



**HAL**  
open science

# Strategic Arms with Side Communication Prevail Over Low-Regret MAB Algorithms

Ahmed Ben Yahmed, Clément Calauzènes, Vianney Perchet

► **To cite this version:**

Ahmed Ben Yahmed, Clément Calauzènes, Vianney Perchet. Strategic Arms with Side Communication Prevail Over Low-Regret MAB Algorithms. ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2024, pp.7435-7439. 10.1109/ICASSP48485.2024.10446895 . hal-04669024

**HAL Id: hal-04669024**

**<https://hal.science/hal-04669024v1>**

Submitted on 29 Aug 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

# STRATEGIC ARMS WITH SIDE COMMUNICATION PREVAIL OVER LOW-REGRET MAB ALGORITHMS

Ahmed Ben Yahmed\*

Clément Calauzènes†

Vianney Perchet‡

## ABSTRACT

In the strategic multi-armed bandit setting, when arms possess perfect information about the player’s behavior, they can establish an equilibrium where: **1.** they retain almost all of their value, **2.** they leave the player with a substantial (linear) regret. This study illustrates that, even if complete information is not publicly available to all arms but is shared among them, it is possible to achieve a similar equilibrium. The primary challenge lies in designing a communication protocol that incentivizes the arms to communicate truthfully.

*Index Terms*— multi-armed bandit, strategic arms, communication through a network, Nash equilibrium.

## 1. INTRODUCTION

The concept of the strategic multi-armed bandit extends the traditional multi-armed bandit (MAB) problem by incorporating the utility aspect of the arms. In this context, arms have the ability to report values that differ from the observed rewards. Formally, we consider a set of  $K$  stochastic arms, each arm  $k$  is characterized by its own reward distribution  $D_k$  with a mean denoted as  $\mu_k = \mathbb{E}[D_k]$ . To maintain clarity while broadening our perspective, we assume an order such that  $1 \geq \mu_1 \geq \mu_2 \geq \dots \geq \mu_K \geq 0$ . During each round  $t$ , the player pulls an arm  $k_t$ . Subsequently, the chosen arm observes a reward  $r_{k_t,t} \sim D_{k_t}$ , and it has the possibility to report a value  $x_{k,t} \neq r_{k,t}$  to the player while retaining  $r_{k,t} - x_{k,t}$  as its own utility. Importantly, only arm  $k_t$  possesses knowledge of the actual observed reward  $r_{k_t,t}$ , whereas the player is only aware of  $x_{k_t,t}$  and remains unaware of the withheld portion. Therefore, the player’s decision is based on the information collected up to time  $t$ , which can be formally encoded in the filtration  $\mathcal{F}_{P,t} = \{k_1, x_{k_1,1}, \dots, k_t, x_{k_t,t}\}$ . We define  $x_k$  (and  $r_k$ ) as the concatenation of reported values (and rewards, respectively) over  $T$  rounds. Accordingly, the utility associated with arm  $k$  can be expressed as:

$$\mathcal{U}_k(x_k, x_{-k}) = \mathbb{E} \left[ \sum_{t=1}^T (r_{k_t,t} - x_{k_t,t}) \cdot \mathbb{1}_{[k_t=k]} \right] \quad (1)$$

This strategic scenario introduces a game-like dynamic that engenders a competition of objectives between the player, who strives to minimize regret (see section 3), and the arms, which are driven by the pursuit of maximizing their utilities. This model encapsulates a diverse array of dynamic agency dilemmas wherein the player selects an arm (agent) to execute a task on his behalf, and the associated cost remains concealed from the player due to his limited domain or market knowledge. Broadly speaking, this model can be viewed as an extension to the multi-agent realm, akin to the principal-agent problem in contract theory [1, 2], albeit with multiple agents in play.

\*ENS Paris Saclay, France (ahmed.ben\_yahmed@ens-paris-saclay.fr)

†CRITEO AI lab, Paris, France (c.calauzenes@criteo.com)

‡CREST, ENSAE, Palaiseau, France / CRITEO AI Lab, Paris, France (vianney.perchet@normalesup.org)

It significantly extends the standard MAB problem, as arms can utilize this reporting mechanism to influence the player’s decisions. For example, arms may opt to report higher values initially to increase their chances of being selected in later rounds. Conversely, they may report lower values at the outset to decrease the reserve price in auctions [3]. Furthermore, our study takes into consideration the existence of side communications among arms, governed by predefined rules. This consideration is motivated by real-world scenarios in which such interactions are prevalent and influential.

