied to sub-Gaussian densities with optimal rates as well.

2. Related Work

Cooperative Decision-Making. Cooperative decision-making for the stochastic multi-armed bandit has recently seen a lot of research interest. Decentralized cooperative estimation has been explored for sub-Gaussian stochastic bandits using a *running consensus* protocol in (Landgren et al., 2016a;b; Martínez-Rubio et al., 2019) and for adversarial bandits (Bar-On & Mansour, 2019; Cesa-Bianchi et al., 2019) using a message-passing protocol. Localized decision-making for sub-Gaussian rewards has also been explored in the work of (Landgren et al., 2018), and a fully-centralized algorithm in (Shahrampour et al., 2017), where all agents select the same action via voting. The stochastic bandit with multiple pulls (Anantharam et al., 1987; Xia et al.) is equivalent to the cooperative multi-armed bandit on a complete \mathcal{G} with a centralized actor (since there are no delays and all agents have the same information $\forall t \in [T]$).

Contrasted to cooperative settings, there is extensive research in competitive settings, where multiple agents compete for arms (Bistritz & Leshem, 2018; Bubeck et al., 2019; Liu & Zhao, 2010a;b;c). For strategic experimentation, Brânzei & Peres (2019) provide an interesting comparison of exploration in cooperative and competitive agents.

A closely-related problem setting is the single-agent *social network* bandit, where a user is picked at random every trial, and the algorithm must infer its contextual mean re-

ward (Cesa-Bianchi et al., 2013; Gentile et al., 2014; 2017; Li et al., 2016), while assuming an underlying *clustering* over the users. This problem setting, while relevant, crucially differs from the one considered herein, since (a) this is a single-agent setting (only one action is taken every round), and (b) there are no delays in the rewards obtained. While a multi-agent variant has been considered (Korda et al., 2016), this work also assumes no delays in communication.

Heavy-Tailed Bandits. Bubeck et al. (2013) first discuss the problem of stochastic bandits with heavy-tailed rewards, and propose the ROBUST-UCB algorithm that uses robust mean estimators to obtain logarithmic regret. Vakili et al. (2013) introduce DSEE, an algorithm that sequences phases of exploration and exploitation to obtain sublinear regret. Thompson Sampling (Thompson, 1933) has been analysed for exponential family bandits (that include Pareto and Weibull heavy-tailed distributions) in the work of Korda et al. (2013), however, these distributions have “lighter” tails owing to the existence of higher order moments. Dubey & Pentland (2019) provide an algorithm for Thompson Sampling for α -stable densities (Borak et al., 2005), at family of heavy-tailed densities typically with infinite variance. Yu et al. (2018) provide a purely exploratory algorithm for best-arm identification for ε -heavy tailed rewards. For the linear bandit, (Medina & Yang, 2016; Shao et al., 2018) provide nearly-optimal algorithms under heavy tails. To our knowledge, this paper is the first to study robust bandit learning in the context of decentralized multi-agent estimation.

3. Preliminaries

Finite-Armed Stochastic Bandit. We consider the family of bandit problems \mathcal{E} for a finite, countable set of actions \mathcal{A} , such that $|\mathcal{A}| = K$. \mathcal{E} is considered to be *unstructured*, i.e. the rewards from each arm are independent of the others.

Definition 2 (Unstructured Bandit Problem). *An environment class of bandit problems \mathcal{E} is unstructured if its action space \mathcal{A} is finite, and there exists a set of distributions \mathcal{M}_a for each $a \in \mathcal{A}$ such that*

$$\mathcal{E} = \{\nu = (P_a : a \in \mathcal{A}) : P_a \in \mathcal{M}_a \forall a \in \mathcal{A}\}.$$

Agents face a common stochastic bandit with K arms. Rewards from arm $k \in [K]$ are drawn from an ε -heavy tailed distribution ν_k with mean μ_k , and known bounds on the $(1 + \varepsilon)$ moments $\mathbb{E}[|X - \mu_k|^{1+\varepsilon}] \leq \rho$ and $\mathbb{E}[|X|^{1+\varepsilon}] \leq u$. The optimal arm is given by $k^* = \arg \max_{k \in [K]} \mu_k$.

Cooperative Problem Setting. We consider M agents communicating via a connected, undirected graph $\mathcal{G} = (V, E)$. Communication is bidirectional, and any message sent from agent v is obtained by agent v' after $d(v, v') - 1$ rounds of the bandit problem, where $d(v, v')$ denotes the length of the shortest path between the agents. Let \mathbf{L} and \mathbf{A}

denote the graph Laplacian and adjacency matrix of \mathcal{G} , and $\mathbf{P} = \mathbf{I}_M - \kappa \cdot d_{\max}^{-1} \mathbf{L}$ is a row stochastic matrix, where \mathbf{I}_M is the identity matrix of order M , $\kappa > 0$ is a constant and $d_{\max} = \max_{m \in \mathcal{G}} \text{degree}(m)$. We assume that the eigenvalues λ_i of \mathbf{P} are ordered such as $\lambda_1 = 1 \geq \dots \geq \lambda_M > -1$. Let A_{mt} denote the action taken by agent m at time t , and X_{mt} denote the corresponding reward. $n_k(T)$ 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. Let the power graph of order γ of \mathcal{G} be given by \mathcal{G}_γ , i.e., \mathcal{G}_γ contains an edge (i, j) if there exists a path of length at most γ in \mathcal{G} between agents i and j . For any agent $v \in V$, let the neighborhood of m in \mathcal{G}_γ be given by $\mathcal{N}_\gamma(v)$. The policy of agent $v \in V$ is given by $(\pi_{v,t})_{t \in [T]}$, and the collective policy is given by $\Pi = (\pi_{v,t})_{v \in V, t \in [T]}$.

