

Heterogeneous Explore-Exploit Strategies on Multi-Star Networks

Udari Madhushani^{ID}, *Graduate Student Member, IEEE*, and Naomi Ehrich Leonard^{ID}, *Fellow, IEEE*

Abstract—We investigate the benefits of heterogeneity in multi-agent explore-exploit decision making where the goal of the agents is to maximize cumulative group reward. To do so we study a class of distributed stochastic bandit problems in which agents communicate over a multi-star network and make sequential choices among options in the same uncertain environment. Typically, in multi-agent bandit problems, agents use homogeneous decision-making strategies. However, group performance can be improved by incorporating heterogeneity into the choices agents make, especially when the network graph is irregular, i.e., when agents have different numbers of neighbors. We design and analyze new heterogeneous explore-exploit strategies, using the multi-star as the model irregular network graph. The key idea is to enable center agents to do more exploring than they would do using the homogeneous strategy, as a means of providing more useful data to the peripheral agents. In the case all agents broadcast their reward values and choices to their neighbors with the same probability, we provide theoretical guarantees that group performance improves under the proposed heterogeneous strategies as compared to under homogeneous strategies. We use numerical simulations to illustrate our results and to validate our theoretical bounds.

Index Terms—Bandit algorithms, distributed learning, heterogeneous strategies.

I. INTRODUCTION

THE INFLUENCE of agent heterogeneity on cooperation in social learning has been a recent focus of research in many fields, including ecology, sociology, and decision theory [1]. Studies on evolutionary human behavior provide evidence that individual differences can be leveraged to enhance collective prosperity [2]. Motivated by applications such as social foraging and multi-robot coordination tasks, we study and design cooperative strategies for a group of agents making sequential explore-exploit decisions in an uncertain environment. The strategies we design incorporate

agent heterogeneity to optimize the performance of the group through collective learning.

Consider a group of agents, each making a sequence of choices among options in an uncertain environment in order to maximize collective payoff. At each time step in the sequence, each agent chooses an option depending on the knowledge it has acquired about the environment up to that time step. Maximizing payoff necessitates striking a balance between making choices that yield high immediate payoff, i.e., exploiting, and making choices that yield high information content and possibly high future payoffs, i.e., exploring. When an agent fails to acquire sufficient information about the environment to make optimal decisions, it must sacrifice exploitation potential in order to explore. However, in the group setting, agents can recover exploitation potential by gaining information through cooperation, i.e., through collective learning.

Sequential decision making in uncertain environments that requires trading off exploitation and exploration is modeled mathematically by the bandit framework [3]. In the multi-armed bandit (MAB) problem, an agent is repeatedly faced with the task of choosing an option from a given set of options. At each time step the agent receives a stochastic reward drawn from a fixed probability distribution associated with the chosen option. The agent's goal is to maximize the cumulative reward by the end of the decision-making process. This requires choosing frequently enough the optimal option, i.e., the option with highest expected reward. In order to meet this requirement, the agent must simultaneously choose options that are known to provide high rewards (exploit) and choose lesser known options (explore) that might potentially provide even higher rewards [4], [5].

Maximizing cumulative reward is equivalent to minimizing cumulative regret, defined as the loss incurred by an agent choosing a sub-optimal option instead of the optimal option. Since the probability distribution associated with each option is fixed, cumulative regret can be minimized by reducing the number of times sub-optimal options are chosen. Performance of the proposed algorithms for this problem is measured using expected cumulative regret. The letter [4] establishes that any efficient policy chooses suboptimal options asymptotically logarithmically in time. The letter [5] proposes an Upper Confidence Bound (UCB) based sampling rule that achieves a logarithmic expected cumulative regret uniformly in time.

The papers [6]–[16] extend to the multi-agent setting and capture different aspects of collective learning. In [6]–[9], agents share their estimates of the expected reward of options

Manuscript received September 1, 2020; revised November 2, 2020; accepted November 17, 2020. Date of publication December 3, 2020; date of current version December 23, 2020. This research was supported in part by ONR Grant N00014-18-1-2873 and Grant N00014-19-1-2556, and in part by the ARO under Grant W911NF-18-1-0325. Recommended by Senior Editor M. Arcak. (*Corresponding author: Udari Madhushani.*)

The authors are with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: udarim@princeton.edu; naomi@princeton.edu).

Digital Object Identifier 10.1109/LCSYS.2020.3042459

2475-1456 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.

with neighbors according to fixed communication structures. The papers [6], [7] use a running consensus algorithm to update estimates and provide graph-structure-dependent performance measures that predict the relative performance of agents and networks. The letter [7] also addresses the case of a constrained reward model in which agents that choose the same option at the same time step receive no reward. The letter [8] proposes an accelerated consensus procedure assuming agents know the spectral gap of the communication graph and designs a decentralized UCB algorithm based on delayed rewards. The letter [9] considers a P2P communication where an agent is only allowed to communicate with two other agents at each time step.

