

A Closer Look at Small-loss Bounds for Bandits with Graph Feedback

Chung-Wei Lee

Haipeng Luo

Mengxiao Zhang

University of Southern California

LEECHUNG@USC.EDU

HAIPENGL@USC.EDU

MENGXIAO.ZHANG@USC.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We study *small-loss* bounds for adversarial multi-armed bandits with graph feedback, that is, adaptive regret bounds that depend on the loss of the best arm or related quantities, instead of the total number of rounds. We derive the first small-loss bound for general strongly observable graphs, resolving an open problem of Lykouris et al. (2018). Specifically, we develop an algorithm with regret $\tilde{O}(\sqrt{\kappa L_*})$ where κ is the clique partition number and L_* is the loss of the best arm, and for the special case of self-aware graphs where every arm has a self-loop, we improve the regret to $\tilde{O}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_*}\})$ where $\alpha \leq \kappa$ is the independence number. Our results significantly improve and extend those by Lykouris et al. (2018) who only consider self-aware undirected graphs.

Furthermore, we also take the first attempt at deriving small-loss bounds for weakly observable graphs. We first prove that no typical small-loss bounds are achievable in this case, and then propose algorithms with alternative small-loss bounds in terms of the loss of some specific subset of arms. A surprising side result is that $\tilde{O}(\sqrt{T})$ regret is achievable even for weakly observable graphs as long as the best arm has a self-loop.

Our algorithms are based on the Online Mirror Descent framework but require a suite of novel techniques that might be of independent interest. Moreover, all our algorithms can be made parameter-free without the knowledge of the environment.

Keywords: multi-armed bandits, feedback graph, small-loss bounds.

1. Introduction

Adversarial multi-armed bandits with graph feedback is an online learning model that generalizes the classic expert problem (Freund and Schapire, 1997) as well as the standard multi-armed bandits problem (Auer et al., 2002). In this model, the learner needs to choose one of the K arms at each round, while simultaneously the adversary decides the loss of each arm. Afterwards, the learner receives feedback based on a graph with the K arms as nodes. Specifically, the learner observes the loss of every arm to which the chosen arm is connected. Clearly, the full-information expert problem corresponds to having a complete feedback graph, while the standard multi-armed bandits problem corresponds to having a feedback graph with only self-loops.

Alon et al. (2015) provided a full characterization of the minimax regret for this problem. Specifically, it was shown that the minimax regret for strongly observable graphs and that for weakly observable graphs are $\tilde{\Theta}(\sqrt{\alpha T})$ and $\tilde{\Theta}(d^{1/3}T^{2/3})$ respectively, where T is the total number of rounds and α and d are the independence number and weak domination number of the feedback graph respectively (see Section 2 for definitions).

Table 1: Main results and comparisons with prior work. T is the number of rounds, $L_\star \leq T$ is the total loss of the best arm, α , κ , and d are the independence, clique partition, and weak domination number respectively. For our results for weakly observable graphs, γ can be any value in $[1/3, 1/2]$, i^\star is the best arm, S is the set of nodes with a self-loop, $L_{i^\star_S}$ is the loss of the best arm in S , $L_{\mathcal{D}}$ is the average loss of nodes in a weakly dominating set, and dependence on other parameters is omitted. See Section 2 for detailed definitions. All our algorithms have parameter-free versions.

Graph Type		Regret		
		Minimax (Alon et al., 2015)	(Lykouris et al., 2018)	Our work
Strongly Observable	General	$\tilde{\Theta}(\sqrt{\alpha T})$	N/A	$\tilde{\mathcal{O}}(\sqrt{(\kappa + 1)L_\star})$
	Special case: self-aware		$\tilde{\mathcal{O}}(\alpha^{1/3} L_\star^{2/3}), \tilde{\mathcal{O}}(\sqrt{\kappa L_\star})$ (undirected graphs only)	$\tilde{\mathcal{O}}(\min\{\sqrt{\kappa L_\star}, \sqrt{\alpha T}\})$
Weakly Observable	General	$\tilde{\Theta}(d^{1/3} T^{2/3})$	N/A	(no $o(L_\star)$ bounds achievable) $\begin{cases} \tilde{\mathcal{O}}(L_{\mathcal{D}}^{1-\gamma}), & \text{if } i^\star \in S, \\ \tilde{\mathcal{O}}(L_{\mathcal{D}}^{(1+\gamma)/2}), & \text{else.} \end{cases}$
	Special case: bipartite	$\tilde{\Theta}(T^{2/3})$		$\begin{cases} \tilde{\mathcal{O}}(\sqrt{L_\star}), & \text{if } i^\star \in S, \\ \tilde{\mathcal{O}}(L_{i^\star_S}^{2/3}), & \text{else.} \end{cases}$

However, it is well-known that more adaptive data-dependent regret bounds are achievable for a wide range of online learning problems. Among them, perhaps the most common one is the so-called *first-order* or *small-loss* bounds, which replaces the dependence on T by the total loss of the best arm $L_\star \leq T$. Such bounds are usually never worse than the worst-case bounds, but could be potentially much smaller if a relatively good arm exists. Achieving small-loss bounds for bandits with graph feedback has been surprisingly challenging. Lykouris et al. (2018) took the first attempt and their algorithms achieve regret $\tilde{\mathcal{O}}(\alpha^{1/3} L_\star^{2/3})$ or $\tilde{\mathcal{O}}(\sqrt{\kappa L_\star})$ (κ is the clique partition number), but *only for self-aware undirected graphs* (self-aware means that every node has a self-loop). It was left as a major open problem whether better and more general small-loss bounds are achievable.

Our work makes a significant step towards a full understanding of small-loss bounds for bandits with a fixed directed feedback graph. Specifically, our contributions are (see also Table 1):

- (Section 3.3) For general strongly observable graphs, we develop an algorithm with regret $\tilde{\mathcal{O}}(\sqrt{(\kappa + 1)L_\star})$. This is the first small-loss bound for the general case, extending the results of (Lykouris et al., 2018) that only hold for self-aware undirected graphs and resolving an open problem therein.
- (Section 3.4) For the special case of self-aware (directed) graphs, we develop an algorithm with regret $\tilde{\mathcal{O}}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\})$, again strictly improving (Lykouris et al., 2018) by providing an extra robustness guarantee (note that $\alpha \leq \kappa$ always holds).
- (Section 4) For weakly observable graphs (where small-loss bounds have not been studied before at all), we prove that *no algorithm can achieve typical small-loss bounds* (such as $o(L_\star)$). Despite this negative result, we develop an algorithm with regret $\tilde{\mathcal{O}}(L_{\mathcal{D}}^{2/3})$ where $L_{\mathcal{D}}$ is the average loss of a weakly dominating set (and dependence on other parameters is omitted for simplicity). More generally, we also achieve different trade-offs between the case when

the best arm has a self-loop and the case without, such as $\tilde{O}(\sqrt{L_{\mathcal{D}}})$ versus $\tilde{O}(L_{\mathcal{D}}^{3/4})$. We further consider a special case with a complete bipartite graph, and show that our algorithm achieves $\tilde{O}(\sqrt{L_{\star}})$ when the best arm has a self-loop and $\tilde{O}(L_{i_S^{\star}}^{2/3})$ otherwise, where $L_{i_S^{\star}}$ is the loss of the best arm with a self-loop. A surprising implication of our result is that $\tilde{O}(\sqrt{T})$ regret is possible even for weakly observable graphs as long as the best arm has a self-loop.

- (Appendix) We provide parameter-free versions of all our algorithms using sophisticated doubling tricks, which we emphasize is highly non-trivial for bandit settings, especially because some of our algorithms consist of a layer structure combining different subroutines.

Our algorithms are based on the well-known Online Mirror Descent framework, but importantly with a suite of different techniques including hybrid regularizers, unconstrained loss shifting trick, increasing learning rates, combining algorithms with partial information, adding correction terms to loss estimators, and their combination in an innovative way. We defer further discussion on the novelty of each component and comparisons with prior work to the description of each algorithm.

Related work. The bandits with graph feedback model was first proposed by (Mannor and Shamir, 2011). Later, Alon et al. (2015, 2017) gave a full characterization of the minimax regret for this problem. There are many follow-ups that consider different variants and extensions of this problem, such as (Kocák et al., 2016; Feng and Loh, 2018; Rangi and Franceschetti, 2019; Arora et al., 2019).

Small-loss bounds have been widely studied in the online learning literature. For the full-information expert problem, the classic Hedge algorithm (Freund and Schapire, 1997) achieves $\tilde{O}(\sqrt{L_{\star}})$ regret already. For the standard multi-armed bandits problem and its variant semi-bandits, there are also several different approaches to achieve $\tilde{O}(\sqrt{L_{\star}})$ regret (Allenberg et al., 2006; Neu, 2015; Foster et al., 2016; Wei and Luo, 2018; Bubeck and Sellke, 2020). Even for the challenging contextual bandits setting (which is in fact a special case of learning with time-varying feedback graphs), it was shown by Allen-Zhu et al. (2018) that $\tilde{O}(\sqrt{L_{\star}})$ regret is also achievable.

The work most related to ours is (Lykouris et al., 2018). As mentioned, we significantly extend their results to more general graphs, including graphs with directed edges, graphs without self-loops, and even weakly observable graphs, and we also improve their bound for self-aware graphs. Our algorithms are also based on very different ideas compared to theirs which are mainly built on the recursive arm freezing technique. We point out that, however, they also studied high probability bounds and time-varying graphs for some cases, which is not the focus of this work.

2. Problem Setup and Notations

Throughout the paper, we denote $\{1, \dots, m\}$ by $[m]$ for some positive integer m . Before the game starts, the adversary decides a sequence of T loss vectors $\ell_1, \dots, \ell_T \in [0, 1]^K$ for some integers $K \geq 2$ and $T \geq 2K$, and a directed feedback graph $G = ([K], E)$ for $E \subseteq [K] \times [K]$ which is fixed and known. Each node in the graph represents one of the K arms, and in this paper we use the terms “arm” and “node” interchangeably. At each round $t \in [T]$, the learner selects an arm $i_t \in [K]$ and incurs loss ℓ_{t, i_t} . At the end of this round, the learner receives feedback according to G . Specifically, the learner observes the loss of arm i for all i such that $i_t \in N^{\text{in}}(i)$, where $N^{\text{in}}(i) \triangleq \{j : (j, i) \in E\}$ is the set of nodes that can observe i . The *regret* with respect to an arm i is defined as $\text{Reg}_i \triangleq \mathbb{E} \left[\sum_{t=1}^T \ell_{t, i_t} - \sum_{t=1}^T \ell_{t, i} \right]$, which is the expected difference between the learner’s total loss and that of arm i (the expectation is with respect to the learner’s internal

randomness). We denote the best arm by $i^* \triangleq \operatorname{argmin}_{i \in [K]} \sum_{t=1}^T \ell_{t,i}$ and define $\operatorname{Reg} \triangleq \operatorname{Reg}_{i^*}$. Alon et al. (2015) show that the minimax regret (in terms of T) for strongly observable graphs and weakly observable graphs (definitions to follow) are $\operatorname{Reg} = \tilde{\Theta}(\sqrt{T})$ and $\operatorname{Reg} = \tilde{\Theta}(T^{2/3})$ respectively.

Our goal is to obtain the so-called *small-loss* regret bounds that could potentially be much smaller than the minimax bounds. Specifically, for an arm i , we denote its total loss by $L_i \triangleq \sum_{t=1}^T \ell_{t,i} \leq T$, and we use the shorthand $L_* \triangleq L_{i^*}$. Our goal is to replace the dependence of T by L_* when bounding Reg , or more generally, to replace T by L_i when bounding Reg_i for each arm i . Below, we introduce some graph-related notions and other notations necessary for discussions.

Observability. A node i is *observable* if $N^{\text{in}}(i) \neq \emptyset$. An observable node is *strongly observable* if either $i \in N^{\text{in}}(i)$ or $N^{\text{in}}(i) = [K] \setminus \{i\}$, and *weakly observable* otherwise. Similarly, a graph is observable if all of its nodes are observable. An observable graph is strongly observable if all nodes are strongly observable, and weakly observable otherwise.

Special cases. We denote by $S \triangleq \{i \in [K] : i \in N^{\text{in}}(i)\}$ the subset of nodes with a self-loop, and by $\bar{S} \triangleq [K] \setminus S$ the subset of nodes without a self-loop. We further use s and \bar{s} to denote $|S|$ and $|\bar{S}|$. A graph is *self-aware* if $S = [K]$, that is, every node has a self-loop, which is a special case of strongly observable graphs. We also consider a special case of weakly observable graphs with $(i, j) \in E$ for every $i \in S$ and every $j \in \bar{S}$, and call it a *directed complete bipartite graph*.¹

Independence sets and cliques. We remind the reader the standard concepts of independence sets and cliques. An independent set \mathcal{I} is a subset of nodes such that for any two distinct nodes $i, j \in \mathcal{I}$, we have $(i, j) \notin E$. The independence number of a graph is the cardinality of its largest independent set. A clique \mathcal{C} is a subset of nodes such that for any two distinct nodes $i, j \in \mathcal{C}$, we have $(i, j) \in E$. A clique partition $\{\mathcal{C}_1, \dots, \mathcal{C}_m\}$ of a graph is a partition of its nodes such that each \mathcal{C}_k in this collection is a clique. The clique partition number of a graph is the cardinality of its smallest clique partition. As in previous works, our bounds for strongly observable graphs depend on the independence number α of G , or the clique partition number κ of the subgraph G_S obtained by restricting G to only the nodes in S . Note that $\alpha \leq \kappa + 1$ always holds.

Some of our algorithms rely on having a clique partition of G_S , which we assume is given, even though it is in general NP-hard to find (Karp, 1972). We emphasize that, however, our algorithms work with any clique partition and the bounds hold with κ replaced by the size of this partition.

Weakly dominating sets. Following (Alon et al., 2015), for a weakly observable graph, we define a *weakly dominating set* \mathcal{D} to be a set of nodes such that all weakly observable nodes can be observed by at least one node in \mathcal{D} . Our bounds for weakly observable graphs depend on the *weak domination number* d of graph G , which is the cardinality of its smallest weakly dominating set. Similarly, we assume that a weakly dominating set of size d is given, but our algorithms work using any weakly dominating set and our bounds hold with d replaced by the size of this set.

Other notations. For a differentiable convex function ψ defined on a convex set Ω , the associated Bregman divergence is defined as $D_\psi(x, y) = \psi(x) - \psi(y) - \langle \nabla \psi(y), x - y \rangle$ for any two points $x, y \in \Omega$. For a positive definite matrix $M \in \mathbb{R}^{K \times K}$, we define norm $\|z\|_M \triangleq \sqrt{z^\top M z}$ for any vector $z \in \mathbb{R}^K$. For two matrices M_1 and M_2 , $M_1 \preceq M_2$ means that $M_2 - M_1$ is positive semi-definite, and for two vectors v_1 and v_2 , $v_1 \preceq v_2$ means that v_1 is coordinate-wise less than

1. Note that unlike the traditional definition of bipartite graphs, here we allow additional edges other than those from S to \bar{S} , as adding more edges only makes the problem easier.

or equal to v_2 . We denote the $(K - 1)$ -dimensional simplex by $\Delta(K)$, the diagonal matrix with v_1, \dots, v_K on the diagonal by $\text{diag}\{v_1, \dots, v_K\}$, and the all-zero and all-one vectors in \mathbb{R}^K by $\mathbf{0}$ and $\mathbf{1}$ respectively. The notation $\tilde{\mathcal{O}}(\cdot)$ hides logarithmic dependence on K and T .

3. Strongly Observable Graphs

In this section, we focus on strongly observable graphs. We first show how to achieve $\text{Reg} = \tilde{\mathcal{O}}(\sqrt{(\kappa + 1)L_\star})$ in general, and then discuss how to further improve it to $\tilde{\mathcal{O}}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\})$ for the special case of self-aware graphs (Section 3.4).

There are three key components/ideas to achieve $\tilde{\mathcal{O}}(\sqrt{(\kappa + 1)L_\star})$ regret. Specifically, starting from the EXP3.G algorithm of (Alon et al., 2015), which is an instance of Online Mirror Descent (OMD) with the entropy regularizer, natural loss estimators for graph feedback, and an additional $\Theta(1/\sqrt{T})$ amount of uniform exploration, we make the following three modifications:

- (Section 3.1) Reduce the amount of uniform exploration to $\Theta(1/T)$ and add a constant amount of log-barrier to the regularizer. We show that this modification maintains the same $\tilde{\mathcal{O}}(\sqrt{\alpha T})$ regret as EXP3.G, but importantly, the smaller amount of uniform exploration is the key for further obtaining small-loss bounds.
- (Section 3.2) Replace the entropy regularizer with the log-barrier regularizer for nodes in S . This leads to a small-loss bound of order $\tilde{\mathcal{O}}(\sqrt{(s + 1)L_\star})$.
- (Section 3.3) Create a clique partition for nodes in S , run a Hedge variant within each clique, and run the algorithm from Section 3.2 treating each clique as a meta-node, which finally improves the regret to $\tilde{\mathcal{O}}(\sqrt{(\kappa + 1)L_\star})$. This part relies on highly nontrivial extensions of techniques from (Agarwal et al., 2017) on combining algorithms in the bandit setting.

3.1. Reducing the Amount of Uniform Exploration

We start by describing the EXP3.G algorithm of (Alon et al., 2015). It maintains a distribution $p_t \in \Delta(K)$ for each time t , and samples i_t according to p_t . Then a standard importance-weighted loss estimator $\hat{\ell}_t$ is constructed such that

$$\hat{\ell}_{t,i} = \frac{\ell_{t,i}}{W_{t,i}} \mathbb{1}\{i_t \in N^{\text{in}}(i)\} \quad \text{where } W_{t,i} \triangleq \sum_{j \in N^{\text{in}}(i)} p_{t,j}. \quad (1)$$

It is clear that $\mathbb{E}[\hat{\ell}_{t,i}] = \ell_{t,i}$, that is, the estimator is unbiased. With such a loss estimator, the distribution is updated according to the classic OMD algorithm:

$$p_{t+1} = \underset{p \in \Omega}{\text{argmin}} \left\langle p, \hat{\ell}_t \right\rangle + D_\psi(p, p_t), \quad \text{with } p_1 = \underset{p \in \Omega}{\text{argmin}} \psi(p). \quad (2)$$

For EXP3.G, $\psi(p) = \frac{1}{\eta} \sum_{i \in [K]} p_i \ln p_i$ is the standard Shannon entropy regularizer with learning rate $\eta \leq 1/2$, and $\Omega = \left\{ p \in \Delta(K) : p_i \geq \frac{2\eta}{K}, \forall i \in [K] \right\}$ is the clipped simplex and enforces $\mathcal{O}(\eta)$ amount of uniform exploration.²

2. In the original EXP3.G algorithm, Ω is the exact simplex $\Delta(K)$ and uniform exploration is enforced by sampling i_t according to p_t with probability $1 - 2\eta$ and according to a uniform distribution with probability 2η . Nevertheless, our slight modification essentially serves the same purpose and makes subsequent discussions easier.

Standard OMD analysis shows that the instantaneous regret of EXP3.G against any $u \in \Delta(K)$ at time t is bounded as $\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \|\hat{\ell}_t\|_{\nabla^{-2}\psi(p_t)}^2$. However, the last term $\|\hat{\ell}_t\|_{\nabla^{-2}\psi(p_t)}^2$ (often called the local-norm term) could be prohibitively large for general strongly observable graphs. The analysis of EXP3.G overcomes this issue via a key *loss shifting* trick. Specifically, it is shown that the following more general bound holds

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 \quad (3)$$

for any $z \leq 1/\eta$, and in particular, with $z = \sum_{i \in S} p_{t,i} \hat{\ell}_{t,i}$, the local-norm term $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2$ is bounded by $\tilde{O}(\eta\alpha)$ in expectation. This choice of z is indeed not larger than $1/\eta$ due to the form of $\hat{\ell}_t$ and *importantly the $\mathcal{O}(\eta)$ amount of uniform exploration*. The rest of the analysis is straightforward and shows that $\text{Reg} = \tilde{O}(\frac{1}{\eta} + \eta\alpha T)$, which is $\tilde{O}(\sqrt{\alpha T})$ by picking $\eta = 1/\sqrt{\alpha T}$.