Definition 3 (Consistent Bandit Policy). *Let Π be any bandit policy, potentially running over multiple agents. Π is **consistent** if, for any suboptimal arm $k \in [K], k \neq k^*$, horizon $T > 0$, one has $\mathbb{E}[N_k(T)] = o(T^a)$ for any $a > 0$.*

Univariate Robust Estimation. Optimal algorithms have been proposed for robust estimation of location in the univariate setting with polynomial running time. The simplest of these is the trimmed mean, that rejects outlying samples based on an upper bound on the moments. Its runtime for N samples obtained sequentially (with changing confidence δ) is $O(N^2)$, which we improve to $O(N \ln N)$ (Algorithm 3).

Definition 4 (Trimmed Mean). *Consider n copies X_1, \dots, X_n of a heavy-tailed random variable X such that $\mathbb{E}[X] = \mu, \mathbb{E}[X^{1+\varepsilon}] \leq u$ for some $\varepsilon \in (0, 1]$. The online trimmed mean, for some $\delta \in (0, 1)$ is defined as*

$$\hat{\mu}_O = \frac{1}{n} \sum_{i=1}^n X_i \mathbf{1} \left\{ |X_i| \leq \left(\frac{ui}{\log \delta^{-1}} \right)^{\frac{1}{1+\varepsilon}} \right\}.$$

Several alternative robust mean estimators exist, such as the median-of-means or Catoni’s estimator (Catoni, 2012). Under stricter tail assumptions, they provide better estimates, however, for simplicity, we continue with the trimmed mean. In the analysis, we assume that a mean estimator exists that achieves the following optimal rate (up to constants).

Assumption 1 (Rate Assumption). *Let X_1, \dots, X_n be n samples of an ε -heavy tailed random variable, where $\varepsilon \in (0, 1]$, and $\mathbb{E}[X] = \mu$. For positive constants c, ρ suppose that there exists a robust estimator $\hat{\mu}(\delta, n)$ such that, with probability at least $1 - \delta$,*

$$|\hat{\mu}(\delta, n) - \mu| \leq 2\rho^{\frac{1}{1+\varepsilon}} \left(\frac{c \log(\delta^{-1})}{n} \right)^{\frac{\varepsilon}{1+\varepsilon}}.$$

Catoni (2012) provides this as the optimal achievable rate under heavy tails, and Bubeck et al. (2013) demonstrate that the trimmed mean achieves this rate (see appendix).

4. Lower Bounds

We now present lower bounds on cooperative decision-making. All full proofs are presented in the appendix for brevity. We consider $|V| = M$ agents communicating over graph \mathcal{G} , with $\text{diameter}(\mathcal{G}) = \gamma_* \ll M$. We first make some (mild) assumptions on the communication protocol.

Assumption 2 (Communication Protocol). *The communication protocol considered follows:*

1. Any agent m is capable of sending a message $\mathbf{q}_m(t)$ to any other agent $m' \in [M]$, which is earliest received at time $t + \min(0, d(m, m') - 1)$.
2. $\mathbf{q}_m(t)$ is a function of the action-reward pairs of agent m , i.e. $\mathbf{q}_m(t) = \mathbf{F}_t(A_{m,1}, X_{m,1}, \dots, A_{m,t}, X_{m,t})$ for any deterministic, bijective and differentiable set of functions $\mathbf{F}_t = (f_{i,t})_{i \in [L]}, f_{i,t} : \mathbb{R}^{2t} \rightarrow \mathbb{R}$.
3. \mathbf{F}_t satisfies $|\det(\mathbf{J}_t)| = \Lambda(m, t)$. Here, $\mathbf{J}_t(\cdot)$ is the Jacobian of \mathbf{F}_t , and Λ is only a function of m and t .

This assumption ensures that (a) information can flow between any two agents, and (b) that the messages are not stochastic and are independent the bandit problem. We can then derive a lower bound on the group regret.

Theorem 1 (Lower Bound). *For any consistent cooperative multi-agent policy $\Pi = (\Pi_t)_{t \in [T]}$ on M agents that satisfies Assumption 2 the following is true.*

$$\liminf_{T \rightarrow \infty} \frac{R_{\mathcal{G}}(T)}{\ln T} \geq \sum_{k: \Delta_k > 0} \frac{\Delta_k}{\mathbb{D}_k^{\text{inf}}}.$$

Here, $\mathbb{D}_k^{\text{inf}} = \inf_{\nu' \in \mathcal{M}_k} \{\mathbb{D}_{\text{KL}}(\nu, \nu') : \mu(\nu') > \mu^*\}$, and $\mathbb{D}_{\text{KL}}(\cdot, \cdot)$ denotes the Kullback-Leibler divergence.

Remark 1. *Theorem 1 does not guarantee an overhead from delayed communication, since it includes protocols that allow information to flow completely through the (connected) network \mathcal{G} , albeit at a delay (which is independent of T). Making stronger assumptions about the connectivity of \mathcal{G} and communication protocol can lead to stronger bounds.*

Remark 2. *This result generalizes that obtained by Anantharam et al. (1987) for a centralized agent with multiple pulls to the case where rewards are obtained after finite delays. This can be understood by considering a complete \mathcal{G} , which is equivalent to having a centralised agent (since there is no difference in information between agents). The comparison with a single agent pulling MT arms (Martínez-Rubio et al., 2019), is therefore an incorrect benchmark.*

For the specific case of $(1 + \varepsilon)$ -heavy tailed rewards, the single-agent lower bound provided by (Bubeck et al., 2013) can be easily extended to the cooperative multi-agent case.

Corollary 1 (Lower Bound on Heavy-Tailed Cooperative Regret). *For any $\Delta \in (0, 1/4)$, there exist $K \geq 2$ distributions ν_1, \dots, ν_K satisfying $\mathbb{E}_{X \sim \nu_k} [|X|^{1+\varepsilon}] \leq u$, and $\mathbb{E}_{X \sim \nu_*} [X] - \mathbb{E}_{X \sim \nu_k} [X] = \Delta \forall k \in [K]$, such that any consistent decentralized policy $\Pi_t = (\pi_{m,t})_{m \in [M], t \in [T]}$ that satisfies Assumption 2 obtains group regret of $\Omega(K \Delta^{-1/\varepsilon} \ln T)$.*

The $O(\Delta^{-1/\varepsilon})$ dependency is unavoidable, as shown in (Bubeck et al., 2013), and it can be matched using robust estimators. The formulation of robust estimators makes averaging-based communication protocols infeasible, such as the *running consensus*, as shown in the following section.

