

An Efficient Algorithm for Cooperative Semi-Bandits

Riccardo Della Vecchia

*Artificial Intelligence Lab, Institute for Data Science & Analytics
Bocconi University, Milano, Italy*

RICCARDO.DELLAVECCHIA@UNIBOCCONI.IT

Tommaso R. Cesari

*Toulouse School of Economics (TSE), Toulouse, France
Artificial and Natural Intelligence Toulouse Institute (ANITI), Toulouse, France*

TOM.CESARI@UNIV-TOULOUSE.FR

Editors: Vitaly Feldman, Katrina Ligett and Sivan Sabato

Abstract

We consider the problem of asynchronous online combinatorial optimization on a network of communicating agents. At each time step, some of the agents are stochastically activated, requested to make a prediction, and the system pays the corresponding loss. Then, neighbors of active agents receive semi-bandit feedback and exchange some succinct local information. The goal is to minimize the network regret, defined as the difference between the cumulative loss of the predictions of active agents and that of the best action in hindsight, selected from a combinatorial decision set. The main challenge in such a context is to control the computational complexity of the resulting algorithm while retaining minimax optimal regret guarantees. We introduce Coop-FTPL, a cooperative version of the well-known Follow The Perturbed Leader algorithm, that implements a new loss estimation procedure generalizing the Geometric Resampling of [Neu and Bartók \(2013\)](#) to our setting. Assuming that the elements of the decision set are k -dimensional binary vectors with at most m non-zero entries and α_1 is the independence number of the network, we show that the expected regret of our algorithm after T time steps is of order $Q\sqrt{mkT \log(k)(k\alpha_1/Q + m)}$, where Q is the total activation probability mass. Furthermore, we prove that this is only $\sqrt{k \log k}$ -away from the best achievable rate and that Coop-FTPL has a state-of-the-art $T^{3/2}$ worst-case computational complexity.

Keywords: follow the perturbed leader, online stochastic mirror descent, regret minimization, multi-agent learning

1. Introduction

Distributed online settings with communication constraints arise naturally in large-scale learning systems. For example, in domains such as finance or online advertising, agents often serve high volumes of prediction requests and have to update their local models in an online fashion. Bandwidth and computational constraints may therefore preclude a central processor from having access to all the observations from all sessions and synchronizing all local models at the same time. With these motivations in mind, we introduce and analyze a new online learning setting in which a network of agents solves efficiently a common nonstochastic combinatorial semi-bandit problem by sharing information only with their network neighbors. At each time step t , some agents v belonging to a communication network \mathcal{G} are asked to make a prediction $x_t(v)$ belonging to a subset \mathcal{A} of $\{0, 1\}^k$ and pay a (linear) loss $\langle x_t(v), \ell_t \rangle$ where $\ell_t \in [0, 1]^k$ is chosen adversarially by an oblivious environment. Then, any such agent v receives the feedback $(x_t(1, v)\ell_t(1), \dots, x_t(k, v)\ell_t(k))$, which

is shared, together with some local information, to its first neighbors in the graph. The goal is to minimize the network regret after T time steps

$$R_T = \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \langle x_t(v), \ell_t \rangle - \sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \langle a, \ell_t \rangle \right], \quad (1)$$

where \mathcal{S}_t is the set of agents v that made a prediction at time t . In words, this is the difference between the cumulative loss of the “active” agents and the loss that they would have incurred had they consistently made the best prediction in hindsight.

For this setting, we design a distributed algorithm that we call Coop-FTPL (Algorithm 1), and prove that its regret is upper bounded by $Q\sqrt{mkT \log(k)(k\alpha_1/Q + m)}$ (Theorem 5), where α_1 is the independence number of the network \mathcal{G} and Q is the sum over all agents of the probability that the agent is active during a time step. Our algorithm employs an estimation technique that we call Cooperative Geometric Resampling (Coop-GR, Algorithm 2). It is an extension of a similar procedure appearing in (Neu and Bartók, 2013) that relies on the fact that the reciprocal of the probability of an event can be estimated by measuring the reoccurrence time. We can leverage this idea in the context of cooperative learning thanks to some statistical properties of the minimum of a family of geometric random variables (see Lemmas 1–3). Our algorithm has a state-of-the-art dependence on time of order $T^{3/2}$ for the worst-case computational complexity (Proposition 6). Moreover, we show with a lower bound (Theorem 7) that no algorithm can achieve a regret smaller than $Q\sqrt{mkT\alpha_1/Q}$ on all cooperative semi-bandit instances. Thus, our Coop-FTPL is at most a multiplicative factor of $\sqrt{k \log k}$ -away from the minimax result.

To the best of our knowledge, ours is the first computationally efficient near-optimal learning algorithm for the problem of cooperative learning with nonstochastic combinatorial bandits, where not all agents are necessarily active at all time steps.

2. Related work and further applications

Single-agent combinatorial bandits find applications in several fields, such as path planning, ranking and matching problems, finding minimum-weight spanning trees, cut sets, and multitask bandits. An efficient algorithm for this setting is Follow-The-Perturbed-Leader (FTPL), which was first proposed by Hannan (1957) and later rediscovered by Kalai and Vempala (2005). Neu and Bartók (2013) show that combining FTPL with a loss estimation procedure called Geometric Resampling (GR) leads to a computationally efficient solution for this problem. More precisely, the solution is efficient given that the offline optimization problem of finding

$$a^* = \operatorname{argmin}_{a \in \mathcal{A}} \langle a, y \rangle, \quad \forall y \in [0, +\infty)^k \quad (2)$$

admits a computationally efficient algorithm. This assumption is minimal, in the sense that if the offline problem in Eq. (2) is hard to approximate, then any algorithm with low regret must also be inefficient.¹ Grötschel et al. (2012) and Lee et al. (2018) give some sufficient conditions for the validity of this assumption. They essentially rely on having an efficient membership oracle for the convex hull $\operatorname{co}(\mathcal{A})$ of \mathcal{A} and an evaluation oracle for the linear function to optimize. Audibert et al. (2014) note that Online Stochastic Mirror Descent (OSMD) or Follow The Regularized Leader

1. A slight relaxation in this direction would be assuming that Eq. (2) can be approximated accurately and efficiently.

(FTRL)-type algorithms can be efficiently implemented by convex programming if the convex hull of the decision set can be described by a polynomial number of constraints. Suehiro et al. (2012) investigate the details of such efficient implementations and design an algorithm with k^6 time-complexity, which might still be unfeasible in practice. Methods based on the exponential weighting of each decision vector can be implemented efficiently only in a handful of special cases—see, e.g., (Koolen et al., 2010) and (Cesa-Bianchi and Lugosi, 2012) for some examples.

The study of cooperative nonstochastic online learning on networks was pioneered by Awerbuch and Kleinberg (2008), who investigated a bandit setting in which the communication graph is a clique, agents belong to clusters characterized by the same loss, and some agents may be non-cooperative. In our multi-agent setting, the end goal is to control the total network regret (1). This objective was already studied by Cesa-Bianchi et al. (2019a) in the full-information case. A similar line of work was pursued by Cesa-Bianchi et al. (2019b), where the authors consider networks of learning agents that cooperate to solve the same nonstochastic bandit problem. In their setting, all agents are simultaneously active at all time steps, and the feedback propagates throughout the network with a maximum delay of d time steps, where d is a parameter of the proposed algorithm. The authors introduce a cooperative version of Exp3 that they call Exp3-COOP with regret of order $\sqrt{(d+1 + K\alpha_d/N)(T \log K)}$ where K is the number of arms in the nonstochastic bandit problem, N is the total number of agents in the network, and α_d is the independence number of the d -th power of the communication network. The case $d = 1$ corresponds to information that arrives with one round of delay and communication limited to first neighbors. In this setting Exp3-COOP has regret of order $\sqrt{(1 + K\alpha_1/N)(T \log K)}$. Thus, our work can be seen as an extension of this setting to the case of combinatorial bandits with stochastic activation of agents. Finally, we point out that if the network consists of a single node, our cooperative setting collapses into a single-agent combinatorial semi-bandit problem. In particular, when the number of arms is k and $m = 1$, this becomes the well-known adversarial multiarmed bandit problem (see (Auer et al., 2002)). Hence, ours is a proper generalization of all the settings mentioned above. These cooperative problems are also studied in stochastic setting (see, e.g., Martínez-Rubio et al. (2019)).

Finally, the reader may wonder what kind of results could be achieved if the agents are activated adversarially rather than stochastically. Cesa-Bianchi et al. (2019a) showed that in this setting no learning can occur, not even in with full-information feedback.

3. Cooperative semi-bandit setting

In this section, we present our cooperative semi-bandit protocol and we introduce all relevant definitions and notation.

We say that $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a *communication network* over N agents if it is an undirected graph over a set \mathcal{V} with cardinality $|\mathcal{V}| = N$, whose elements we refer to as *agents*. Without loss of generality, we assume that $\mathcal{V} = \{1, \dots, N\}$. For any agent $v \in \mathcal{V}$, we denote by $\mathcal{N}(v)$ the set of agents containing v and the neighborhood $\{w \in \mathcal{V} : (v, w) \in \mathcal{E}\}$. We say that α_1 is the *independence number* of the network \mathcal{G} if is the largest cardinality of an *independent set* of \mathcal{G} , where an independent set of \mathcal{G} is a subset of agents, no two of which are neighbors.