## 1.1. Related Work

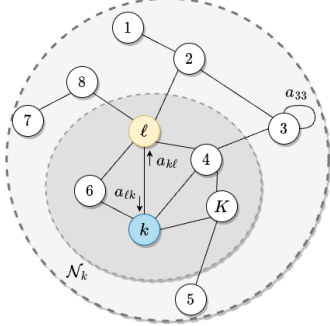
Previous studies, such as [4, 5, 6], have already examined scenarios involving connected and communicating arms. In these scenarios, pulling an arm  $k$  at time step  $t$  not only provides information about arm  $k$  itself but also reveals information about some related arms. A typical example of such a situation is advertising on social networks, where a decision-maker targets individual users of an online platform with promotions, hoping to maximize purchases. However, in this context, the arms (i.e users) are connected and capable of communication. Conceptually, pulling an arm  $k$  triggers instantaneous communication through the arms, revealing aggregated information about all arms related to  $k$  and itself [4]. The concept of a strategic arm in the MAB setting was first introduced in the groundbreaking work by [7]. This notion highlights the challenge of dealing with arms that are not limited to providing their true rewards but can instead manipulate the player to maximize their own utilities. To achieve this, authors present an equilibrium strategy for such arms especially when  $\mu_1 - \mu_2 \leq \frac{\mu_1}{K}$ , enabling them to leave the player with only a minimal reward. This strategic approach poses a significant obstacle for any low-regret algorithm employed by the player.

## 1.2. Contributions

This study addresses the limitation of [7] where arms need full information to be able to have an equilibrium strategy in which they can extract (almost) all the value. The new strategy introduced here, which includes a proper communication protocol, enables them to achieve equilibrium for full surplus extraction and prevents the player from generating high revenue, regardless of the low-regret MAB algorithm the player chooses. We build upon the strategy presented in [7], which serves as a foundation for our work. The main challenge in constructing an equilibrium strategy that doesn’t rely on public knowledge of the entire history is to guarantee that the communication scheme doesn’t encourage arms to convey false information. Indeed, from a game-theoretic perspective, both the report  $x_{k_t,t}$  to the player and the information they share with other arms are components of strategic behavior. We support our claims with theoretical analysis and experimental results.

## 2. MODELING COMMUNICATION

We assume that the arms are interconnected and form a network modeled by a graph. This setting is widely used in distributed learning and multi-agent communication [8] as it provides a stronger privacy protection and reduced risk of communication bottleneck. Arms will be linked by some graph topology (see Fig. 1) that allows neighboring arms to share information. At any point in time, no single arm will have access to all the information that is available across the graph. Additionally, all types of message passing occur simultaneously throughout the graph. This means that all information at all arms is updated instantaneously and equivalently at the same time (i.e the network has a synchronized clock).



**Fig. 1.** Arms are linked by a graph topology and can share information over their shared edges. The neighborhood of an arm is the collection of all arms linked to it. The neighborhood of arm  $k$  is marked by the highlighted area and denoted by  $\mathcal{N}_k$ . We assign a pair of nonnegative scaling weights,  $\{a_{k\ell}, a_{\ell k}\}$ , to the edge connecting  $k$  and  $\ell$ . The scalar  $a_{\ell k}$  will be used by arm  $k$  to scale information it receives from arm  $\ell$ ; this scaling can be interpreted as a measure of confidence [9].

## 3. LOW-REGRET MAB ALGORITHMS

The efficiency of MAB algorithms is most commonly compared with the notion of regret. Regret measures the cumulative loss incurred over  $T$  rounds by choosing arm  $k_t$  at time  $t$  instead of the best arm  $k^*$ . Ideally, we aim for a low-regret, which implies convergence to the best arm, at least asymptotically. Formally, let  $k_t$  be the strategic arm selected by the player at round  $t$  using algorithm  $A$ . The selected arm observes a reward  $r_{k_t,t}$  that we consider in this work in  $[0, 1]$ . Then, it reports, in an adversarial fashion, a value  $x_{k_t,t}$ . The regret of the algorithm  $A$  is the random variable:

$$\mathcal{R}(A) = \max_k \sum_{t=1}^T x_{k,t} - \sum_{t=1}^T x_{k_t,t} \quad (2)$$