5. The Limits of Running Consensus

Under the consensus protocol, agents maintain an estimate of values of interest, which they average with their neighbors every round. The protocol stores $2K$ opinion vectors $\hat{\mathbf{s}}_k(t) = (\hat{s}_k^v(t))_{v \in V}, k \in [K]$ and $\hat{\mathbf{n}}_k(t) = (\hat{n}_k^v(t))_{v \in V}, k \in [K]$, that are updated as follows.

$$\hat{\mathbf{s}}_k(t) = \mathbf{P}(\hat{\mathbf{s}}_k(t-1) + \mathbf{r}_k(t) \circ \zeta_k(t)). \quad (2)$$

$$\hat{\mathbf{n}}_k(t) = \mathbf{P}(\hat{\mathbf{n}}_k(t-1) + \zeta_k(t)). \quad (3)$$

Here $\hat{\mathbf{s}}_k(t)$ is a vector of reward sums for arm k for each agent, $\zeta_k(t)$ is a vector of indicators of whether the agents pulled arm k at time t , and $\mathbf{r}_k(t)$ is the vector of rewards obtained by the agents from arm k . Using this, any agent $v \in V$ computes the empirical mean of each arm k .

$$\hat{\mu}_k^{(v)}(t) = \hat{s}_k^{(v)}(t) / \hat{n}_k^{(v)}(t). \quad (4)$$

When $\varepsilon = 1$, i.e. the reward distributions have finite variance, we can design a UCB algorithm **CONSENSUS-UCB** (Landgren et al., 2016a), where each agent chooses the arm that maximizes the following UCB.

$$A_{v,t} = \arg \max_{k \in [K]} \left\{ \hat{\mu}_k^{(v)}(t) + \sqrt{\frac{6\rho t^{2/3}}{|V|} \left(\frac{\hat{n}_k^v(t) + \epsilon^k}{\hat{n}_k^v(t)^2} \right)} \right\}. \quad (5)$$

Theorem 2. *The CONSENSUS-UCB algorithm obtains a group regret of $O\left((1 + h(\mathcal{G}))KT^{\frac{2}{3}}\right)$ after T trials, where $h(\mathcal{G})$ is, for constants a_{pj} that only depend on \mathcal{G} ,*

$$h(\mathcal{G}) = \sum_{p=1}^{|V|} \sum_{j=2}^{|V|} \frac{|\lambda_p \lambda_j|}{1 - |\lambda_p \lambda_j|} a_{pj}.$$

A full description of this algorithm is included in the appendix. Since we are utilizing the empirical mean for CONSENSUS-UCB, the UCB utilized cannot be made tighter, suggesting that the algorithm is suboptimal in T . More importantly, it can be expected that any algorithm that

utilizes the empirical mean (running consensus converges to the empirical mean as $T \rightarrow \infty$ (Aysal & Barner, 2010)) is suboptimal in T owing to the suboptimality of the empirical mean itself (see appendix).

6. Message-Passing Cooperative UCB

In the message-passing protocol, agents $v \in V$ communicate via messages $\mathbf{q}_v(t) = \langle v, t, A_{v,t}, X_{v,t} \rangle$. This message is first sent to its neighbors in \mathcal{G} , and it is subsequently forwarded by any agent that receives it until time $t + \gamma$, after which it is discarded. $0 \leq \gamma \leq \text{diameter}(\mathcal{G})$ is therefore the communication *density*, where lower values of γ imply less communication in the network.

Let $Q_v(t)$ denote the set of incoming messages received by agent v at instant t . During any trial, the agent first pulls an arm, and creates the message $\mathbf{q}_v(t)$. It then processes all messages in $Q_v(t)$, and updates its beliefs as per any bandit algorithm. Finally, it discards all messages older than $t - \gamma$ and forwards all remaining messages in $Q_v(t) \cup \{\mathbf{q}_v(t)\}$ to all its neighbors in \mathcal{G} . This protocol has been used in distributed optimization (Moallemi, 2007), non-stochastic bandit settings (Bar-On & Mansour, 2019; Cesa-Bianchi et al., 2019) and asynchronous online learning (Suomela, 2013). This protocol satisfies Assumption 2 with $\gamma = \text{diameter}(\mathcal{G})$.

6.1. Decentralized Algorithm

In the decentralized setting, each agent acts independently, i.e., there is no centralized controller that dictates actions. In this setting, each agent v maintains a set $S_k^v(t)$ of rewards obtained from arm k , which it updates at each trial from its own pulls and incoming messages. Then it computes the robust mean of $S_k^v(t)$ via the estimator $\hat{\mu}(|S_k^v(t)|, \delta)$. Using Assumption 1, it then estimate a UCB for each arm mean, and selects the arm with the largest UCB (Algorithm 1).

Theorem 3. *The group regret for Algorithm 1 when run with parameter γ and mean estimator $\hat{\mu}(n, \delta)$ that satisfies Assumption 1 with constants c and ρ satisfies:*

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

Here, $C > 0$ is a constant independent of T, K, M , and $\bar{\chi}(\cdot)$ refers to the clique number.