To obtain small-loss bounds, one clear obstacle in EXP3.G is the uniform exploration, which contributes to $\mathcal{O}(\eta T) = \mathcal{O}(\sqrt{T})$ regret already by the above optimal choice of η . Our first step is thus to get rid of this large amount uniform exploration, and we take an approach that is completely different from (Lykouris et al., 2018). Specifically, noting that the key reason to have uniform exploration is the constraint $z \leq 1/\eta$ in the loss shifting trick Eq. (3), our key idea is to *remove this constraint completely*, which turns out to be possible if the regularizer contains a constant amount of log-barrier (similar to (Bubeck et al., 2018; Zheng et al., 2019)), as shown in the following lemma.

Lemma 1 (Unconstrained Loss Shifting) *Let $p_{t+1} = \arg\min_{p \in \Omega} \langle p, \hat{\ell}_t \rangle + D_\psi(p, p_t)$, for $\Omega \subseteq \Delta(K)$ and $\psi : \Omega \rightarrow \mathbb{R}$. Suppose (a) $0 \leq \hat{\ell}_{t,i} \leq \max\{\frac{1}{p_{t,i}}, \frac{1}{1-p_{t,i}}\}$, $\forall i \in [K]$, (b) $\nabla^{-2}\psi(p) \preceq 4\nabla^{-2}\psi(q)$ holds when $p \preceq 2q$, (c) $\nabla^2\psi(p) \succeq \text{diag}\{64K/p_1^2, \dots, 64K/p_K^2\}$, $\forall p \in \Omega$. Then we have*

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 8 \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2. \quad (4)$$

Condition (a) is clearly satisfied for strongly observable graphs if $\hat{\ell}_t$ is defined by Eq. (1) since $W_{t,i} \geq p_{t,i}$ for $i \in S$ and $W_{t,i} = 1 - p_{t,i}$ for $i \notin S$. Condition (b) is trivially satisfied for all common regularizers for the simplex such as Shannon entropy, Tsallis entropy, log-barrier, and any of their combinations. Finally, to ensure Condition (c), one only needs to include a log-barrier component $c \sum_{i \in [K]} \ln \frac{1}{p_i}$ for $c \geq 64K$ in the regularizer, whose Hessian is exactly $\text{diag}\{c/p_1^2, \dots, c/p_K^2\}$. This inspires us to propose the following hybrid regularizer

$$\psi(p) = \frac{1}{\eta} \sum_{i \in [K]} p_i \ln p_i + c \sum_{i \in [K]} \ln \frac{1}{p_i}, \quad (5)$$

and we prove the following theorem.

Theorem 2 *The OMD update Eq. (2) with $\Omega = \{p \in \Delta(K) : p_i \geq \frac{1}{T}, \forall i \in [K]\}$, $\hat{\ell}_t$ defined in Eq. (1), and ψ defined in Eq. (5) for $c = 64K$ ensures $\text{Reg} \leq \tilde{O}(\frac{1}{\eta} + \eta\alpha T + K^2)$ for any strongly observable graph. Choosing $\eta = 1/\sqrt{\alpha T}$, we have $\text{Reg} = \tilde{O}(\sqrt{\alpha T} + K^2)$.*

Note that we still enforce a small $1/T$ amount of uniform exploration, which is important for a technical lemma (Alon et al., 2015, Lemma 5), but this only contributes $\mathcal{O}(K)$ regret. Also, adding the log-barrier leads to a small $\mathcal{O}(K^2)$ overhead in the Bregman divergence term $D_\psi(u, p_1)$, but only makes the local-norm term smaller and thus $\min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2$ is still of order $\tilde{O}(\eta\alpha)$ in expectation. The proof of Theorem 2 is now straightforward and is deferred to Appendix A.

3.2. $\tilde{\mathcal{O}}(\sqrt{(s+1)L_\star})$ Regret Bound

Having solved the uniform exploration issue, we now discuss our first attempt to obtain small-loss bounds for strongly observable graphs. Inspired by the fact that for multi-armed bandits, that is, the case where E contains all the self-loops but nothing else, OMD with the log-barrier regularizer achieves a small-loss bound (Foster et al., 2016), we propose to replace the entropy regularizer with the log-barrier for all nodes in S , while keeping the hybrid regularizer Eq. (5) for nodes in \bar{S} :

$$\psi(p) = \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{p_i} + \frac{1}{\eta} \sum_{i \in \bar{S}} p_i \ln p_i + c \sum_{i \in \bar{S}} \ln \frac{1}{p_i}. \quad (6)$$

We note that it is important not to also use $1/\eta$ amount of log-barrier for nodes in \bar{S} , since this leads to an overhead of K/η for the Bregman divergence term and thus at best gives $\tilde{\mathcal{O}}(\sqrt{KL_\star})$ regret. We prove the following theorem for our proposed algorithm.

Theorem 3 *OMD Eq. (2) with $\Omega = \{p \in \Delta(K) : p_i \geq \frac{1}{T}, \forall i \in [K]\}$, $\hat{\ell}_t$ defined in Eq. (1), and ψ defined in Eq. (6) for $c = 64K$ and $\eta \leq \frac{1}{64K}$ ensures $\text{Reg} = \mathcal{O}\left(\frac{s \ln T + \ln K}{\eta} + \eta L_\star + K^2 \ln T\right)$ for any strongly observable graph. Choosing $\eta = \min\left\{\sqrt{\frac{s+1}{L_\star}}, \frac{1}{64K}\right\}$ gives $\tilde{\mathcal{O}}\left(\sqrt{(s+1)L_\star} + K^2\right)$.³*

While the algorithmic idea is straightforward, the main challenge in the analysis is to deal with the nodes in \bar{S} . Specifically, for the particular choices of η and c , we know that the conditions of Lemma 1 hold, and the key is thus again to bound the local-norm term $\min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2$. Simply taking $z = 0$ or the previous choice $z = \sum_{i \in \bar{S}} p_{t,i} \hat{\ell}_{t,i}$ from (Alon et al., 2015) does not give the desired bound. Instead, we propose a novel shift: $z = \hat{\ell}_{t,i_0}$ for some $i_0 \in \bar{S}$ with $p_{t,i_0} \geq 1/2$, or $z = 0$ if no such i_0 exists. Direct calculations then show $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 = \mathcal{O}(\eta \langle p_t, \hat{\ell}_t \rangle)$. In expectation, the local-norm term is thus related to the loss of the algorithm $\langle p_t, \ell_t \rangle$, and it is well-known that this is enough for obtaining small-loss bounds. For the complete proof, see Appendix B.

3.3. $\tilde{\mathcal{O}}(\sqrt{(\kappa+1)L_\star})$ Regret Bound

Finally, we discuss how to further improve our bound to $\tilde{\mathcal{O}}(\sqrt{(\kappa+1)L_\star})$. Let $\mathcal{C}_1, \dots, \mathcal{C}_\kappa$ be a clique partition of G_S (recall that G_S is G restricted to S). Essentially, we compress each clique as one meta-node, and together with nodes from \bar{S} , this creates a meta-graph with $\beta \triangleq \kappa + \bar{s}$ nodes, which can be seen as a “low-resolution” version of G . We index these meta-nodes as $1, \dots, \kappa$, and without loss of generality we assume that the original indices of nodes in \bar{S} are $\kappa + 1, \dots, \kappa + \bar{s}$, so that $[\beta]$ is the set of nodes for this meta-graph, and $[\kappa]$ is the set of nodes with a self-loop (taking the same role as S in the original graph). If we were actually dealing with a problem with this meta-graph, running the algorithm from Theorem 3 would thus give $\tilde{\mathcal{O}}(\sqrt{(\kappa+1)L_\star})$ regret.

To solve the original problem, however, we need to specify what to do when a meta-node is selected. Note that within a meta-node, we are essentially facing the classic expert problem with full-information (Freund and Schapire, 1997) since nodes are all connected with each other. A natural idea is thus to run an expert algorithm with a small-loss bound within each clique, and when a

3. In fact, the s dependence can be improved to the number of nodes that are not observed by every other nodes (by using log-barrier only for these nodes). However, we simplify the presentation with a looser bound since this improvement does not help improve the final main result in Section 3.3.

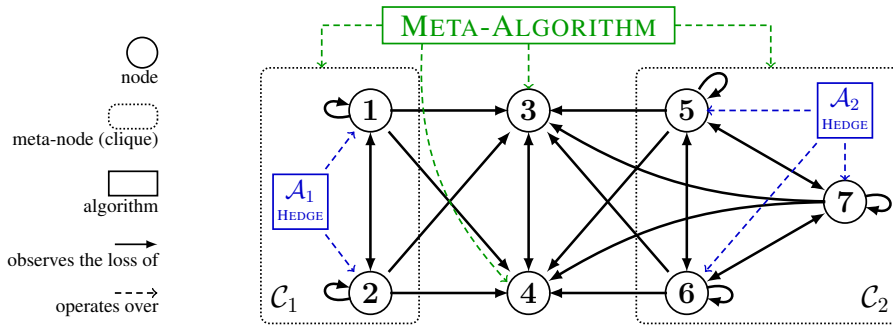


Figure 1: An illustration of Algorithm 1 for a graph with 7 nodes. Here, we have $S = \{1, 2, 5, 6, 7\}$, $\bar{S} = \{3, 4\}$, and $\kappa = 2$ with $C_1 = \{1, 2\}$ and $C_2 = \{5, 6, 7\}$ being a minimum clique partition of G_S . The meta-algorithm operates over nodes 3 and 4, and also the two cliques, each with a Hedge instance running inside.

clique is selected, follow the suggestion of the corresponding expert algorithm. We choose to use an adaptive version of Hedge (Freund and Schapire, 1997) as the expert algorithm, with details deferred to Algorithm 3 in Appendix C. Importantly, the regret of Hedge has only logarithmic dependence on the number of nodes and thus does not defeat the purpose of obtaining $\tilde{O}(\sqrt{(\kappa + 1)L_*})$ regret.

Figure 1 illustrates the main idea of our algorithm, and Algorithm 1 shows the complete pseudocode. We use i to index a node in the original graph and j to index a node in the meta-graph. Note that nodes from \bar{S} appear in both graphs so they are indexed by either i or j , depending on the context. The κ instances of Hedge are denoted by A_1, \dots, A_κ , where A_j only operates over nodes in C_j . For notational convenience, however, we require A_j to output at time t a distribution $\tilde{p}_t^{(j)} \in \Delta(K)$ over all nodes with the constraint $\tilde{p}_{t,i}^{(j)} = 0$ for all $i \notin C_j$ (Line 1), and we also feed the estimated losses of all nodes to A_j (Line 6), even though it ignores those outside C_j .

The algorithm maintains a distribution $p_t \in \Delta(\beta)$ for the meta-graph, updated using the algorithm from Theorem 3 (Line 8). The only difference is that we use a time-varying regularizer defined in Eq. (7), where the learning rate $\eta_{t,j}$ for meta-node $j \in [\kappa]$ is time-varying and also different for different j (all starting from $\eta_{1,j} = \eta$; more explanation to follow). At the beginning of time t , the algorithm first samples $j_t \in p_t$. If j_t happens to be a node in \bar{S} , we play $i_t = j_t$; otherwise, we sample i_t from the distribution received from A_{j_t} . See Line 2 and Line 3.

After playing arm i_t and receiving loss feedback, we construct loss estimator $\tilde{\ell}_t \in \mathbb{R}^K$ for nodes in G (Line 5) and estimator $\hat{\ell}_t \in \mathbb{R}^\beta$ for nodes in the meta-graph (Line 7). The estimator for nodes in \bar{S} , for either G or the meta-graph, is exactly the same as the standard one described in Eq. (1). The estimator for a node $i \in S$ also essentially follows Eq. (1), except that we ignore all edges that point to i but are not from those nodes in the same clique, so the probability of observing i is $p_{t,j}$ for j being the index of the clique to which i belongs.⁴ Finally, the estimator for a meta-node $j \in [\kappa]$ is simply $\langle \tilde{p}_t^{(j)}, \tilde{\ell}_t \rangle$, which is the estimated loss of the corresponding Hedge A_j .

While the idea of combining algorithms in such a two-level hierarchy is natural and straightforward, doing it in a partial-information setting is notoriously challenging and requires extra tech-

4. One could also follow exactly Eq. (1) to construct more complicated estimators, but it does not lead to a better bound.

Algorithm 1 Algorithm for General Strongly Observable Graphs

Input: Feedback graph G and a clique partition $\{\mathcal{C}_1, \dots, \mathcal{C}_\kappa\}$ of G_S , parameters η, c .

Define: $\beta = \kappa + \bar{s}$ and $\Omega = \{p \in \Delta(\beta) : p_j \geq \frac{1}{T}, \forall j \in [\beta]\}$.

Initialize: $p_1 = \operatorname{argmin}_{p \in \Omega} \psi_1(p)$ (see Eq. (7)), $\eta_{1,j} = \eta, \rho_{1,j} = 2\kappa, \forall j \in [\kappa]$.

Initialize: an instance \mathcal{A}_j of adaptive Hedge (Algorithm 3) with nodes in $\mathcal{C}_j, \forall j \in [\kappa]$.

for $t = 1, 2, \dots, T$ **do**

- 1 **For each** $j \in [\kappa]$, receive $\tilde{p}_t^{(j)} \in \{p \in \Delta(K) : p_{t,i} = 0, \forall i \notin \mathcal{C}_j\}$ from \mathcal{A}_j .
 - 2 Sample $j_t \sim p_t$.
 - 3 **if** $j_t \in [\kappa]$ **then** draw $i_t \sim \tilde{p}_t^{(j_t)}$; **else** let $i_t = j_t$.
 - 4 Pull arm i_t and receive feedback $\ell_{t,i}$ for all i such that $i_t \in N^{\text{in}}(i)$.
 - 5 Construct estimator $\tilde{\ell}_t \in \mathbb{R}^K$ such that $\tilde{\ell}_{t,i} = \begin{cases} \frac{\ell_{t,i}}{p_{t,j_t}} \mathbb{1}\{j_t \in [\kappa], i \in \mathcal{C}_{j_t}\}, & \text{for } i \in S, \\ \frac{\ell_{t,i}}{1-p_{t,i}} \mathbb{1}\{i \neq i_t\}, & \text{for } i \in \bar{S}. \end{cases}$
 - 6 **For each** $j \in [\kappa]$, feed $\tilde{\ell}_t$ to \mathcal{A}_j .
 - 7 Construct estimator $\hat{\ell}_t \in \mathbb{R}^\beta$ for meta-nodes such that $\hat{\ell}_{t,j} = \begin{cases} \langle \tilde{p}_t^{(j)}, \tilde{\ell}_t \rangle, & \text{for } j \in [\kappa], \\ \tilde{\ell}_{t,j}, & \text{for } j \in \bar{S}. \end{cases}$
 - 8 Compute $p_{t+1} = \operatorname{argmin}_{p \in \Omega} \left\{ \langle p, \hat{\ell}_t \rangle + D_{\psi_t}(p, p_t) \right\}$ where

$$\psi_t(p) = \sum_{j \in [\kappa]} \frac{1}{\eta_{t,j}} \ln \frac{1}{p_j} + \frac{1}{\eta} \sum_{j \in \bar{S}} p_j \ln p_j + c \sum_{j \in \bar{S}} \ln \frac{1}{p_j}. \quad (7)$$
 - 9 **for** $j \in [\kappa]$ **do**
 - if** $\frac{1}{p_{t+1,j}} > \rho_{t,j}$ **then** set $\rho_{t+1,j} = \frac{2}{p_{t+1,j}}, \eta_{t+1,j} = \eta_{t,j} e^{\frac{1}{\ln T}}$.
 - else** set $\rho_{t+1,j} = \rho_{t,j}, \eta_{t+1,j} = \eta_{t,j}$.
-

niques, as explained in detail in (Agarwal et al., 2017). In a word, the difficulty is that the Hedge algorithms do not always receive feedback and thus do not yield the usual regret bound as one would get for a full-information problem. To address this issue, we apply the increasing learning rate technique from (Agarwal et al., 2017). Specifically, we maintain a threshold $\rho_{t,j}$ for each time t and each meta-node j (starting from $\rho_{1,j} = 2\kappa$). Every time after the OMD update, if $p_{t+1,j}$ becomes too small and $1/p_{t+1,j}$ exceeds the threshold $\rho_{t,j}$, we increase the learning rate for j by a factor of $e^{\frac{1}{\ln T}}$ and set the new threshold to be $\rho_{t+1,j} = 2/p_{t+1,j}$ (Line 9). The high-level idea behind this technique is that when the probability of picking a clique is small and thus the corresponding Hedge receives little feedback, increasing the corresponding learning rate allows its faster recovery, should a node in the clique become the best node later. For more intuition, we refer the reader to (Agarwal et al., 2017). We are now ready to show the main theorem of this section.

Theorem 4 Algorithm 1 with $c = 64\beta$ and $\eta \leq \eta_{\max} \triangleq \min \left\{ \frac{1}{64\beta}, \frac{1}{1000(\ln T) \ln^2(KT)} \right\}$ guarantees:

$$\text{Reg} \leq \tilde{O} \left(\frac{\kappa \ln T + \ln K}{\eta} + \eta L_\star + \beta^2 \ln T \right)$$

for any strongly observable graph. Choosing $\eta = \min \left\{ \eta_{max}, \sqrt{\frac{\kappa+1}{L_\star}} \right\}$ gives $\tilde{O} \left(\sqrt{(\kappa+1)L_\star} + \beta^2 \right)$.

Our algorithm strictly improves over the GREEN-IX-GRAPH algorithm of (Lykouris et al., 2018) (in terms of expected regret), which achieves the same bound but only for undirected self-aware graphs. The proof of Theorem 4 is deferred to Appendix C. Moreover, in Appendix C.3 we also provide an adaptive version of our algorithm with a sophisticated doubling trick, which achieves the same bound but without the need of knowing L_\star for learning rate tuning. We remark that doing doubling trick in such a two-level structure and with partial information is highly non-trivial.

3.4. $\tilde{O}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\})$ Regret Bound for Self-Aware Graphs

Although our bound $\tilde{O}(\sqrt{(\kappa+1)L_\star})$ could be much smaller than the worst-case bound $\tilde{O}(\sqrt{\alpha T})$, it is not always better since $\alpha \leq \kappa + 1$. To remedy this drawback, we propose another algorithm with $\tilde{O}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\})$ regret for the special case of self-aware graphs. The high-level idea is to first run an algorithm with $\tilde{O}(\sqrt{\kappa L_\star})$ regret while keeping an estimate of L_\star , and when we are confident that $\tilde{O}(\sqrt{\alpha T})$ is the better bound, switch to any algorithm with $\tilde{O}(\sqrt{\alpha T})$ regret. For the first part, using Algorithm 1 would create some technical issues and we are unable to analyze it unfortunately (otherwise we could have dealt with general strongly observable graphs). Instead, we introduce a new algorithm using a similar clipping technique of (Lykouris et al., 2018).

We emphasize that the key challenge here is that the algorithm has to be adaptive in the sense that it does not need the knowledge of L_\star — otherwise, simply comparing the two bounds and running the corresponding algorithm with the better bound solves the problem already. We again design a sophisticated doubling trick to resolve this issue. All details are included in Appendix D.

4. Weakly Observable Graphs

In this section, we consider small-loss bounds for weakly observable graphs, which have not been studied before. Recall that the minimax regret bound in this case is $\tilde{O}(d^{1/3}T^{2/3})$ where d is the weak domination number. The most natural small-loss bound one would hope for is therefore $\tilde{O}(d^{1/3}L_\star^{2/3})$. However, it turns out that this is not achievable, and in fact, no typical small-loss bounds are achievable for any weakly observable graph, as we prove in the following theorem.

Theorem 5 *For any weakly observable graph and any algorithm \mathcal{A} (without the knowledge of L_\star), if \mathcal{A} guarantees $\tilde{O}(1)$ regret when $L_\star = 0$ (ignoring dependence on K), then there exists a sequence of loss vectors such that the regret of \mathcal{A} is $\Omega(T^{1-\epsilon})$ for any $\epsilon \in (0, 1/3)$.*

Note that this precludes small-loss bounds such as $\tilde{O}(L_\star^{2/3})$, or even $\tilde{O}(\min\{L_\star^a, T^{2/3}\})$ for any $a > 0$. The proof crucially relies on the fact that for a weakly observable graph, one can always find a pair of nodes u and v such that neither of them can observe u . See Appendix E for details.