We study the following cooperative online combinatorial optimization protocol. Initially, hidden from the agents, the environment draws a sequence of subsets $\mathcal{S}_1, \mathcal{S}_2, \dots \subset \mathcal{V}$ of agents, that we call *active*, and a sequence of *loss vectors* $\ell_1, \ell_2, \dots \in \mathbb{R}^k$. We assume that each agent v has a probability $q(v)$ of being activated, which need only be known by v . The set of active agents \mathcal{S}_t

at time $t \in \{1, 2, \dots\}$ is then determined by drawing, for each agent $v \in \mathcal{V}$, a Bernoulli random variable $X_t(v)$ with bias $q(v)$, independently of the past, and \mathcal{S}_t consists exclusively of agents $v \in \mathcal{V}$ for which $X_t(v) = 1$. The *decision set* is a subset \mathcal{A} of $\{a \in \{0, 1\}^k : \sum_{i=1}^k a_i \leq m\}$, for some $m \in \{1, \dots, k\}$.

For each time step $t \in \{1, 2, \dots\}$:

1. each active agent $v \in \mathcal{S}_t$ predicts with $x_t(v) \in \mathcal{A}$ (possibly drawn at random);
2. each neighbor $v \in \mathcal{N}(w)$ of an active agent $w \in \mathcal{S}_t$ receives the feedback

$$f_t(w) := (x_t(1, w)\ell_t(1), \dots, x_t(k, w)\ell_t(k)) ; \quad (3)$$

3. each agent $v \in \bigcup_{w \in \mathcal{S}_t} \mathcal{N}(w)$ receives some local information from its neighbors in $\mathcal{N}(v)$;
4. the system incurs the loss $\sum_{v \in \mathcal{S}_t} \langle x_t(v), \ell_t \rangle$.

The goal is to minimize the expected *network regret* as a function of the *time horizon* T , defined by

$$R_T := \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \langle x_t(v), \ell_t \rangle - \sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \langle a, \ell_t \rangle \right], \quad (4)$$

where the expectation is taken with respect to the draws of $\mathcal{S}_1, \dots, \mathcal{S}_T$ and (possibly) the randomization of the learners. In the next sections we will also denote by \mathbb{P}_t the probability conditioned to the history up to and including round $t - 1$, and by \mathbb{E}_t the corresponding expectation.

The nature of the local information exchanged by neighbors of active agents will be clarified in the next section. In short, they share succinct representations of the current state of their local prediction model.

4. Coop-FTPL and upper bound

In this section we introduce and analyze our efficient Coop-FTPL algorithm for cooperative online combinatorial optimization.

4.1. The Coop-FTPL algorithm

Coop-FTPL (Algorithm 1) takes as input a decision set $\mathcal{A} \subset \{0, 1\}^k$, a time horizon $T \in \mathbb{N}$, a learning rate $\eta > 0$, a truncation parameter β , and an exploration distribution ζ . At each time step t , all active agents v make a FTPL prediction $x_t(v)$ (line 4) with an i.i.d. perturbation sampled from ζ (line 3), then they receive some feedback $f_t(v)$ and share it with their first neighbors (line 6). Afterwards, each agent who received some feedback this round, requests from its neighbors w a vector $K_t(w)$ of geometric random samples (line 7) which is efficiently computed by Algorithm 2 and will be described in detail later. With these geometric samples, each agent v computes an estimated loss $\hat{\ell}_t(v)$ (line 8) and updates the cumulative loss estimate $\hat{L}_t(v)$ (line 9). This estimator is described in detail in Section 4.3.

4.2. Reduction to OSMD

Before describing $K_t(w)$ and $\hat{\ell}_t(v)$, we make a connection between FTPL and the Online Stochastic Mirror Descent algorithm (OSMD)² that will be crucial for our analysis.³

2. For a brief overview of some key convex analysis and OSMD facts, see Appendices A and B

3. For a similar approach in the single-agent case, see (Lattimore and Szepesvári, 2020).

Algorithm 1: Follow the perturbed leader for cooperative semi-bandits (Coop-FTPL)

Input: decision set \mathcal{A} , horizon T , learning rate η , truncation parameter β , exploration pdf ζ

Initialization: $\widehat{L}_0 = 0 \in \mathbb{R}^k$

```

1 for each time  $t = 1, 2, \dots$  do
2   for each active agent  $v \in \mathcal{S}_t$  do                                     // active agents
3     sample  $Z_t(v) \sim \zeta$ , independently of the past
4     make the prediction  $x_t(v) = \operatorname{argmax}_{a \in \mathcal{A}} \langle a, Z_t(v) - \eta \widehat{L}_{t-1}(v) \rangle$ 
5   for each agent  $v \in \bigcup_{w \in \mathcal{S}_t} \mathcal{N}(w)$  do                               // neighbors of active agents
6     receive the feedback  $f_t(w)$  (Eq. (3)) from each active neighbor  $w \in \mathcal{N}(v) \cap \mathcal{S}_t$ 
7     receive  $K_t(i, w)$  ( $\forall i \in \{1, \dots, k\}$ ) from each neighbor  $w \in \mathcal{N}(v)$  using Algorithm 2
8     compute  $\widehat{\ell}_t(i, v) = \ell_t(i) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{K_t(i, w)\}$ ,  $\forall i \in \{1, \dots, k\}$  (Eq. (6)–(7))
9     update  $\widehat{L}_t(v) = \widehat{L}_{t-1}(v) + \widehat{\ell}_t(v)$ 

```

Algorithm 2: Geometric resampling for cooperative semi-bandits (Coop-GR)

Input: time step t , component i , agent w , truncation parameter β , exploration pdf ζ

```

1 for  $s = 1, 2, \dots$  do
2   sample  $z'_s \sim \zeta$ , independently of the past
3   let  $x'_s$  be the  $i$ -th component of the vector  $\operatorname{argmax}_{a \in \mathcal{A}} \langle a, z'_s - \eta \widehat{L}_{t-1}(w) \rangle$ 
4   sample  $y'_s$  from a Bernoulli distribution with parameter  $q(w)$ , independently of the past
5   if ( $x'_s = 1$  and  $y'_s = 1$ ) or  $s = \beta$ , break
6 return  $K_t(i, w) = s$ 

```

Fix any time step t and an agent v . As we mentioned above, if v is active, it makes the following FTPL prediction (line 4)

$$x_t(v) = \operatorname{argmin}_{a \in \mathcal{A}} \langle a, \eta \widehat{L}_{t-1}(v) - Z_t(v) \rangle ,$$

where $Z_t(v) \in \mathbb{R}^k$ is sampled i.i.d. from ζ (the random perturbations introduce the exploration, which for an appropriate choice of ζ is sufficient to guarantee small regret). On the other hand, given a Legendre potential F with $\operatorname{dom}(\nabla F) = \operatorname{int}(\operatorname{co}(\mathcal{A}))$, an OSMD algorithm makes the prediction

$$\bar{x}_t(v) = \operatorname{argmin}_{x \in \operatorname{co}(\mathcal{A})} \left(\langle x, \eta \widehat{\ell}_{t-1}(v) \rangle + \mathcal{B}_F(x, \bar{x}_{t-1}(v)) \right) ,$$

where \mathcal{B}_F is the Bregman divergence induced by F and $\operatorname{co}(\mathcal{A})$ is the convex hull of \mathcal{A} . Using the fact that $\operatorname{dom}(\nabla F) = \operatorname{int}(\operatorname{co}(\mathcal{A}))$, the argmin above can be computed in a standard way by studying when the gradient of its argument is equal to zero, and proceeding inductively, we obtain the two identities $\nabla F(\bar{x}_t(v)) = \nabla F(\bar{x}_{t-1}(v)) - \eta \widehat{\ell}_{t-1}(v) = -\eta \widehat{L}_{t-1}(v)$. By duality this implies that $\bar{x}_t(v) = \nabla F^*(-\eta \widehat{L}_{t-1}(v))$. We now want to relate $x_t(v)$ and $\bar{x}_t(v)$ so that

$$\bar{x}_t(v) = \mathbb{E}_t[x_t(v)] = \mathbb{E}_t \left[\operatorname{argmin}_{a \in \mathcal{A}} \langle a, \eta \widehat{L}_{t-1}(v) - Z_t(v) \rangle \right] , \quad (5)$$

where the conditional expectation \mathbb{E}_t (given the history up to time $t-1$) is taken with respect to $Z_t(v)$. Thus, in order to view FTPL as an instance of OSMD, it suffices to find a Legendre potential F with $\operatorname{dom}(\nabla F) = \operatorname{int}(\operatorname{co}(\mathcal{A}))$ such that $\nabla F^*(-\eta \widehat{L}_{t-1}(v)) = \mathbb{E}_t[\operatorname{argmax}_{a \in \mathcal{A}} \langle a, Z_t(v) - \eta \widehat{L}_{t-1}(v) \rangle]$. In order to satisfy this condition, we need that for any $x \in \mathbb{R}^k$, the Fenchel conjugate F^* of F enjoys $\nabla F^*(x) = \int_{\mathbb{R}^k} \operatorname{argmax}_{a \in \mathcal{A}} \langle a, z - x \rangle \zeta(z) dz$. Then, we define $h(x) := \operatorname{argmax}_{a \in \mathcal{A}} \langle a, x \rangle$ for any $x \in \mathbb{R}^k$, where $h(x)$ is chosen to be an arbitrary maximizer if multiple maximizers exist. From convex analysis, if the convex hull $\operatorname{co}(\mathcal{A})$ of \mathcal{A} had a smooth boundary, then the support function $x \mapsto \phi(x) := \max_{a \in \operatorname{co}(\mathcal{A})} \langle a, x \rangle = \max_{a \in \mathcal{A}} \langle a, x \rangle$, of $\operatorname{co}(\mathcal{A})$ would satisfy $\nabla \phi(x) = h(x)$. For combinatorial bandits, $\operatorname{co}(\mathcal{A})$ is non-smooth, but, being ζ a density with respect to Lebesgue measure, one can prove (see, e.g., [Lattimore and Szepesvári \(2020\)](#)) that $\nabla \int_{\mathbb{R}^k} \phi(x+z) \zeta(z) dz = \int_{\mathbb{R}^k} h(x+z) \zeta(z) dz$, for all $x \in \mathbb{R}^k$. This shows that FTPL can be interpreted as OSMD with a potential F defined implicitly by its Fenchel conjugate

$$F^*(x) := \int_{\mathbb{R}^k} \phi(x+z) \zeta(z) dz , \quad \forall x \in \mathbb{R}^k .$$