Proof (sketch). We first bound the regret in each clique \mathcal{C} within the clique covering \mathcal{C}_γ of \mathcal{G}_γ . This is done by noticing that the upper confidence bound for any arm at a selected t deviates by a constant amount between agents based on the

Algorithm 1 DECENTRALIZED MP-UCB

```

1: Input: Arms  $k \in [K]$ , parameters  $\varepsilon, c, \rho$ , estimator  $\hat{\mu}(n, \delta)$ 
2:  $S_k^v \leftarrow \phi \forall k \in [K], Q_v(t) \leftarrow \phi, \forall v \in V.$ 
3: for each iteration  $t \in [T]$  do
4:   for each agent  $v \in V$  do
5:     if  $t \leq K$  then
6:        $A_{m,t} \leftarrow t.$ 
7:     else
8:       for Arm  $k \in [K]$  do
9:          $\hat{\mu}_k^{(v)} \leftarrow \hat{\mu}(S_k^v, 1/t^2).$ 
10:         $\text{UCB}_k^{(v)}(t) \leftarrow \rho^{\frac{1}{1+\varepsilon}} \left( \frac{2c \ln t}{|S_k^v|} \right)^{\frac{\varepsilon}{1+\varepsilon}}.$ 
11:       end for
12:        $A_{v,t} \leftarrow \arg \max_{k \in [K]} \{ \hat{\mu}_k^{(v)}(t) + \text{UCB}_k^{(v)}(t) \}.$ 
13:     end if
14:      $X_{v,t} \leftarrow \text{PULL}(A_{v,t}).$ 
15:      $S_{A_{v,t}}^v \leftarrow S_{A_{v,t}}^v \cup \{X_{v,t}\}$ 
16:      $Q_v(t) \leftarrow Q_v(t) \cup \{ \langle v, t, A_{v,t}, X_{v,t} \rangle \}.$ 
17:     for each neighbor  $v'$  in  $\mathcal{N}_1(v)$  do
18:        $\text{SENDMESSAGES}(v, v', Q_v(t)).$ 
19:     end for
20:   end for
21:   for each agent  $v \in V$  do
22:      $Q_v(t+1) \leftarrow \phi.$ 
23:     for each neighbor  $v'$  in  $\mathcal{N}_1(v)$  do
24:        $Q' \leftarrow \text{RECEIVEMESSAGES}(v', v)$ 
25:        $Q_v(t+1) \leftarrow Q_v(t+1) \cup Q'.$ 
26:     end for
27:     for  $\langle v', t', a', x' \rangle \in Q_v(t+1)$  do
28:       if  $v' \in \text{CLIQUE}(v, \mathcal{G}_\gamma)$  then
29:          $S_{a'}^v \leftarrow S_{a'}^v \cup \{x'\}.$ 
30:       end if
31:     end for
32:   end for
33: end for

```

number of times each agent has pulled an arm. By bounding this deviation, we obtain a relationship between the confidence bound of each arm for each agent within the clique \mathcal{C} . Next, we bound the probability of pulling a suboptimal arm within the clique \mathcal{C} using the previous result. Summing over the clique cover \mathcal{C}_γ delivers the final form of the result. The complete proof is included in the appendix for brevity.

Remark 3. *Communication density determines the group regret dependence or “cooperation” in Algorithm 1. When $\gamma = \text{diameter}(\mathcal{G})$, $\bar{\chi}(\mathcal{G}_\gamma) = 1$, and we incur optimal group regret $O(K\Delta^{-1/\varepsilon} \ln T)$, and also satisfies both assumptions of Assumption 2. However, when $\gamma = 0$, i.e. agents do not communicate, regret is $O(|V|K\Delta^{-1/\varepsilon} \ln T)$.*

Each agent in Algorithm 1 utilizes observations only from its own clique in \mathcal{G}_γ to make decisions, effectively partitioning \mathcal{G} . When \mathcal{G} is sparse (e.g., small-world networks (Barabasi, 2005)), the clique number of the graph \mathcal{G}_γ can be large. In this case, a centralized variant can provide lower regret.

Algorithm 2 CENTRALIZED MP-UCB

```

1: Input: Same as Algorithm 1.
2: Set  $S_k^v \leftarrow \phi \forall k \in [K], Q_v(t) \leftarrow \phi, A_v^* \leftarrow 1$ , for all  $v \in V$ .
3: for each iteration  $t \in [T]$  do
4:   for each agent  $v \in V$  do
5:     if  $t \leq K$  then
6:        $A_{v,t} \leftarrow t$ .
7:     else if  $v \in V'$  or  $t \leq d(v, l(v))$  then
8:       Run lines 8-12 of Algorithm 1.
9:     else
10:       $A_{v,t} \leftarrow A_v^*$ .
11:    end if
12:    Run lines 14-19 of Algorithm 1.
13:  end for
14:  for each agent  $v \in V$  do
15:    Run lines 22-26 of Algorithm 1.
16:    for  $\langle v', t', a', x' \rangle \in Q_v(t+1)$  do
17:       $S_{a'}^v \leftarrow S_{a'}^v \cup \{x'\}$ .
18:    end for
19:     $A_v^* = \text{CHOOSELASTACTION}(\cup_k S_k^v(t+1))$ .
20:  end for
21: end for

```

6.2. Centralized Algorithm

In the centralized setting, we present a version of the “follow-the-leader” strategy. Here, the agents are partitioned into “leaders” and “followers”. The leader agents follow the same procedure identically to Algorithm 1, and the follower agents simply copy the most recent action they have observed of their associated leader. We now describe how the graph \mathcal{G} is partitioned into leaders and followers.

Definition 5 (Maximal Weighted Independent Set). *An independent set of a graph $\mathcal{G} = (V, E)$ is a set of vertices $V' \subseteq V$ such that no two vertices in V' are connected. A maximal independent set V^* is the largest independent set in \mathcal{G} , and the independence number $\alpha(\mathcal{G}) = |V^*|$. For a vertex-weighted graph, a maximal weighted independent set $V_w^* \subseteq V$ is the maximal independent set such that the sum of weights for all vertices in V_w^* is the largest possible.*

We select the leaders as the members of a maximal independent set $V' \subseteq V$ of \mathcal{G}_γ . For each follower agent $v \in V \setminus V'$ we assign a leader $l(v)$ to it such that (a) there is an edge between v and $l(v)$ in \mathcal{G}_γ , and (b) $l(v)$ has maximum degree in $V' \cap \mathcal{N}_\gamma(v)$, i.e. $l(v) \in V'$ such that $l(v) = \arg \max_{v' \in V' \cap \mathcal{N}_\gamma(v)} \text{degree}(v')$. It is trivial to demonstrate that each agent will either be a leader node, or be connected to a leader (see appendix). Algorithm 2 describes this algorithm particularly from its differences with the decentralized version.