Despite this negative result, nevertheless, we provide alternative first-order regret bounds in terms of the loss of some specific subset of nodes. Specifically, ignoring dependence on other parameters, for the special case of directed complete bipartite graphs, we obtain regret $\tilde{O}(\sqrt{L_\star})$ when $i^\star \in S$ and $\tilde{O}(L_{i^\star_S}^{2/3})$ otherwise, where $i^\star_S = \operatorname{argmin}_{i \in S} \sum_{t=1}^T \ell_{t,i}$ is the best node with a self-loop. Moreover, for general weakly observable graphs, we achieve regret $\tilde{O}(\sqrt{L_{\mathcal{D}}})$ when $i^\star \in S$ and $\tilde{O}(L_{\mathcal{D}}^{3/4})$ otherwise (or other different trade-offs between the two cases), where \mathcal{D} is a weakly dominating set and $L_{\mathcal{D}} \triangleq \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \sum_{t=1}^T \ell_{t,i}$ is the average total loss of nodes in \mathcal{D} .

Algorithm 2 Algorithm for Weakly Observable Graphs

Input: Feedback graph G , decision set Ω , parameter $\eta \leq \frac{1}{5}$ and $\bar{\eta}$.

Define: hybrid regularizer $\psi(p) = \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{p_i} + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} p_i \ln p_i$.

Initialize: p_1 is such that $p_{1,i} = \frac{1}{2s}$ for $i \in S$ and $p_{1,i} = \frac{1}{2\bar{s}}$ for $i \in \bar{S}$.

for $t = 1, 2, \dots, T$ **do**

- 1 Play arm $i_t \sim p_t$ and receive feedback $\ell_{t,i}$ for all i such that $i_t \in N^{\text{in}}(i)$.
 - 2 Construct estimator $\hat{\ell}_t$ such that $\hat{\ell}_{t,i} = \frac{\ell_{t,i}}{W_{t,i}} \mathbb{1}\{i_t \in N^{\text{in}}(i)\}$ where $W_{t,i} = \sum_{j \in N^{\text{in}}(i)} p_{t,j}$.
 - 3 Construct correction term a_t such that $a_{t,i} = \begin{cases} 2\eta p_{t,i} \hat{\ell}_{t,i}^2, & \text{for } i \in S, \\ 2\bar{\eta} \hat{\ell}_{t,i}^2, & \text{for } i \in \bar{S}. \end{cases}$
 - 4 Compute $p_{t+1} = \operatorname{argmin}_{p \in \Omega} \{\langle p, \hat{\ell}_t + a_t \rangle + D_\psi(p, p_t)\}$.
-

Our algorithm is summarized in [Algorithm 2](#). The key algorithmic idea is inspired by the work of ([Wei and Luo, 2018](#)). They show that a variant of OMD with certain correction terms added to the loss estimators leads to a regret bound on Reg_i where the typical local-norm term $\|\hat{\ell}_t\|_{\nabla^{-2}\psi(p_t)}^2$ is replaced by a term that is *only in terms of the information of arm i* . For our problem this is the key to achieve different orders of regret for different arms. More specifically, the algorithm performs a standard OMD update, except that $\hat{\ell}_t$ is replaced by $\hat{\ell}_t + a_t$ for some correction terms a_t ([Line 4](#)).

Before specifying our correction term, we first describe the regularizer. Similar to the algorithms for strongly observable graphs, we again use a hybrid regularizer $\psi(p) = \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{p_i} + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} p_i \ln p_i$, that is, log-barrier for nodes in S and entropy for nodes in \bar{S} . Note that we do not enforce a small amount of log-barrier for every node as in [Section 3](#) (reasons to follow). Also note that the learning rate for nodes in S and \bar{S} are different (η and $\bar{\eta}$ respectively), which is also crucial for getting different orders of regret for different nodes. In light of this specific choice of regularizer, our correction term a_t is defined as in [Line 3](#), because $\eta p_{t,i} \hat{\ell}_{t,i}^2$ is the typical correction term for log-barrier ([Wei and Luo, 2018](#)), and on the other hand $\bar{\eta} \hat{\ell}_{t,i}^2$ is the typical one for entropy ([Steinhardt and Liang, 2014](#)). Mixing these two correction terms is novel as far as we know.

The estimator $\hat{\ell}_t$ is constructed exactly by [Eq. \(1\)](#), and it remains to specify the decision set $\Omega \subseteq \Delta(K)$, which is different for different cases and will be discussed separately. In both cases, the decision set is such that some relatively large amount of uniform exploration is enforced over a subset of nodes, which is also the reason why the small amount of log-barrier is not needed anymore. Similar to the analysis of ([Wei and Luo, 2018](#)), we prove the following lemma (see [Appendix E](#)).

Lemma 6 *Algorithm 2 ensures $\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \langle u, a_t \rangle$ for all $u \in \Omega$, as long as Ω is a subset of $\{p \in \Delta(K) : \sum_{j \in N^{\text{in}}(i)} p_{t,j} \geq 5\bar{\eta}, \forall i \in \bar{S}\}$.*

Naturally, to compete with node i , we let u almost concentrate on i , in which case $\langle u, a_t \rangle$ is roughly $a_{t,i}$ (only in terms of i ; key difference compared to [Eq. \(4\)](#)). To understand why this is useful, consider the case when $i \in S$ so $a_{t,i}$ is $\eta p_{t,i} \hat{\ell}_{t,i}^2$. By the construction of the loss estimator, the latter is bounded by $\eta \ell_{t,i}$ in expectation, which is the key of getting $\tilde{O}(\sqrt{T})$ regret in this case.

Directed Complete Bipartite Graphs. For the special case of directed complete bipartite graphs, we take $\Omega = \{p \in \Delta(K) : \sum_{i \in S} p_i \geq \sqrt{\bar{\eta}}\}$, which ensures that every node in \bar{S} is observed with

probability at least $\sqrt{\bar{\eta}}$ (by the definition of directed complete bipartite graphs). This, however, unavoidably introduces dependence on $L_{i_S^*}$ when bounding Reg_i for $i \in \bar{S}$, as shown below.

Theorem 7 *For any directed complete bipartite graph, Algorithm 2 with $\eta \leq \frac{1}{5}$, $\bar{\eta} \leq \frac{1}{25}$ and $\Omega = \{p \in \Delta(K) : \sum_{i \in S} p_i \geq \sqrt{\bar{\eta}}\}$ guarantees:*

$$\text{Reg}_i \leq \begin{cases} \frac{s \ln T}{\eta} + 2\eta L_i + 2s, & \text{for } i \in S, \\ \frac{2s \ln T}{\eta} + \frac{2 \ln K}{\bar{\eta}} + 2\sqrt{\bar{\eta}} L_{i_S^*} + 2\sqrt{\bar{\eta}} L_i + 2s, & \text{for } i \in \bar{S}, \end{cases}$$

where $i_S^* = \text{argmin}_{i \in S} \sum_{t=1}^T \ell_{t,i}$. Choosing $\eta = \min \left\{ \sqrt{\frac{s}{L_{i_S^*}}}, \frac{1}{5} \right\}$ and $\bar{\eta} = \min \left\{ L_{i_S^*}^{-2/3}, \frac{1}{25} \right\}$, we have: $\text{Reg}_i = \tilde{O}(\sqrt{sL_i} + s)$ for $i \in S$ and $\text{Reg}_i = \tilde{O}\left(L_{i_S^*}^{2/3} + \sqrt{sL_{i_S^*}} + s\right)$ for $i \in \bar{S}$.

Note that even though Alon et al. (2015) show that the worst-case regret of any weakly observable graph is $\Omega(T^{2/3})$, our result indicates that for directed complete bipartite graphs, one can in fact achieve much better regret of order $\tilde{O}(\sqrt{T})$ when the best node has a self-loop, while still maintaining the worst-case regret $\tilde{O}(T^{2/3})$. Moreover, in the former case, we can even achieve a typical small-loss bound, while in the latter case, the regret could be better than $\tilde{O}(T^{2/3})$ as long as the best node in S has sublinear total loss. In Appendix E.3, we also provide an adaptive version of the algorithm without the need of knowing L_i or $L_{i_S^*}$ to tune learning rates (while maintaining the same bound), which requires a nontrivial combination of a clipping technique and doubling trick.

General Case. For general weakly observable graphs, following similar ideas of forcing the algorithm to observe nodes in \bar{S} with a large enough probability, we take $\Omega = \{p \in \Delta(K) : p_i \geq \delta, \forall i \in \mathcal{D}\}$ where \mathcal{D} is a minimum weakly dominating set with size d and δ is some parameter. By definition, this ensures that each node in \bar{S} is observed with probability at least δ . However, this also introduces dependence on $L_{\mathcal{D}}$ for Reg_i , even when $i \in S$, as shown in the following theorem.

Theorem 8 *For any weakly observable graph, Algorithm 2 with $\frac{1}{T} \leq \delta \leq \min \left\{ \frac{1}{125}, \frac{1}{4s}, \frac{1}{4d} \right\}$, $\eta \leq \frac{1}{25}$, $\bar{\eta} \leq \delta^{\frac{4}{3}}$, and $\Omega = \{p \in \Delta(K) : p_i \geq \delta, \forall i \in \mathcal{D}\}$ guarantees:*

$$\text{Reg}_i \leq \begin{cases} \frac{2s \ln T}{\eta} + 2\eta L_i + 2\delta d L_{\mathcal{D}} + 2s, & \text{for } i \in S, \\ \frac{s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + \frac{2\bar{\eta} L_i}{\delta} + 2\delta d L_{\mathcal{D}} + 2s, & \text{for } i \in \bar{S}. \end{cases}$$

For any $\gamma \in [\frac{1}{3}, \frac{1}{2}]$, setting $\delta = \min \left\{ \frac{1}{125}, \frac{1}{4s}, \frac{1}{4d}, L_{\mathcal{D}}^{-\gamma} \right\}$, $\eta = \min \left\{ \sqrt{\frac{1}{L_{\mathcal{D}}}}, \frac{1}{25} \right\}$, and $\bar{\eta} = \min \left\{ \sqrt{\frac{\delta}{L_{\mathcal{D}}}}, \delta^{\frac{4}{3}} \right\}$ gives $\text{Reg} = \tilde{O}\left(L_{\mathcal{D}}^{1-\gamma}\right)$ if $i^* \in S$ and $\tilde{O}\left(L_{\mathcal{D}}^{(1+\gamma)/2}\right)$ otherwise (ignoring dependence on s and d).

Note that unlike the special case of directed complete bipartite graphs, we face a trade-off here when setting the parameters, due to the extra parameter δ that appears in both cases ($i \in S$ or $i \in \bar{S}$). For example, when picking $\gamma = \frac{1}{3}$, we achieve $\text{Reg} = \tilde{O}(L_{\mathcal{D}}^{2/3})$ always, better than the worst-case bound as long as $L_{\mathcal{D}}$ is sublinear. On the other hand, picking $\gamma = \frac{1}{2}$, we achieve $\text{Reg} = \tilde{O}(\sqrt{L_{\mathcal{D}}})$ when $i^* \in S$ and $\tilde{O}(L_{\mathcal{D}}^{3/4})$ otherwise. Once again, we provide an adaptive version in Appendix E.5.

Acknowledgments

The authors are supported by NSF Award IIS-1755781 and thank Sébastien Bubeck, Akshay Krishnamurthy, Thodoris Lykouris, and Chen-Yu Wei for helpful discussions.

References

- Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corraling a band of bandit algorithms. In *Proceedings of the 2017 Conference on Learning Theory*, 2017.
- Zeyuan Allen-Zhu, Sébastien Bubeck, and Yuanzhi Li. Make the minority great again: First-order regret bound for contextual bandits. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Chamy Allenberg, Peter Auer, László Györfi, and György Ottucsák. Hannan consistency in on-line learning in case of unbounded losses under partial monitoring. In *International Conference on Algorithmic Learning Theory*, 2006.
- Noga Alon, Nicolò Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *Proceedings of The 28th Conference on Learning Theory*, 2015.
- Noga Alon, Nicolò Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 2017.
- Raman Arora, Teodor Vanislavov Marinov, and Mehryar Mohri. Bandits with feedback graphs and switching costs. In *Advances in Neural Information Processing Systems 32*, 2019.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Sébastien Bubeck and Mark Sellke. First-order regret analysis of thompson sampling. In *International Conference on Algorithmic Learning Theory*, 2020.
- Sébastien Bubeck, Michael Cohen, and Yuanzhi Li. Sparsity, variance and curvature in multi-armed bandits. In *Proceedings of Algorithmic Learning Theory*, 2018.
- Sébastien Bubeck, Yuanzhi Li, Haipeng Luo, and Chen-Yu Wei. Improved path-length regret bounds for bandits. In *Proceedings of the 32nd Conference On Learning Theory*, 2019.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Zhili Feng and Po-Ling Loh. Online learning with graph-structured feedback against adaptive adversaries. In *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018.
- Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems 29*, 2016.

- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 1997.
- Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- Tomáš Kocák, Gergely Neu, and Michal Valko. Online learning with noisy side observations. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, 2016.
- Thodoris Lykouris, Karthik Sridharan, and Éva Tardos. Small-loss bounds for online learning with partial information. In *Proceedings of the 31st Conference On Learning Theory*, 2018.
- Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. In *Advances in Neural Information Processing Systems 24*, 2011.
- Gergely Neu. First-order regret bounds for combinatorial semi-bandits. In *Proceedings of The 28th Conference on Learning Theory*, 2015.
- Anshuka Rangi and Massimo Franceschetti. Online learning with feedback graphs and switching costs. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 2019.
- Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In *International Conference on Machine Learning*, pages 1593–1601, 2014.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Proceedings of the 31st Conference On Learning Theory*, 2018.
- Kai Zheng, Haipeng Luo, Ilias Diakonikolas, and Liwei Wang. Equipping experts/bandits with long-term memory. In *Advances in Neural Information Processing Systems 32*, 2019.

Conclusions and Open Problems

In this work, we provide various new results on small-loss bounds for bandits with a directed feedback graph (either strongly observable or weakly observable), making a significant step towards a full understanding of this problem.

One clear open question is whether one can achieve $\tilde{O}(\sqrt{\alpha L_\star})$ regret for strongly observable graphs, which would be a strict improvement over the minimax regret $\tilde{O}(\sqrt{\alpha T})$. Note again that our bound $\tilde{O}(\sqrt{(\kappa+1)L_\star})$ is weaker since $\alpha \leq \kappa+1$ always holds. Achieving this ideal bound appears to require new ideas. Even for the special case of self-aware graphs, the problem remains challenging and the closest result is the bound $\tilde{O}(\alpha^{1/3} L_\star^{2/3})$ by [Lykouris et al. \(2018\)](#). Another future direction is to generalize our results to time-varying feedback graphs, which also appears to require new ideas.

Appendix A. Proofs for Section 3.1

In this section, we prove [Lemma 1](#) and [Theorem 2](#). To prove [Lemma 1](#), we first show the following auxiliary lemma, which states that the OMD update enjoys multiplicative stability under certain conditions.

Lemma 9 *Let $p_{t+1} = \operatorname{argmin}_{p \in \Omega} \langle p, \hat{\ell}_t \rangle + D_\psi(p, p_t)$ for $\Omega \subseteq \Delta(K)$ and $\psi : \Omega \rightarrow \mathbb{R}$ such that $\nabla^2 \psi(p) \succeq \operatorname{diag} \left\{ \frac{C_1}{p_i^2}, \dots, \frac{C_1}{p_K^2} \right\}$ for some $C_1 \geq 9$ and $\nabla^{-2} \psi(p) \preceq 4 \nabla^{-2} \psi(q)$ as long as $p \preceq 2q$. If there exists $z \in \mathbb{R}$ such that $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} \leq \frac{1}{8}$, then we have $\frac{1}{2} p_t \preceq p_{t+1} \preceq 2 p_t$.*

Proof The proof follows similar ideas of recent work such as ([Wei and Luo, 2018](#)) or ([Bubeck et al., 2019](#)). Let $F_t(p) \triangleq \langle p, \hat{\ell}_t - z \cdot \mathbf{1} \rangle + D_\psi(p, p_t)$, for z satisfying the condition $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} \leq \frac{1}{8}$. As we only shift each entry of the loss estimator by a constant, according to the algorithm, we have $p_{t+1} = \operatorname{argmin}_{p \in \Omega} F_t(p)$. We first show that $F_t(p') \geq F_t(p_t)$ for any $p' \in \Omega$ such that $\|p' - p_t\|_{\nabla^2 \psi(p_t)} = 1$. We start by applying Taylor expansion:

$$\begin{aligned} F_t(p') &= F_t(p_t) + \nabla F_t(p_t)^\top (p' - p_t) + \frac{1}{2} (p' - p_t)^\top \nabla^2 F_t(\xi) (p' - p_t) \\ &= F_t(p_t) + \left(\hat{\ell}_t - z \cdot \mathbf{1} \right)^\top (p' - p_t) + \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\xi)}^2 \\ &\geq F_t(p_t) - \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} \|p' - p_t\|_{\nabla^2 \psi(p_t)} + \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\xi)}^2 \\ &= F_t(p_t) - \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} + \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\xi)}^2, \end{aligned}$$

where the inequality is by Hölder's inequality and ξ is some point on the line segment between p_t and p' . By the condition $\nabla^2 \psi(p) \succeq \operatorname{diag} \left\{ \frac{9}{p_i^2}, \dots, \frac{9}{p_K^2} \right\}$, we have $1 = \|p' - p_t\|_{\nabla^2 \psi(p_t)}^2 \geq 9 \sum_{i \in [K]} \frac{(p'_i - p_{t,i})^2}{p_{t,i}^2}$, which implies $\frac{|p'_i - p_{t,i}|}{p_{t,i}} \leq \frac{1}{3}$ for all $i \in [K]$. Therefore, we have $\xi \preceq \frac{4}{3} p_t \preceq 2 p_t$, which leads to $\nabla^2 \psi(\xi) \succeq \frac{1}{4} \nabla^2 \psi(p_t)$ according to the assumption. Plugging it into the previous inequality, we have

$$F_t(p') - F_t(p_t) \geq -\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} + \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\xi)}^2 \geq -\frac{1}{8} + \frac{1}{8} = 0.$$

Therefore, according to the optimality of p_{t+1} and the convexity of F_t , we have $\|p_{t+1} - p_t\|_{\nabla^2 \psi(p_t)} \leq 1$. Following the previous analysis, we further have:

$$1 \geq \|p_{t+1} - p_t\|_{\nabla^2 \psi(p_t)}^2 \geq 9 \sum_{i \in [K]} \frac{(p_{t+1,i} - p_{t,i})^2}{p_{t,i}^2} \geq 9 \frac{(p_{t+1,j} - p_{t,j})^2}{p_{t,j}^2}, \forall j \in [K].$$

So we conclude $p_{t+1,i} \in [\frac{2}{3}p_{t,i}, \frac{4}{3}p_{t,i}] \subseteq [\frac{1}{2}p_{t,i}, 2p_{t,i}]$ for all $i \in [K]$, finishing the proof. \blacksquare

The next lemma further shows that the condition $\exists z, \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} \leq \frac{1}{8}$ is easily satisfied as long as $0 \leq \hat{\ell}_{t,i} \leq \max\left\{\frac{1}{p_{t,i}}, \frac{1}{1-p_{t,i}}\right\}$ for all i .