Thus, recalling (5), we can think of the update $\bar{x}_t(v)$ of OSMD as the average of a random component-wise draw $\bar{x}_t(i, v) = \sum_{a \in \mathcal{A}} P_t(a, v) a(i)$ (for all $i \in \{1, \dots, k\}$), with respect to a distribution $P_t(v)$ on \mathcal{A} defined in terms of the distribution of Z_t , as

$$P_t(a, v) = \mathbb{P}_t \left[h(Z_t(v) - \eta \widehat{L}_{t-1}(v)) = a \right] , \quad \forall a \in \mathcal{A} ,$$

where \mathbb{P}_t is the probability conditioned of the history up to time $t-1$.

4.3. An efficient estimator

For the understanding of the definitions and analyses of $K_t(w)$ and $\widehat{\ell}_t(v)$, we introduce three useful lemmas on geometric distributions. We defer their proofs to Appendix C.

Lemma 1 *Let Y_1, \dots, Y_m be m independent random variables such that each Y_j has a geometric distribution with parameter $p_j \in [0, 1]$. Then, the random variable $Z := \min_{j \in \{1, \dots, m\}} Y_j$ has a geometric distribution with parameter $1 - \prod_{j=1}^m (1 - p_j)$.*

Lemma 2 *Let G be a geometric random variable with parameter $q \in (0, 1]$ and $\beta > 0$. Then, the expectation of the random variable $\min\{G, \beta\}$ satisfies $\mathbb{E}[\min\{G, \beta\}] = (1 - (1 - q)^\beta)/q$.*

Lemma 3 *For all $v \in \mathcal{V}$, fix two arbitrary numbers $p_1(v), p_2(v) \in [0, 1]$. Consider a collection $\{X_s(v), Y_s(v)\}_{s \in \mathbb{N}, v \in \mathcal{V}}$ of independent Bernoulli random variables such that $\mathbb{E}[X_s(v)] = p_1(v)$ and $\mathbb{E}[Y_s(v)] = p_2(v)$ for any $s \in \mathbb{N}$ and all $v \in \mathcal{V}$. Then, the random variables $\{G(v)\}_{v \in \mathcal{V}}$ defined for all $v \in \mathcal{V}$ by $G(v) := \inf\{s \in \mathbb{N} : X_s(v) Y_s(v) = 1\}$ are all independent and they have a geometric distribution with parameter $p_1(v) p_2(v)$.*

Fix now any time step t , agent v , and component $i \in \{1, \dots, k\}$. The loss estimator $\widehat{\ell}_t(i, v)$ depends on the algorithmic definition of $K_t(i, w)$ in Algorithm 2, where $w \in \mathcal{N}(v)$. By Lemma 3, we have that for any w , conditionally on the history up to time $t - 1$, the random variable $K_t(i, w)$, has a truncated geometric distribution with success probability equal to $\bar{x}_t(i, w)q(w)$ and truncation parameter β . The idea is that using $K_t(i, w)$ as an estimator we can reconstruct the inverse of the probability that agent v observes i at time t . The truncation parameter β is not needed for the analysis, but it is used just to optimize the computational efficiency of the algorithm.⁴

The loss estimator of v is then defined as

$$\widehat{\ell}_t(i, v) := \ell_t(i) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{K_t(i, w)\}, \quad (6)$$

where

$$B_t(i, v) = \mathbb{I}\{\exists w \in \mathcal{N}(v) : w \in \mathcal{S}_t, x_t(i, w) = 1\}, \quad K_t(i, w) = \min\{G_t(i, w), \beta\}, \quad (7)$$

and given the history up to time $t - 1$, for each $i \in \{1, \dots, k\}$, the family $\{G_t(i, w)\}_{w \in \mathcal{V}}$ consists of independent geometric random variables with parameter $\bar{x}_t(i, w)q(w)$. Note that the geometric random variables $G_t(i, w)$ are actually never computed by Algorithm 2 which efficiently computes only their truncations $K_t(i, w)$, with truncation parameter β . Nevertheless, as it will be apparent later, they are a useful tool for the theoretical analysis. Note that, by Eq. (5), we have

$$\mathbb{P}_t[x_t(i, w) = 1] = \mathbb{E}_t[x_t(i, w)] = \bar{x}_t(i, w),$$

4. Previously known cooperative algorithms for limited feedback settings need to exchange at least two real numbers: the probability according to which predictions are drawn and the loss. Instead of the probability, we only need to pass the integer $K_t(i, w)$, which requires at most $\log_2(\beta)$ bits (order of $\log(T)$, when tuned). Note also that for the loss, we could exchange an approximation l_t of ℓ_t using only $n = O(\log(T))$ bits. Indeed one can show that Lemma 4, in this case, is true replacing ℓ_t with l_t in the first equality. Everything else works the same in the proof of Theorem 5 up an extra (negligible) $m2^{-n}T$ term.

therefore

$$\bar{B}_t(i, v) := \mathbb{E}_t[B_t(i, v)] = 1 - \prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) = \frac{1}{\mathbb{E}_t[\min_{w \in \mathcal{N}(v)} G_t(i, w)]},$$

where the last identity follows by Lemma 1. Moreover from Lemmas 1 and 2, we have

$$\mathbb{E}_t \left[\min_{w \in \mathcal{N}(v)} K_t(i, w) \right] = \frac{1 - \prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w))^\beta}{\bar{B}_t(i, v)}.$$

The following key lemma gives an upper bound on the expected estimated loss.

Lemma 4 *For any time t , component i , agents v , and truncation parameter β , the expectation of the loss estimator in (6), given the history up to time $t - 1$, satisfies*

$$\mathbb{E}_t \left[\widehat{\ell}_t(i, v) \right] = \ell_t(i) \left(1 - \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right) \leq \ell_t(i).$$

Proof Using the fact that, conditioned on the history up to time $t - 1$, the random variable $\min_{w \in \mathcal{N}(v)} G_t(i, w)$ has a geometric distribution with parameter $\bar{B}_t(i, v)$ (Lemmas 1-3), we get

$$\begin{aligned} & \mathbb{E}_t \left[\widehat{\ell}_t(i, v) \right] \\ &= \mathbb{E}_t \left[\ell_t(i) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{ \min\{G_t(i, w), \beta\} \} \right] = \mathbb{E}_t \left[\ell_t(i) B_t(i, v) \min \left\{ \min_{w \in \mathcal{N}(v)} G_t(i, w), \beta \right\} \right] \\ &= \ell_t(i) \mathbb{E}_t[B_t(i, v)] \mathbb{E}_t \left[\min \left\{ \min_{w \in \mathcal{N}(v)} G_t(i, w), \beta \right\} \right] = \ell_t(i) \bar{B}_t(i, v) \frac{(1 - (1 - \bar{B}_t(i, v))^\beta)}{\bar{B}_t(i, v)} \\ &= \ell_t(i) \left(1 - (1 - \bar{B}_t(i, v))^\beta \right) = \ell_t(i) \left(1 - \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right), \end{aligned}$$

where we plugged in the definition of $\bar{B}_t(i, v)$ in the last equation. From the fact that $\bar{x}_t(i, w) q(w) \in [0, 1]$ and $\beta > 0$ follows that $\mathbb{E}_t \left[\widehat{\ell}_t(i, v) \right] \leq \ell_t(i)$. \blacksquare

4.4. Analysis

We can finally state our upper bound on the regret of Coop-FTPL. The key idea is to apply OSMD techniques to our FTPL algorithm, as explained in Section 4.2. The proof proceeds by splitting the regret of each agent in the network in three terms. The first two are treated with standard techniques; the first one depends on the diameter of $\text{co}(\mathcal{A})$ with respect to the regularizer F and the second one on the Bregman divergence of consecutive updates. The last term is related to the bias of the estimator and is analyzed leveraging the lemmas on geometric distributions introduced in Section 4.3. Then, this terms are summed, each with a weight corresponding to their probabilities of being activated, and this sum is controlled using results that relate a sum of weights over the nodes of a graph of with the independence number of the graph.

Theorem 5 *If ζ is the Laplace density $z \mapsto \zeta(z) = 2^{-k} \exp(-\|z\|_1)$, $\eta > 0$, and $0 < \beta \leq 1/(\eta k)$, then the regret of our Coop-FTPL algorithm (Algorithm 1) satisfies*

$$R_T \leq \frac{3Qm \log k}{\eta} + 3Q\eta k T \left(\frac{k}{Q} \alpha_1 + m \right) + \frac{\alpha_1 k T}{\beta}.$$

In particular, tuning the parameters η, β as follows

$$\beta = \left\lfloor \frac{1}{k\eta} \right\rfloor \quad \text{and} \quad \eta = \sqrt{\frac{3m \log(k)}{5kT \left(\frac{k}{Q} \alpha_1 + m \right)}}, \quad \text{where} \quad Q = \sum_{v \in \mathcal{V}} q(v), \quad (8)$$

yields

$$R_T \leq 2Q \sqrt{15mkT \log(k) \left(\frac{k}{Q} \alpha_1 + m \right)}.$$

We now present a detailed sketch of the proof of our main result (full proof in Appendix D).