The papers [10]–[16] consider the case in which agents share reward values and choices with neighbors. In [10]–[12], agents use stochastic communication structures that depend on the decision-making process. In [10], each agent observes rewards and actions of its neighbors when it is exploring. In [11], each agent instead broadcasts its rewards and actions to its neighbors when it is exploring. In [12], at each time step, agents decide either to sample an option or to broadcast the last obtained reward to the entire group.

The setup in our earlier paper [13] is closest to that in the present paper: agents observe reward values and actions of their neighbors defined by a network graph that changes in time according to probabilistic edge weights. An underlying fixed network graph is given, and each agent k observes its neighbors with probability p_k . The communication structure is independent of the decision-making process.

The papers [6]–[13] consider homogeneous protocols, whereas the papers [14]–[16] consider protocols where some agents (followers) copy actions of others (leaders). In [14], followers observe rewards and choices of their neighbors. In [15] one leader explores and estimates the mean reward of options, while all other agents choose the option with highest estimated mean per the leader. The letter [16] proposes the FYL algorithm, which uses a deterministic communication protocol and exploits degree heterogeneity of the communication network graph. FYL outperforms our algorithm when $p_k = p = 1$; however, our algorithm provides a method to exploit agent heterogeneity when agents share information with probability $0 < p < 1$.

When communication among agents is defined by an irregular network graph, e.g., some agents serve as information hubs, group performance can be improved by using heterogeneous explore-exploit strategies. To understand this, consider an environment with unconstrained resources. Then, agents can only influence the decisions of one another through the information they share, and the structure of interactions that defines neighbors, i.e., who is sharing information with whom, strongly affects the quality and quantity of information received by each individual.

We consider the case that all agents broadcast their instantaneous rewards and actions to their neighbors with probability p . This communication protocol is motivated by real-world applications in which estimates of mean rewards or the sum of collected rewards, which rely on the history of choices and rewards, are deliberately not disclosed to protect privacy [17]. For example, in user targeted recommender systems [18] (or clinical trials [19]), sharing user (patient) history of choices can reveal sensitive information about users (patients). Even

when an agent is broadcasting only its current rewards and actions to neighbors, an adversarial agent can listen to the broadcasts and access the history of choices made by the agent. To reduce such privacy leakage we consider agents that broadcast instantaneous rewards and actions probabilistically. Further, if communication failures are possible, then having agents broadcast only current rewards and actions avoids problems associated with agents losing track of what information has and has not been received by neighbors. In this context, $1 - p$ represents the probability of communication failure.

In irregular and centralized networks like the multi-star, center agents have more neighbors and thus receive more information than peripheral agents. This leads to an imbalanced exploitation potential across the group [1], [20], and group performance degrades with increasing number of peripheral agents. We investigate improving group performance by leveraging heterogeneity in the exploitation potential of agents. To do so we propose heterogeneous explore-exploit strategies that require center agents to explore more and thus increase the exploitation potential of peripheral agents.

The multi-star network models recommender systems, where there are many small servers, assigned to different regions, that each make sequential recommendations based on user feedback and communicate only with a large central server. Performance can be improved by using the central server to suggest more exploratory recommendations which allows the system to gather more information about user preferences. Probabilistic communication accounts for random communication failures between servers.

This letter is organized as follows. In Section II we provide the problem formulation and notation. Section III presents the proposed algorithm and intuition. We analyze performance of the proposed algorithm in Section IV and provide improved theoretical bounds for the expected cumulative group regret. In Section V we show numerical simulations to illustrate and validate the theoretical results. We conclude in Section VI.

II. PROBLEM FORMULATION

In this section we present the problem formulation and relevant mathematical notations. Consider a group of K agents, each faced with the same N -armed bandit problem for T time steps. At each time step $t \in \{1, \dots, T\}$, each agent chooses an option and receives a stochastic reward associated with the chosen option. Let X_i be a sub-Gaussian random variable that denotes the *reward associated with option* $i \in \{1, \dots, N\}$. Sub-Gaussian rewards include widely used distributions such as Bernoulli, Gaussian, and bounded rewards. Define $\mu_i = \mathbb{E}(X_i)$ and σ_i^2 as the *expected reward* and *variance proxy* associated with option i , respectively. Let $i^* = \arg_i \max\{\mu_1, \dots, \mu_N\}$ be the *optimal option* with highest expected reward. Define $\Delta_i = \mu_{i^*} - \mu_i$ as the *expected reward gap* between option i^* and option i .

Let $G(\mathcal{V}, \mathcal{E})$ be a fixed undirected network graph that defines the structure of the interactions between agents. This captures the inherent hard communication constraints of the system. Here \mathcal{V} is a set of K vertices such that each vertex corresponds to an agent. Each edge $e(k, j) \in \mathcal{E}$ in the graph denotes that agent k and agent j are *neighbors*. At each time step, each agent broadcasts its reward value and action to its neighbors

with *broadcasting probability* p . Let $\mathbb{I}_{\{k,j\}}^t$ be the indicator random variable that takes value 1 if agent k receives information from agent j at time t and 0 otherwise. Then, for every time t , $\mathbb{E}(\mathbb{I}_{\{k,j\}}^t) = p$, $\forall k, j$ such that $e(k, j) \in \mathcal{E}$, and $\mathbb{E}(\mathbb{I}_{\{k,j\}}^t) = 0$ otherwise. We define $\mathbb{I}_{\{k,k\}}^t = 1$, $\forall k, t$.