Lemma 10 *If $0 \leq \hat{\ell}_{t,i} \leq \max\left\{\frac{1}{p_{t,i}}, \frac{1}{1-p_{t,i}}\right\}$ for all $i \in [K]$, under the same conditions of Lemma 9 with $C_1 = 64K$, there exists $z \in \mathbb{R}$ such that $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)} \leq \frac{1}{8}$.*

Proof If $p_{t,i} \leq \frac{1}{2}$ for all $i \in [K]$, then we have $p_{t,i} \hat{\ell}_{t,i} \leq \max\left\{1, \frac{p_{t,i}}{1-p_{t,i}}\right\} \leq 1$ for all $i \in [K]$. In this case, $z = 0$ satisfies:

$$\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2 \leq \sum_{i \in [K]} \frac{p_{t,i}^2 \hat{\ell}_{t,i}^2}{C_1} \leq \frac{K}{C_1} = \frac{1}{64}.$$

On the other hand, if there is one node $i_{t,0}$ such that $p_{t,i_{t,0}} > \frac{1}{2}$, then $p_{t,i} \leq \frac{1}{2}$ and $p_{t,i} \hat{\ell}_{t,i} \leq 1$ must be true for all $i \neq i_{t,0}$. In this case picking $z = \hat{\ell}_{t,i_{t,0}}$ gives the following bound on $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}$:

$$\begin{aligned} \|\hat{\ell}_t - \hat{\ell}_{t,i_{t,0}} \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2 &\leq \frac{1}{C_1} \sum_{i \neq i_{t,0}} p_{t,i}^2 (\hat{\ell}_{t,i} - \hat{\ell}_{t,i_{t,0}})^2 \\ &\leq \frac{1}{C_1} \sum_{i \neq i_{t,0}} \left(p_{t,i}^2 \hat{\ell}_{t,i}^2 + p_{t,i}^2 \hat{\ell}_{t,i_{t,0}}^2 \right) \\ &\leq \frac{(K-1)}{C_1} + \frac{(1-p_{t,i_{t,0}})^2 \hat{\ell}_{t,i_{t,0}}^2}{C_1} \quad (\sum_{i \neq i_{t,0}} p_{t,i}^2 \leq (1-p_{t,i_{t,0}})^2) \\ &\leq \frac{1}{64}. \quad (0 \leq \hat{\ell}_{t,i_{t,0}} \leq \frac{1}{1-p_{t,i_{t,0}}}) \end{aligned}$$

Combining the two cases finishes the proof. \blacksquare

Now we are ready to prove Lemma 1.

Proof of Lemma 1. For any time step t and any $z \in \mathbb{R}$, we first follow standard Online Mirror Descent analysis and show

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 2\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(\xi)}^2 \quad (8)$$

for some ξ on the line segment of p_t and p_{t+1} . Define $F_t(p) \triangleq \langle p, \hat{\ell}_t - z \cdot \mathbf{1} \rangle + D_\psi(p, p_t)$. As we only shift each entry of the loss estimator by a constant, according to the algorithm, we have

$p_{t+1} = \operatorname{argmin}_{p \in \Omega} F_t(p)$ and thus by Taylor expansion, it holds for some ξ on the line segment of p_t and p_{t+1} that

$$\begin{aligned} F_t(p_t) - F_t(p_{t+1}) &= \nabla F_t(p_{t+1})^\top (p_t - p_{t+1}) + \frac{1}{2} (p_t - p_{t+1})^\top \nabla^2 F_t(\xi) (p_t - p_{t+1}) \\ &\geq \frac{1}{2} \|p_t - p_{t+1}\|_{\nabla^2 \psi(\xi)}^2. \end{aligned} \quad (9)$$

On the other hand, by the non-negativity of Bregman divergence and Hölder's inequality, we have

$$\begin{aligned} F_t(p_t) - F_t(p_{t+1}) &= \left\langle p_t - p_{t+1}, \hat{\ell}_t - z \cdot \mathbf{1} \right\rangle - D_\psi(p_{t+1}, p_t) \\ &\leq \left\langle p_t - p_{t+1}, \hat{\ell}_t - z \cdot \mathbf{1} \right\rangle \\ &\leq \|p_t - p_{t+1}\|_{\nabla^2 \psi(\xi)} \cdot \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(\xi)}. \end{aligned} \quad (10)$$

Combining Eq. (9) and Eq. (10), we have $\|p_t - p_{t+1}\|_{\nabla^2 \psi(\xi)} \leq 2 \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(\xi)}$. Furthermore, standard analysis of Online Mirror Descent (see e.g. (Wei and Luo, 2018, Lemma 6)) shows

$$\left\langle p_t - u, \hat{\ell}_t \right\rangle = \left\langle p_t - u, \hat{\ell}_t - z \cdot \mathbf{1} \right\rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \left\langle p_t - p_{t+1}, \hat{\ell}_t - z \cdot \mathbf{1} \right\rangle,$$

which proves Eq. (8) after applying Hölder's inequality again and the previous conclusion $\|p_t - p_{t+1}\|_{\nabla^2 \psi(\xi)} \leq 2 \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(\xi)}$.

Finally, according to Lemma 10, we know that the conditions of Lemma 9 hold, and thus multiplicative stability $\frac{1}{2} p_t \preceq p_{t+1} \preceq 2 p_t$ holds, implying $\xi \preceq 2 p_t$. By Condition (b) of the lemma statement, we have $\nabla^{-2} \psi(\xi) \preceq 4 \nabla^{-2} \psi(p_t)$, which shows $\|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(\xi)}^2 \leq 4 \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2$ and completes the proof as z is arbitrary. \blacksquare

Finally, we prove Theorem 2.

Proof of Theorem 2. According to the choice of c and the construction of loss estimators, the conditions of Lemma 1 hold and we have

$$\left\langle p_t - u, \hat{\ell}_t \right\rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 8 \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2.$$

To bound the local-norm term $\min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2$, one could follow the analysis of (Alon et al., 2015). However, to be consistent with other proofs in this work, we provide a different analysis based on a novel loss shift (that is critical for all other proofs). Specifically, we consider two cases. First, if $p_{t,i} < \frac{1}{2}$ holds for all $i \in \bar{S}$, then we relax the local-norm term by taking $z = 0$:

$$\begin{aligned} \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2 &\leq \|\hat{\ell}_t\|_{\nabla^{-2} \psi(p_t)}^2 = \sum_{i \in [K]} \frac{1}{c/p_{t,i}^2 + 1/\eta p_{t,i}} \hat{\ell}_{t,i}^2 \\ &\leq \sum_{i \in [K]} \eta p_{t,i} \hat{\ell}_{t,i}^2 \leq \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i}^2 + 2 \sum_{i \in \bar{S}} \eta p_{t,i} \hat{\ell}_{t,i}, \end{aligned}$$

where the last step is because $\hat{\ell}_{t,i} \leq \frac{1}{1-p_{t,i}} \leq 2$ for $i \in \bar{S}$. On the other hand, if there exists $i_{t,0} \in \bar{S}$ such that $p_{t,i_{t,0}} \geq \frac{1}{2}$, then we take $z = \hat{\ell}_{t,i_{t,0}}$ and arrive at:

$$\min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2 \leq \|\hat{\ell}_t - \hat{\ell}_{t,i_{t,0}} \cdot \mathbf{1}\|_{\nabla^{-2} \psi(p_t)}^2$$

$$\begin{aligned}
 &\leq \sum_{i \neq i_{t,0}} \eta p_{t,i} \left(\hat{\ell}_{t,i} - \hat{\ell}_{t,i_{t,0}} \right)^2 \\
 &\leq \sum_{i \neq i_{t,0}} \eta p_{t,i} \left(\hat{\ell}_{t,i}^2 + \hat{\ell}_{t,i_{t,0}}^2 \right) \\
 &= \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i}^2 + \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i}^2 + \sum_{i \in [K], i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i_{t,0}}^2 \\
 &\leq \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i}^2 + 2 \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i} + \sum_{i \in [K], i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i_{t,0}}^2 \\
 &= \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i}^2 + 2 \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i} + \eta (1 - p_{t,i_{t,0}}) \hat{\ell}_{t,i_{t,0}}^2 \\
 &\leq \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i}^2 + 2 \sum_{i \in \bar{S}} \eta p_{t,i} \hat{\ell}_{t,i},
 \end{aligned}$$

where the second to last inequality is because $\hat{\ell}_{t,i} \leq \frac{1}{1-p_{t,i}} \leq 2$ for $i \in \bar{S} \setminus \{i_{t,0}\}$ and the final inequality is because $(1 - p_{t,i_{t,0}}) \hat{\ell}_{t,i_{t,0}} \leq \frac{1-p_{t,i_{t,0}}}{1-p_{t,i_{t,0}}} = 1 \leq 2p_{t,i_{t,0}}$. Therefore, combining the two cases we have shown:

$$\left\langle p_t - u, \hat{\ell}_t \right\rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 8\eta \sum_{i \in S} p_{t,i} \hat{\ell}_{t,i}^2 + 16\eta \sum_{i \in \bar{S}} p_{t,i} \hat{\ell}_{t,i}.$$

Summing over t and telescoping, we further have:

$$\begin{aligned}
 \sum_{t=1}^T \left\langle p_t - u, \hat{\ell}_t \right\rangle &\leq D_\psi(u, p_1) + 8\eta \sum_{t=1}^T \sum_{i \in S} p_{t,i} \hat{\ell}_{t,i}^2 + 16\eta \sum_{t=1}^T \sum_{i \in \bar{S}} p_{t,i} \hat{\ell}_{t,i} \\
 &\leq D_\psi(u, p_1) + 8\eta \sum_{t=1}^T \sum_{i \in S} \frac{p_{t,i}}{W_{t,i}} \hat{\ell}_{t,i} + 16\eta \sum_{t=1}^T \sum_{i \in \bar{S}} p_{t,i} \hat{\ell}_{t,i}.
 \end{aligned}$$

We choose $u = (1 - \frac{K}{T}) e_{i^*} + \frac{1}{T} \cdot \mathbf{1}$. By the optimality of p_1 , we bound the Bregman divergence term as:

$$\begin{aligned}
 D_\psi(u, p_1) &\leq \psi(u) - \psi(p_1) \\
 &\leq \frac{1}{\eta} \sum_{i \in [K]} p_{1,i} \ln \frac{1}{p_{1,i}} + c \sum_{i \in [K]} \ln \frac{1}{u_i} \leq \frac{\ln K}{\eta} + cK \ln T.
 \end{aligned}$$

Comparing u and e_{i^*} , we bound $\sum_{t=1}^T \left\langle p_t - e_{i^*}, \hat{\ell}_t \right\rangle$ by

$$\frac{\ln K}{\eta} + cK \ln T + 8\eta \sum_{t=1}^T \sum_{i \in S} \frac{p_{t,i}}{W_{t,i}} \hat{\ell}_{t,i} + 16\eta \sum_{t=1}^T \sum_{i \in \bar{S}} p_{t,i} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [K]} \hat{\ell}_{t,i}.$$

Taking expectation over both sides, we arrive at:

$$\text{Reg} \leq \frac{\ln K}{\eta} + cK \ln T + \mathbb{E} \left[8\eta \sum_{t=1}^T \sum_{i \in S} \frac{p_{t,i}}{W_{t,i}} \ell_{t,i} + 16\eta \sum_{t=1}^T \sum_{i \in \bar{S}} p_{t,i} \ell_{t,i} \right] + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [K]} \ell_{t,i}$$

$$\leq \frac{\ln K}{\eta} + cK \ln T + 32\eta\alpha T \ln \left(\frac{4KT}{\alpha} \right) + 16\eta T + K = \tilde{O} \left(\frac{1}{\eta} + \eta\alpha T + K^2 \right),$$

where the last inequality uses the fact $\ell_{t,i} \leq 1$ and also a graph-theoretic lemma (Alon et al., 2015, Lemma 5) which asserts $\sum_{i \in S} \frac{p_{t,i}}{W_{t,i}} \leq 4\alpha \ln \left(\frac{4KT}{\alpha} \right)$. This finishes the proof. \blacksquare

Appendix B. Proofs for Section 3.2

We prove Theorem 3 in this section.

Proof of Theorem 3. Similar to the proof of Theorem 2, the conditions of Lemma 1 hold and we have

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 8 \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2.$$

Once again, we bound the local-norm term by considering two cases separately.

(i). If $p_{t,i} < \frac{1}{2}$ holds for all $i \in \bar{S}$, then choosing $z = 0$ we have:

$$\begin{aligned} \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 &\leq \|\hat{\ell}_t\|_{\nabla^{-2}\psi(p_t)}^2 \\ &= \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 + \sum_{i \in \bar{S}} \frac{1}{c/p_{t,i}^2 + 1/\eta p_{t,i}} \hat{\ell}_{t,i}^2 \\ &\leq \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 + \sum_{i \in \bar{S}} \eta p_{t,i} \hat{\ell}_{t,i}^2 \\ &\leq \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i} + 2 \sum_{i \in \bar{S}} \eta p_{t,i} \hat{\ell}_{t,i} \\ &\leq 2\eta \langle p_t, \hat{\ell}_t \rangle. \end{aligned}$$

The third inequality is because $p_{t,i} \hat{\ell}_{t,i} \leq 1$ for $i \in S$ and $\hat{\ell}_{t,i} \leq \frac{1}{1-p_{t,i}} \leq 2$ for $i \in \bar{S}$.

(ii). If there exists $\exists i_{t,0} \in \bar{S}$ such that $p_{t,i_{t,0}} \geq \frac{1}{2}$, the choosing $z = \hat{\ell}_{t,i_{t,0}}$ we have:

$$\begin{aligned} \min_{z \in \mathbb{R}} \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 &\leq \|\hat{\ell}_t - \hat{\ell}_{t,i_{t,0}} \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 \\ &= \sum_{i \in S} \eta p_{t,i}^2 (\hat{\ell}_{t,i} - \hat{\ell}_{t,i_{t,0}})^2 + \sum_{i \in \bar{S}, i \neq i_{t,0}} \frac{1}{c/p_{t,i}^2 + 1/\eta p_{t,i}} (\hat{\ell}_{t,i} - \hat{\ell}_{t,i_{t,0}})^2 \\ &\leq \sum_{i \in S} \eta p_{t,i}^2 (\hat{\ell}_{t,i}^2 + \hat{\ell}_{t,i_{t,0}}^2) + \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} (\hat{\ell}_{t,i}^2 + \hat{\ell}_{t,i_{t,0}}^2) \\ &= \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 + \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i}^2 + \sum_{i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i_{t,0}}^2 \\ &= \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 + \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i}^2 + \eta (1 - p_{t,i_{t,0}}) \hat{\ell}_{t,i_{t,0}}^2 \\ &\leq \sum_{i \in S} \eta p_{t,i} \hat{\ell}_{t,i} + 2 \sum_{i \in \bar{S}, i \neq i_{t,0}} \eta p_{t,i} \hat{\ell}_{t,i} + 2\eta p_{t,i_{t,0}} \hat{\ell}_{t,i_{t,0}} \end{aligned}$$

$$\leq 2\eta \langle p_t, \hat{\ell}_t \rangle.$$

The second to last inequality is because $p_{t,i} \hat{\ell}_{t,i} \leq 1$ for $i \in S$, $\hat{\ell}_{t,i} \leq \frac{1}{1-p_{t,i}} \leq 2$ for $i \in \bar{S} \setminus \{i_{t,0}\}$, and $(1-p_{t,i_{t,0}}) \hat{\ell}_{t,i_{t,0}} \leq 1 \leq 2p_{t,i_{t,0}}$.

Combining the two cases, we have

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + 16\eta \langle p_t, \hat{\ell}_t \rangle,$$

and summing over t and telescoping, we further have

$$\sum_{t=1}^T \langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_1) + 16\eta \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle.$$

We choose $u = (1 - \frac{K}{T}) e_{i^*} + \frac{1}{T} \cdot \mathbf{1}$ and calculate the Bregman divergence term as

$$\begin{aligned} D_\psi(u, p_1) &\leq \psi(u) - \psi(p_1) && \text{(by optimality of } p_1) \\ &\leq \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{u_i} + \frac{1}{\eta} \sum_{i \in \bar{S}} p_{1,i} \ln \frac{1}{p_{1,i}} + c \sum_{i \in \bar{S}} \ln \frac{1}{u_i} \leq \frac{s \ln T}{\eta} + \frac{\ln K}{\eta} + cK \ln T. \end{aligned}$$

Comparing the difference between u and e_{i^*} and rearranging, we arrive at:

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \leq \frac{1}{1-16\eta} \left(\frac{s \ln T + \ln K}{\eta} + cK \ln T + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [K]} \hat{\ell}_{t,i} + 16\eta \sum_{t=1}^T \hat{\ell}_{t,i^*} \right).$$

Taking expectation on both sides shows

$$\begin{aligned} \text{Reg} &\leq \frac{1}{1-16\eta} \left(\frac{s \ln T + \ln K}{\eta} + cK \ln T + K + 16\eta L_\star \right) \\ &= \mathcal{O} \left(\frac{s \ln T + \ln K}{\eta} + \eta L_\star + K^2 \ln T \right), \end{aligned}$$

finishing the proof. ■

Appendix C. Omitted details for Section 3.3

In this section, we provide omitted details for Section 3.3, including the adaptive Hedge subroutine used in Algorithm 1 and its regret bound (Appendix C.1), the proof of Theorem 4 (Appendix C.2), and an adaptive version of Algorithm 1 and its analysis (Appendix C.3).

C.1. Hedge with Adaptive Learning Rates

We first provide details of the Hedge variant used in Algorithm 1. Algorithm 3 shows the complete pseudocode. Note that as described in Section 3.3, each Hedge instance only operates over a subset of arms, denoted by C as an input of the algorithm. At each time t , the algorithm proposes a distribution \tilde{p}_t , and then receives a loss vector $\ell_t \in \mathbb{R}_+^K$.

Algorithm 3 Hedge with Adaptive Learning Rates

Input: The number of arms K , the set of active arms $C \subseteq [K]$.

Define: $\Omega = \left\{ p \in \Delta(K) : p_i \geq \frac{1}{|C|T}, \forall i \in C, \text{ and } p_i = 0, \forall i \notin C \right\}$

Initialize: \tilde{p}_1 is the uniform distribution over C .

for $t = 1, 2, \dots, T$ **do**

1 Propose distribution \tilde{p}_t .

2 Receive feedback $\tilde{\ell}_t \in \mathbb{R}_+^K$.

3 Compute $\tilde{p}_{t+1} = \operatorname{argmin}_{p \in \Omega} \left\{ \langle p, \tilde{\ell}_t \rangle + D_{\psi_t}(p, \tilde{p}_t) \right\}$, where

$$\psi_t(p) = \frac{1}{\eta_t} \sum_{i \in [K]} p_i \ln p_i, \quad \text{with } \eta_t = \sqrt{\frac{1}{1 + \sum_{\tau=1}^t \sum_{i=1}^K \tilde{p}_{\tau,i} \tilde{\ell}_{\tau,i}^2}}.$$

Vanilla Hedge is simply OMD with the entropy regularizer over the simplex. Our variant makes the following two modifications. First, the decision set Ω is restricted to a subset of simplex so that zero probability is assigned to arms outside C and at least $\frac{1}{|C|T}$ probability is assigned to each arm in C for exploration purpose. Second, we apply an adaptive time-varying learning rate as specified in [Line 3](#). This adaptive learning rate schedule ensures an adaptive regret bound (which is important for our analysis), as shown in the following lemma.

Lemma 11 *Algorithm 3 ensures that for any $i \in C$, we have*

$$\sum_{t=1}^T \langle \tilde{p}_t - e_i, \tilde{\ell}_t \rangle \leq 25\rho \ln^2(KT) + 10 \ln(KT) \sqrt{\rho \sum_{t=1}^T \tilde{\ell}_{t,i}}, \quad (11)$$

where $\rho = \max \left\{ 1, \max_{t \in [T], i \in C} \tilde{\ell}_{t,i} \right\}$.