Proof sketch For the sake of convenience, we define the expected individual regret of an agent $v \in \mathcal{V}$ in the network with respect to a fixed action $a \in \mathcal{A}$ by

$$R_T(a, v) := \mathbb{E} \left[\sum_{t=1}^T \langle x_t(v), \ell_t \rangle - \sum_{t=1}^T \langle a, \ell_t \rangle \right],$$

where the expectation is taken with respect to the internal randomization of the agent, but not to its activation probability $q(v)$. With this definition the total regret on the network in Eq. (4) can be decomposed as

$$\begin{aligned} R_T &= \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \left(\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right) \right] = \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{v \in \mathcal{S}_t} \left(\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right) \right] \right] \\ &= \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{V}} q(v) \mathbb{E}_t \left[\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right] \right] = \max_{a \in \mathcal{A}} \sum_{v \in \mathcal{V}} q(v) R_T(a, v). \end{aligned} \quad (9)$$

The proof then proceeds by isolating the bias in the loss estimators. For each $a \in \mathcal{A}$ we have

$$R_T(a, v) = \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \hat{\ell}_t(v) \rangle \right] + \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \ell_t - \hat{\ell}_t(v) \rangle \right].$$

Exploiting the analogy that we established between FTPL and OSMD, we begin by using the standard bound for the regret of OSMD in the first term of the previous equation. For the reader's convenience, we restate it in Appendix B, Theorem 15. This leads to

$$R_T(a, v) \leq \underbrace{\frac{F(\bar{x}_1(v)) - F(a)}{\eta}}_{(I)} + \underbrace{\mathbb{E} \left[\frac{1}{\eta} \sum_{t=1}^T \mathcal{B}_F(\bar{x}_t(v), \bar{x}_{t+1}(v)) \right]}_{(II)} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \ell_t - \hat{\ell}_t(v) \rangle \right]}_{(III)}.$$

The three terms are studied separately and in detail in Appendix D. Here, we provide a sketch of the bounds.

For the first term (I), we use the fact that the regularizer F satisfies, for all $a \in \mathcal{A}$,

$$F(a) \geq -m(1 + \log(k)), \quad (10)$$

which follows by the definition of F , the properties of the perturbation distribution, and the fact that $\|a\|_1 \leq m$ for any $a \in \mathcal{A}$. One can also show that $F(a) \leq 0$ for all $a \in \mathcal{A}$, and this, combined with the previous equation, leads to

$$(I) \leq \frac{m(1 + \log k)}{\eta}.$$

For the second term (II), we have

$$\begin{aligned} \mathcal{B}_F(\bar{x}_t(v), \bar{x}_{t+1}(v)) &= \mathcal{B}_{F^*}(\nabla F(\bar{x}_{t+1}(v)), \nabla F(\bar{x}_t(v))) \\ &= \mathcal{B}_{F^*}\left(-\eta \widehat{L}_{t-1}(v) - \eta \widehat{\ell}_t(v), -\eta \widehat{L}_{t-1}(v)\right) \\ &= \frac{\eta^2}{2} \left\| \widehat{\ell}_t(v) \right\|_{\nabla^2 F^*(\xi(v))}^2, \end{aligned} \quad (11)$$

where the first equality is a standard property of Bregmann divergence (see Theorem 13 in Appendix A), the second follows from the definitions of the updates and the last by Taylor's theorem, where $\xi(v) = -\eta \widehat{L}_{t-1}(v) - \alpha \eta \widehat{\ell}_t(v)$, for some $\alpha \in [0, 1]$. The estimation of the entries of the Hessian are non trivial (but tedious); the interested reader can find them in Appendix D. Exploiting our assumption that $\beta \leq 1/(\eta k)$, we get, for all $i, j \in \{1, \dots, k\}$,

$$\nabla^2 F^*(\xi(v))_{ij} \leq e \bar{x}_t(i, v).$$

Plugging this estimate in Eq. (11) yields

$$\begin{aligned} \frac{\eta^2}{2} \left\| \widehat{\ell}_t(v) \right\|_{\nabla^2 F^*(\xi(v))}^2 &= \frac{\eta^2}{2} \sum_{i=1}^k \sum_{j=1}^k \nabla^2 F^*(\xi(v))_{i,j} \widehat{\ell}_t(i, v) \widehat{\ell}_t(j, v) \\ &\leq \frac{\eta^2 e}{2} \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i, v) \widehat{\ell}_t(i, v) \widehat{\ell}_t(j, v) \\ &\leq \frac{\eta^2 e}{2} \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i, v) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{G_t(i, w)\} B_t(j, v) \min_{w \in \mathcal{N}(v)} \{G_t(j, w)\}, \end{aligned}$$

where the last inequality follows by neglecting the truncation with β . Hence multiplying (II) by $q(v)$ and summing over $v \in \mathcal{V}$ yields

$$\begin{aligned} \sum_{v \in \mathcal{V}} q(v) \mathbb{E} \left[\frac{\eta}{2} \sum_{t=1}^T \left\| \widehat{\ell}_t(v) \right\|_{\nabla^2 F^*(\xi(v))}^2 \right] &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\left\| \widehat{\ell}_t(v) \right\|_{\nabla^2 F^*(\xi(v))}^2 \right] \right] \\ &\leq \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i,j=1}^k \bar{x}_t(i, v) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{G_t(i, w)\} B_t(j, v) \min_{w \in \mathcal{N}(v)} \{G_t(j, w)\} \right] \right], \end{aligned}$$

which is rewritten as

$$\begin{aligned}
 & \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i,j=1}^k \bar{x}_t(i,v) B_t(i,v) \min_{w \in \mathcal{N}(v)} \{G_t(i,w)\} B_t(j,v) \min_{w \in \mathcal{N}(v)} \{G_t(j,w)\} \right] \right] \\
 &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i,v) B_t(i,v) \tilde{G}_t(i,v) B_t(j,v) \tilde{G}_t(j,v) \right] \right] \\
 &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i,v) \mathbb{E}_t[B_t(i,v) B_t(j,v)] \mathbb{E}_t[\tilde{G}_t(i,v)] \mathbb{E}_t[\tilde{G}_t(j,v)] \right] =: (\star),
 \end{aligned}$$

where in the first equality we defined $\tilde{G}_t(i,v) = \min_{w \in \mathcal{N}(v)} \{G_t(i,w)\}$ and, analogously, $\tilde{G}_t(j,v) = \min_{w \in \mathcal{N}(v)} \{G_t(j,w)\}$, while the second follows by the conditional independence of the three terms $(B_t(i,v), B_t(j,v))$, $\tilde{G}_t(i,v)$, and $\tilde{G}_t(j,v)$ given the history up to time $t-1$. Furthermore, making use of Lemmas 1–3 and upper bounding, we get:

$$\begin{aligned}
 (\star) &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i=1}^k \sum_{j=1}^k \frac{\bar{x}_t(i,v)}{\bar{B}_t(i,v)} B_t(i,v) \frac{B_t(j,v)}{\bar{B}_t(j,v)} \right] \right] \\
 &\leq \frac{\eta e k}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{v \in \mathcal{V}} \frac{\bar{x}_t(i,v) q(v)}{\bar{B}_t(i,v)} \right] \leq \frac{\eta e k T}{2(1-e^{-1})} (k\alpha_1 + mQ),
 \end{aligned}$$

where the first equality uses the expected value of the geometric random variables \tilde{G} , the first inequality is obtained neglecting the indicator function $B_t(i,v)$ and taking the conditional expectation of $B_t(j,v)$, and the last inequality follows by a known upper bound involving independence numbers appearing, for example in [Cesa-Bianchi et al. \(2019a,b\)](#). For the sake of convenience, we add this result to Appendix E, Lemma 17. We now consider the last term (III). Since $\ell_t \geq \mathbb{E}_t[\hat{\ell}_t(v)]$ by Lemma 4, we have

$$\begin{aligned}
 \text{(III)} &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\langle \bar{x}_t(v) - a, \ell_t - \hat{\ell}_t(v) \rangle \right] \right] \leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\langle \bar{x}_t(v), \ell_t - \hat{\ell}_t(v) \rangle \right] \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \ell_t(i) \bar{x}_t(i,v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i,w) q(w)) \right)^\beta \right].
 \end{aligned}$$