We follow the  $(\rho, \delta)$ -low-regret definition given in [7], i.e an algorithm  $A$  is a  $(\rho, \delta)$ -low-regret for the MAB problem if with probability  $1 - \rho$ ,

$$\mathcal{R}(A) \leq \delta \quad (3)$$

The majority of MAB algorithms, particularly in the adversarial scenario – which is relevant here, considering that the reporting protocol is widely regarded as an adversarial setting – assign a probability  $p_{k,t}$  to each arm  $k \in \{1, \dots, K\}$  to be pulled at round  $t$ . The algorithm then selects the arm according to these probabilities. The probability  $p_{k,t}$  is primarily determined by the parameters of the algorithm  $A$ , an intrinsic information  $I_k$  of arm  $k$ , and the average

value of this information  $\tilde{I}$  across all the arms. Formally it can be seen as the following function:

$$\Pr : \{\text{set of parameters}\} \times \mathbb{R} \times \mathbb{R} \rightarrow [0, 1] \quad (4)$$

$$(A, I_{k,t}, \tilde{I}_t) \mapsto p_{k,t} \quad (5)$$

For instance, in the EXP3 algorithm [10], to compute the probability  $p_{k,t}$  of choosing arm  $k$  at time  $t$ , we require the parameter  $\gamma$ , which determines the trade-off between exploration and exploitation. Additionally, we need the weighted estimated reward which is the intrinsic information  $I_{k,t}$  associated with arm  $k$  at time  $t$ , and the sum of these rewards over all arms, denoted as  $K\tilde{I}_t$ .

## 4. ARMS' STRATEGY THAT PREVAILS OVER LOW-REGRET MAB ALGORITHMS

In this section, we extend the work presented in [7] and introduce a strategy that enables strategic arms with restricted communication to reach an  $\epsilon$ -Nash equilibrium while providing only a marginal utility to the player. The Strategy 1 adopts market sharing techniques, where arms select their actions in a way that ensures they are chosen an equal number of times. As a result, the player receives only minimal revenue, as he is unable to commit to selecting the best arm consistently, deviating from the traditional bandit setting. This strategy does not assume that the arms possess prior knowledge of their own distributions or the history of selected arms. In other words, their respective information available at time  $t$  is less complex compared to the setting presented in [7]. The strategy is presented as follow: let  $A$  be the low-regret MAB algorithm used by the player,  $N_k(t)$  be the number of times arm  $k$  has been pulled up to time  $t$  and  $\mathbb{A} = [a_{k\ell}]$  the combination matrix describing the communication graph topology. Set  $B = 7\sqrt{KT\delta}$  and  $\theta = \sqrt{\frac{K\delta}{T}}$ . Then the strategy that the arms shall use is Strategy 1.

---

### Strategy 1: Equilibrium strategy

---

```

1 for  $t = 1, \dots, T$  do
2   Update  $I_{k,t}$  the intrinsic information available for all
   arm  $k$  at time  $t$ . Initiate  $\tilde{I}_{k,0} = I_{k,t}, \forall k$ .
3   for  $n = 1, \dots, \tau$  do
4     for  $k = 1, \dots, K$  do
5        $\tilde{I}_{k,n} = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \tilde{I}_{\ell, n-1}$ 
6     end
7   end
8   for  $k = 1, \dots, K$  do
9     If at any time  $s \leq t$  in the past  $N_k(s) < \frac{s}{K} - B$ 
       then arm  $k$  defects and offers its full value
        $x_{k,t} = r_{k,t}$ .
10    Else arm  $k$  computes the probability
        $\hat{p}_{k,t} = \Pr(A, I_k, \tilde{I}_{k,\tau})$  and offers
        $x_{k,t} = \theta(1 - \hat{p}_{k,t})$ .
11  end
12 end