Let d_k be the *degree* (number of neighbors) of agent k and $d_{avg} = \frac{1}{K} \sum_{k=1}^K d_k$ be the *average degree of the network*. Let d_k^{avg} be the *average degree of neighbors of agent k* : $d_k^{avg} = \frac{1}{d_k} \sum_{e(k,j) \in \mathcal{E}} d_j$.

We focus on *multi-star* graphs defined as follows. Let there be m center agents and $K - m$ peripheral agents. Without loss of generality let each agent k , $k \leq m$, be a *center agent*. All center agents are neighbors of one another, i.e., $e(k, j) \in \mathcal{E}$, $\forall k, j \leq m$, and a center agent's degree d_k is at least $m - 1$. Each *peripheral agent* k , $k > m$, has exactly one neighbor ($d_k = 1$), and the neighbor is a center agent. To reduce complexity, we assume the graph is symmetric, which implies that all center agents have the same number of neighbors. Thus $K - m$ is an integer multiple of m . If $K > 2$ and $m < K$, the multi-star graph is *irregular*, i.e., the degree of center agents differs from the degree of peripheral agents. Let d_{cen} be the degree of each center agent. Then, $d_{cen} = \frac{K-m}{m} + m - 1$. When $m = 1$ the graph is a star, the most irregular multi-star graph. When $m = K$, there are no peripheral agents and the graph is all-to-all and thus regular.

Let φ_t^k be a random variable that denotes the *option chosen by agent $k \in \{1, \dots, K\}$ at time $t \in \{1, \dots, T\}$* . Let $\mathbb{I}_{\{\varphi_t^k=i\}}$ be an indicator random variable that takes value 1 if agent k chooses option i at time t and 0 otherwise. Let $n_i^k(t)$ be the *total number of times agent k chooses option i until time t* and let $N_i^k(t)$ be the *total number of times agent k observes option i until time t* . The total number of observations is the sum of the number of samples taken from option i by agent k and the number of broadcasts on option i by its neighbors:

$$n_i^k(t) = \sum_{\tau=1}^t \mathbb{I}_{\{\varphi_\tau^k=i\}}, \quad N_i^k(t) = \sum_{\tau=1}^t \sum_{j=1}^K \mathbb{I}_{\{\varphi_\tau^j=i\}} \mathbb{I}_{\{k,j\}}^\tau. \quad (1)$$

Let $\widehat{\mu}_i^k(t)$ denote the estimate of expected reward of agent k for option i at time t . Then, $\widehat{\mu}_i^k(t) = \frac{S_i^k(t)}{N_i^k(t)}$, where $S_i^k(t) = \sum_{\tau=1}^t \sum_{j=1}^K X_i \mathbb{I}_{\{\varphi_\tau^j=i\}} \mathbb{I}_{\{k,j\}}^\tau$.

Expected regret is defined as the expected loss suffered by agents by sampling sub-optimal options. Let $R(t)$ be the *cumulative group regret* at time t . Then *expected cumulative group regret* can be computed as

$$\mathbb{E}(R(t)) = \sum_{i=1}^N \sum_{k=1}^K \Delta_i \mathbb{E}(n_i^k(t)). \quad (2)$$

III. ALGORITHM

To realize the goal of maximizing cumulative group reward, agents should minimize the number of times they sample sub-optimal options. Each agent employs an agent-based strategy that captures the trade-off between exploring and exploiting by constructing an objective function that strikes a balance between the estimation of the expected reward and the uncertainty associated with the estimate [5].

Since center agents have more neighbors they are more likely to obtain a high number of observations. This reduces

the uncertainty associated with their estimate of the expected reward of options. Thus, identifying the optimal option requires less exploring, which increases their exploitation potential. Since peripheral agents only have one neighbor they are more likely to obtain a low number of observations. Thus, identifying the optimal option requires more exploring, which decreases their exploitation potential. Further, since center agents do less exploring, the usefulness of the information they broadcast is reduced, also decreasing the peripheral agents' exploitation potential. Accordingly, homogeneous sampling rules in irregular, multi-star networks lead to imbalanced exploitation potential across the group and thus degraded group performance.

To improve group performance, we propose heterogeneous explore-exploit strategies that regulate exploitation potential across the network. When center agents are more exploratory their performance degrades, but the usefulness of the information they broadcast increases and so the performance of peripheral agents improves. When there are more peripheral agents than center agents, and broadcasting probability p is sufficiently high, the performance improvement obtained by peripheral agents outweighs the performance degradation incurred by center agents, and group performance increases. If p is too small, for example, when broadcasting is costly or risky, center agents do not broadcast enough information to benefit peripheral agents. Thus it doesn't pay for center agents to increase their exploration. Indeed, when $p = 0$ all agents have the same exploitation potential.