Multiplying (III) by $q(v)$ and summing over the agents, we now upper bound $\ell_t(i)$ with 1 and use the facts that $1 - x \leq e^{-x}$ for all $x \in [0, 1]$ and $e^{-y} \leq 1/y$ for all $y > 0$, to obtain

$$\begin{aligned}
 & \sum_{v \in \mathcal{V}} q(v) \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \ell_t(i) \bar{x}_t(i, v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\
 & \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{v \in \mathcal{V}} \bar{x}_t(i, v) q(v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\
 & = \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \bar{x}_t(i, v) q(v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\
 & \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \bar{x}_t(i, v) q(v) \exp \left(-\beta \sum_{w \in \mathcal{N}(v)} \bar{x}_t(i, w) q(w) \right) \right] \\
 & \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \frac{\bar{x}_t(i, v) q(v)}{\beta \sum_{w \in \mathcal{N}(v)} \bar{x}_t(i, w) q(w)} \right] \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \frac{\alpha_1}{\beta} \right] = \frac{\alpha_1 k T}{\beta}
 \end{aligned}$$

where the last inequality follows by a known upper bound involving independence numbers appearing, for example in (Alon et al., 2017, Lemma 10). For the sake of convenience, we add this result to Appendix E, Lemma 16. Putting everything together and recalling that $\beta = \lfloor 1/(k\eta) \rfloor \geq 1/(2k\eta)$, we can finally conclude that

$$\begin{aligned}
 R_T & \leq Q \frac{m(1 + \log(k))}{\eta} + Q \frac{\eta k T}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + \frac{\alpha_1 k T}{\beta} \\
 & \leq Q \frac{m(1 + \log(k))}{\eta} + Q \frac{\eta k T}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + 2\eta \alpha_1 k^2 T \\
 & = Q \frac{m(1 + \log(k))}{\eta} + \eta Q k T \left(\frac{e}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + 2\alpha_1 \frac{k}{Q} \right) \\
 & \leq Q \frac{m(1 + \log(k))}{\eta} + 5\eta Q k T \left(\frac{k}{Q} \alpha_1 + m \right) \\
 & \leq 2Q \sqrt{15mkT \log(k)} \left(\frac{k}{Q} \alpha_1 + m \right).
 \end{aligned}$$

■

We conclude this section by discussing the computational complexity of our Coop-FTPL algorithm. The next result shows that the total number of elementary operations performed by Coop-FTPL over T time-steps scales with $T^{3/2}$ in the worst-case. To the best of our knowledge, no known algorithm attains a lower worst-case computational complexity.

Proposition 6 *If the optimization problem (2) can be solved with at most $c \in \mathbb{N}$ elementary operations, the worst-case computational complexity $\gamma_{\text{Coop-FTPL}}$ of each agent $v \in \mathcal{V}$ running our Coop-FTPL algorithm with the optimal tuning (8) for T rounds is*

$$\gamma_{\text{Coop-FTPL}} = \mathcal{O}\left(T^{3/2} c \sqrt{\frac{\alpha_1/Q + 1}{m}}\right).$$

Proof The result follows immediately by noting that the number of elementary operations performed by each agent v at each time step t is at most

$$c(\beta + 1) \leq c\left(\frac{1}{k\eta} + 1\right) = c\left(\frac{1}{k} \sqrt{\frac{5kT(k\alpha_1/Q + m)}{2m \log k}} + 1\right) = c\left(\sqrt{\frac{5T(\alpha_1/Q + m/k)}{2m \log k}} + 1\right).$$

■

5. Lower bound

In this section we show that no cooperative semi-bandit algorithm can beat the $Q\sqrt{mkT\alpha_1/Q}$ rate. The key idea for constructing the lower bound is simple: if the activation probabilities $q(v)$ are non-zero only for agents v belonging to an independent set with cardinality α_1 , then the problem is reduced to α_1 independent instances of single-agent semi-bandits, whose minimax rate is known.

Theorem 7 *For each communication network \mathcal{G} with independence number α_1 there exist cooperative semi-bandit instances for which the regret of any learning algorithm satisfies*

$$R_T = \Omega(Q\sqrt{mkT\alpha_1/Q}).$$

Proof Let $\mathcal{W} = \{w_1, \dots, w_{\alpha_1}\} \subset \mathcal{V}$ be an independent set with cardinality α_1 . Furthermore, let $q \in (0, 1]$ be a positive probability and for all agents $v \in \mathcal{V}$, let

$$q(v) = q\mathbb{I}\{v \in \mathcal{W}\}.$$

In words, only agents belonging to an independent set with largest cardinality are activated (with positive probability), and all with the same probability. Thus, only agents in \mathcal{W} contribute to the expected regret and their total mass $Q = \sum_{v \in \mathcal{V}} q(v)$ is equal to $\alpha_1 q$. Moreover, note that being non-adjacent, agents in \mathcal{W} never exchange any information. Each agent $w \in \mathcal{W}$ is therefore running an independent single-agent online linear optimization problem with semi-bandit feedback for an average of qT rounds. Since for single-agent semi-bandits, the worst-case lower bound on the regret after T' time steps is known to be $\Omega(\sqrt{mkT'})$ (see, e.g., Audibert et al. (2014); Lattimore et al. (2018)) and the cardinality of \mathcal{W} is α_1 , the regret of any cooperative semi-bandit algorithm run on this instance satisfies

$$R_T = \Omega(\alpha_1 \sqrt{mkqT}) = \Omega(\alpha_1 q \sqrt{mkT/q}) = \Omega(Q\sqrt{mkT\alpha_1/Q}),$$

where we used $Q = \alpha_1 q$. This concludes the proof. ■

In the previous section we showed that the expected regret of our Coop-FTPL algorithm can always be upper bounded by $Q\sqrt{mkT \log(k)(k\alpha_1/Q + m)}$ (ignoring constants). Thus, Theorem 7 shows that, up to the additive m term inside the rightmost bracket, the regret of Coop-FTPL is at most $\sqrt{k \log k}$ -away from the minimax optimal rate.

6. Conclusions and open problems

Motivated by spatially distributed large-scale learning systems, we introduced a new cooperative setting for adversarial semi-bandits in which only some of the agents are active at any given time step. We designed and analyzed an efficient algorithm that we called Coop-FTPL for which we proved near-optimal regret guarantees with state-of-the-art computational complexity costs. Our analysis relies on the fact that agents are aware of their activation probabilities, and they have some prior knowledge about the connectivity of the graph. Two interesting new lines of research are investigating if either of these assumptions could be lifted while retaining low regret and good computational complexity. In particular, removing the need for prior knowledge of the independence number would represent a significant theoretical and practical improvement, given that computing α_1 is NP-hard in the worst-case. It is however unclear if existing techniques that address this problem in similar settings (e.g., Cesa-Bianchi et al. (2019b)) would yield any results in our general case. We believe that entirely new ideas will be required to deal with this issue. We leave these intriguing problems open for future work.

Acknowledgments

This project has received funding from the French “Investing for the Future – PIA3” program under the Grant agreement ANITI ANR-19-PI3A-0004.⁵

References

- Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 46(6):1785–1826, 2017.
- Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2014.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Baruch Awerbuch and Robert Kleinberg. Competitive collaborative learning. *Journal of Computer and System Sciences*, 74(8):1271–1288, 2008.
- Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Nicolò Cesa-Bianchi, Tommaso R Cesari, and Claire Monteleoni. Cooperative online learning: Keeping your neighbors updated. *arXiv preprint arXiv:1901.08082*, 2019a.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Delay and cooperation in nonstochastic bandits. *The Journal of Machine Learning Research*, 20(1):613–650, 2019b.
- Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

5. <https://aniti.univ-toulouse.fr/>

- James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Wouter M Koolen, Manfred K Warmuth, Jyrki Kivinen, et al. Hedging structured concepts. In *COLT*, pages 93–105. Citeseer, 2010.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Tor Lattimore, Branislav Kveton, Shuai Li, and Csaba Szepesvari. Toprank: A practical algorithm for online stochastic ranking. In *Advances in Neural Information Processing Systems*, pages 3945–3954, 2018.
- Yin Tat Lee, Aaron Sidford, and Santosh S Vempala. Efficient convex optimization with membership oracles. In *Conference On Learning Theory*, pages 1292–1294. PMLR, 2018.
- David Martínez-Rubio, Varun Kanade, and Patrick Rebeschini. Decentralized cooperative stochastic bandits. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 4529–4540. Curran Associates, Inc., 2019.
- Gergely Neu and Gábor Bartók. An efficient algorithm for learning with semi-bandit feedback. In *International Conference on Algorithmic Learning Theory*, pages 234–248. Springer, 2013.
- Daiki Suehiro, Kohei Hatano, Shuji Kijima, Eiji Takimoto, and Kiyohito Nagano. Online prediction under submodular constraints. In *International Conference on Algorithmic Learning Theory*, pages 260–274. Springer, 2012.

Appendix A. Legendre functions and Fenchel conjugates

In this section, we briefly recall a few known definitions and facts in convex analysis.

Definition 8 (Interior, boundary, and convex hull) For any subset E of \mathbb{R}^k , we denote its topological interior by $\text{int}(E)$, its boundary by ∂E , and its convex hull by $\text{co}(E)$.

Definition 9 (Effective domain) The effective domain of a convex function $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ is

$$\text{dom}(F) := \{x \in \mathbb{R}^k : F(x) < +\infty\} . \quad (12)$$

With a slight abuse of notation, we will denote with the same symbol f a convex function $f: \rightarrow \mathbb{R} \cup \{+\infty\}$ and its restriction $\tilde{f}: \text{dom}(f) \rightarrow \mathbb{R}$ to its effective domain.

Definition 10 (Legendre function) A convex function $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ is Legendre if

1. $\text{int}(\text{dom}(F))$ is non-empty;
2. F is differentiable and strictly convex on $\text{int}(\text{dom}(F))$;
3. for all $x_0 \in \partial[\text{int}(\text{dom}(F))]$, if $x \in \text{int}(\text{dom}(F))$, $x \rightarrow x_0$, then $\|\nabla F(x)\|_2 \rightarrow +\infty$.