Proof Let $q_{t+1,i} = \tilde{p}_{t,i} \exp(-\eta_t \tilde{\ell}_{t,i})$. One can verify $\tilde{p}_{t+1} = \operatorname{argmin}_{p \in \Omega} D_{\psi_t}(p, q_{t+1})$ and also for any $u \in \Omega$,

$$\begin{aligned} \langle \tilde{p}_t - u, \tilde{\ell}_t \rangle &= D_{\psi_t}(u, \tilde{p}_t) - D_{\psi_t}(u, q_{t+1}) + D_{\psi_t}(\tilde{p}_t, q_{t+1}) \\ &\leq D_{\psi_t}(u, \tilde{p}_t) - D_{\psi_t}(u, \tilde{p}_{t+1}) + D_{\psi_t}(\tilde{p}_t, q_{t+1}), \end{aligned}$$

where the second step uses the generalized Pythagorean theorem. On the other hand, using the fact $\exp(-x) \leq 1 - x + x^2$ for any $x \geq 0$, we also have

$$\begin{aligned} D_{\psi_t}(\tilde{p}_t, q_{t+1}) &= \frac{1}{\eta_t} \sum_{i \in [K]} \left(\tilde{p}_{t,i} \ln \frac{\tilde{p}_{t,i}}{q_{t+1,i}} + q_{t+1,i} - \tilde{p}_{t,i} \right) \\ &= \frac{1}{\eta_t} \sum_{i \in [K]} \tilde{p}_{t,i} \left(\exp(-\eta_t \tilde{\ell}_{t,i}) - 1 + \eta_t \tilde{\ell}_{t,i} \right) \leq \eta_t \sum_{i \in [K]} \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2. \end{aligned}$$

Summing over t we have shown

$$\begin{aligned}
 \sum_{t=1}^T \langle \tilde{p}_t - u, \tilde{\ell}_t \rangle &\leq \sum_{t=1}^T (D_{\psi_t}(u, \tilde{p}_t) - D_{\psi_t}(u, \tilde{p}_{t+1})) + \sum_{t=1}^T \eta_t \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2 \\
 &\leq \text{KL}(u|\tilde{p}_1) + \sum_{t=1}^{T-1} \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \text{KL}(u|\tilde{p}_{t+1}) + \sum_{t=1}^T \eta_t \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2 \\
 &\leq \text{KL}(u|\tilde{p}_1) + \frac{\max_{p \in \Omega} \text{KL}(u|p)}{\eta_T} + \sum_{t=1}^T \eta_t \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2 \\
 &\leq \ln K + \frac{\ln(KT)}{\eta_T} + \sum_{t=1}^T \sum_{i=1}^K \frac{\tilde{p}_{t,i} \tilde{\ell}_{t,i}^2}{\sqrt{1 + \sum_{\tau=1}^t \sum_{i=1}^K \tilde{p}_{\tau,i} \tilde{\ell}_{\tau,i}^2}} \\
 &\leq \ln K + \frac{\ln(KT)}{\eta_T} + \int_0^{\sum_{t=1}^T \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2} \frac{1}{\sqrt{x+1}} dx \\
 &\leq \ln K + \frac{\ln(KT)}{\eta_T} + 2 \sqrt{1 + \sum_{t=1}^T \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2} \\
 &= \ln K + (\ln(KT) + 2) \sqrt{1 + \sum_{t=1}^T \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2}.
 \end{aligned}$$

Now choosing $u = (1 - \frac{1}{T}) e_i + \frac{1}{|C|T} \cdot \mathbf{1}_C \in \Omega$ where $\mathbf{1}_C$ is the vector with one for coordinates in C and zero otherwise, we have

$$\begin{aligned}
 \sum_{t=1}^T \langle \tilde{p}_t - e_i, \tilde{\ell}_t \rangle &\leq \ln K + (\ln(KT) + 2) \sqrt{1 + \sum_{t=1}^T \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2} + \frac{1}{|C|T} \sum_{t=1}^T \sum_{i \in C} \tilde{\ell}_{t,i} \\
 &\leq 4 \ln(KT) + 3 \ln(KT) \sqrt{\sum_{t=1}^T \sum_{i=1}^K \tilde{p}_{t,i} \tilde{\ell}_{t,i}^2} + \frac{1}{|C|T} \sum_{t=1}^T \sum_{i \in C} \tilde{\ell}_{t,i} \\
 &\leq 4 \ln(KT) + 3 \ln(KT) \sqrt{\rho \sum_{t=1}^T \langle \tilde{p}_t, \tilde{\ell}_t \rangle} + \rho.
 \end{aligned}$$

Let $\tilde{L}_T \triangleq \sum_{t=1}^T \langle \tilde{p}_t, \tilde{\ell}_t \rangle$ and $\tilde{L}_{T,i} \triangleq \sum_{t=1}^T \tilde{\ell}_{t,i}$. By solving the quadratic inequality, we have

$$\begin{aligned}
 \sqrt{\tilde{L}_T} &\leq \frac{3 \ln(KT) \sqrt{\rho} + \sqrt{9 \ln^2(KT) \rho + 4 \cdot (4 \ln(KT) + \rho + \tilde{L}_{T,i})}}{2} \\
 &\leq 5 \ln(KT) \sqrt{\rho} + \sqrt{\tilde{L}_{T,i}}.
 \end{aligned}$$

Finally, squaring both sides proves $\tilde{L}_T - \tilde{L}_{T,i} \leq 25 \ln^2(KT) \rho + 10 \ln(KT) \sqrt{\rho \tilde{L}_{T,i}}$. ■

C.2. Proofs of [Theorem 4](#)

To prove [Theorem 4](#), we combine the regret bounds of the meta-algorithm and the Hedge subroutine. For the former, we prove the following lemma, which combines the result of [Theorem 3](#) and the effect of the increasing learning rate schedule proposed in ([Agarwal et al., 2017](#)), leading to an important negative regret term.

Lemma 12 *Algorithm 1* with $c = 64\beta$ and $\eta \leq \frac{1}{64\beta}$ ensures that for any $j \in [\beta]$,

$$\begin{aligned} \sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle &\leq \mathcal{O} \left(\frac{\kappa \ln T + \ln K}{\eta} + \beta^2 \ln T \right) + 80\eta \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle \\ &\quad + \frac{1}{T} \sum_{t=1}^T \sum_{j \in [\beta]} \hat{\ell}_{t,j} - \frac{\rho_{T,j}}{20\eta \ln T} \mathbb{1}\{j \in [\kappa]\}. \end{aligned} \quad (12)$$

Proof We first show that according to our increasing learning rate schedule, the final learning rate is upper bounded by a constant times the original learning rate. Fix a node $j \in [\kappa]$. Let n_j be such that $\eta_{T,j} = \sigma^{n_j} \eta_{1,j}$ with $\sigma = e^{\frac{1}{\ln T}}$, where we assume $n_j \geq 1$ (the case $n_j = 0$ is trivial as one will see). Let t_1, \dots, t_{n_j} be the rounds in which the learning rate update is executed for node j . Since $\frac{1}{p_{t_{n_j}+1,j}} > \rho_{t_{n_j},j} > 2\rho_{t_{n_j-1},j} > \dots > 2^{n_j-1} \rho_{1,j} = 2^{n_j} \kappa$ and $\frac{1}{p_{t_{n_j}+1,j}} \leq T$, we have $n_j \leq \log_2 T$. Therefore, we have $\eta_{T,j} \leq \sigma^{\log_2 T} \eta_{1,j} \leq 5\eta_{1,j} = 5\eta$.

Next, according to our choice of c and η , the conditions of [Lemma 1](#) hold and we have

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_{\psi_t}(u, p_t) - D_{\psi_t}(u, p_{t+1}) + 8 \min_z \|\hat{\ell}_t - z\|_{\nabla^{-2}\psi_t(p_t)}^2.$$

We consider the Bregman divergence terms and choose $u = (1 - \frac{\kappa}{T}) e_j + \frac{1}{T} \mathbf{1}$. If $j \in [\kappa]$, then with $h(y) = y - 1 - \ln y$ we have

$$\begin{aligned} \sum_{t=1}^T D_{\psi_t}(u, p_t) - D_{\psi_t}(u, p_{t+1}) &\leq D_{\psi_1}(u, p_1) + \sum_{t=1}^{T-1} (D_{\psi_{t+1}}(u, p_{t+1}) - D_{\psi_t}(u, p_{t+1})) \\ &\leq D_{\psi_1}(u, p_1) + \left(\frac{1}{\eta_{t_{n_j}+1,j}} - \frac{1}{\eta_{t_{n_j},j}} \right) h \left(\frac{u_j}{p_{t_{n_j}+1,j}} \right) \\ &= D_{\psi_1}(u, p_1) + \frac{1-\sigma}{\sigma^{n_j} \eta} h \left(\frac{u_j}{p_{t_{n_j}+1,j}} \right) \\ &\leq \frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T - \frac{1}{5\eta \ln T} h \left(\frac{u_j}{p_{t_{n_j}+1,j}} \right), \end{aligned}$$

where we use the facts $1 - \sigma \leq -\frac{1}{\ln T}$ and $\sigma^{n_j} \leq 5$ as shown earlier, and also the exact same analysis of bounding $D_{\psi_1}(u, p_1)$ as in the proof of [Theorem 3](#). Note that $\frac{u_j}{p_{t_{n_j}+1,j}} \geq \frac{1}{2p_{t_{n_j}+1,j}} \geq 2^{n_j-1} \kappa \geq 1$. Combining the facts that $h(y)$ is increasing when $y \geq 1$ and $\rho_{T,j} = \frac{2}{p_{t_{n_j}+1,j}} \leq 2T$, we have:

$$h \left(\frac{u_j}{p_{t_{n_j}+1,j}} \right) \geq h \left(\frac{1}{2p_{t_{n_j}+1,j}} \right) = \frac{\rho_{T,j}}{4} - 1 - \ln \left(\frac{\rho_{T,j}}{4} \right) \geq \frac{\rho_{T,j}}{4} - 2 \ln T.$$

We have thus shown when $j \in [\kappa]$,

$$\sum_{t=1}^T D_{\psi_t}(u, p_t) - D_{\psi_t}(u, p_{t+1}) \leq \mathcal{O}\left(\frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T\right) - \frac{\rho_{T,j}}{20\eta \ln T}$$

On the other hand, if $j \in \bar{S}$, then we have by the monotonicity of learning rates

$$\begin{aligned} \sum_{t=1}^T D_{\psi_t}(u, p_t) - D_{\psi_t}(u, p_{t+1}) &\leq D_{\psi_1}(u, p_1) + \sum_{t=1}^{T-1} (D_{\psi_{t+1}}(u, p_{t+1}) - D_{\psi_t}(u, p_{t+1})) \\ &\leq D_{\psi_1}(u, p_1) \leq \frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T. \end{aligned}$$

It remains to deal with the local-norm term $\min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi_t(p_t)}^2$. Following the exact analysis in the proof of [Theorem 3](#) and the fact $\eta_{t,j} \leq 5\eta$ for all $t \in [T]$ and $j \in [\kappa]$, we have:

$$\min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi_t(p_t)}^2 \leq 5 \min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi_1(p_t)}^2 \leq 10\eta \langle p_t, \hat{\ell}_t \rangle.$$

Combining the bounds for the Bregman divergence terms and the local-norm term, and accounting for the difference between u and e_j complete the proof. \blacksquare

We are now ready to prove [Theorem 4](#).

Proof of [Theorem 4](#). The main idea of the proof is as follows. When $i^* \in \bar{S}$, the regret is exactly $\mathbb{E} \left[\sum_{t=1}^T \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \right]$. Therefore, [Lemma 12](#) already provides the small-loss bound guarantee by rearranging the terms and taking expectation on both sides. When $i^* \in C_j$, according to our loss estimator construction, the regret is exactly the regret of the meta-algorithm plus the regret of \mathcal{A}_j , and we apply [Lemma 12](#) and [Lemma 11](#) to bound each of these two parts and importantly use the negative term from [Eq. \(12\)](#) to cancel the corresponding terms in [Eq. \(11\)](#).

Formally, when $i^* \in \bar{S}$, we apply [Lemma 12](#) with $j = i^*$ and rearrange terms to arrive at

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \leq \mathcal{O}\left(\frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T + \eta \sum_{t=1}^T \hat{\ell}_{t,i^*} + \frac{1}{T} \sum_{t=1}^T \sum_{j \in [\beta]} \hat{\ell}_{t,j}\right).$$

Note that in this case $\mathbb{E} \left[\hat{\ell}_{t,i^*} \right] = \mathbb{E} \left[\frac{\ell_{t,i^*}}{1-p_{t,i^*}} \mathbb{1}\{i_t \neq i^*\} \right] = \ell_{t,i^*}$ for all $t \in [T]$. Thus, taking expectation shows

$$\text{Reg} = \mathcal{O}\left(\frac{\kappa \ln T + \ln K}{\eta} + \beta^2 \ln T + \eta L_\star\right).$$

On the other hand, when $i^* \in C_j$ for some $j \in [\kappa]$, we decompose the regret as

$$\text{Reg} = \mathbb{E} \left[\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \langle e_{i^*}, \tilde{\ell}_t \rangle \right] = \mathbb{E} \left[\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle \right] + \mathbb{E} \left[\sum_{t=1}^T \langle \tilde{p}_t^{(j)} - e_{i^*}, \tilde{\ell}_t \rangle \right].$$

Here, in the first equality, we use the facts

$$\begin{aligned} \mathbb{E} \left[\langle p_t, \hat{\ell}_t \rangle \right] &= \mathbb{E} \left[\sum_{j \in [\kappa]} p_{t,j} \sum_{i \in \mathcal{C}_j} \tilde{p}_{t,i}^{(j)} \frac{\ell_{t,i}}{p_{t,j}} \mathbb{1}\{j = j_t\} + \sum_{i \in \bar{S}} p_{t,i} \frac{\ell_{t,i}}{1 - p_{t,i}} \mathbb{1}\{i \neq i_t\} \right] \\ &= \mathbb{E} \left[\sum_{j \in [\kappa]} p_{t,j} \sum_{i \in \mathcal{C}_j} \tilde{p}_{t,i}^{(j)} \ell_{t,i} + \sum_{i \in \bar{S}} p_{t,i} \ell_{t,i} \right] \\ &= \mathbb{E} [\ell_{t,i_t}] \end{aligned}$$

and $\mathbb{E} [\tilde{\ell}_{t,i^*}] = \mathbb{E} \left[\frac{\ell_{t,i^*}}{p_{t,j}} \mathbb{1}\{j = j_t\} \right] = \ell_{t,i^*}$; and the second equality is directly by the definition of $\hat{\ell}_{t,j}$ for $j \in [\kappa]$.

For the first part of the decomposition, we apply [Lemma 12](#) directly; for the second part, noting that the scale of $\tilde{\ell}_t^{(j)}$ for all $t \in [T]$ is no more than $\rho_{T,j}$, according to [Lemma 11](#) we have:

$$\sum_{t=1}^T \langle \tilde{p}_t^{(j)} - e_{i^*}, \tilde{\ell}_t \rangle \leq 25\rho_{T,j} \ln^2(KT) + 10 \ln(KT) \sqrt{\rho_{T,j} \sum_{t=1}^T \tilde{\ell}_{t,i^*}}.$$

Combining the two gives

$$\begin{aligned} &\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \langle e_{i^*}, \tilde{\ell}_t \rangle \\ &\leq \mathcal{O} \left(\frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T \right) + 80\eta \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle + \frac{1}{T} \sum_{t=1}^T \sum_{j \in [\beta]} \hat{\ell}_{t,j} \\ &\quad - \frac{\rho_{T,j}}{40\eta \ln T} + 25\rho_{T,j} \ln^2(KT) - \frac{\rho_{T,j}}{40\eta \ln T} + 10 \ln(KT) \sqrt{\rho_{T,j} \sum_{t=1}^T \tilde{\ell}_{t,i^*}} \\ &\leq \mathcal{O} \left(\frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T \right) + 80\eta \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle + \frac{1}{T} \sum_{t=1}^T \sum_{j \in [\beta]} \hat{\ell}_{t,j} \\ &\quad + 1000\eta(\ln T) \ln^2(KT) \sum_{t=1}^T \tilde{\ell}_{t,i^*}, \end{aligned} \tag{13}$$

where the second inequality is by the fact $-ax + \sqrt{bx} \leq \frac{b}{4a}$ for $a, b > 0$ and also the condition $\eta \leq \frac{1}{1000(\ln T) \ln^2(KT)}$. By rearranging we have:

$$\begin{aligned} &\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \langle e_{i^*}, \tilde{\ell}_t \rangle \\ &\leq \mathcal{O} \left(\frac{\kappa \ln T + \ln K}{\eta} + c\beta \ln T + \eta(\ln T) \ln^2(KT) \sum_{t=1}^T \tilde{\ell}_{t,i^*} \right) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i}. \end{aligned} \tag{14}$$

Taking expectation on both sides finishes the proof. \blacksquare

Algorithm 4 Adaptive Version of [Algorithm 1](#)

Input: Feedback graph G and a clique partition $\{\mathcal{C}_1, \dots, \mathcal{C}_\kappa\}$ of G_S , parameter η and c .

Define: $\beta = \kappa + \bar{s}$ and $\Omega = \{p \in \Delta(\beta) : p_j \geq \frac{1}{T}, \forall j \in [\beta]\}$.

for $\lambda = 1, 2, \dots$ **do**

$T_\lambda = t - 1, \eta_{t,j} = \eta, \rho_{t,j} = 2\kappa, \forall j \in [\kappa], p_t = \operatorname{argmin}_{p \in \Omega} \psi_t(p)$ (ψ_t defined in [Eq. \(7\)](#))

Create an instance \mathcal{A}_j of adaptive Hedge ([Algorithm 3](#)) with nodes in $\mathcal{C}_j, \forall j \in [\kappa]$.

while $t \leq T$ **do**

Execute [Line 1](#) to [Line 9](#) of [Algorithm 1](#).

if $\frac{\kappa+1}{\eta} \leq \eta \sum_{\tau=T_\lambda+1}^t \langle p_t, \hat{\ell}_t \rangle$ **then**

$\eta \leftarrow \frac{\eta}{2}, t \leftarrow t + 1.$

Break

$t \leftarrow t + 1.$

C.3. Adaptive Version of [Algorithm 1](#)

In this section, we provide [Algorithm 4](#), an adaptive version of [Algorithm 1](#) with a doubling trick to remove the need of tuning the learning rate η in terms of L_\star . The algorithm mostly follows [Algorithm 1](#), starting from a relatively large value of η . The key difference is that at the end of each round, we check if condition $\frac{\kappa+1}{\eta} \leq \eta \sum_{\tau=T_\lambda+1}^t \langle p_t, \hat{\ell}_t \rangle$ holds, where $T_\lambda + 1$ is the time step of the most recent reset. If the condition holds, it implies that the current learning rate η is not small enough, and we thus halve the learning rate, and at the same time reset the algorithm, which includes resetting the parameters $\eta_{t,j}, \rho_{t,j}$, and the distribution p_t , as well as resetting the Hedge instances.

Below we prove that [Algorithm 4](#) achieves the same regret bound as [Algorithm 1](#) without knowing L_\star . The main difficulty of the doubling trick analysis is that $\langle p_t, \hat{\ell}_t \rangle$ is not well bounded when the graph is not self-aware, which is not the case in prior works such as [Wei and Luo \(2018\)](#). We resolve this issue by again utilizing the negative regret term from the increasing learning rate schedule.

Theorem 13 [Algorithm 4](#) with $c = 64\beta$ and $\eta = \frac{1}{2000(\ln T) \ln^2(KT) + 80\kappa \ln T}$ guarantees

$$\operatorname{Reg} = \tilde{O}\left(\sqrt{(\kappa+1)L_\star} + \beta^2\right).$$

Proof We call the time steps between two resets an epoch (indexed by λ) and let η_λ be the value of η during epoch λ so that $\eta_\lambda = 2^{1-\lambda}\eta_1$. Also let λ^\star be the index of the last epoch. For notational convenience, define

$$\widehat{\operatorname{Reg}} \triangleq \begin{cases} \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \tilde{\ell}_{t,i^\star}, & i^\star \in S. \\ \sum_{t=1}^T \langle p_t - e_{i^\star}, \hat{\ell}_t \rangle, & i^\star \in \bar{S}. \end{cases}$$

Note that $\operatorname{Reg} = \mathbb{E}[\widehat{\operatorname{Reg}}]$. We will first prove the following

$$\widehat{\operatorname{Reg}} \leq \sum_{\lambda=1}^{\lambda^\star} \tilde{O}\left(\frac{\kappa+1}{\eta_\lambda} + \beta^2\right) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i}. \quad (15)$$

To show this, consider the regret in each epoch λ . When $i^* \in \bar{S}$, we have:

$$\begin{aligned}
 & \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \\
 & \leq \mathcal{O} \left(c\beta \ln T + \frac{\kappa \ln T + \ln K}{\eta_\lambda} \right) + 40\eta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i} \\
 & \leq \tilde{\mathcal{O}} \left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda} \right) + 80\eta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \langle p_t, \hat{\ell}_t \rangle + \frac{5}{8} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i} \\
 & \leq \tilde{\mathcal{O}} \left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda} \right) + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i}.
 \end{aligned}$$

Here, the first inequality is according to the analysis of [Lemma 12](#). In the second inequality, we bound $\min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2$ by $2 \langle p_t, \hat{\ell}_t \rangle$ for $t = T_\lambda + 1, \dots, T_{\lambda+1} - 1$ by the same analysis of [Theorem 3](#), and bound the same term for $t = T_{\lambda+1}$ by $\frac{1}{64}$ by [Lemma 10](#). The final inequality is because the reset condition does not hold for $t = T_{\lambda+1} - 1$. Summing over the epochs proves [Eq. \(15\)](#) for the first case.