Using this intuition, we propose the following heterogeneous sampling rules. Assume that variance proxy σ_i^2 for each option i is known to all agents.

Definition 1 (Heterogeneous Sampling Rules): The sampling rule $\{\varphi_t^k\}_1^T$ of agent k at time $t \in \{1, \dots, T\}$ is

$$\mathbb{I}_{\{\varphi_{t+1}^k=i\}} = \begin{cases} 1, & i = \arg \max\{Q_1^k(t), \dots, Q_N^k(t)\} \\ 0, & \text{o.w.} \end{cases}$$

with

$$Q_i^k(t) = \widehat{\mu}_i^k(t) + C_i^k(t) \quad (3)$$

$$C_i^k(t) = \sigma_i \sqrt{\frac{2(1 + \alpha_k)(\xi + 1) \log t}{N_i^k(t)}} \quad (4)$$

where $\xi > 1$ and

$$\alpha_k = \begin{cases} \frac{p^{1-p}(d_k - d_k^{avg})}{d_k}, & k \leq m \\ 0, & k > m. \end{cases} \quad (5)$$

$C_i^k(t)$ in (4) represents *agent k 's uncertainty in its estimated mean of option i* , and Definition 1 implies that for any agent k , when $C_i^k(t)$ is high, agent k will more likely explore. By (4), $C_i^k(t)$ can be high when $N_i^k(t)$, the number of agent k 's observations of option i , is low, i.e., when option i is under-sampled. $C_i^k(t)$ can also be high when *agent k 's exploration bias $\alpha_k > 0$* is high.

By (5), $\alpha_k \neq 0$ only for center agents. Since peripheral agents have one center agent neighbor, $d_k^{avg} \leq d_k$ and thus $\alpha_k \geq 0$ for every center agent $k \leq m$. In fact, $\alpha_k \geq 0$ is designed to grow with increasing irregularity: in the regular case (all-to-all) when $m = K$, $\alpha_k = 0$, and in the most irregular case (star) when $m = 1$, $d_1 = K - 1$ and $d_1^{avg} = 1$ so $(d_1 - d_1^{avg})/d_1 = (K - 2)/(K - 1)$. Further, α_k grows

with p according to the factor p^{1-p} , which grows rapidly for intermediate values of p and is large (i.e., saturates to 1) only when center agents are broadcasting their reward values and actions with sufficiently high probability p .

Definition 2: To get the corresponding homogeneous sampling rules let $\alpha_k = 0, \forall k$, in Definition 1. Heterogeneous and homogeneous rules for peripheral agents are the same.

By design, the heterogeneous rules of Definition 1 drive center agents to explore more than the corresponding homogeneous rules and only when it benefits group performance.

IV. PERFORMANCE ANALYSIS

In this section we analyze the performance of the heterogeneous sampling rules of Definition 1. Using an approach similar to [5] with a few key modifications, we upper bound the expected cumulative group regret $\mathbb{E}(R(T))$. We show that the bound is lower than the upper bound in the case of the corresponding homogeneous sampling rules, and so we can conclude that the designed heterogeneous strategies provide better group performance than the homogeneous strategies.

By (2), we upper bound $\mathbb{E}(R(T))$ if we upper bound $\sum_{k=1}^K \mathbb{E}(n_i^k(T))$, where $n_i^k(T)$ is the number of times agent k samples sub-optimal option i until time T . By Definition 1, agent k chooses sub-optimal option i at time t if $Q_i^k(t) \geq Q_{i^*}^k(t)$. Then, $n_i^k(t) = \sum_{\tau=1}^t \mathbb{I}_{\{\varphi_i^k = i\}} \leq \sum_{\tau=1}^t \mathbb{I}_{\{Q_i^k(\tau) \geq Q_{i^*}^k(\tau)\}}$. For each option i and agent k let $\{\eta_i^k(t)\}_1^T$ be a sequence of nonnegative nondecreasing functions. Then,

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}(n_i^k(T)) &\leq \sum_{k=1}^K \sum_{t=1}^T \mathbb{E}(\mathbb{I}_{\{\varphi_i^k = i\}}, N_i^k(t) \leq \eta_i^k(t)) \\ &\quad + \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(Q_i^k(t) \geq Q_{i^*}^k(t), N_i^k(t) > \eta_i^k(t)). \end{aligned} \quad (6)$$

It remains to upper bound the right hand side of (6) and we do so in two steps. First, we upper bound the second summation term of (6) as follows. From (3) we have

$$\begin{aligned} \{Q_i^k(t) \geq Q_{i^*}^k(t)\} &\subseteq \{\mu_{i^*} < \mu_i + 2C_i^k(t)\} \\ &\cup \{\widehat{\mu}_{i^*}^k(t) \leq \mu_{i^*} - C_{i^*}^k(t)\} \cup \{\widehat{\mu}_i^k(t) \geq \mu_i + C_i^k(t)\}. \end{aligned} \quad (7)$$