```

---

This strategy consists of two parts. In the first part, each arm updates its information individually. For example, if an arm is pulled, it updates its reward information based on the received reward. If it is not chosen, the information remains the same as in the previous time step  $t - 1$ . Then, iteratively, the arms use the communication

scheme to compute a local estimate, denoted as  $\tilde{I}_{k,\tau}$ , of the true average information  $\tilde{I}_t$ . During the second part, each arm  $k$  uses its estimated average to compute the probability  $\hat{p}_{k,t}$  of being selected by the player. The arm then adjusts its offer based on this probability. If  $\hat{p}_{k,t}$  is high, the arm sets a low value for  $x_{k,t}$ , allowing other arms to have a higher chance of being chosen. On the other hand, if  $\hat{p}_{k,t}$  is low, the arm sets a high value for  $x_{k,t}$ , increasing its own chance of being chosen since other arms are more likely to be selected. By following this approach, the arms achieve an equilibrium in terms of market sharing, which results in lower revenue for the player. The parameter  $\theta$  in  $x_{k,t}$  is introduced to facilitate the theoretical analysis of the strategy. In the next section, we will proceed with the theoretical analysis, where we will demonstrate that if the arms follow this strategy, they will reach an equilibrium, discouraging any defection.

## 5. THEORETICAL ANALYSIS

To facilitate the analysis, we suppose that  $K \leq \frac{T^{\frac{1}{3}}}{\log(T)}$ ,  $\rho \leq \frac{1}{T^2}$  and  $\delta \geq \sqrt{T \log(T)}$ .

### 5.1. Reliable approximation of pulling probabilities

First, we will demonstrate the utility of the communication steps and how, after a sufficient number of iterations  $\tau$ , the local values  $\tilde{I}_{k,\tau}$  serve as accurate estimates of the true average  $\tilde{I}_t$  which allows each arm  $k$  to compute an accurate estimate  $\tilde{p}_{k,t}$  of  $p_{k,t}$ . To do so, we start by introducing some assumptions that are commonly used in the literature [8, 11, 9, 12].

**Assumption 1 (Doubly-stochastic combination matrix).** *The combination matrix  $\mathbb{A} = [a_{\ell k}]$  representing the graph topology is doubly-stochastic and symmetric. This means that the matrix has non-negative elements and satisfies:*

$$\mathbb{A} \mathbb{1}_K = \mathbb{1}_K^1, \mathbb{A}^\top = \mathbb{A} \quad (6)$$

We also assume that the matrix  $\mathbb{A}$  is primitive. This implies that there exist paths, in both directions, between any two distinct nodes with nonzero scaling weights. Additionally, there is at least one non-trivial self-loop present, meaning that  $a_{kk} > 0$  for at least one node  $k$ .

By applying the Perron-Frobenius theorem, Assumption 1 states that the mixing rate  $\lambda$  of the combination matrix (i.e., the spectral radius of  $\mathbb{A} - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top$ ) is strictly less than 1:

$$\lambda < 1 \quad (7)$$

**Assumption 2 (Lipschitz mapping).** *The mapping  $\tilde{I} \mapsto \Pr(\cdot, \cdot, \tilde{I})$  is Lipschitz, namely:  $\exists L \in \mathbb{R}_+$  such that  $\forall \tilde{I}, \tilde{I}'$ :*

$$\|\Pr(\cdot, \cdot, \tilde{I}) - \Pr(\cdot, \cdot, \tilde{I}')\| \leq L \|\tilde{I} - \tilde{I}'\| \quad (8)$$

Assumption 2 is valid since we are considering a finite horizon  $T$ .

**Theorem 1 (Network disagreement).** *Under Assumption 1, the network disagreement between the true average  $\tilde{I}_t$  and the local estimates  $\tilde{I}_{k,\tau}$  converges to zero.*

$$\frac{1}{K} \sum_{k=1}^K \|\tilde{I}_{k,\tau} - \tilde{I}_t\|^2 \leq \alpha_t \lambda^{2\tau} \xrightarrow{\tau \rightarrow +\infty} 0 \text{ with } \alpha_t > 0 \quad (9)$$

<sup>1</sup> $\mathbb{1}_K$  is a vector of length  $K$  consisting of ones.