Definition 11 (Fenchel conjugate) Let $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function. The Fenchel conjugate F^* of F is defined as the function

$$F^*: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$$

$$z \mapsto F^*(z) := \sup_{x \in \mathbb{R}^k} (\langle x, z \rangle - F(x)) .$$

Definition 12 (Bregman divergence) Let $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ a convex function with non-empty $\text{int}(\text{dom}(F))$ that is differentiable on $\text{int}(\text{dom}(F))$. The Bregman divergence induced by F is

$$\mathcal{B}_F: \mathbb{R}^k \times \text{int}(\text{dom}(F)) \rightarrow \mathbb{R} \cup \{+\infty\}$$

$$(x, y) \mapsto \mathcal{B}_F(x, y) := F(x) - F(y) - \langle \nabla F(y), x - y \rangle .$$

The following results are taken from (Lattimore and Szepesvári, 2020, Theorem 26.6 and Corollary 26.8).

Theorem 13 Let $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be a Legendre function. Then:

1. the Fenchel conjugate F^* of F is Legendre;
2. $\nabla F: \text{int}(\text{dom}(F)) \rightarrow \text{int}(\text{dom}(F^*))$ is bijective with inverse $(\nabla F)^{-1} = \nabla F^*$;
3. $\mathcal{B}_F(x, y) = \mathcal{B}_{F^*}(\nabla F(y), \nabla F(x))$, for all $x, y \in \text{int}(\text{dom}(f))$.

Corollary 14 If $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ is a Legendre function and $x \in \arg\min_{x \in \text{dom}(F)} F(x)$, then $x \in \text{int}(\text{dom}(F))$.

Appendix B. Online Stochastic Mirror Descent (OSMD)

In this section, we briefly recall the standard Online Stochastic Mirror Descent algorithm (OSMD) (Algorithm 3) and its analysis.

For an overview on some basic convex analysis definitions and results, we refer the reviewer to the previous Appendix A. For a convex function $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ that is differentiable on the non-empty interior $\text{int}(\text{dom}(F)) \neq \emptyset$ of its effective domain $\text{dom}(F)$, we denote by \mathcal{B}_F the Bregman divergence induced by F (Definition 12). Following the existing convention, we refer to the input function F of OSMD as a *potential*. Furthermore, given a measure P on a subset of \mathbb{R}^k , we say that a vector $x \in \mathbb{R}^k$ is the *mean of the measure P* if x is the component-wise expectation of a \mathbb{R}^k -valued random variable with distribution P . For any time step $t \in \{1, 2, \dots\}$, we denote by \mathbb{E}_t the expectation conditioned to the history up to and including round $t - 1$.

It is known that since $\text{co}(\mathcal{A})$ is convex and compact, $\text{int}(\text{dom}(F)) \cap \text{co}(\mathcal{A}) \neq \emptyset$, and F is Legendre, then, all the argmin's exist in Algorithm 3 and $\bar{x}_t \in \text{int}(\text{dom}(F)) \cap \text{co}(\mathcal{A})$ for all $t \in \{1, 2, \dots\}$ (see, e.g., (Lattimore and Szepesvári, 2020, Exercise 28.2)).

The following result is taken from (Lattimore and Szepesvári, 2020, Theorem 28.10) and gives an upper bound on the regret of OSMD.

Theorem 15 Suppose that OSMD (Algorithm 3) is run with input F, \mathcal{A}, η . If the estimates $\widehat{\ell}_t$ computed at line 4 satisfy $\mathbb{E}_t[\widehat{\ell}_t] = \ell_t$ for all $t \in \{1, 2, \dots\}$, then, for all $x \in \text{co}(\mathcal{A})$,

$$\mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t, \ell_t \rangle - \sum_{t=1}^T \langle x, \ell_t \rangle \right] \leq \mathbb{E} \left[\frac{F(x) - F(\bar{x}_1)}{\eta} + \sum_{t=1}^T \langle \bar{x}_t - \bar{x}_{t+1}, \widehat{\ell}_t \rangle - \frac{1}{\eta} \sum_{t=1}^T \mathcal{B}_F(\bar{x}_{t+1}, \bar{x}_t) \right] .$$

Algorithm 3: Online Stochastic Mirror Descent (OSMD)

Input: Legendre potential $F: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$, compact action set $\mathcal{A} \subset \mathbb{R}^k$ with $\text{int}(\text{dom}(F)) \cap \text{co}(\mathcal{A}) \neq \emptyset$, learning rate $\eta > 0$

Initialization: $\bar{x}_1 = \text{argmin}_{x \in \text{dom}(F) \cap \text{co}(\mathcal{A})} F(x)$

- 1 **for** $t = 1, 2, \dots$ **do**
 - 2 choose a measure P_t on \mathcal{A} with mean \bar{x}_t
 - 3 make a prediction x_t drawn from \mathcal{A} according to P_t and suffer the loss $\langle x_t, \ell_t \rangle$
 - 4 compute an estimate $\hat{\ell}_t$ of the loss vector ℓ_t
 - 5 make the update: $\bar{x}_{t+1} = \text{argmin}_{x \in \text{dom}(F) \cap \text{co}(\mathcal{A})} \eta \langle x, \hat{\ell}_t \rangle + \mathcal{B}_F(x, \bar{x}_t)$
-

Furthermore, letting

$$\tilde{x}_{t+1} = \text{argmin}_{x \in \text{dom}(F)} \eta \langle x, \hat{\ell}_t \rangle + \mathcal{B}_F(x, \bar{x}_t)$$

and assuming that $-\eta \hat{\ell}_t + \nabla F(x) \in \nabla F(\text{dom}(F))$ for all $x \in \text{co}(\mathcal{A})$ almost surely, then

$$\sup_{x \in \text{co}(\mathcal{A})} \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t, \ell_t \rangle - \sum_{t=1}^T \langle x, \ell_t \rangle \right] \leq \frac{\text{diam}_F(\text{co}(\mathcal{A}))}{\eta} + \frac{1}{\eta} \sum_{t=1}^T \mathbb{E} [\mathcal{B}_F(\bar{x}_t, \tilde{x}_{t+1})],$$

where $\text{diam}_F(\text{co}(\mathcal{A})) := \sup_{x, y \in \text{co}(\mathcal{A})} (F(x) - F(y))$ is the diameter of $\text{co}(\mathcal{A})$ with respect to F .

Appendix C. Proofs of lemmas on geometric distributions

In this section we provide all missing proofs on geometric distributions that we stated in Section 4.

Proof of Lemma 1 For all $j \in \{1, \dots, m\}$, the cumulative distribution function (c.d.f.) of Y_j is given, for all $n \in \mathbb{N}$, by

$$\mathbb{P}[Y_j \leq n] = p_j \sum_{i=1}^n (1 - p_j)^{i-1} = 1 - (1 - p_j)^n.$$

The c.d.f. of Z is given, for all $n \in \mathbb{N}$, by

$$\begin{aligned} \mathbb{P}[Z \leq n] &= \mathbb{P} \left[\min_{j \in \{1, \dots, m\}} Y_j \leq n \right] = 1 - \prod_{j=1}^m \mathbb{P}[Y_j > n] = 1 - \prod_{j=1}^m (1 - \mathbb{P}[Y_j \leq n]) \\ &= 1 - \prod_{j=1}^m (1 - (1 - (1 - p_j)^n)) = 1 - \left(\prod_{j=1}^m (1 - p_j) \right)^n \\ &= 1 - \left(1 - \left[1 - \prod_{j=1}^m (1 - p_j) \right] \right)^n, \end{aligned}$$

and this is the c.d.f. of a geometric random variable with parameter $(1 - \prod_{j=1}^m (1 - p_j))$. ■

Proof of Lemma 2 By elementary calculations,

$$\begin{aligned}
 \mathbb{E}[\min\{G, \beta\}] &= \sum_{n=1}^{\infty} n(1-q)^{n-1}q - \sum_{n=\beta}^{\infty} (n-\beta)(1-q)^{n-1}q \\
 &= \sum_{n=1}^{\infty} n(1-q)^{n-1}q - (1-q)^{\beta} \sum_{n=\beta}^{\infty} (n-\beta)(1-q)^{n-\beta-1}q \\
 &= \left(1 - (1-q)^{\beta}\right) \sum_{n=1}^{\infty} n(1-q)^{n-1}q = \frac{\left(1 - (1-q)^{\beta}\right)}{q}.
 \end{aligned}$$

■

Proof of Lemma 3 The proof follows immediately from the fact that $\{X_s(v) Y_s(v)\}_{s \in \mathcal{I}, v \in \mathcal{V}}$ is a collection of independent Bernoulli random variables with expectation $\mathbb{E}[X_s(v) Y_s(v)] = p_1(v) p_2(v)$ for any $s \in \mathbb{N}$ and all $v \in \mathcal{V}$. ■

Appendix D. Proof of Theorem 5

In this section, we present a complete proof of Theorem 5.