For all k let

$$\eta_i^k(t) = (1 + \alpha_k)\eta_i(t), \quad \eta_i(t) = \frac{8\sigma_i^2(\xi + 1)\log t}{\Delta_i^2}. \quad (8)$$

Then, by (4), $\{\mu_{i^*} < \mu_i + 2C_i^k(t)\} \cap \{N_i^k(t) > \eta_i^k(t)\} = \emptyset$ where \emptyset is the empty set. Using (7) we obtain

$$\begin{aligned} \mathbb{P}(Q_i^k(t) \geq Q_{i^*}^k(t), N_i^k(t) > \eta_i^k(t)) \\ \leq \mathbb{P}(\widehat{\mu}_{i^*}^k(t) \leq \mu_{i^*} - C_{i^*}^k(t)) + \mathbb{P}(\widehat{\mu}_i^k(t) \geq \mu_i + C_i^k(t)). \end{aligned} \quad (9)$$

To upper bound the right hand side of (9) we use the tail probability bound provided in the following lemma.

Lemma 1: For any $\xi > 1$, some $\zeta > 1$ and for $\sigma_i > 0$ in the uncertainty $C_i^k(t)$ given by (4), we get

$$\mathbb{P}(|\widehat{\mu}_i^k(t) - \mu_i| > C_i^k(t)) \leq \frac{1}{\log \zeta} \frac{\log((1 + d_k)t)}{t^{(\xi+1)(1+\alpha_k)}}.$$

Proof: From Theorem 1 in this letter [13] we have for some $\zeta > 1$ and for $\sigma_i > 0$ there exists a $\vartheta_k > 0$ such that

$$\mathbb{P}\left(\widehat{\mu}_i^k(t) - \mu_i > \sqrt{\frac{\vartheta_k}{N_i^k(t)}}\right) \leq \frac{\nu \log((d_k + 1)T)}{\exp(2\kappa \vartheta_k)}$$

where, $\nu = \frac{1}{\log \zeta}$, $\kappa = \frac{1}{\sigma_i^2(\zeta^{\frac{1}{4}} + \zeta^{-\frac{1}{4}})^2}$. Since $\alpha_k \geq 0, \forall k$, we can use $\vartheta_k = 2\sigma_i^2(1 + \alpha_k)(\xi + 1)\log t$ to get the statement of the lemma. ■

Using the statement of Lemma 1 in (9),

$$\mathbb{P}(Q_i^k(t) \geq Q_{i^*}^k(t), N_i^k(t) > \eta_i^k(t)) \leq \frac{2}{\log \zeta} \frac{\log((1 + d_k)t)}{t^{(\xi+1)(1+\alpha_k)}}. \quad (10)$$

Summing the right hand side of (10) over t we get

$$\begin{aligned} \sum_{t=1}^T \frac{\log((1 + d_k)t)}{t^{(\xi+1)(1+\alpha_k)}} &\leq \log(1 + d_k) \\ &\quad + \frac{\log(1 + d_k)(\xi \alpha_k + \xi + \alpha_k) + 1}{(\xi \alpha_k + \xi + \alpha_k)^2}. \end{aligned} \quad (11)$$

Since log is concave, substituting (11) into (10) we get

$$\begin{aligned} \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(Q_i^k(t) \geq Q_{i^*}^k(t), N_i^k(t) > \eta_i^k(t)) \\ \leq \frac{2K}{\log \zeta} \log(1 + d_{avg}) \\ + \frac{2}{\log \zeta} \sum_{k=1}^K \frac{\log(1 + d_k)(\xi \alpha_k + \xi + \alpha_k) + 1}{(\xi \alpha_k + \xi + \alpha_k)^2}, \end{aligned} \quad (12)$$

which upper bounds the second summation of (6).

Next, we upper bound the first summation term of (6) as follows. Since we restrict to symmetric graphs where all center agents have the same number and type of neighbors, $\alpha_k = \alpha, \forall k \leq m$. Then, by (8) we have $\eta_i^k(t) = (1 + \alpha)\eta_i(t), \forall k \leq m$, and $\eta_i^k(t) = \eta_i(t), \forall k > m$. Let $[x]^+ = \max\{x, 0\}$.

Lemma 2: Let G be a symmetric multi-star graph with m center agents and $K - m$ peripheral agents. Let $\{\eta_i^k(t)\}_1^T$ be the sequence of nonnegative nondecreasing functions given by (8). Then with some high probability $1 - \delta(p)$

$$\begin{aligned} \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\mathbb{I}_{\{\varphi_i^k = i\}} = 1, N_i^k(t) \leq \eta_i^k(t)) &\leq (K - m)\eta_i(T) \\ &\quad + \frac{m}{1 + p(m - 1)} \left[1 - p \frac{K - m}{m}\right]^+ (1 + \alpha)\eta_i(T). \end{aligned}$$

Proof: Recall the definitions of $n_i^k(t)$ and $N_i^k(t)$ in (1). Since the communication structure is independent of the decision-making process $\forall k$,

$$\mathbb{E}(n_i^k(t)) + p \sum_{e(k,j) \in \mathcal{E}} \mathbb{E}(n_i^j(t)) = \mathbb{E}(N_i^k(t)). \quad (13)$$