**Proof:** for generality we suppose that the dimension of variables is  $M$ , i.e.  $\dim(I_{k,t}) = \dim(\tilde{I}_t) = \dim(\tilde{I}_{k,t}) = M$ . We begin by defining the following variable that aggregates the local variables of each arm into a single variable:

$$\tilde{\mathcal{I}}_n \triangleq \text{col}\{\tilde{I}_{1,n}, \dots, \tilde{I}_{K,n}\} \quad (10)$$

$$\mathcal{A} \triangleq \mathbb{A} \otimes \mathbb{1}_M^2 \quad (11)$$

Where  $\tilde{\mathcal{I}}_n$  is a vector of length  $K \times M$ , obtained by vertically concatenating vectors enclosed in brackets. We express the update in the communication scheme using a more concise notation:

$$\tilde{\mathcal{I}}_n = \mathcal{A}^\top \tilde{\mathcal{I}}_{n-1} \quad (12)$$

$$\Rightarrow \left( \frac{1}{K} \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_n \stackrel{(3)}{=} \left( \frac{1}{K} \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_{n-1} \quad (13)$$

It should be noted that due to the nature of the combination matrix  $\mathbb{A}$ , the true average  $\tilde{I}_t$  is equivalent to the average of the variables  $\tilde{I}_{k,n}$  for any given  $n$ . So we write:

$$\tilde{I}_t = \frac{1}{K} \sum_{k=1}^K \tilde{I}_{k,n} = \left( \frac{1}{K} \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_n \quad (14)$$

and we define the extended average as a vector of length  $K \times M$ :

$$\tilde{\mathcal{I}}_t \triangleq \mathbb{1}_K \otimes \tilde{I}_t = \left( \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_n \quad (15)$$

We get:

$$\tilde{\mathcal{I}}_n - \tilde{\mathcal{I}}_t = \left( \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_{n-1} \quad (16)$$

$$\stackrel{(4)}{=} \left( \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \left( \mathbb{1}_M - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \tilde{\mathcal{I}}_{n-1} \quad (17)$$

$$= \left( \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) (\tilde{\mathcal{I}}_{n-1} - \tilde{\mathcal{I}}_t) \quad (18)$$

Taking the square norm:

$$\|\tilde{\mathcal{I}}_n - \tilde{\mathcal{I}}_t\|^2 = \left\| \left( \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) (\tilde{\mathcal{I}}_{n-1} - \tilde{\mathcal{I}}_t) \right\|^2 \quad (19)$$

$$\leq \lambda^2 \|\tilde{\mathcal{I}}_{n-1} - \tilde{\mathcal{I}}_t\|^2 \quad (20)$$

Iterating from  $\tau$  to 0:

$$\|\tilde{\mathcal{I}}_\tau - \tilde{\mathcal{I}}_t\|^2 \leq \lambda^{2\tau} \|\tilde{\mathcal{I}}_0 - \tilde{\mathcal{I}}_t\|^2 \quad (21)$$

Taking  $\alpha_t = \frac{\|\tilde{\mathcal{I}}_0 - \tilde{\mathcal{I}}_t\|^2}{K}$  finishes the proof.

**Corollary 1.** *Under Assumption 2 and using Theorem 1, we have:*

$$|p_{k,t} - \hat{p}_{k,t}| = |\Pr(A, I_{k,t}, \tilde{I}_t) - \Pr(A, I_{k,t}, \tilde{I}_{k,t})| \quad (22)$$

$$\leq L \sqrt{\alpha_t} \lambda^\tau \xrightarrow{\tau \rightarrow +\infty} 0 \quad (23)$$

Therefore, we have demonstrated that after a sufficient number of iterations  $\tau$ ,  $\hat{p}_{k,t}$  provides a reliable approximation for  $p_{k,t}$ . This allows the arms to calibrate their rewards as if they have access to the complete information available to the player.

<sup>2</sup> $\mathbb{1}_M$  refers to the identity matrix of dimension  $M$ , and  $\otimes$  represents the Kronecker product.