Proof of Theorem 5 For the sake of convenience, we define the expected individual regret of an agent $v \in \mathcal{V}$ in the network with respect to a fixed action $a \in \mathcal{A}$ by

$$R_T(a, v) := \mathbb{E} \left[\sum_{t=1}^T \langle x_t(v), \ell_t \rangle - \sum_{t=1}^T \langle a, \ell_t \rangle \right],$$

where the expectation is taken with respect to the internal randomization of the agent, but not its activation probability $q(v)$. With this definition the total regret on the network in Eq. (4) can be decomposed as

$$\begin{aligned}
 R_T &= \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{S}_t} \left(\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right) \right] = \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{v \in \mathcal{S}_t} \left(\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right) \right] \right] \\
 &= \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \sum_{v \in \mathcal{V}} q(v) \mathbb{E}_t \left[\langle x_t(v), \ell_t \rangle - \langle a, \ell_t \rangle \right] \right] = \max_{a \in \mathcal{A}} \sum_{v \in \mathcal{V}} q(v) R_T(a, v). \tag{13}
 \end{aligned}$$

The proof then proceeds by isolating the bias in the loss estimators. For each $a \in \mathcal{A}$ we get

$$\begin{aligned}
 R_T(a, v) &= \mathbb{E} \left[\sum_{t=1}^T \langle x_t(v) - a, \ell_t \rangle \right] = \mathbb{E} \left[\mathbb{E}_t \left[\sum_{t=1}^T \langle x_t(v) - a, \ell_t \rangle \right] \right] = \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \ell_t \rangle \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \widehat{\ell}_t(v) \rangle \right] + \mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \ell_t - \widehat{\ell}_t(v) \rangle \right] \\
 &\leq \underbrace{\frac{F(\bar{x}_1(v)) - F(a)}{\eta}}_{\text{(I)}} + \underbrace{\mathbb{E} \left[\frac{1}{\eta} \sum_{t=1}^T \mathcal{B}_F(\bar{x}_t(v), \bar{x}_{t+1}(v)) \right]}_{\text{(II)}} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T \langle \bar{x}_t(v) - a, \ell_t - \widehat{\ell}_t(v) \rangle \right]}_{\text{(III)}}
 \end{aligned}$$

where the inequality follows by the standard analysis of OMD. We bound the three terms separately. For the first term (I), we have

$$\begin{aligned}
 F(a) &= \sup_{x \in \mathbb{R}^k} (\langle a, x \rangle - F^*(x)) = \sup_{x \in \mathbb{R}^k} \left(\langle a, x \rangle - \mathbb{E} \left[\max_{a' \in \mathcal{A}} \langle a', x + Z \rangle \right] \right) \\
 &\geq -\mathbb{E} \left[\max_{a' \in \mathcal{A}} \langle a', Z \rangle \right] \geq -m \mathbb{E}[\|Z\|_\infty] = -m \sum_{i=1}^k \frac{1}{i} \geq -m(1 + \log(k)), \quad (14)
 \end{aligned}$$

where the first inequality follows by choosing $x = 0$, the second follows from Hölder's inequality and $\|a\|_1 \leq m$ for any $a \in \mathcal{A}$, and the last equality is Exercise 30.4 in [Lattimore and Szepesvári \(2020\)](#). It follows that

$$F(\bar{x}_1(v)) - F(a) \leq m(1 + \log(k)),$$

where we use the fact that $F(a) \leq 0$ for all $a \in \mathcal{A}$ and this follows from the first line of Eq. (14) by the convexity of the maximum of a finite number of linear functions, using Jensen's inequality and the fact that the random variable Z is centered. Thus

$$\text{(I)} \leq \frac{m(1 + \log k)}{\eta}.$$

We now study the second term (II). We have

$$\begin{aligned}
 \mathcal{B}_F(\bar{x}_t(v), \bar{x}_{t+1}(v)) &= \mathcal{B}_{F^*}(\nabla F(\bar{x}_{t+1}(v)), \nabla F(\bar{x}_t(v))) \\
 &= \mathcal{B}_{F^*}(-\eta \widehat{L}_{t-1}(v) - \eta \widehat{\ell}_t(v), -\eta \widehat{L}_{t-1}(v)) \\
 &= \frac{\eta^2}{2} \left\| \widehat{\ell}_t(v) \right\|_{\nabla^2 F^*(\xi(v))}^2, \quad (15)
 \end{aligned}$$

where the first equality is a standard property of Bregmann divergence, the second follows from the definitions of the updates and the last by Taylor's theorem, where $\xi(v) = -\eta \widehat{L}_{t-1}(v) - \alpha \eta \widehat{\ell}_t(v)$, for some $\alpha \in [0, 1]$. To calculate the Hessian we use a change of variable to avoid applying the gradient

to the (possibly) non-differentiable argmax and we get:

$$\begin{aligned}
 \nabla^2 F^*(x) &= \nabla(\nabla F^*(x)) = \nabla \mathbb{E}[h(x+Z)] = \nabla \int_{\mathbb{R}^k} h(x+z)\zeta(z)dz \\
 &= \nabla \int_{\mathbb{R}^k} h(u)\zeta(u-x)du = \int_{\mathbb{R}^k} h(u)(\nabla \zeta(u-x))^\top du \\
 &= \int_{\mathbb{R}^k} h(u)\text{sign}(u-x)^\top \zeta(u-x)du = \int_{\mathbb{R}^k} h(x+z)\text{sign}(z)^\top \zeta(z)dz
 \end{aligned}$$

Using the definition of $\xi(v)$ and the fact that $h(x)$ is nonnegative,

$$\begin{aligned}
 \nabla^2 F^*(\xi(v))_{ij} &= \int_{\mathbb{R}^k} h(\xi(v)+z)_i \text{sign}(z)_j \zeta(z)dz \\
 &\leq \int_{\mathbb{R}^k} h(\xi(v)+z)_i \zeta(z)dz \\
 &= \int_{\mathbb{R}^k} h\left(z - \eta \widehat{L}_{t-1}(v) - \alpha \eta \widehat{\ell}_t(v)\right)_i \zeta(z)dz \\
 &= \int_{\mathbb{R}^k} h\left(u - \eta \widehat{L}_{t-1}(v)\right)_i \zeta\left(u + \alpha \eta \widehat{\ell}_t(v)\right)du \\
 &\leq \exp\left(\left\|\alpha \eta \widehat{\ell}_t(v)\right\|_1\right) \int_{\mathbb{R}^k} h\left(u - \eta \widehat{L}_{t-1}(v)\right)_i \zeta(u)du \\
 &\leq \exp\left(\alpha \eta \sum_{i=1}^k B_t(i, v) \beta\right) \bar{x}_t(i, v) \\
 &\leq \exp(\alpha \eta k \beta) \bar{x}_t(i, v) \\
 &\leq e \bar{x}_t(i, v)
 \end{aligned}$$

where the last inequality follows by $\alpha \leq 1$ and $\beta \leq 1/(\eta k)$. Plugging this estimate in Eq. (15) yields

$$\begin{aligned}
 \frac{\eta^2}{2} \left\|\widehat{\ell}_t(v)\right\|_{\nabla^2 F^*(\xi(v))}^2 &= \frac{\eta^2}{2} \sum_{i=1}^k \sum_{j=1}^k \nabla^2 F^*(\xi(v))_{i,j} \widehat{\ell}_t(i, v) \widehat{\ell}_t(j, v) \\
 &\leq \frac{\eta^2 e}{2} \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i, v) \widehat{\ell}_t(i, v) \widehat{\ell}_t(j, v) \\
 &\leq \frac{\eta^2 e}{2} \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i, v) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{G_t(i, w)\} B_t(j, v) \min_{w \in \mathcal{N}(v)} \{G_t(j, w)\},
 \end{aligned}$$

where the last inequality follows by neglecting the truncation with β . Hence multiplying (II) by $q(v)$ and summing over $v \in \mathcal{V}$ yields

$$\begin{aligned}
 \sum_{v \in \mathcal{V}} q(v) \mathbb{E} \left[\frac{\eta}{2} \sum_{t=1}^T \left\|\widehat{\ell}_t(v)\right\|_{\nabla^2 F^*(\xi(v))}^2 \right] &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\left\|\widehat{\ell}_t(v)\right\|_{\nabla^2 F^*(\xi(v))}^2 \right] \right] \\
 &\leq \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i,j=1}^k \bar{x}_t(i, v) B_t(i, v) \min_{w \in \mathcal{N}(v)} \{G_t(i, w)\} B_t(j, v) \min_{w \in \mathcal{N}(v)} \{G_t(j, w)\} \right] \right],
 \end{aligned}$$

which is rewritten as

$$\begin{aligned}
 & \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i,j=1}^k \bar{x}_t(i,v) B_t(i,v) \min_{w \in \mathcal{N}(v)} \{G_t(i,w)\} B_t(j,v) \min_{w \in \mathcal{N}(v)} \{G_t(j,w)\} \right] \right] \\
 &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i,v) B_t(i,v) \tilde{G}_t(i,v) B_t(j,v) \tilde{G}_t(j,v) \right] \right] \\
 &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{j=1}^k \bar{x}_t(i,v) \mathbb{E}_t[B_t(i,v) B_t(j,v)] \mathbb{E}_t[\tilde{G}_t(i,v)] \mathbb{E}_t[\tilde{G}_t(j,v)] \right] =: (\star),
 \end{aligned}$$

where in the first equality we defined $\tilde{G}_t(i,v) = \min_{w \in \mathcal{N}(v)} \{G_t(i,w)\}$ and, analogously, $\tilde{G}_t(j,v) = \min_{w \in \mathcal{N}(v)} \{G_t(j,w)\}$, while the second follows by the conditional independence of the three terms $(B_t(i,v), B_t(j,v))$, $\tilde{G}_t(i,v)$, and $\tilde{G}_t(j,v)$ given the history up to time $t-1$. Furthermore, making use of Lemmas 1–3, we get