When $i^* \in \mathcal{C}_j$ for some $j \in [\kappa]$, similar to the previous analysis and the derivation of [Eq. \(13\)](#), we have:

$$\begin{aligned}
 & \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \tilde{\ell}_{t,i^*} \\
 & \leq \mathcal{O} \left(c\beta \ln T + \frac{\kappa \ln T + \ln K}{\eta_\lambda} \right) + 40\eta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \min_z \|\hat{\ell}_t - z \cdot \mathbf{1}\|_{\nabla^{-2}\psi(p_t)}^2 + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i} \\
 & \quad - \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T} - \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T} + 25\rho_{T_{\lambda+1},j} \ln^2(KT) - \frac{\rho_{T_{\lambda+1},j}}{40\eta_\lambda \ln T} + 10 \ln(KT) \sqrt{\rho_{T_{\lambda+1},j} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \tilde{\ell}_{t,i^*}} \\
 & \leq \tilde{\mathcal{O}} \left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda} \right) - \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T} + 1000\eta_\lambda (\ln T) \ln^2(KT) \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \tilde{\ell}_{t,i^*} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i},
 \end{aligned}$$

where the second inequality uses the fact $\eta_\lambda \leq \eta_1 \leq \frac{1}{2000 \ln T \ln^2(KT)}$ and the AM-GM inequality. By rearranging terms, we have

$$\begin{aligned}
 & (1 + 1000\eta (\ln T) \ln^2(KT)) \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \tilde{\ell}_{t,i^*} \right) \\
 & \leq \tilde{\mathcal{O}} \left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda} \right) - \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T} + 1000\eta_\lambda (\ln T) \ln^2(KT) \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i}
 \end{aligned}$$

$$\leq \tilde{\mathcal{O}}\left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda}\right) + \frac{1}{2} \langle p_{T_{\lambda+1}}, \hat{\ell}_{T_{\lambda+1}} \rangle - \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T} + \frac{1}{T} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i},$$

where the second inequality is again because the condition $\frac{\kappa+1}{\eta_\lambda} \leq \eta_\lambda \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}-1} \langle p_t, \hat{\ell}_t \rangle$ does not hold and $\eta_\lambda \leq \eta_1 \leq \frac{1}{2000(\ln T) \ln^2(KT)}$. Now if for all $i \in \bar{S}$, $p_{T_{\lambda+1},i} < \frac{1}{2}$, then we have

$$\langle p_{T_{\lambda+1}}, \hat{\ell}_{T_{\lambda+1}} \rangle \leq \sum_{j' \in [\kappa]} p_{T_{\lambda+1},j'} \frac{1}{p_{T_{\lambda+1},j'}} + \sum_{i \in \bar{S}} \frac{p_{T_{\lambda+1},i}}{1 - p_{T_{\lambda+1},i}} \leq \beta.$$

Otherwise, we have exactly one $i_0 \in \bar{S}$ such that $p_{T_{\lambda+1},i_0} > \frac{1}{2}$ and then we have

$$\langle p_{T_{\lambda+1}}, \hat{\ell}_{T_{\lambda+1}} \rangle \leq \sum_{j' \in [\kappa]} p_{T_{\lambda+1},j'} \frac{1}{p_{T_{\lambda+1},j'}} + \sum_{i \in \bar{S}} \frac{p_{T_{\lambda+1},i}}{1 - p_{T_{\lambda+1},i}} \leq \kappa + \frac{1}{1 - p_{T_{\lambda+1},i_0}} \leq \kappa + \rho_{T_{\lambda+1},j}.$$

The last inequality is because $1 - p_{T_{\lambda+1},i_0} \geq p_{T_{\lambda+1},j}$. Therefore, as $\eta_\lambda \leq \eta_1 = \frac{1}{2000 \ln T \ln^2(KT) + 80\kappa \ln T}$, we always have

$$\langle p_{T_{\lambda+1}}, \hat{\ell}_{T_{\lambda+1}} \rangle \leq \frac{\rho_{T_{\lambda+1},j}}{80\eta_\lambda \ln T}.$$

We have thus shown

$$\sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \tilde{\ell}_{t,i^*} \leq \tilde{\mathcal{O}}\left(\beta^2 + \frac{\kappa + 1}{\eta_\lambda}\right) + \frac{1}{T} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in [\beta]} \hat{\ell}_{t,i}.$$

Summing up the regret from epoch 1 to λ^* gives [Eq. \(15\)](#).

Next, using the definition of η_λ , we further have

$$\widehat{\text{Reg}} \leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}}\left(\frac{\kappa + 1}{\eta_\lambda} + \beta^2\right) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i} \leq \tilde{\mathcal{O}}\left(\frac{\kappa + 1}{\eta_{\lambda^*}} + \beta^2 \lambda^*\right) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i}. \quad (16)$$

When $\lambda^* = 1$, direct calculation gives $\widehat{\text{Reg}} \leq \tilde{\mathcal{O}}(\beta^2) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i}$. On the other hand, if $\lambda^* \geq 2$, consider the time step at the end of epoch $\lambda^* - 1$. Using the reset condition, we have:

$$(\kappa + 1) \left(\frac{2^{\lambda^* - 2}}{\eta_1}\right)^2 = (\kappa + 1) \frac{1}{\eta_{\lambda^* - 1}^2} \leq \sum_{t=T_{\lambda^* - 1} + 1}^{T_{\lambda^*}} \langle p_t, \hat{\ell}_t \rangle \leq \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle \leq T^2.$$

So $\lambda^* = \mathcal{O}(\ln T)$, $\left(\frac{\kappa+1}{\eta_{\lambda^*}}\right)^2 = \tilde{\mathcal{O}}\left((\kappa + 1) \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle\right)$. Plugging these into [Eq. \(16\)](#), we have

$$\widehat{\text{Reg}} \leq \tilde{\mathcal{O}}\left(\sqrt{(\kappa + 1) \left(\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle\right) + \beta^2}\right) + \frac{1}{T} \sum_{t=1}^T \sum_{i \in [\beta]} \hat{\ell}_{t,i},$$

which also holds for the case $\lambda^* = 1$. Finally, taking expectation on both sides gives:

$$\text{Reg} = \mathbb{E} \left[\widehat{\text{Reg}} \right] \leq \tilde{\mathcal{O}} \left(\mathbb{E} \left[\sqrt{(\kappa + 1) \left(\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle\right) + \beta^2} \right] \right)$$

$$\leq \tilde{\mathcal{O}} \left(\sqrt{(\kappa + 1) \left(\mathbb{E} \left[\sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle \right] \right)} + \beta^2 \right) \leq \tilde{\mathcal{O}} \left(\sqrt{(\kappa + 1) \left(\mathbb{E} \left[\sum_{t=1}^T \langle p_t, \ell_t \rangle \right] \right)} + \beta^2 \right).$$

Solving the quadratic inequality, we obtain the regret bound $\text{Reg} \leq \tilde{\mathcal{O}} \left(\sqrt{(\kappa + 1)L_*} + \beta^2 \right)$. \blacksquare

Appendix D. Omitted details for Section 3.4

Algorithm 5 An Algorithm with Regret $\tilde{\mathcal{O}}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_*}\})$ for Self-aware Graphs

Input: A clique partition $\{\mathcal{C}_1, \dots, \mathcal{C}_\kappa\}$ of G_S , parameter η_{init} and ϵ .

Define: $\Omega = \{p \in \Delta(K) : p_i \geq \frac{1}{T}, \forall i \in [K]\}$.

```

1  for  $m = 1, 2, \dots, \log_2 T$  do
2       $\eta = \eta_{\text{init}}$ .
3      for  $\lambda = 1, 2, \dots$  do
4           $p_t = \frac{1}{K} \cdot \mathbf{1}, T_\lambda = t - 1$ .
5          while  $t \leq T$  do
6              Pull arm  $i_t \sim p_t$  and receive feedback  $\ell_{t,i}$  for all  $i$  such that  $i_t \in N^{\text{in}}(i)$ .
7              Construct estimator  $\hat{\ell}_t \in \mathbb{R}^K$  such that  $\hat{\ell}_{t,i} = \begin{cases} \frac{\ell_{t,i} \cdot \mathbb{1}\{i_t \in N^{\text{in}}(i)\}}{\sum_{j \in N^{\text{in}}(i)} p_{t,j}}, & \text{if } p_{t,i} > 0. \\ 0, & \text{if } p_{t,i} = 0. \end{cases}$ 
8              Compute  $\hat{p}_{t+1} = \operatorname{argmin}_{p \in \Delta_K} \left\{ \langle p, \hat{\ell}_t \rangle + D_\psi(p, \hat{p}_t) \right\}$ , where
                    
$$\psi(p) = \frac{1}{\eta} \sum_{i \in [K]} p_i \ln p_i.$$

9               $p_t = \hat{p}_t$ .
10             for  $j = 1, 2, \dots, \kappa$  do
11                 if  $\sum_{i \in \mathcal{C}_j} p_{t+1,i} \leq \epsilon$  then
12                     for  $i \in \mathcal{C}_j$  do
13                          $p_{t+1,i} = 0$ .
14             Renormalize  $p_{t+1}$  such that  $p_{t+1} \in \Delta(K)$ .
15             if  $\frac{1}{\eta} \leq 4\eta\kappa \min_{i \in [K]} \left\{ \sum_{\tau=T_\lambda+1}^t \hat{\ell}_{\tau,i} \right\}$  then
16                  $\eta \leftarrow \frac{\eta}{2}, \epsilon = \max\{2\eta, \frac{1}{T}\}, t \leftarrow t + 1$ .
17                 if  $\eta \leq \sqrt{\frac{1}{\alpha T}}$  then
18                     Jump to Line 1.
19                 Jump to Line 3.
20                 break.
21 Run the algorithm from Theorem 2 (from scratch) for the rest of the game.
    
```

In this section, we discuss how to obtain $\tilde{O}(\min\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\})$ regret for self-aware graphs. [Algorithm 5](#) shows the complete pseudocode. It consists of two stages. The first stage is when $m \leq \log_2 T$ and runs yet another parameter-free algorithm with $\tilde{O}(\sqrt{\kappa L_\star})$ regret for self-aware graphs (for technical reasons we are not able to use [Algorithm 1](#) directly here). The second stage exactly runs the algorithm mentioned in [Section 3.1 \(Theorem 2\)](#), which achieves $\tilde{O}(\sqrt{\alpha T})$ regret.

The first stage mainly follows the clipping idea of ([Allenberg et al., 2006](#)) for OMD with entropy regularizer. At each round t , after updating the distribution with OMD, we clip the probability of a clique to zero if it has low probability to be chosen ([Line 11](#) and [Line 13](#)), and normalize the distribution after clipping ([Line 14](#)). Then, a doubling trick is introduced: once the condition $\frac{1}{\eta} \leq 4\eta\kappa \min_{i \in [K]} \{\sum_{\tau=T_\lambda+1}^t \hat{\ell}_{\tau,i}\}$ holds ([Line 15](#)), we halve the learning rate and reset ([Line 19](#)). We point out that thanks to the clipping trick, the term $\min_{i \in [K]} \{\sum_{\tau=T_\lambda+1}^t \hat{\ell}_{\tau,i}\}$ is nicely bounded, which is crucial for the doubling trick analysis. We say that we start a new *epoch* in this case.

Furthermore, once the learning rate is smaller than $\sqrt{1/\alpha T}$ ([Line 17](#)), we jump to [Line 1](#) and reset it to the initial learning rate $\eta = \eta_{\text{init}}$ in [Line 2](#). We say that we start a new *meta-epoch* in this case. After having $\log_2 T$ meta-epochs, the algorithm is confident that $\sqrt{\alpha T} \leq \tilde{O}(\sqrt{\kappa L_\star})$, and thus switches to the second stage and runs the algorithm introduced in [Section 3.1](#) with regret $\tilde{O}(\sqrt{\alpha T})$. We point out that in fact any algorithm with $\tilde{O}(\sqrt{\alpha T})$ regret is acceptable for this second stage.

The guarantee of this algorithm is shown below.

Theorem 14 *Algorithm 5 with $\eta_{\text{init}} = \frac{1}{4\kappa}$ and $\epsilon = \frac{1}{2\kappa}$ guarantees*

$$\text{Reg} = \tilde{O}\left(\min\left\{\sqrt{\alpha T}, \sqrt{\kappa L_\star}\right\} + K^2\right).$$

To prove the theorem, we first prove a bound on the regret within an epoch λ , where $\eta \leq \frac{1}{4\kappa}$ and $\epsilon = \max\{2\eta, \frac{1}{T}\}$ are fixed.

Lemma 15 *Fix a meta-epoch and consider an epoch λ in the first stage. Algorithm 5 guarantees*

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} = \tilde{O}\left(\frac{1}{\eta} + \eta\kappa \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \kappa + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}\right).$$

Proof According to the analysis of online mirror descent with entropy regularizer ([Cesa-Bianchi and Lugosi, 2006](#)), we have

$$\begin{aligned} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} &\leq \frac{\ln K}{\eta} + \eta \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i=1}^K \hat{p}_{t,i} \hat{\ell}_{t,i}^2 \\ &\leq \frac{\ln K}{\eta} + \eta \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{j \in [K]} \sum_{i \in \mathcal{C}_j} \frac{\hat{p}_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i}. \end{aligned}$$

As we do clipping for each clique, if $p_{t,i} > 0$, then $1 - \kappa\epsilon \leq \frac{\hat{p}_{t,i}}{p_{t,i}} \leq 1$ and otherwise $\hat{\ell}_{t,i} = 0$. Therefore, we have

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} (1 - \kappa\epsilon) \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \leq \frac{\ln K}{\eta} + \eta \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{j \in [K]} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i} \quad (17)$$

Now consider a fixed clique \mathcal{C}_j that is not clipped at time t . We have $\frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} = \frac{\hat{p}_{t,i}}{\sum_{i' \in \mathcal{C}_j} \hat{p}_{t,i'}} = \frac{\exp(-\eta \sum_{\tau=1}^{t-1} \hat{\ell}_{\tau,i})}{\sum_{i' \in \mathcal{C}_j} \exp(-\eta \sum_{\tau=1}^{t-1} \hat{\ell}_{\tau,i'})}$, which can be considered as a probability distribution generated by online mirror descent with entropy regularizer inside the clique. Therefore for any clique \mathcal{C}_j , $j \in [\kappa]$, we have

$$\begin{aligned} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i} - \min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} &\leq \frac{\ln K}{\eta} + \eta \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i}^2 \\ &\leq \frac{\ln K}{\eta} + \frac{\eta}{\epsilon} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i} \\ &\leq \frac{\ln K}{\eta} + \frac{1}{2} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i}. \end{aligned}$$

The second inequality is because $\hat{\ell}_{t,i} \leq \frac{1}{\epsilon}$ for all $i \in \mathcal{C}_j$, $j \in [\kappa]$ and the third inequality is because $\epsilon \geq 2\eta$. Rearranging the terms, we have

$$\sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \sum_{i \in \mathcal{C}_j} \frac{p_{t,i}}{\sum_{i' \in \mathcal{C}_j} p_{t,i'}} \hat{\ell}_{t,i} \leq \frac{2 \ln K}{\eta} + 2 \min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i}.$$

Therefore, by combining with Eq. (17), we have

$$\begin{aligned} (1 - \kappa\epsilon) \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} &\leq \frac{\ln K}{\eta} + \eta \sum_{j \in [\kappa]} \left(\frac{2 \ln K}{\eta} + 2 \min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right) \\ &\leq \frac{\ln K}{\eta} + 2\kappa \ln K + 2\eta \sum_{j \in [\kappa]} \min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i}. \quad (18) \end{aligned}$$

Furthermore, let T_{i_0} be the last round such that $\hat{\ell}_{t,i_0} > 0$ for some $i_0 \in \mathcal{C}_j$. Note that T_{i_0} is also the last round such that $\hat{\ell}_{t,i'} > 0$ for all $i' \in \mathcal{C}_j$. Then we have for any $i' \in [K]$:

$$\begin{aligned} \epsilon &\leq \sum_{i \in \mathcal{C}_j} \hat{p}_{T_{i_0},i} = \frac{\sum_{i \in \mathcal{C}_j} \exp(-\eta \sum_{\tau=T(s)+1}^{T_{i_0}-1} \hat{\ell}_{\tau,i})}{\sum_{i=1}^K \exp(-\eta \sum_{\tau=T(s)+1}^{T_{i_0}-1} \hat{\ell}_{\tau,i})} \\ &\leq \frac{\sum_{i \in \mathcal{C}_j} \exp(-\eta \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{\eta}{\epsilon})}{\sum_{i=1}^K \exp(-\eta \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i})} \\ &\leq \frac{|\mathcal{C}_j| \exp(-\eta \min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{\eta}{\epsilon})}{\exp(-\eta \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i'})}, \end{aligned}$$

where the second inequality is because $\hat{\ell}_{T_{i_0},i} \leq \frac{1}{\epsilon}$ for all $i \in \mathcal{C}_j$ and the fact $\hat{\ell}_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [K]$. Therefore, by rearranging terms, for any $i' \in [K]$ and any $j \in [\kappa]$, we have

$$\min_{i \in \mathcal{C}_j} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i} \leq \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i'} + \frac{1}{\epsilon} + \frac{1}{\eta} \ln \frac{K}{\epsilon}.$$

Combining Eq. (18) and choosing $i' = \operatorname{argmin}_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}$ further show

$$\begin{aligned}
 (1 - \kappa\epsilon) \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} &\leq \frac{\ln K}{\eta} + 2\kappa \ln K + 2\eta\kappa \left(\min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{1}{\epsilon} + \frac{1}{\eta} \ln \frac{K}{\epsilon} \right) \\
 &\leq \frac{\ln K}{\eta} + 2\eta\kappa \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + 3\kappa \ln K + 2\kappa \ln(KT) \\
 &\leq \frac{\ln K}{\eta} + 2\eta\kappa \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + 5\kappa \ln(KT),
 \end{aligned}$$

where the second inequality is because $\epsilon = \max\{2\eta, \frac{1}{T}\}$. Finally, rearranging terms again shows

$$\begin{aligned}
 &\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \\
 &\leq \frac{1}{1 - \kappa\epsilon} \left(\frac{\ln K}{\eta} + \kappa(2\eta + \epsilon) \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + 5\kappa \ln(KT) \right) \\
 &\leq \frac{2 \ln K}{\eta} + 8\eta\kappa \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + 10\kappa \ln(KT) + \frac{2K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}.
 \end{aligned}$$

The second inequality is because $\eta \leq \frac{1}{4\kappa}$, which means $\kappa\epsilon \leq \max\{\frac{\kappa}{T}, 2\eta\kappa\} \leq \frac{1}{2}$, as $T \geq 2K$. \blacksquare

Next we bound the regret within a meta-epoch. In the remaining of the section, we use \mathcal{T}_m to denote the set of rounds in meta-epoch m .