<sup>3</sup> $\left( \frac{1}{K} \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \mathcal{A}^\top = \frac{1}{K} \mathbb{1}_K^\top \otimes \mathbb{1}_M$

<sup>4</sup> $\left( \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) \left( \mathbb{1}_M - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M \right) = \mathcal{A}^\top - \frac{1}{K} \mathbb{1}_K \mathbb{1}_K^\top \otimes \mathbb{1}_M$

## 5.2. Equilibrium resulting from Strategy 1

Following Strategy 1, arms won't defect and will achieve a market sharing situation where each arm is pulled approximately an equal number of times. This renders the utilized low-regret MAB algorithm inefficient. This observation is formalized as follows:

**Theorem 2.** *If arms use Strategy 1, then with high probability  $(1 - \frac{3}{T})$ ,  $N_k(t) \geq \frac{t}{K} - B, \forall t \in [T], k \in [K]$  and they will be in an  $O(\sqrt{KT\delta})$ -Nash equilibrium.*

**Proof sketch**<sup>5</sup>: if arms faithfully adhere to the Strategy 1 denoted as  $S^*$ , then by employing Corollary 1 and similar arguments as in [7], we can demonstrate that with high probability  $(1 - \frac{3}{T})$ ,  $N_k(t) \geq \frac{t}{K} - B$  for all  $t \in [T]$  and  $k \in [K]$ . This implies that arms do not defect, and step 9 of the strategy is never executed. To establish equilibrium, we will evaluate the utility of arm  $k$  when all arms are adhering to  $S^*$ , while the player employs a low-regret MAB algorithm. We show that:

$$U_k(S_k^*, S_{-k}^*) \geq \frac{\mu_k T}{K} - O(\sqrt{KT\delta}) \quad (24)$$

On the other hand, if arm  $k$  plays any strategy  $S$  other than  $S^*$ , we can demonstrate that:

$$U_k(S_k, S_{-k}^*) \leq \frac{\mu_k T}{K} + O(\sqrt{KT\delta}) \quad (25)$$

Therefore, for all strategies  $S$  different from  $S^*$ , we can derive the following inequality:

$$U_k(S_k, S_{-k}^*) - U_k(S_k^*, S_{-k}^*) \leq O(\sqrt{KT\delta}) \quad (26)$$

showing that  $(S_1^*, \dots, S_K^*)$  is an  $O(\sqrt{KT\delta})$ -Nash equilibrium for all arms. Therefore, it becomes evident that the equilibrium is primarily determined by the number of times each arm is pulled. At equilibrium, we observe that each arm is pulled approximately the same number of times and receives in average  $\frac{\mu_k}{K}$  per round. If arms choose to deviate from this strategy by dishonestly reporting either *their values to the player* or *the values communicated to their neighbors*, one of two scenarios will unfold. In the first scenario, this deviation will not impact the number of times each arm is pulled, thus failing to activate the defection step 9. Consequently, the utility of the arms remains unaffected, and the equilibrium remains intact. In the second case, the defection step 9 is triggered, resulting in arm 1, which has the highest real mean value, emerging as the winner. In this case, it gains at maximum  $\mu_1 - \mu_2$  per round. However, it's important to note that  $\mu_1 - \mu_2 \leq \frac{\mu_1}{K}$  by assumption. Given this condition, there is no incentive for arms to deviate from the strategy, as the potential gain from defection is less than what they can achieve by adhering to the equilibrium strategy.

This equilibrium proves detrimental to the player, resulting in constrained revenue, irrespective of the low-regret algorithm employed:

**Corollary 2.** *If arms follow Strategy 1, the player gets at most  $O(\sqrt{KT\delta})$  revenue.*

**Proof:** given that playing according to Strategy 1 implies that with high probability  $\forall t \in [T], k \in [K], N_k(t) \geq \frac{t}{K} - B$ , arms won't defect and the player gets  $T\theta = O(\sqrt{KT\delta})$ . In the case of

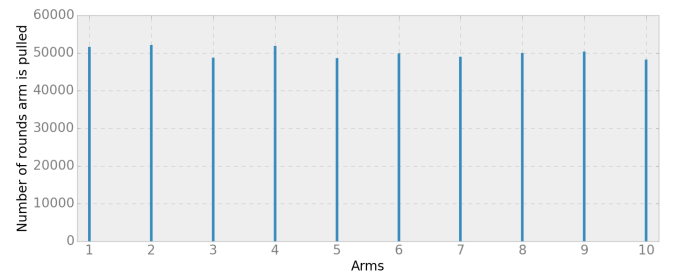
the low probable event the player will get at most  $T$ . So the player revenue is:

$$\text{Player-revenue} \leq (1 - \frac{3}{T})O(\sqrt{KT\delta}) + \frac{3}{T}T \quad (27)$$