Since $N_i^k(t)$ is a nonnegative random variable, $N_i^k(t) \leq \eta_i^k(t) \implies \mathbb{E}(N_i^k(t)) \leq \eta_i^k(t)$. Thus, from (13), for all k ,

$$\begin{aligned} \mathbb{E}(n_i^k(t), N_i^k(t) \leq \eta_i^k(t)) \\ + p \sum_{e(k,j) \in \mathcal{E}} \mathbb{E}(n_i^j(t), N_i^k(t) \leq \eta_i^k(t)) \leq \eta_i^k(t). \end{aligned} \quad (14)$$

To upper bound $\sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\mathbb{I}_{\{\varphi_i^k=i\}} = 1, N_i^k(t) \leq \eta_i^k(t))$ we maximize $\sum_{k=1}^K \mathbb{E}(n_i^k(t))$ subject to the constraint given by (14). This is the linear programming optimization problem: maximize $\sum_{k=1}^K \mathbb{E}(n_i^k(t), N_i^k(t) \leq \eta_i^k(t))$ subject to (14) and $\mathbb{E}(n_i^k(t), N_i^k(t) \leq \eta_i^k(t)) \geq 0$ for all k . For $p = 1$ the solution is the sum of $\eta_i^k(t)$ over the maximal independent set of G , which for a multi-star graph is the set of peripheral agents $k \geq m + 1$. Thus, for general p we have

$$\begin{aligned} \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\mathbb{I}_{\{\varphi_i^k=i\}} = 1, N_i^k(t) \leq \eta_i^k(t)) &\leq \sum_{k=m+1}^K \eta_i^k(T) \\ &+ \sum_{k=1}^m \frac{1}{1+p(m-1)} \left[1 - p \frac{K-m}{m} \right]^+ \eta_i^k(T), \end{aligned}$$

and the statement of the lemma follows. \blacksquare

This concludes upper bounding the first summation of (6).

Theorem 1: Consider a distributed stochastic bandit problem with N options, K agents, and T time steps. Let communication graph G be a symmetric multi-star graph with m center agents and $K - m$ peripheral agents. If all agents sample according to the heterogeneous sampling rules defined in Definition 1, with some high probability $1 - \delta(p)$, the expected cumulative group regret satisfies

$$\begin{aligned} \mathbb{E}(R(T)) &\leq c_1(K, m, \alpha, p) \sum_{i=1}^N \frac{8\sigma_i^2(\xi + 1) \log T}{\Delta_i} \\ &+ \frac{2}{\log \zeta} \sum_{i=1}^N \Delta_i \left(K \log(1 + d_{avg}) + (K - m) \frac{\xi \log 2 + 1}{\xi^2} \right. \\ &\quad \left. + m \frac{\log(1 + d_{cen})(\xi \alpha + \xi + \alpha) + 1}{(\xi \alpha + \xi + \alpha)^2} \right), \\ c_1(K, m, \alpha, p) &= K - m + \frac{m(1 + \alpha)}{1 + p(m - 1)} \left[1 - p \frac{K - m}{m} \right]^+, \end{aligned}$$

where Δ_i is the expected reward gap between the options i^* and i , σ_i^2 is the variance proxy, and $\xi, \zeta > 1$.

Proof: Result follows from (2), (6), (12) and Lemma 2. \blacksquare

Remark 1: Recall that under the corresponding homogeneous sampling rules we have $\alpha_k = 0, \forall k$. Thus, we can recover the expected cumulative group regret bound for the homogeneous sampling rules as follows:

$$\begin{aligned} \mathbb{E}(R(T)) &\leq c_2(K, m, p) \sum_{i=1}^N \frac{8\sigma_i^2(\xi + 1) \log T}{\Delta_i} \\ &+ \frac{2}{\log \zeta} \sum_{i=1}^N \Delta_i \left(K \log(1 + d_{avg}) + (K - m) \frac{\xi \log 2 + 1}{\xi^2} \right. \\ &\quad \left. + m \frac{\log(1 + d_{cen})\xi + 1}{\xi^2} \right), \end{aligned}$$

$$c_2(K, m, p) = K - m + \frac{m}{1+p(m-1)} [1 - p \frac{K-m}{m}]^+.$$

When the network graph has a large enough ratio of peripheral agents to center agents and a sufficiently high broadcasting probability p , i.e., $p(K-m)/m > 1$, we have $[1 - p \frac{K-m}{m}]^+ = 0$, which implies $c_1 = c_2 = K - m$. And since $\alpha > 0$ we have

$$\frac{\log(1 + d_{cen})(\xi \alpha + \xi + \alpha) + 1}{(\xi \alpha + \xi + \alpha)^2} < \frac{\log(1 + d_{cen})\xi + 1}{\xi^2}.$$

Plugging these results into the bounds of Theorem 1 and Remark 1, we see that the heterogeneous sampling rules provide a lower theoretical regret bound than the corresponding homogeneous sampling rules, which implies that the heterogeneous sampling rules provide better group performance than the homogeneous sampling rules.