Theorem 4. *Algorithm 2 run with parameters γ, c, ρ obtains the following group regret (where $\alpha(\cdot)$ denotes the independence number).*

$$R_{\mathcal{G}}(T) = O\left(\alpha(\mathcal{G}_\gamma) \left(\sum_{k: \Delta_k > 0} \Delta_k^{-1/\varepsilon}\right) \ln T\right).$$

Proof (sketch). The key idea is to partition \mathcal{G}_γ into non-overlapping sets given by V' and to note that $R_v \leq R_{l(v)} + \gamma$ for any $v \in V \setminus V'$. Then, we can bound the number of times any element in V' selects an arm until time t as a function of its neighborhood in \mathcal{G}_γ . Using this bound, we can then create an UCB to bound the probability of pulling a suboptimal arm for any agent $v \in V'$, and collectively bound the group regret of the entire neighborhood. Summing over $v \in V'$ delivers the final result, since V' forms a vertex cover in \mathcal{G}_γ . The complete proof is available in the appendix.

Since $\alpha(\mathcal{G}) \leq \bar{\chi}(\mathcal{G})$ for any graph \mathcal{G} , the centralized version of the MP-UCB algorithm obtains regret strictly no worse compared to the decentralized version. We are aware that the set of leader nodes must form a maximal independent set in \mathcal{G}_γ , however, for large graphs there may be multiple maximal independent sets present, and selecting a suboptimal independent set can increase group regret. Agents present more “centrally” may be a better choice as leaders, compared to “peripheral” agents. Our choice of independent set is motivated by the following result.

Corollary 2. *For agent $v \in \mathcal{G}$, let v^* denote its corresponding leader agent ($v = v^*$ for leaders), and $F(v^*)$ denote the corresponding set of follower agents for v^* (including v^*). The following holds for the regret $R_v(T)$.*

$$R_v(T) = O\left(\frac{K \ln T}{|F(v^*)| \Delta_{\min}^{1/\varepsilon}}\right).$$

We see that, intuitively, even for agents that themselves are not well-connected, as long as they are connected to a well-connected leader (with large $|F(v^*)|$), the individual regret will be low. By this result, we select the weight assigned to any agent v as its degree in \mathcal{G}_γ , since, asymptotically performance depends on $(|F(v^*)|^{-1})$. A few additional remarks can be made, inspired by Bar-On & Mansour (2019).

Remark 4. *The average regret from Algorithm 2 is $O((\alpha(\mathcal{G}_\gamma)/|V|)K \ln T)$, i.e. optimal when $\gamma = \text{diam}(\mathcal{G})$. When $\gamma = \sqrt{K}$, Algorithm 2 can obtain a per-agent regret of $O(\Delta_*^{-1/\varepsilon} \sqrt{K} \ln T)$. This can be shown following the procedure in Bar-On & Mansour (2019), by noticing that when \mathcal{G} is connected, $\alpha(\mathcal{G}_\gamma) \leq \lceil 2|V|/(\gamma + 2) \rceil$. Also note that we need only \sqrt{K} leaders at most to obtain this regret. When $\gamma = \text{diam}(\mathcal{G})$, then, only 1 arbitrarily chosen leader can deliver optimal regret, regardless of its position in \mathcal{G} .*

6.3. Additional Optimizations

$O(K)$ Per-round Communication. We now demonstrate that communicating additional information beyond just action-reward pairs can significantly improve performance, and obtain optimal regret *without knowledge of \mathcal{G}* . In this case, the message $\mathbf{q}_v(t)$ is augmented as follows.

$$\mathbf{q}_v(t) = \langle v, t, A_{v,t}, X_{v,t}, \hat{\boldsymbol{\mu}}^{(v)}(t), \mathbf{N}^v(t) \rangle. \quad (6)$$

Where $\hat{\boldsymbol{\mu}}^{(v)}(t) = (\hat{\mu}_k^{(v)}(t))_{k \in [K]}$ are the robust mean estimates used by agent m to make decisions at time t , and $\mathbf{N}^v(t) = (|S_k^v(t)|)_{k \in [K]}$ is the vector containing the number of reward samples possessed by agent v until time t . Each agent v also maintains a set W of the most recent $(\hat{\boldsymbol{\mu}}^{(v')}(t), \mathbf{N}^{v'}(t))$ for each $v' \in \mathcal{N}_\gamma(v) \cup \{v\}$, which they update with each message received from agent v' . At any instant, the agent chooses, for each arm k , the corresponding $\hat{\mu}_k^{v^*}(t)$ and $N^{v^*}(t)$ in W with the largest $N^{v^*}(t)$ (and “tightest” UCB) to construct its upper confidence bound. The full algorithm (KMP-UCB) is described in the appendix.

Theorem 5. KMP-UCB obtains group regret $R_{\mathcal{G}}(T)$ of $O(\alpha(\mathcal{G}_\gamma)K\Delta^{-1/\varepsilon} \ln T)$ over any connected graph \mathcal{G} .

Proof (sketch). We first note that there will be an independent set of agents $V' \in V$ that has, at any given t , the largest set of observations within their neighborhoods. Since at trial $t + \text{diameter}(\mathcal{G}_\gamma)$, any other agent will either use the confidence estimates of $v \in V'$, or will have better estimates (from more samples). This provides us a technique to lower bound the number of times the entire group of agents V will pull an arm at any time t in terms of the pulls of $v \in V'$, and then construct a UCB for each arm from it. We then proceed by the standard UCB technique for a single-agent, and use concentration of the robust mean to derive regret for each $v \in V'$. Finally, summing over v gives us the desired result.

Contrasted to the regret bound of $O(\sqrt{|V|}\alpha(\mathcal{G})TK \ln K)$ obtained by Cesa-Bianchi et al. (2019) for the nonstochastic case (where communication is also $O(K)$ per agent), our algorithm obtains lower group regret in the stochastic case. Additionally, this implies a $O(\ln |V|)$ improvement over the previous bound in the stochastic case (Martínez-Rubio et al., 2019).