$$\begin{aligned}
 (\star) &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i=1}^k \sum_{j=1}^k \frac{\bar{x}_t(i,v)}{\bar{B}_t(i,v)} B_t(i,v) \frac{B_t(j,v)}{\bar{B}_t(j,v)} \right] \right] \\
 &\leq \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\sum_{i=1}^k \sum_{j=1}^k \frac{\bar{x}_t(i,v)}{\bar{B}_t(i,v)} \frac{B_t(j,v)}{\bar{B}_t(j,v)} \right] \right] \\
 &= \sum_{v \in \mathcal{V}} q(v) \frac{\eta e}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{j=1}^k \frac{\bar{x}_t(i,v)}{\bar{B}_t(i,v)} \frac{\bar{B}_t(j,v)}{\bar{B}_t(j,v)} \right] \\
 &= \frac{\eta e k}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{v \in \mathcal{V}} \frac{\bar{x}_t(i,v) q(v)}{\bar{B}_t(i,v)} \right] \\
 &\leq \frac{\eta e k}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \left(\frac{1}{1-e^{-1}} \left(\alpha_1 + \sum_{v \in \mathcal{V}} \bar{x}_t(i,v) q(v) \right) \right) \right] \\
 &= \frac{\eta e k}{2} \mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{1-e^{-1}} \left(k \alpha_1 + \sum_{v \in \mathcal{V}} \sum_{i=1}^k \bar{x}_t(i,v) q(v) \right) \right) \right] \\
 &\leq \frac{\eta e k T}{2(1-e^{-1})} (k \alpha_1 + m Q),
 \end{aligned}$$

where the first equality uses the expected value of the geometric random variables \tilde{G} , the first inequality is obtained neglecting the indicator function $B_t(i,v)$, the following equality uses the expected value of the geometric random variables B_t , the second inequality follows by Lemma 17.

We now consider the last term (III). Since $\ell_t \geq \mathbb{E}[\widehat{\ell}_t(v)]$, from Lemma 4, we have

$$\begin{aligned} \text{(III)} &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\left\langle \bar{x}_t(v) - a, \ell_t - \widehat{\ell}_t(v) \right\rangle \right] \right] \leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t \left[\left\langle \bar{x}_t(v), \ell_t - \widehat{\ell}_t(v) \right\rangle \right] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \ell_t(i) \bar{x}_t(i, v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right]. \end{aligned}$$

Multiplying (III) by $q(v)$ and summing over the agents, we can now upper bound $\ell_t(i)$ with 1, then we use facts that $1 - x \leq e^{-x}$ for $x \in [0, 1]$ and that $e^{-y} \leq 1/y$ for all $y > 0$, to obtain

$$\begin{aligned} &\sum_{v \in \mathcal{V}} q(v) \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \ell_t(i) \bar{x}_t(i, v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{v \in \mathcal{V}} \bar{x}_t(i, v) q(v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \bar{x}_t(i, v) q(v) \left(\prod_{w \in \mathcal{N}(v)} (1 - \bar{x}_t(i, w) q(w)) \right)^\beta \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \bar{x}_t(i, v) q(v) \exp \left(-\beta \sum_{w \in \mathcal{N}(v)} \bar{x}_t(i, w) q(w) \right) \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \sum_{\substack{v \in \mathcal{V} \\ \bar{x}_t(i, v) q(v) > 0}} \frac{\bar{x}_t(i, v) q(v)}{\beta \sum_{w \in \mathcal{N}(v)} \bar{x}_t(i, w) q(w)} \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^k \frac{\alpha_1}{\beta} \right] = \frac{\alpha_1 k T}{\beta} \end{aligned}$$

where in the last inequality follows by Lemma 16. Putting all together and recalling that $\beta = \lfloor \frac{1}{k\eta} \rfloor \geq \frac{1}{2k\eta}$, we can conclude that for every $a \in \mathcal{A}$, thanks to Eq. (13), we have

$$\begin{aligned}
 R_T &\leq Q \frac{m(1 + \log(k))}{\eta} + Q \frac{\eta e k T}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + \frac{\alpha_1 k T}{\beta} \\
 &\leq Q \frac{m(1 + \log(k))}{\eta} + Q \frac{\eta e k T}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + 2\eta \alpha_1 k^2 T \\
 &= Q \frac{m(1 + \log(k))}{\eta} + \eta Q k T \left(\frac{e}{2(1 - e^{-1})} \left(\frac{k}{Q} \alpha_1 + m \right) + 2\alpha_1 \frac{k}{Q} \right) \\
 &\leq Q \frac{m(1 + \log(k))}{\eta} + 5\eta Q k T \left(\frac{k}{Q} \alpha_1 + m \right) \\
 &\leq 2Q \sqrt{15mkT \log(k)} \left(\frac{k}{Q} \alpha_1 + m \right).
 \end{aligned}$$

■

Appendix E. Bounds on independence numbers

The two following lemmas provide upper bounds of sums of weights over nodes of a graph expressed in terms of its independence number.

Lemma 16 *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with indecence number α_1 , $q(v) \geq 0$, and $Q(v) = \sum_{w \in \mathcal{N}(v)} q(w) > 0$ for all $v \in \mathcal{V}$. Then*

$$\sum_{v \in \mathcal{V}} \frac{q(v)}{Q(v)} \leq \alpha_1$$

Proof Initialize $V_1 = \mathcal{V}$, fix $w_1 \in \operatorname{argmin}_{w \in V_1} Q(w)$, and denote $V_2 = \mathcal{V} \setminus \mathcal{N}(w_1)$. For $k \geq 2$ fix $w_k \in \operatorname{argmin}_{w \in V_k} Q(w)$ and shrink $V_{k+1} = V_k \setminus \mathcal{N}(w_k)$ until $V_{k+1} = \emptyset$. Since \mathcal{G} is undirected $w_k \notin \bigcup_{s=1}^{k-1} \mathcal{N}(w_s)$, therefore the number m of times that an action can be picked this way is upper bounded by α_1 . Denoting $\mathcal{N}'(w_k) = V_k \cap \mathcal{N}(w_k)$ this implies

$$\begin{aligned}
 \sum_{v \in \mathcal{V}} \frac{q(v)}{Q(v)} &= \sum_{k=1}^m \sum_{v \in \mathcal{N}'(w_k)} \frac{q(v)}{Q(v)} \leq \sum_{k=1}^m \sum_{v \in \mathcal{N}'(w_k)} \frac{q(v)}{Q(w_k)} \\
 &\leq \sum_{k=1}^m \frac{\sum_{v \in \mathcal{N}(w_k)} q(v)}{Q(w_k)} = m \leq \alpha_1
 \end{aligned}$$

concluding the proof. ■

Lemma 17 *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with independence number α_1 . For each $v \in \mathcal{V}$, let $\mathcal{N}(v)$ be the neighborhood of node v (including v itself), and $p(1, v), \dots, p(k, v) \geq 0$. Then, for all $i \in \{1, \dots, k\}$,*

$$\sum_{v \in \mathcal{V}} \frac{p(i, v)}{q(i, v)} \leq \frac{1}{1 - e^{-1}} \left(\alpha_1 + \sum_{v \in \mathcal{V}} p(i, v) \right) \quad \text{where} \quad q(i, v) = 1 - \prod_{w \in \mathcal{N}(v)} (1 - p(i, w)).$$

Proof Fix $i \in \{1, \dots, k\}$ and set for brevity $P(i, v) = \sum_{w \in \mathcal{N}(v)} p(i, w)$. We can write

$$\sum_{v \in \mathcal{V}} \frac{p(i, v)}{q(i, v)} = \underbrace{\sum_{v \in \mathcal{V}: P(i, v) \geq 1} \frac{p(i, v)}{q(i, v)}}_{\text{(I)}} + \underbrace{\sum_{v \in \mathcal{V}: P(i, v) < 1} \frac{p(i, v)}{q(i, v)}}_{\text{(II)}},$$

and proceed by upper bounding the two terms (I) and (II) separately. Let $r(v)$ be the cardinality of $\mathcal{N}(v)$. We have, for any given $v \in \mathcal{V}$,

$$\min \left\{ q(i, v) : \sum_{w \in \mathcal{N}(v)} p(i, w) \geq 1 \right\} = 1 - \left(1 - \frac{1}{r(v)} \right)^{r(v)} \geq 1 - e^{-1}.$$

The equality is due to the fact that the minimum is achieved when $p(i, w) = \frac{1}{r(v)}$ for all $w \in \mathcal{N}(v)$, and the inequality comes from $r(v) \geq 1$ (for, $v \in \mathcal{N}(v)$). Hence

$$\text{(I)} \leq \sum_{v \in \mathcal{V}: P(i, v) \geq 1} \frac{p(i, v)}{1 - e^{-1}} \leq \sum_{v \in \mathcal{V}} \frac{p(i, v)}{1 - e^{-1}}.$$

As for (II), using the inequality $1 - x \leq e^{-x}$, $x \in [0, 1]$, with $x = p(i, w)$, we can write

$$q(i, v) \geq 1 - \exp \left(- \sum_{w \in \mathcal{N}(v)} p(i, w) \right) = 1 - \exp(-P(i, v)).$$

In turn, because $P(i, v) < 1$ in terms (II), we can use the inequality $1 - e^{-x} \geq (1 - e^{-1})x$, holding when $x \in [0, 1]$, with $x = P(i, v)$, thereby concluding that

$$q(i, v) \geq (1 - e^{-1})P(i, v)$$

Thus

$$\text{(II)} \leq \sum_{v \in \mathcal{V}: P(i, v) < 1} \frac{p(i, v)}{(1 - e^{-1})P(i, v)} \leq \frac{1}{1 - e^{-1}} \sum_{v \in \mathcal{V}} \frac{p(i, v)}{P(i, v)} \leq \frac{\alpha_1}{1 - e^{-1}},$$

where in the last step we used Lemma 16. ■