Remark 2: Our bounds hold for sub-exponential reward distributions, where X_i is a sub-exponential random variable with mean μ_i and parameters (σ_i^2, b) with $b \leq \frac{\sigma_i}{2\sqrt{2(\xi+1)\log T}}$.

V. SIMULATION RESULTS

In this section we provide numerical simulations to illustrate results and validate theoretical bounds. For all simulations, we consider 10 options ($N = 10$) with Gaussian reward distributions. Expected reward for the optimal option is $\mu_{i^*} = 11$ and for all sub-optimal options $i \neq i^*$ is $\mu_i = 10$. We let variance associated with all options i be $\sigma_i^2 = 1$. Because the expected reward gaps $\Delta_i = 1, i \neq i^*$, are equal to the variances $\sigma_i^2 = 1$, it is a challenging problem to distinguish the optimal option from the sub-optimal options. For all simulations, we consider 1000 time steps ($T = 1000$) and use 1000 Monte Carlo simulations with $\xi = 1.01$.

We show simulation results for performance of a group of $K = 36$ agents that communicate over two different symmetric multi-star graphs and use the heterogeneous sampling rules of Definition 1. We compare to the case when agents use the corresponding homogeneous sampling rules of Definition 2. The first multi-star graph has $m = 2$ center agents and $K - m = 34$ peripheral agents, with each center agent communicating with 17 peripheral agents and the other center agent. The second multi-star graph has $m = 3$ center agents and $K - m = 33$ peripheral agents, with each center agent communicating with 11 peripheral agents and the other center agents. In each case, center agents are interchangeable and peripheral agents are interchangeable, so the average performance of a center (peripheral) agent is the same as the individual performance of a center (peripheral) agent.

Figure 1 shows how average expected cumulative group regret varies with broadcasting probability p for agents using the heterogeneous rules (dotted) and homogeneous rules (solid). Regret is inversely related to performance: lower group regret implies higher group performance. Results are plotted on the left for the graph with 2 center agents and on the right for the graph with 3 center agents. When $p = 0$ there is no communication at all. So when p becomes even just a little positive and agents learn about options from their neighbors, regret falls, i.e., group performance rises.

In the case of the homogeneous rules, as p increases through intermediate values, center agents do less and less exploring and the usefulness of the information received by peripheral agents decreases. This leads to increased regret for peripheral agents, and the group overall, and thus degraded group performance. When p approaches 1, center agents receive sufficient information from their peripheral neighbors such that their improved performance outweighs the degraded performance of peripheral agents. This leads to a final decrease in group regret and increase in group performance.

The improvement in performance provided by the heterogeneous rules relative to the homogeneous rules, as predicted by Theorem 1 and Remark 1, can be clearly seen in Figure 1




Fig. 1. Average expected cumulative group regret for $K = 36$ agents at time $t = 1000$ as a function of broadcasting probability p with communication over a symmetric multi-star graph. Left: 2 center and 34 peripheral agents. Right: 3 center and 33 peripheral agents. Dotted lines and solid line shows average regret when agents use heterogeneous and homogeneous sampling rules, respectively.

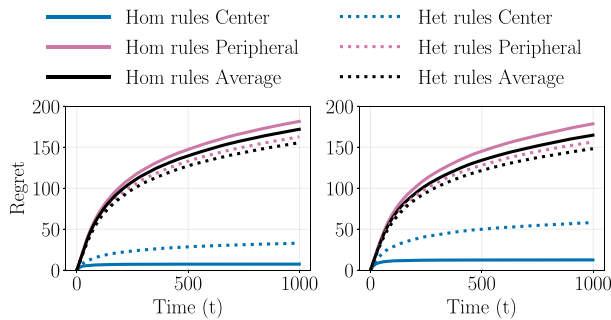


Fig. 2. Expected cumulative regret of center agent, peripheral agent, and average agent for $K = 36$ agents as a function of time t for $p = 0.8$ and the same two symmetric multi-star graphs as in Figure 1: 2 center agents (left) and 3 center agents (right) where agents use heterogeneous (dotted) and homogeneous (solid) sampling rules.

by observing how much lower the dotted regret curve is than the solid regret curve. The growth in regret in the homogeneous case, as p increases through intermediate values, is reduced in the heterogeneous case. This is because, by design, center agents are biased toward more exploring, which improves the information that peripheral agents receive. The group performance increase that comes, as p increases further, occurs in the heterogeneous case well before p approaches 1.

The influence of irregularity of the graph can be observed in Figure 1 by comparing the left plot (2 center agents and more irregular) to the right plot (3 center agents and less irregular). The results suggest that performance is higher with more center agents, i.e., with greater regularity in the graph.