Online Estimation of Trimmed Mean. The trimmed mean estimator requires selecting a sample X_i at time t only if $|X_i| \leq (2ui \ln t)^{1/(1+\varepsilon)}$ (Definition 4). This implies that the i^{th} reward sample an agent has will be selected at the smallest time t such that $(|X_i|^{1+\varepsilon}/i) \leq 2u \ln t$. When T is known, we can utilize a binary search tree to make an update to the robust mean $O(\ln t)$ instead of $O(t)$ at time t . We outline this procedure in Algorithm 3.

Algorithm 3 assumes that for any t , a new set of observations O_t is available, which it incorporates into the robust mean with $O(\ln t)$ per sample (instead of typically recomputing the mean for each t). The complexity stems from the binary search, assuming the dictionary lookup is $O(1)$.

7. Experiments

Our primary contributions are in leveraging cooperation to accelerate overall decision-making, and the most interesting aspects of this study pertain to how graph structures,

Algorithm 3 ONLINE TRIMMED MEAN ESTIMATOR

```

1: Input:  $u, T$ .
2: Create dictionary  $D$  of size  $T$ , where  $D(t) = \phi \forall t \in [T]$ .
3: Create BST  $B$  with entries  $((2u \ln t)^{1/(1+\varepsilon)})_{t \in [T]}$ .
4:  $\hat{S}_O \leftarrow 0, n \leftarrow 0$ 
5: for  $t \in [T]$  do
6:    $O_t \leftarrow \text{OBSERVATIONS}(t)$ .
7:   for  $x_t \in O_t$  do
8:      $n \leftarrow n + 1$ 
9:      $i_t \leftarrow \max(t, \text{SEARCH}(B, (|x_t|^{1+\varepsilon}/n)))$ .
10:     $D(i_t) \leftarrow D(i_t) \cup \{x_t\}$ .
11:   end for
12:   for  $x \in D(t)$  do
13:      $\hat{S}_O \leftarrow \hat{S}_O + x$ .
14:   end for
15:    $\hat{\mu}_O(t) \leftarrow \hat{S}_O/n$ .
16: end for
    
```

scalability, heavy tails and decentralized vs. centralized estimation affect the group regret. To this end, we analyse these aspects in our experimental setup, and relegate other comparisons (Δ_k , number of arms, etc.) to the appendix.

Reward Distributions. We conduct experiments using α -stable densities (Lévy, 1925), that admit finite moments only of order $< \alpha \leq 2$, and we consider α -stable densities where $\alpha \geq 1$. The α -stable family includes several widely used distributions, such as Gaussian ($\alpha = 2$, only light-tailed density), Lévy ($\alpha = 0.5$) and Cauchy ($\alpha=1$). The primary advantage of this density is that α can be adjusted to alter the heaviness of the reward distribution ($\alpha > 1$).

Graph Partitioning. For Algorithm 2, we require computing the maximal weighted independent set of \mathcal{G} . This problem is NP-Hard for arbitrary \mathcal{G} , and difficult to approximate. We use the approximate algorithm presented in (Lucas, 2014) that uses the QUBO (Glover & Kochenberger, 2018) solver.

Experiment 1: Random Graphs. We set $K = 5$, $\alpha = 1.9$ for the standard α -stable density, and sample arm means randomly from the interval $[0, 1]$ for each arm every experiment. We then construct random graphs on 200 agents from the Erdos-Renyi (ER) ($p = 0.7$) and Barabasi-Albert (BA) ($m = 5$) random graph families, and compare all three of our algorithms (using the trimmed mean estimator, with $\gamma = \text{diam}(\mathcal{G})/2$) with the CONSENSUS-UCB and single-agent ROBUST-UCB (Bubeck et al., 2013) algorithms. We compare the group regret $R_{\mathcal{G}}(T)$ vs. T , averaged over 100 random graphs and bandit instances. The results for Erdos-Renyi graphs (Figure (1A)) and Barabasi-Albert graphs (Figure (1B)) demonstrate that while our algorithms outperform the baselines (in the order dictated by regret bounds), the gain is larger for the former. We attribute this to the network connectivity, i.e., since Barabasi-Albert graphs have “hubs”, the clique number $\bar{\chi}(\mathcal{G})$ for these graphs is larger.