Lemma 16 *For any meta-epoch m in the first stage, Algorithm 5 guarantees*

$$\sum_{t \in \mathcal{T}_m} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \leq \tilde{\mathcal{O}} \left(\min \left\{ \sqrt{\alpha T}, \sqrt{\kappa \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}} \right\} + \kappa + \frac{K}{T} \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \right).$$

Proof Let $\eta_\lambda = 2^{1-\lambda}\eta_1$ and $\epsilon_\lambda = \max\{2\eta_\lambda, 1/T\}$ be the value of η and ϵ during epoch λ and let λ^* be the index of the last epoch. Consider the regret in epoch λ . Using Lemma 15, we know that

$$\begin{aligned}
 &\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \\
 &= \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \eta_\lambda \kappa \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \kappa + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right) \\
 &= \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \frac{\eta_\lambda \kappa}{\epsilon_\lambda} + \kappa + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right) \quad (\text{the update rule and } \hat{\ell}_{T_{\lambda+1},i} \leq \frac{1}{\epsilon_\lambda} \text{ for all } i \in [K])
 \end{aligned}$$

$$\begin{aligned}
 &= \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \kappa + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right) && (\epsilon_\lambda = \max\{\frac{1}{T}, 2\eta_\lambda\}) \\
 &= \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right). && (\eta_\lambda \leq \eta_1 = \frac{1}{4\kappa})
 \end{aligned}$$

Taking a summation from $\lambda = 1, 2, \dots, \lambda^*$, we have

$$\begin{aligned}
 \sum_{t \in \mathcal{T}_m} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} &\leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \frac{K}{T} \min_{i \in [K]} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\frac{1}{\eta_{\lambda^*}} + \frac{K}{T} \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \right). \tag{19}
 \end{aligned}$$

Now we come to bound $\frac{1}{\eta_{\lambda^*}}$. If $\lambda^* = 1$, we have $\sum_{t \in \mathcal{T}_m} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \leq \tilde{\mathcal{O}}(K) + \frac{2K}{T} \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}$. If $\lambda^* \geq 2$, consider the last round of epoch $\lambda^* - 1$. According to the update rule of η , we have

$$\frac{1}{\eta_{\lambda^*-1}^2} \leq 4\kappa \min_{i \in [K]} \sum_{t=T_{\lambda^*-1}+1}^{T_{\lambda^*}} \hat{\ell}_{t,i} \leq 4\kappa \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}.$$

Therefore, we have $\frac{1}{\eta_{\lambda^*}} \leq \sqrt{16\kappa \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}}$. In addition, note that $\eta_{\lambda^*} \geq \sqrt{\frac{1}{\alpha T}}$ by [Line 17](#) of the algorithm. So $\frac{1}{\eta_{\lambda^*}} = \tilde{\mathcal{O}} \left(\min \left\{ \sqrt{\alpha T}, \sqrt{\kappa \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}} \right\} + \kappa \right)$. Plugging this into [Eq. \(19\)](#) completes the proof. \blacksquare

Now we are ready to prove [Theorem 14](#).

Proof of Theorem 14 We first show that with high probability, the algorithm does not enter the second stage if $\alpha T \geq 64\kappa L_*$. Let $m^* \leq \log_2 T$ be the number of meta-epochs executed in the first stage. Define random variables as follows:

$$R_m \triangleq \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i}, \quad m \in [m^*].$$

As we reset all the parameters for each meta-epoch, we have

$$\mathbb{E} [R_m \mid R_1, \dots, R_{m-1}] \leq \mathbb{E} \left[\min_{i \in [K]} \sum_{t=1}^T \hat{\ell}_{t,i} \right] \leq \min_{i \in [K]} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right] = L_*.$$

On the other hand, according to [Line 17](#), the learning rate η used at the last round of each meta-epoch m is at most $\sqrt{\frac{4}{\alpha T}}$ before being halved. Combining with [Line 15](#), we have $\alpha T \leq \frac{4}{\eta^2} \leq 16\kappa R_m$. Now suppose $\alpha T \geq 64\kappa L_*$. Using Markov inequality, we have

$$\text{Prob} \left[R_m \geq \frac{\alpha T}{16\kappa} \mid R_1, \dots, R_{m-1} \right] \leq \frac{16\kappa L_*}{\alpha T} \leq \frac{1}{4}, \quad \forall m \in [m^*].$$

Therefore, the probability that the algorithm reaches the second stage when $\alpha T \geq 64\kappa L_\star$ is upper bounded as follows

$$\begin{aligned} \text{Prob} [\text{Algorithm 5 reaches the second stage}] &\leq \prod_{m=1}^{\log_2 T} \text{Prob} \left[R_m \geq \frac{\alpha T}{16\kappa} \mid R_1, \dots, R_{m-1} \right] \\ &\leq \left(\frac{1}{4} \right)^{\log_2 T} = \frac{1}{T^2}. \end{aligned}$$

This shows that it is unlikely to reach the second stage when $\alpha T \geq 64\kappa L_\star$. Now consider the expected regret when $\alpha T \geq 64\kappa L_\star$. With probability $1 - \frac{1}{T^2}$, the regret only comes from the first stage. According to [Lemma 16](#), we have

$$\begin{aligned} \text{Reg} &\leq \mathbb{E} \left[\sum_{m=1}^{m^\star} \left(\sum_{t \in \mathcal{T}_m} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \right) \right] + \frac{1}{T^2} \cdot T \\ &\leq \tilde{\mathcal{O}} \left(\mathbb{E} \left[\min \left\{ \sqrt{\alpha T}, \sqrt{\kappa \min_{i \in [K]} \sum_{t=1}^T \hat{\ell}_{t,i}} \right\} \right] + K \right) \\ &\leq \tilde{\mathcal{O}} \left(\min \left\{ \sqrt{\alpha T}, \sqrt{\kappa L_\star} \right\} + K \right). \end{aligned}$$

On the other hand, when $\alpha T \leq 64\kappa L_\star$, the regret is bounded by the sum of the worst-case regret from both stages. According to [Theorem 2](#) and [Lemma 16](#), we arrive at

$$\begin{aligned} \text{Reg} &\leq \mathbb{E} \left[\sum_{m=1}^{m^\star} \left(\sum_{t \in \mathcal{T}_m} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in [K]} \sum_{t \in \mathcal{T}_m} \hat{\ell}_{t,i} \right) \right] + \tilde{\mathcal{O}} \left(\sqrt{\alpha T} + K^2 \right) \\ &\leq \tilde{\mathcal{O}} \left(\sqrt{\alpha T} + K^2 \right) = \tilde{\mathcal{O}} \left(\min \left\{ \sqrt{\alpha T}, \sqrt{\kappa L_\star} \right\} + K^2 \right). \end{aligned}$$

Combining the two cases completes the proof. ■

Appendix E. Omitted Details for [Section 4](#)

In this section, we provide omitted details for [Section 4](#), including the proof of [Theorem 5](#) ([Appendix E.1](#)), the proof of [Theorem 7](#) ([Appendix E.2](#)), an adaptive version of [Algorithm 2](#) for directed complete bipartite graphs and its analysis ([Appendix E.3](#)), the proof of [Theorem 8](#) ([Appendix E.4](#)), and an adaptive version of [Algorithm 2](#) for general weakly observable graphs and its analysis ([Appendix E.5](#)). For notational convenience, we use $\mathbf{1}_U$ to denote a vector in \mathbb{R}^K whose i -th coordinate is 1 if $i \in U$ and 0 otherwise. We also define $h(x) \triangleq x - 1 - \ln x$ so that for a hybrid regularizer of the form $\psi(p) = \frac{1}{\eta} \sum_{i \in \mathcal{S}} \ln \frac{1}{p_i} + \frac{1}{\eta} \sum_{i \in \bar{\mathcal{S}}} p_i \ln p_i$, its associated Bregman divergence is $D_\psi(p, q) = \frac{1}{\eta} \sum_{i \in \mathcal{S}} h(p_i/q_i) + \frac{1}{\eta} \sum_{i \in \bar{\mathcal{S}}} p_i \ln(p_i/q_i)$.

E.1. Proof of [Theorem 5](#)

Proof Fix any $\epsilon \in (0, 1/3)$. Since the feedback graph is weakly observable, there must exist two nodes u and v , such that u does not have a self-loop and v cannot observe u . Consider the

following environment: $\forall t \in [T], \ell_{t,u} = 0, \ell_{t,v} = T^{-b}$ for some b such that $\epsilon < b < 1 - 2\epsilon$, and $\ell_{t,w} = 1, \forall w \neq u, v$. We call nodes other than u and v “bad arms”. Note that the loss of each node is constant over time and $L_\star = 0$.

Since \mathcal{A} achieves $\tilde{\mathcal{O}}(1)$ regret when $L_\star = 0$, the expected number of times that bad arms are selected by \mathcal{A} (denoted by \mathcal{N}_{bad}) is at most $\mathcal{O}(T^\epsilon)$. Similarly, the expected number of times that v is selected by \mathcal{A} (denoted by \mathcal{N}_v) is at most $\mathcal{O}(T^{b+\epsilon})$. The condition $\mathcal{N}_{\text{bad}} = \mathcal{O}(T^\epsilon)$ implies that we can find a interval of length $\frac{T}{2\mathcal{N}_{\text{bad}}+1} = \Omega(T^{1-\epsilon})$ such that \mathcal{A} selects bad arms less than $\frac{1}{2}$ times in expectation in this interval. Therefore, by Markov inequality, we have with probability $\frac{1}{2}$, \mathcal{A} does not select bad arms at all in this interval.

Now consider creating another environment by switching the loss of u to 1 only in this interval. Note that with probability $\frac{1}{2}$, \mathcal{A} cannot notice the change because u 's loss is not revealed if none of the bad arms is selected. Since $\mathcal{N}_v = \mathcal{O}(T^{b+\epsilon})$, we conclude that \mathcal{A} suffers expected loss $\Omega(T^{1-\epsilon} - \mathcal{N}_v) = \Omega(T^{1-\epsilon})$ in this interval using the condition $b < 1 - 2\epsilon$. Moreover, v becomes the best arm in this new environment and $L_\star = T^{1-b}$. Therefore, \mathcal{A} suffers expected regret $\Omega(T^{1-\epsilon} - T^{1-b}) = \Omega(T^{1-\epsilon})$ since $\epsilon < b$, which completes the proof. \blacksquare

E.2. Proofs for Theorem 7

We first prove Lemma 6, which will be useful for the proofs of Theorem 7 and Theorem 8.

Proof of Lemma 6 Let $\tilde{p}_{t+1} = \operatorname{argmin}_{p \in \mathbb{R}_+^K} \left\{ \langle p, \hat{\ell}_t + a_t \rangle + D_\psi(p, p_t) \right\}$. One can verify that $p_{t+1} = \operatorname{argmin}_{p \in \Omega} \{ D_\psi(p, \tilde{p}_{t+1}) \}$ and for any $u \in \Omega$, we have

$$\begin{aligned} & \langle p_t - u, \hat{\ell}_t + a_t \rangle \\ &= D_\psi(u, p_t) - D_\psi(u, \tilde{p}_{t+1}) + D_\psi(p_t, \tilde{p}_{t+1}) \\ &\leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + D_\psi(p_t, \tilde{p}_{t+1}) \\ &= D_\psi(u, p_t) - D_\psi(u, p_{t+1}) \\ &\quad + \frac{1}{\eta} \sum_{i \in S} \left(\frac{p_{t,i}}{\tilde{p}_{t+1,i}} - 1 - \ln \left(\frac{p_{t,i}}{\tilde{p}_{t+1,i}} \right) \right) + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} \left(p_{t,i} \ln \frac{p_{t,i}}{\tilde{p}_{t+1,i}} - (p_{t,i} - \tilde{p}_{t+1,i}) \right), \end{aligned}$$

where the second inequality is by the generalized Pythagorean theorem. According to the choice of ψ , one can also obtain the close-form of \tilde{p}_{t+1} , which satisfies

$$\frac{p_{t,i}}{\tilde{p}_{t+1,i}} = 1 + \eta p_{t,i} (\hat{\ell}_{t,i} + a_{t,i}), \text{ for } i \in S; \quad \ln \left(\frac{p_{t,i}}{\tilde{p}_{t+1,i}} \right) = \bar{\eta} (\hat{\ell}_{t,i} + a_{t,i}), \text{ for } i \in \bar{S}.$$

Plugging this into the previous inequality shows

$$\begin{aligned} & \langle p_t - u, \hat{\ell}_t + a_t \rangle \\ &\leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \frac{1}{\eta} \sum_{i \in S} \left(\eta p_{t,i} (\hat{\ell}_{t,i} + a_{t,i}) - \ln \left(1 + \eta p_{t,i} (\hat{\ell}_{t,i} + a_{t,i}) \right) \right) \\ &\quad + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} \left(\bar{\eta} p_{t,i} (\hat{\ell}_{t,i} + a_{t,i}) - p_{t,i} + p_{t,i} \exp \left(-\bar{\eta} (\hat{\ell}_{t,i} + a_{t,i}) \right) \right) \end{aligned}$$

Using the facts $\ln(1+x) \geq x - x^2$ and $\exp(-x) \leq 1 - x + x^2$ for any $x \geq 0$, and realizing $\hat{\ell}_{t,i} + a_{t,i} \geq 0$ holds, we further have

$$\langle p_t - u, \hat{\ell}_t + a_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \sum_{i \in S} \eta p_{t,i}^2 (\hat{\ell}_{t,i} + a_{t,i})^2 + \sum_{i \in \bar{S}} \bar{\eta} p_{t,i} (\hat{\ell}_{t,i} + a_{t,i})^2. \quad (20)$$

Next we use the conditions for η and $\bar{\eta}$ and the concrete form of a_t to bound the last two terms as

$$\begin{aligned} & \sum_{i \in S} \eta p_{t,i}^2 (\hat{\ell}_{t,i} + a_{t,i})^2 + \sum_{i \in \bar{S}} \bar{\eta} p_{t,i} (\hat{\ell}_{t,i} + a_{t,i})^2 \\ &= \sum_{i \in S} \eta p_{t,i}^2 (\hat{\ell}_{t,i} + 2\eta p_{t,i} \hat{\ell}_{t,i}^2)^2 + \sum_{i \in \bar{S}} \bar{\eta} p_{t,i} (\hat{\ell}_{t,i} + 2\bar{\eta} \hat{\ell}_{t,i}^2)^2 \\ &= \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 (1 + 2\eta p_{t,i} \hat{\ell}_{t,i})^2 + \sum_{i \in \bar{S}} \bar{\eta} p_{t,i} \hat{\ell}_{t,i}^2 (1 + 2\bar{\eta} \hat{\ell}_{t,i})^2 \\ &\leq \sum_{i \in S} \eta p_{t,i}^2 \hat{\ell}_{t,i}^2 (1 + 2\eta)^2 + \sum_{i \in \bar{S}} \bar{\eta} p_{t,i} \hat{\ell}_{t,i}^2 \left(1 + \frac{2}{5}\right)^2 \quad (p_{t,i} \hat{\ell}_{t,i} \leq 1, i \in S; \bar{\eta} \hat{\ell}_{t,i} \leq \frac{1}{5}, i \in \bar{S}) \\ &\leq \langle p_t, a_t \rangle. \quad (\eta \leq \frac{1}{5}) \end{aligned}$$

The proof is completed by plugging the inequality above into [Eq. \(20\)](#) and rearranging terms. \blacksquare

Proof of Theorem 7 The condition of [Lemma 6](#) holds since according to the definition of $\bar{\eta}$ and Ω , we have $\frac{\bar{\eta}}{W_{t,i}} \leq \frac{\bar{\eta}}{\sqrt{\bar{\eta}}} = \sqrt{\bar{\eta}} \leq \frac{1}{5}$, $\forall t \in [T]$ and $i \in \bar{S}$. Thus, by [Lemma 6](#), we have for any $u \in \Omega$

$$\langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_t) - D_\psi(u, p_{t+1}) + \langle u, a_t \rangle.$$

Summing over $t \in [T]$, we have:

$$\sum_{t=1}^T \langle p_t - u, \hat{\ell}_t \rangle \leq D_\psi(u, p_1) + \sum_{t=1}^T \langle u, a_t \rangle = \frac{1}{\eta} \sum_{i \in S} h\left(\frac{u_i}{p_{1,i}}\right) + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} u_i \ln \frac{u_i}{p_{1,i}} + \sum_{t=1}^T \langle u, a_t \rangle.$$

When comparing to a node $i \in S$, we set $u = \frac{1}{T} \cdot \mathbf{1}_S + (1 - \frac{s}{T}) \cdot e_i \in \Omega$. Using the definition of p_1 , we have

$$\begin{aligned} D_\psi(u, p_1) &= \frac{1}{\eta} \sum_{i \in S} h\left(\frac{u_i}{p_{1,i}}\right) + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} u_i \ln \frac{u_i}{p_{1,i}} \\ &= \frac{1}{\eta} \sum_{j \neq i, j \in S} h\left(\frac{2s}{T}\right) + \frac{1}{\eta} h\left(2s \left(1 - \frac{s-1}{T}\right)\right) \\ &\leq \frac{s-1}{\eta} \left(\frac{2s}{T} - 1 - \ln \frac{2s}{T}\right) + \frac{1}{\eta} h(2s) \\ &= \frac{s-1}{\eta} \left(\frac{2s}{T} - 1 - \ln \frac{2s}{T}\right) + \frac{1}{\eta} (2s - 1 - \ln 2s) \\ &\leq \frac{(s-1) \ln T}{\eta} + \frac{2s-1-s \ln 2s}{\eta} \leq \frac{s \ln T}{\eta}, \end{aligned}$$

The first inequality is because $h(y)$ is increasing when $y \geq 1$ and $T \geq 2K \geq 2s$. The second inequality is also because $T \geq 2s$ and the last one is because $2s - 1 - s \ln 2s \leq 1 \leq \ln T$ for all $s > 0$. Therefore, we have

$$\begin{aligned} \sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle &\leq \frac{s \ln T}{\eta} + 2\eta \sum_{t=1}^T p_{t,i} \hat{\ell}_{t,i}^2 + \frac{2\eta}{T} \sum_{t=1}^T \sum_{j \in S} p_{t,j} \hat{\ell}_{t,j}^2 + \frac{1}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j} \\ &\leq \frac{s \ln T}{\eta} + 2\eta \sum_{t=1}^T \hat{\ell}_{t,i} + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j}, \end{aligned} \quad (21)$$

where the second inequality is because $p_{t,i} \hat{\ell}_{t,i} \leq 1$ for all $i \in S$ and $\eta \leq \frac{1}{5} \leq \frac{1}{2}$. When comparing with node $i \in \bar{S}$, let $\hat{i}_S^* = \operatorname{argmin}_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}$ and we choose $u = \sqrt{\bar{\eta}} \cdot e_{\hat{i}_S^*} + \frac{1}{T} \cdot \mathbf{1}_S + (1 - \frac{s}{T} - \sqrt{\bar{\eta}}) e_i \in \Omega$. According to the choice of p_1 , we bound the Bregman divergence term as

$$\begin{aligned} D_\psi(u, p_1) &= \frac{1}{\eta} \sum_{i \in S} h\left(\frac{u_i}{p_{1,i}}\right) + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} u_i \ln \frac{u_i}{p_{1,i}} \\ &= \frac{1}{\eta} \sum_{j \neq \hat{i}_S^*, j \in S} h\left(\frac{2s}{T}\right) + \frac{1}{\eta} h\left(2s \left(\frac{1}{T} + \sqrt{\bar{\eta}}\right)\right) + \frac{1}{\bar{\eta}} \left(1 - \frac{s}{T} - \sqrt{\bar{\eta}}\right) \ln\left(2\bar{s} \left(1 - \frac{s}{T} - \sqrt{\bar{\eta}}\right)\right) \\ &\leq \frac{s}{\eta} \left(\frac{2s}{T} - 1 - \ln\left(\frac{2s}{T}\right)\right) + \frac{1}{\eta} \left(\frac{2s}{T} + 2s\sqrt{\bar{\eta}} - 1 - \ln\left(\frac{2s}{T} + 2s\sqrt{\bar{\eta}}\right)\right) + \frac{1}{\bar{\eta}} \left(1 - \frac{s}{T} - \sqrt{\bar{\eta}}\right) \ln(2\bar{s}) \\ &\leq \frac{s \ln T}{\eta} + \frac{1}{\eta} (2s\sqrt{\bar{\eta}} + \ln T) + \frac{1}{\bar{\eta}} \ln(2\bar{s}) \\ &\leq \frac{s \ln T + s + \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} \leq \frac{2s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}}. \end{aligned}$$

Here, the second inequality is because $T \geq 2K \geq 2s$; the third inequality is because $\sqrt{\bar{\eta}} \leq \frac{1}{5} \leq \frac{1}{2}$; and the last inequality is because $\ln T \geq 1$ and $s \geq 1$. Therefore, we have

$$\begin{aligned} \sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle &\leq \frac{2s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + 2\bar{\eta} \sum_{t=1}^T \hat{\ell}_{t,i}^2 + \frac{2\eta}{T} \sum_{t=1}^T \sum_{j \in S} p_{t,j} \hat{\ell}_{t,j}^2 + 2\eta\sqrt{\bar{\eta}} \sum_{t=1}^T p_{t,\hat{i}_S^*} \hat{\ell}_{t,\hat{i}_S^*}^2 \\ &\quad + \sqrt{\bar{\eta}} \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j} \\ &\leq \frac{2s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + 2\sqrt{\bar{\eta}} \sum_{t=1}^T \hat{\ell}_{t,i} + 2\sqrt{\bar{\eta}} \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i} + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j}, \end{aligned} \quad (22)$$

where the second inequality is because $\hat{\ell}_{t,i} \leq \frac{1}{\sqrt{\bar{\eta}}}$ for nodes $i \in \bar{S}$, $p_{t,i} \hat{\ell}_{t,i} \leq 1$ for all $i \in S$, $t \in [T]$ and $\eta \leq \frac{1}{5} \leq \frac{1}{2}$.