Figure 2 shows expected cumulative regret as a function of time t for center (blue), peripheral (pink), and average (black) agents, when $p = 0.8$ and agents use the heterogeneous rules (dotted) and homogeneous rules (solid). Results are plotted on the left for the graph with 2 center agents and on the right for the graph with 3 center agents. It can be observed that, as predicted for the heterogeneous rules, the peripheral agent performance increases and the center agent performance decreases, such that group performance (as represented by the average agent) improves. Further, a comparison of left and right plots suggests that group performance improves with more center agents (more regularity).

VI. CONCLUSION

We have designed and analyzed new heterogeneous rules for how a group of agents that share information over a network

should sample an uncertain environment to maximize group reward. We consider communication networks defined by symmetric multi-star graphs, since these exemplify realistic settings. Using the multi-armed bandit problem as the explore-exploit framework, we show how sampling rules for center agents that favor exploring over exploiting make the information that center agents broadcast to their neighbors more useful, thereby increasing the total reward accumulated by the group.

Our analysis and design advance understanding of the role that heterogeneity does and can play in collective decision making. And our demonstration that heterogeneity can be leveraged to improve the performance of a cooperative multi-agent system suggests that further investigation is warranted.

REFERENCES

- [1] X. Li, Y. Geng, C. Shen, and L. Shi, "The influence of heterogeneous learning ability on the evolution of cooperation," *Sci. Rep.*, vol. 9, no. 1, pp. 1–7, 2019.
- [2] M. A. Amaral, L. Wardil, M. Perc, and J. K. L. da Silva, "Evolutionary mixed games in structured populations: Cooperation and the benefits of heterogeneity," *Phys. Rev. E*, vol. 93, no. 4, 2016, Art. no. 042304.
- [3] H. Robbins, "Some aspects of the sequential design of experiments," *Bull. Amer. Math. Soc.*, vol. 58, no. 5, pp. 527–535, 1952.
- [4] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Adv. Appl. Math.*, vol. 6, no. 1, pp. 4–22, Mar. 1985.
- [5] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Mach. Learn.*, vol. 47, nos. 2–3, pp. 235–256, 2002.
- [6] P. Landgren, V. Srivastava, and N. E. Leonard, "Distributed cooperative decision-making in multiarmed bandits: Frequentist and Bayesian algorithms," in *Proc. IEEE CDC*, 2016, pp. 167–172.
- [7] P. Landgren, V. Srivastava, and N. E. Leonard, "Distributed cooperative decision making in multi-agent multi-armed bandits," 2020. [Online]. Available: arXiv:2003.01312.
- [8] D. Martínez-Rubio, V. Kanade, and P. Rebeschini, "Decentralized cooperative stochastic bandits," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 4531–4542.
- [9] B. Szörényi, R. Busa-Fekete, I. Hegedűs, R. Ormándi, M. Jelasity, and B. Kégl, "Gossip-based distributed stochastic bandit algorithms," in *Proc. ICML*, 2013, pp. 19–27.
- [10] U. Madhushani and N. E. Leonard, "A dynamic observation strategy for multi-agent multi-armed bandit problem," in *Eur. Control Conf. (ECC)*, 2020, pp. 1677–1682.
- [11] U. Madhushani and N. E. Leonard, "Distributed learning: Sequential decision making in resource-constrained environments," 2020. [Online]. Available: arXiv:2004.06171.
- [12] M. Chakraborty, K. Y. P. Chua, S. Das, and B. Juba, "Coordinated versus decentralized exploration in multi-agent multi-armed bandits," in *Proc. IJCAI*, 2017, pp. 164–170.
- [13] U. Madhushani and N. E. Leonard, "Heterogeneous stochastic interactions for multiple agents in a multi-armed bandit problem," in *Proc. Eur. Control Conf. (ECC)*, 2019, pp. 3502–3507.
- [14] P. Landgren, V. Srivastava, and N. E. Leonard, "Social imitation in cooperative multiarmed bandits: Partition-based algorithms with strictly local information," in *Proc. IEEE CDC*, 2018, pp. 5239–5244.
- [15] P.-A. Wang, A. Proutiere, K. Ariu, Y. Jedra, and A. Russo, "Optimal algorithms for multiplayer multi-armed bandits," in *Proc. Int. Conf. Artif. Intell. Stat.*, 2020, pp. 4120–4129.
- [16] R. K. Kolla, K. Jagannathan, and A. Gopalan, "Collaborative learning of stochastic bandits over a social network," *IEEE/ACM Trans. Netw.*, vol. 26, no. 4, pp. 1782–1795, Aug. 2018.
- [17] R. Féraud, R. Alami, and R. Laroche, "Decentralized exploration in multi-armed bandits," 2018. [Online]. Available: arXiv:1811.07763.
- [18] R. Warlop, A. Lazaric, and J. Mary, "Fighting boredom in recommender systems with linear reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 1757–1768.
- [19] A. C. Y. Tossou and C. Dimitrakakis, "Algorithms for differentially private multi-armed bandits," in *Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 2087–2093.
- [20] C. Sueur, J.-L. Deneubourg, and O. Petit, "From social network (centralized vs. decentralized) to collective decision-making (unshared vs. shared consensus)," *PLoS one*, vol. 7, no. 2, 2012, Art. no. e32566.