Experiment 2: Real-World Networks. We select the p2p-

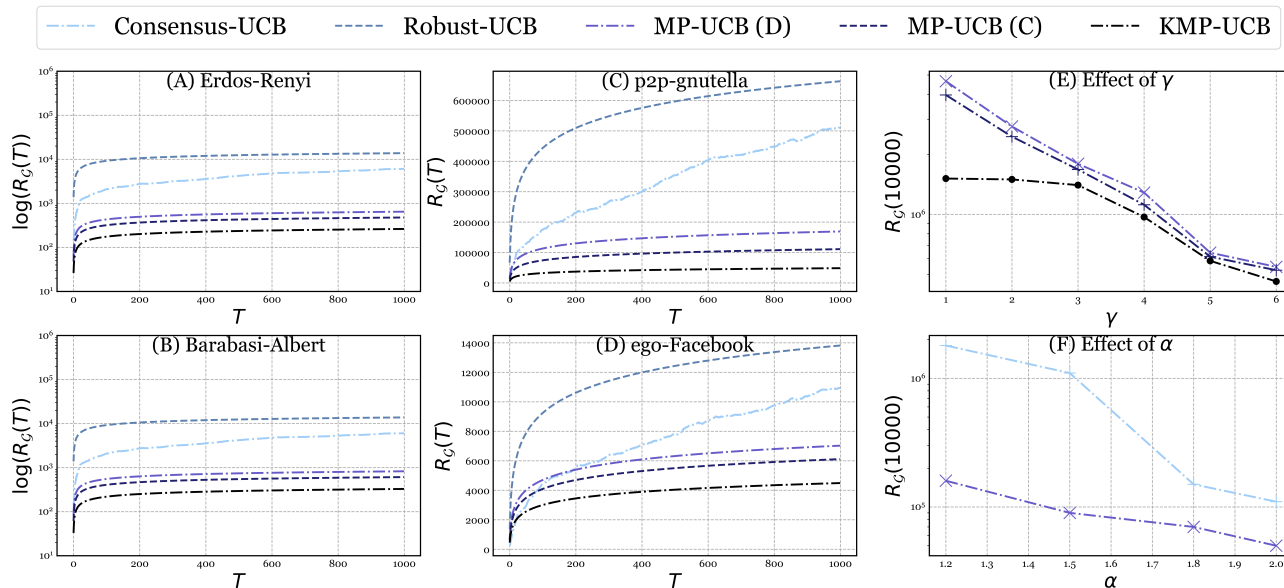


Figure 1. Experimental benchmarks, where each experiment is averaged over 100 trials. Figures (A) and (B) compare performance on samples of random graphs; (C) and (D) compare performance on two classes of real-world networks, and (E) and (F) are ablations.

Gnutella04 (Figure 1C) and ego-Facebook (Figure 1D) network structures from the SNAP repository (Leskovec & Sosič, 2016) to experiment with in the real-world setting. For both experiments, we sample subgraphs of 500 nodes, and use these subgraphs. A common misconception is to compare our distributed *multi-agent* problem with the *social network clustering* problem (Gentile et al., 2014; Li et al., 2016), which is more scalable since it is *single-agent* (i.e., one action chosen per trial). These networks are chosen because they represent two diverse situations cooperative decision-making can be applicable in – social networks and peer-to-peer communication networks. In both cases, we observe a similar trend. The gains are larger in the p2p-Gnutella case since ego-Facebook is dense (with fewer nodes), hence CONSENSUS-UCB performs better as well.

Experiment 3: Effect of γ and α . As ablation experiments, we investigate the effect of communication density γ (Figure 1E) and tail parameter α (Figure 1F) on the group regret. For both experiments, we construct random graphs on 200 agents from the Erdos-Renyi ($p = 0.7$) family. We compare the group regret at $T = 10000$ trials as a function of γ , and α , respectively. First, we observe that communication density has a significant effect on all but the KMT-UCB algorithms. Next, we see that CONSENSUS-UCB progressively gets worse as the tail gets heavier (i.e., $\alpha \rightarrow 1^+$).

8. Conclusion

In this paper, we presented a treatment of cooperative bandit estimation under heavy tails. We provided the first asymp-

otic lower bound on cooperative estimation that holds for arbitrary graphs \mathcal{G} and a wide variety of communication protocols. We present the first robust cooperative estimation algorithms that can all provide optimal regret, even without knowledge of \mathcal{G} . We support our bounds via experiments over random graphs as well. However, our work leaves several open questions in *robust* multi-agent decision-making.

First, we note that our best algorithm provides an asymptotic group regret of $O(\alpha(\mathcal{G}_\gamma)K \ln T)$, which is similar to the results obtained in the non-stochastic case (Cesa-Bianchi et al., 2019; Martínez-Rubio et al., 2019). The $\alpha(\mathcal{G}_\gamma)$ overhead can be attributed to the fact that information does not flow completely through the network (cf. Assumption 2a). This leads us to believe that tighter lower bounds can be obtained by taking this aspect of the communication protocol into account. Moreover, in realistic settings, messages incur stochasticity, i.e. they can be dropped at random, or propagate with varying delay γ . This line of work has been studied in the single-agent setting (Pike-Burke et al., 2017; Vernade et al., 2018), however the problem becomes more challenging when multiple agents interact simultaneously.

The extension of our setting to the contextual case is not trivial. Robust single-agent estimation for linear bandits is a difficult problem from both the algorithmic and computational point of view, since statistically optimal multivariate estimators require exponential time to compute (Lugosi & Mendelson, 2019). Furthermore, delay creates a $\sqrt{\gamma}$ scaling of the regret (Neu et al., 2010), which is amplified in the multi-agent setting. Addressing such scenarios is a difficult but crucial next step in this line of research.