$$\Rightarrow \text{Player-revenue} \leq O(\sqrt{KT\delta}) \quad (28)$$

## 6. EXPERIMENT

In this section we test Strategy 1 against an adapted version of EXP3.P [13] as it is done in [4] to take into consideration the existence of the side communication. It is a  $(\rho, O(\sqrt{T \log(K\rho^{-1})}))$ -low-regret MAB algorithm. The intrinsic information for arm  $k$  is the exponential of its weighted estimated cumulative gain. We create a random Erdos-Rényi graph over  $K = 10$  nodes, where each pair of nodes are linked independently with probability  $p = 0.6$ . Arms are modeled as Bernoulli random variables. For most arms, mean is set at 0.4, while three specific arms have different means: 0.8, 0.85, and 0.9. The combination matrix  $\mathbb{A}$  is generated using the Metropolis rule [14]. We run the experiment for  $T = 5 \cdot 10^5$  rounds and set  $\tau = 50$ .



**Fig. 2.** The number of rounds each arm is pulled by the player when arms are using Strategy 1.

**Table 1.** Summary of the numerical results.

$T$	$\delta$	Experimental total revenue	$\sqrt{KT\delta}$
$5 \cdot 10^5$	2778	105169	117838

Fig. 2 confirms our claim that by following Strategy 1 against a low-regret algorithm, arms will reach an equilibrium. This is evident from the balanced number of rounds the arms are pulled, in contrast to the ordinary MAB setting where the algorithm tends to favor choosing the best arm much more frequently than the others. Table 1 further supports our second claim that this strategy leaves the player with a cumulative reward less than  $O(\sqrt{KT\delta})$ .

## 7. CONCLUSION

In scenarios involving repeated interactions, converting a single-step collusion scenario into an equilibrium within the cumulative game necessitates the ability of each participant to identify instances where others deviate from the collusive behavior. This study illustrates that even when not all historical information is publicly accessible, the arms can implement a communication strategy enabling each of them to detect deviations, whether they involve falsifying player reports or manipulating shared information. The established equilibrium surpasses any low-regret MAB algorithm, resulting in reduced player revenues. Future research can focus on developing mechanisms that encompass not only traditional sequential learning in the classical MAB style but also integrate incentive mechanisms to effectively address the challenges highlighted in this paper.

<sup>5</sup>The detailed proof is omitted due to space limitations.

## 8. REFERENCES

- [1] Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, September 2013. (Cited on 1)
- [2] Jean-Jacques Laffont and David Martimort. The theory of incentives: The principal-agent model. pages i–vi, 2002. (Cited on 1)
- [3] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 11 2013. (Cited on 1)
- [4] Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, page 684–692, Red Hook, NY, USA, 2011. Curran Associates Inc. (Cited on 1, 4)
- [5] Stéphane Caron, Branislav Kveton, Marc Lelarge, and Smriti Bhagat. Leveraging side observations in stochastic bandits. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, August 14-18, 2012*, pages 142–151. AUAI Press, 2012. (Cited on 1)
- [6] Swapna Buccapatnam, Atilla Eryilmaz, and Ness B. Shroff. Stochastic bandits with side observations on networks. In *Measurement and Modeling of Computer Systems*, 2014. (Cited on 1)
- [7] Mark Braverman, Jieming Mao, Jon Schneider, and S. Matthew Weinberg. Multi-armed bandit problems with strategic arms. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 383–416. PMLR, 25–28 Jun 2019. (Cited on 1, 2, 4)
- [8] Ali H. Sayed. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning*, 7(4-5):311–801, 2014. (Cited on 2, 3)
- [9] Ali H. Sayed. Adaptive networks. *Proceedings of the IEEE*, 102(4):460–497, 2014. (Cited on 2, 3)
- [10] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002. (Cited on 2)
- [11] Kun Yuan, Qing Ling, and Wotao Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016. (Cited on 3)
- [12] Mert Kayaalp, Stefan Vlaski, and Ali Sayed. Dif-MAML: Decentralized multi-agent meta-learning. *IEEE Open Journal of Signal Processing*, 3:71–93, 2022. (Cited on 3)
- [13] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5, 04 2012. (Cited on 4)
- [14] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, 21(6):1087–1092, June 1953. (Cited on 4)