References

- Agrawal, S. and Goyal, N. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pp. 39–1, 2012.
- Anantharam, V., Varaiya, P., and Walrand, J. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part i: iid rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976, 1987.
- Aysal, T. C. and Barner, K. E. Convergence of consensus models with stochastic disturbances. *IEEE Transactions on Information Theory*, 56(8):4101–4113, 2010.
- Bar-On, Y. and Mansour, Y. Individual regret in cooperative nonstochastic multi-armed bandits. *arXiv preprint arXiv:1907.03346*, 2019.
- Barabasi, A.-L. The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207, 2005.
- Bistritz, I. and Leshem, A. Distributed multi-player bandits—a game of thrones approach. In *Advances in Neural Information Processing Systems*, pp. 7222–7232, 2018.
- Borak, S., Härdle, W., and Weron, R. Stable distributions. In *Statistical tools for finance and insurance*, pp. 21–44. Springer, 2005.
- Brânzei, S. and Peres, Y. Multiplayer bandit learning, from competition to cooperation. *arXiv preprint arXiv:1908.01135*, 2019.
- Bubeck, S., Cesa-Bianchi, N., and Lugosi, G. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.
- Bubeck, S., Li, Y., Peres, Y., and Sellke, M. Non-stochastic multi-player multi-armed bandits: Optimal rate with collision information, sublinear without. *arXiv preprint arXiv:1904.12233*, 2019.
- Catoni, O. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l’IHP Probabilités et statistiques*, volume 48, pp. 1148–1185, 2012.
- Cesa-Bianchi, N., Gentile, C., and Zappella, G. A gang of bandits. In *Advances in Neural Information Processing Systems*, pp. 737–745, 2013.
- Cesa-Bianchi, N., Gentile, C., and Mansour, Y. Delay and cooperation in nonstochastic bandits. *The Journal of Machine Learning Research*, 20(1):613–650, 2019.
- Crovella, M. E., Taqqu, M. S., and Bestavros, A. Heavy-tailed probability distributions in the world wide web. *A practical guide to heavy tails*, 1:3–26, 1998.
- De Vany, A. S., Lee, C., et al. *Information cascades in multi-agent models*. University of California, Irvine, Department of Economics, 1999.
- Dubey, A. and Pentland, A. S. Thompson sampling on symmetric alpha-stable bandits. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pp. 5715–5721. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/792. URL <https://doi.org/10.24963/ijcai.2019/792>.
- Eom, Y.-H. and Jo, H.-H. Tail-scope: Using friends to estimate heavy tails of degree distributions in large-scale complex networks. *Scientific reports*, 5:09752, 2015.
- Gentile, C., Li, S., and Zappella, G. Online clustering of bandits. In *International Conference on Machine Learning*, pp. 757–765, 2014.
- Gentile, C., Li, S., Kar, P., Karatzoglou, A., Zappella, G., and Etrue, E. On context-dependent clustering of bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1253–1262. JMLR.org, 2017.
- Glover, F. and Kochenberger, G. A tutorial on formulating qubo models. *arXiv preprint arXiv:1811.11538*, 2018.
- Hernandez-Campos, F., Marron, J., Samorodnitsky, G., and Smith, F. D. Variable heavy tails in internet traffic. *Performance Evaluation*, 58(2-3):261–284, 2004.
- Kononov, A. Information cascades and experts institutions. *New Institutional Economics: Research Methods and Research Methods and Tools*, pp. 66, 2010.
- Korda, N., Kaufmann, E., and Munos, R. Thompson sampling for 1-dimensional exponential family bandits. In *Advances in Neural Information Processing Systems*, pp. 1448–1456, 2013.
- Korda, N., Szörényi, B., and Shuai, L. Distributed clustering of linear bandits in peer to peer networks. In *Journal of machine learning research workshop and conference proceedings*, volume 48, pp. 1301–1309. International Machine Learning Society, 2016.
- Landgren, P., Srivastava, V., and Leonard, N. E. On distributed cooperative decision-making in multiarmed bandits. In *2016 European Control Conference (ECC)*, pp. 243–248. IEEE, 2016a.
- Landgren, P., Srivastava, V., and Leonard, N. E. Distributed cooperative decision-making in multiarmed bandits: Frequentist and bayesian algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 167–172. IEEE, 2016b.

- Landgren, P., Srivastava, V., and Leonard, N. E. Social imitation in cooperative multiarmed bandits: Partition-based algorithms with strictly local information. In *2018 IEEE Conference on Decision and Control (CDC)*, pp. 5239–5244. IEEE, 2018.
- Leskovec, J. and Sosič, R. Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.
- Lévy, P. Calcul des probabilités, vol. 9. *Gauthier-Villars Paris*, 1925.
- Li, S., Karatzoglou, A., and Gentile, C. Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 539–548, 2016.
- Linial, N. Locality in distributed graph algorithms. *SIAM Journal on computing*, 21(1):193–201, 1992.
- Liu, K. and Zhao, Q. Decentralized multi-armed bandit with multiple distributed players. In *2010 Information Theory and Applications Workshop (ITA)*, pp. 1–10. IEEE, 2010a.
- Liu, K. and Zhao, Q. Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681, 2010b.
- Liu, K. and Zhao, Q. Distributed learning in cognitive radio networks: Multi-armed bandit with distributed multiple players. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3010–3013. IEEE, 2010c.
- Lucas, A. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.
- Lugosi, G. and Mendelson, S. Robust multivariate mean estimation: the optimality of trimmed mean. *arXiv preprint arXiv:1907.11391*, 2019.
- Martínez-Rubio, D., Kanade, V., and Rebeschini, P. Decentralized cooperative stochastic multi-armed bandits. In *Advances in Neural Information Processing Systems*, 2019.
- McCulloch, C. E. and Neuhaus, J. M. Prediction of random effects in linear and generalized linear models under model misspecification. *Biometrics*, 67(1):270–279, 2011.
- Medina, A. M. and Yang, S. No-regret algorithms for heavy-tailed linear bandits. In *International Conference on Machine Learning*, pp. 1642–1650, 2016.
- Moallemi, C. C. *A message-passing paradigm for optimization*, volume 68. 2007.
- Neu, G., Antos, A., György, A., and Szepesvári, C. Online markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems*, pp. 1804–1812, 2010.
- Pike-Burke, C., Agrawal, S., Szepesvári, C., and Grunewalder, S. Bandits with delayed, aggregated anonymous feedback. *arXiv preprint arXiv:1709.06853*, 2017.
- Shahrampour, S., Rakhlin, A., and Jadbabaie, A. Multi-armed bandits in multi-agent networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2786–2790. IEEE, 2017.
- Shao, H., Yu, X., King, I., and Lyu, M. R. Almost optimal algorithms for linear stochastic bandits with heavy-tailed payoffs. In *Advances in Neural Information Processing Systems*, pp. 8430–8439, 2018.
- Suomela, J. Survey of local algorithms. 45(2):24, 2013.
- Thadakamalla, H., Raghavan, U. N., Kumara, S., and Albert, R. Survivability of multiagent-based supply networks: a topological perspective. *IEEE Intelligent Systems*, 19(5): 24–31, 2004.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Vakili, S., Liu, K., and Zhao, Q. Deterministic sequencing of exploration and exploitation for multi-armed bandit problems. *IEEE Journal of Selected Topics in Signal Processing*, 7(5):759–767, 2013.
- Vernade, C., Carpentier, A., Zappella, G., Ermis, B., and Brueckner, M. Contextual bandits under delayed feedback. *arXiv preprint arXiv:1807.02089*, 2018.
- Xia, Y., Qin, T., Ma, W., Yu, N., and Liu, T.-Y. Budgeted multi-armed bandits with multiple plays.
- Yu, X., Shao, H., Lyu, M. R., and King, I. Pure exploration of multi-armed bandits with heavy-tailed payoffs. In *Conference on Uncertainty in Artificial Intelligence*, pp. 937–946, 2018.