Now we take expectation over both sides. If $i \in S$, then

$$\text{Reg}_i = \mathbb{E} \left[\sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle \right] \leq \frac{s \ln T}{\eta} + 2\eta L_i + 2s.$$

If $i \in \bar{S}$, then

$$\begin{aligned} \text{Reg}_i &= \mathbb{E} \left[\sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle \right] \leq \frac{2s \ln T}{\eta} + \frac{2 \ln K}{\bar{\eta}} + 2\sqrt{\bar{\eta}} L_i + 2\sqrt{\bar{\eta}} \mathbb{E} \left[\min_{i \in \bar{S}} \sum_{t=1}^T \hat{\ell}_{t,i} \right] + 2s \\ &\leq \frac{2s \ln T}{\eta} + \frac{2 \ln K}{\bar{\eta}} + 2\sqrt{\bar{\eta}} L_i + 2\sqrt{\bar{\eta}} \min_{i \in \bar{S}} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right] + 2s \\ &= \frac{2s \ln T}{\eta} + \frac{2 \ln K}{\bar{\eta}} + 2\sqrt{\bar{\eta}} L_i + 2\sqrt{\bar{\eta}} L_{i_S^*} + 2s. \end{aligned}$$

The first inequality is because $\ln(2\bar{s}) \leq \ln(2K) \leq 2 \ln K$ as $K \geq 2$, and the second inequality is because of Jensen's inequality. Finally, choosing $\eta = \min\{\sqrt{s/L_{i_S^*}}, 1/5\}$, $\bar{\eta} = \min\{L_{i_S^*}^{-2/3}, 1/25\}$, we have for $i \in S$

$$\text{Reg}_i \leq \text{Reg}_{i_S^*} \leq \mathcal{O} \left(s \ln T + \sqrt{s L_{i_S^*}} \right) \leq \tilde{\mathcal{O}} \left(\sqrt{s L_i} + s \right).$$

For $i \in \bar{S}$, if $L_i \leq L_{i_S^*}$, we have

$$\text{Reg}_i \leq \frac{2s \ln T}{\eta} + \frac{2 \ln K}{\bar{\eta}} + 4\sqrt{\bar{\eta}} L_{i_S^*} + 2s \leq \tilde{\mathcal{O}} \left(L_{i_S^*}^{2/3} + \sqrt{s L_{i_S^*}} + s \right).$$

Otherwise, we have

$$\text{Reg}_i \leq \text{Reg}_{i_S^*} \leq \tilde{\mathcal{O}} \left(\sqrt{s L_{i_S^*}} + s \right).$$

Combining the above results finishes the proof. \blacksquare

E.3. Adaptive Version of Algorithm 2 for Directed Complete Bipartite Graphs

In order to make Algorithm 2 parameter-free, one may consider directly applying doubling trick. However, one technical issue comes from analyzing the last round before each restart where the loss estimator might be too large. To address this issue, we combine Algorithm 2 and the clipping technique, together with a double trick.

The full algorithm is described in Algorithm 6. Similar to previous doubling trick algorithms, we start from some large η and $\bar{\eta}$, run the procedures of Algorithm 2 (Line 5 to Line 8) and reduce the learning rate when the accumulated estimated loss is too large (Line 10 and Line 11). The key difference is that we again follow the clipping idea of (Allenberg et al., 2006): after computing \hat{p}_{t+1} through OMD, we do clipping with threshold μ and renormalization for nodes in S (Line 9). In this way, $\hat{\ell}_{t,i}$ defined in Line 6 is well upper bounded for all $i \in [K]$, $t \in [T]$, which is crucial for the doubling trick analysis. Formally, we prove the following theorem.

Algorithm 6 Adaptive Version of [Algorithm 2](#) for Directed Complete Bipartite Graphs

Input: Feedback graph G and parameter $\eta \leq \frac{1}{5}$.

Initialize: p_1 is such that $p_{1,i} = \frac{1}{2s}$ for $i \in S$ and $p_{1,i} = \frac{1}{2\bar{s}}$ for $i \in \bar{S}$.

for $\lambda = 1, 2, \dots$ **do**

```

1    $p_t = p_1, \bar{\eta} = s^{-\frac{2}{3}}\eta^{\frac{4}{3}}, T_\lambda = t - 1.$ 
2   Define decision set  $\Omega = \{p \in \Delta(K) : \sum_{i \in S} p_i \geq \sqrt{\bar{\eta}}\}.$ 
3   Define hybrid regularizer  $\psi(p) = \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{p_i} + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} p_i \ln p_i.$ 
4   while  $t \leq T$  do
5       Play arm  $i_t \sim p_t$  and receive feedback  $\ell_{t,i}$  for all  $i$  such that  $i_t \in N^{\text{in}}(i).$ 
6       Construct estimator  $\hat{\ell}_t$  such that  $\hat{\ell}_{t,i} = \begin{cases} 0, & p_{t,i} = 0, \\ \frac{\ell_{t,i}}{W_{t,i}} \cdot \mathbb{1}\{i_t \in N^{\text{in}}(i)\}, & p_{t,i} > 0, \end{cases}$ 
       where  $W_{t,i} = \sum_{j \in N^{\text{in}}(i)} p_{t,j}.$ 
7       Construct correction term  $a_t$  such that  $a_{t,i} = \begin{cases} 2\eta p_{t,i} \hat{\ell}_{t,i}^2, & \text{for } i \in S, \\ 2\bar{\eta} \hat{\ell}_{t,i}^2, & \text{for } i \in \bar{S}. \end{cases}$ 
8       Compute  $\hat{p}_{t+1} = \operatorname{argmin}_{p \in \Omega} \left\{ \langle p, \hat{\ell}_t + a_t \rangle + D_\psi(p, \hat{p}_t) \right\}.$ 
9       Construct  $p_{t+1}$  as follows, where  $\mu = \frac{\eta\sqrt{\bar{\eta}}}{s}$ :
           
$$p_{t+1,i} = \begin{cases} \frac{\hat{p}_{t+1,i} \mathbb{1}\{\hat{p}_{t+1,i} \geq \mu\}}{\sum_{i' \in S} \hat{p}_{t+1,i'} \mathbb{1}\{\hat{p}_{t+1,i'} \geq \mu\}} \cdot \sum_{i' \in S} \hat{p}_{t+1,i'}, & \text{if } i \in S, \\ \hat{p}_{t+1,i}, & \text{if } i \in \bar{S}. \end{cases}$$

10      if  $\frac{s}{\eta} \leq \eta \min_{i \in S} \sum_{\tau=T_\lambda+1}^t \hat{\ell}_{\tau,i}$  then
11           $\eta \leftarrow \eta/2, t \leftarrow t + 1.$ 
          Break.
           $t \leftarrow t + 1.$ 

```

Theorem 17 [Algorithm 6](#) with $\eta = \min\{\frac{1}{5}, \frac{1}{s}\}$ guarantees for any directed complete bipartite graph:

$$\text{Reg} = \begin{cases} \tilde{O}\left(\sqrt{sL_{i_S^*}} + s^2\right), & \text{if } i^* \in S. \\ \tilde{O}\left(L_{i_S^*}^{2/3} + \sqrt{sL_{i_S^*}} + s^2\right), & \text{if } i^* \in \bar{S}. \end{cases}$$

Proof We call the time steps between two resets an epoch (indexed by λ) and let $\eta_\lambda, \bar{\eta}_\lambda,$ and μ_λ be the value of $\eta, \bar{\eta},$ and μ during epoch λ so that $\eta_\lambda = 2^{1-\lambda}\eta_1, \bar{\eta}_\lambda = s^{-2/3}\eta_\lambda^{4/3},$ and $\mu_\lambda = \frac{\eta_\lambda\sqrt{\bar{\eta}_\lambda}}{s}.$ Also let λ^* be the index of the last epoch. As we only do clipping restricted on the nodes in $S,$ all nodes in \bar{S} can still be observed with probability greater than zero. Therefore, $\hat{\ell}_{t,i}$ is still unbiased for any node $i \in \bar{S}$ and we have $\mathbb{E}[\ell_{t,i}] = \mathbb{E}[\langle p_t, \hat{\ell}_t \rangle].$ In addition, in each epoch $\lambda,$ as the clipping threshold is $\mu_\lambda \leq \frac{\sqrt{\bar{\eta}_\lambda}}{s},$ at least one node in S will survive and we have $1 \leq \frac{p_{t,i_t}}{\hat{p}_{t,i_t}} \leq \frac{1}{1 - \frac{s\mu_\lambda}{\sqrt{\bar{\eta}_\lambda}}} = \frac{1}{1 - \eta_\lambda}.$

Now we consider the regret in epoch λ . We will prove that

$$\frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \begin{cases} \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right), & \text{if } i^* \in \bar{S}. \end{cases} \quad (23)$$

When $i^* \in S$, we have

$$\begin{aligned} & \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ & \leq \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \eta_\lambda \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \eta_\lambda \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \hat{\ell}_{t,i} + \eta_\lambda \max_{i \in S} \hat{\ell}_{T_{\lambda+1},i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + s^{\frac{4}{3}} \eta_\lambda^{-\frac{2}{3}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right). \end{aligned}$$

The second inequality is derived by rearranging terms in Eq. (21). The fourth inequality is because $\frac{s}{\eta_\lambda} \leq \eta_\lambda \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \hat{\ell}_{t,i}$ does not hold and $\hat{\ell}_{T_{\lambda+1},i} \leq \frac{1}{\mu_\lambda}$ holds for all $i \in [K]$. The last inequality is because $\eta_\lambda \leq \eta_1 \leq \frac{1}{s}$. On the other hand, if $i^* \in \bar{S}$, we have

$$\begin{aligned} & \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + (\sqrt{\eta_\lambda} + \eta_\lambda) \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} + \sqrt{\eta_\lambda} \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right). \end{aligned}$$

which is also derived by rearranging terms in Eq. (22). Then we consider the following two cases.

If $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}$, we have

$$\begin{aligned} & \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ & \leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + (\sqrt{\eta_\lambda} + \eta_\lambda) \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \end{aligned}$$

$$\begin{aligned}
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + (\sqrt{\eta_\lambda} + \eta_\lambda) \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \hat{\ell}_{t,i} + (\sqrt{\eta_\lambda} + \eta_\lambda) \max_{i \in S} \hat{\ell}_{T_{\lambda+1},i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + (\sqrt{\eta_\lambda} + \eta_\lambda) \max_{i \in S} \hat{\ell}_{T_{\lambda+1},i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + \frac{s}{\eta_\lambda} + s^{\frac{4}{3}} \eta_\lambda^{-\frac{2}{3}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &= \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right).
 \end{aligned}$$

Here, the third inequality is because $\frac{s}{\eta_\lambda} \leq \eta_\lambda \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \hat{\ell}_{t,i}$ does not hold, which also implies that $\frac{1}{\eta_\lambda} \leq \eta_\lambda \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}-1} \hat{\ell}_{t,i}$ does not hold; the fourth inequality is also $\hat{\ell}_{T_{\lambda+1},i} \leq \frac{1}{\mu_\lambda}$ for all $i \in S$; and the last inequality is because $\eta_\lambda \leq \eta_1 \leq \frac{1}{s}$.

On the other hand, if $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \geq \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}$, then based on previous results, we have

$$\begin{aligned}
 \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} &\leq \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \\
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right).
 \end{aligned}$$

Combining the two cases, we finish proving [Eq. \(23\)](#). Now we sum up the regret over all epochs $\lambda = 1, 2, \dots, \lambda^*$. For $i^* \in S$, we have

$$\begin{aligned}
 \sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \right) &\leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_{\lambda^*}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &= \tilde{\mathcal{O}} \left(\frac{2^{\lambda^*} s}{\eta_1} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \tag{24}
 \end{aligned}$$

For $i^* \in \bar{S}$, we have

$$\begin{aligned}
 \sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \right) &\leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}} \left(\frac{s}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\frac{s}{\eta_{\lambda^*}} + \frac{1}{\bar{\eta}_{\lambda^*}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right)
 \end{aligned}$$

$$= \tilde{\mathcal{O}} \left(\frac{2^{\lambda^*} s}{\eta_1} + \frac{2^{\frac{4}{3}\lambda^*}}{\bar{\eta}_1} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right). \quad (25)$$

Below we show that λ^* is well upper bounded. When $\lambda^* \geq 2$, consider the last time step of epoch $\lambda^* - 1$, we have

$$\frac{2^{2\lambda^*-2} s}{\eta_1^2} = \frac{s}{\eta_{\lambda^*-1}^2} \leq \min_{i \in S} \sum_{t=T_{\lambda^*-1}+1}^{T_{\lambda^*}} \hat{\ell}_{t,i} \leq \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}.$$

Therefore, we know that $2^{\lambda^*} \leq 2\eta_1 \sqrt{\frac{\min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}}{s}}$. Plugging this into Eq. (24), we have for $i^* \in S$,

$$\begin{aligned} \sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \right) &\leq \tilde{\mathcal{O}} \left(\frac{2^{\lambda^*} s}{\eta_1} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ &\leq \tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right). \end{aligned}$$

On the other hand, for $i^* \in \bar{S}$, plugging $2^{\lambda^*} \leq 2\eta_1 \sqrt{\frac{\min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}}{s}}$ into Eq. (25) gives

$$\begin{aligned} \sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \right) \\ \leq \tilde{\mathcal{O}} \left(\frac{2^{\lambda^*} s}{\eta_1} + \frac{2^{\frac{4}{3}\lambda^*}}{\bar{\eta}_1} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ \leq \tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + \left(\min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i} \right)^{\frac{2}{3}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right). \end{aligned}$$

Combining with the case $\lambda^* = 1$, we have the following result

$$\begin{aligned} \frac{1}{1-\eta_\lambda} \left(\sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle \right) - \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ \leq \begin{cases} \tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + s^2 + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + \left(\min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i} \right)^{\frac{2}{3}} + s^2 + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right), & \text{if } i^* \in \bar{S}. \end{cases} \end{aligned}$$

Now we take the expectation over both sides. First, for the left hand side, we have

$$\mathbb{E} \left[\sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_{\lambda+1}}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \right) \right]$$

$$\begin{aligned}
 &= \mathbb{E} \left[\sum_{\lambda=1}^{\lambda^*} \left(\frac{1}{1-\eta_\lambda} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \ell_{t,i_t} \right) \right] + \mathbb{E} \left[\sum_{t=1}^T \ell_{t,i_t} - \hat{\ell}_{t,i^*} \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \frac{1}{1-\eta_{\lambda_t}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \ell_{t,i_t} \right] + \mathbb{E} \left[\sum_{t=1}^T \ell_{t,i_t} - \ell_{t,i^*} \right] \\
 &= \sum_{t=1}^T \mathbb{E}_{\lambda_t} \left[\mathbb{E}_{i_t|\lambda_t} \left[\frac{1}{1-\eta_{\lambda_t}} \langle \hat{p}_t, \hat{\ell}_t \rangle - \ell_{t,i_t} \right] \right] + \mathbb{E} \left[\sum_{t=1}^T \ell_{t,i_t} - \ell_{t,i^*} \right] \\
 &\geq \mathbb{E} \left[\sum_{t=1}^T \ell_{t,i_t} - \ell_{t,i^*} \right].
 \end{aligned}$$

Here, λ_t represents the epoch that time t belongs to. The last inequality is because of the fact that whether t is in epoch λ or not is independent of what action is realized in time t and $1 \leq \frac{p_{t,i_t}}{\hat{p}_{t,i_t}} \leq \frac{1}{1-\eta_\lambda}$.

Next we consider the right hand side. For $i^* \in S$, we have

$$\begin{aligned}
 \mathbb{E} \left[\tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + s^2 + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \right] &\leq \tilde{\mathcal{O}} \left(\left[\sqrt{s \min_{i \in S} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right]} \right] + s^2 \right) \\
 &= \tilde{\mathcal{O}} \left(\sqrt{sL_{i_S^*}} + s^2 \right).
 \end{aligned}$$

where we use Jensen's inequality. For $i^* \in \bar{S}$, we have

$$\begin{aligned}
 &\mathbb{E} \left[\tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i}} + \left(\min_{i \in S} \sum_{t=1}^T \hat{\ell}_{t,i} \right)^{\frac{2}{3}} + s^2 + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \right] \\
 &\leq \tilde{\mathcal{O}} \left(\sqrt{s \min_{i \in S} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right]} + \left(\min_{i \in S} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right] \right)^{\frac{2}{3}} + s^2 \right) \\
 &= \tilde{\mathcal{O}} \left(\sqrt{sL_{i_S^*}} + L_{i_S^*}^{2/3} + s^2 \right).
 \end{aligned}$$

Finally combining the results above proves the theorem statement:

$$\text{Reg} = \begin{cases} \tilde{\mathcal{O}} \left(\sqrt{sL_{i_S^*}} + s^2 \right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}} \left(L_{i_S^*}^{2/3} + \sqrt{sL_{i_S^*}} + s^2 \right), & \text{if } i^* \in \bar{S}. \end{cases}$$

■

E.4. Proof of Theorem 8

Proof The condition of Lemma 6 holds since according to the choice of $\bar{\eta}$ and δ , we have $\frac{\bar{\eta}}{W_{t,i}} \leq \frac{\bar{\eta}}{\delta} \leq \delta^{1/3} \leq \frac{1}{5}$, $\forall t \in [T]$, $i \in \bar{S}$. Therefore, we know that for any $u \in \Omega$,

$$\sum_{t=1}^T \langle p_t - u, \hat{\ell}_t \rangle \leq \frac{1}{\eta} \sum_{i \in S} h \left(\frac{u_i}{p_{1,i}} \right) + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} \left(u_i \ln \frac{u_i}{p_{1,i}} \right) + \sum_{t=1}^T \langle u, a_t \rangle.$$

Set $u = \frac{1}{T} \cdot \mathbf{1}_{S \setminus \mathcal{D}} + \delta \cdot \mathbf{1}_{\mathcal{D}} + \left(1 - d\delta - \frac{|S \setminus \mathcal{D}|}{T}\right) \cdot e_i$. When comparing with $i \in S$, we have

$$\begin{aligned}
 & \sum_{j \in S} h\left(\frac{u_j}{p_{1,j}}\right) + \sum_{j \in \bar{S}} \left(u_j \ln \frac{u_j}{p_{1,j}}\right) \\
 &= \sum_{j \neq i, j \in S \setminus \mathcal{D}} \left(\frac{2s}{T} - 1 - \ln \frac{2s}{T}\right) + \sum_{j \neq i, j \in S \cap \mathcal{D}} (2s\delta - 1 - \ln 2s\delta) \\
 & \quad + \left(2s \left(1 - d\delta - \frac{|S \setminus \mathcal{D}|}{T}\right) - 1 - \ln 2s \left(1 - d\delta - \frac{|S \setminus \mathcal{D}|}{T}\right)\right) + |\bar{S} \cap \mathcal{D}| (\delta \ln(2\bar{s}\delta)) \\
 & \leq (s-1) \ln \frac{T}{2s} + 2s - 1 - \ln \frac{s}{2} \quad (T \geq 2K \geq 2s \text{ and } \frac{1}{T} \leq \delta \leq \min\{\frac{1}{4d}, \frac{1}{4s}\}) \\
 & = (s-1) \ln T - s \ln 2s + 2s + \ln 4 - 1 \\
 & \leq 2s \ln T. \quad (-s \ln 2s + 2s + \ln 4 - 1 \leq 2 \leq (s+1) \ln T)
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 \sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle & \leq \frac{2s \ln T}{\eta} + 2\eta \sum_{t=1}^T p_{t,i} \hat{\ell}_{t,i}^2 + \delta \sum_{t=1}^T \left(2\eta \sum_{j \in S \cap \mathcal{D}} p_{t,j} \hat{\ell}_{t,j}^2 + 2\bar{\eta} \sum_{j \in \bar{S} \cap \mathcal{D}} \hat{\ell}_{t,j}^2\right) \\
 & \quad + \frac{2\eta}{T} \sum_{t=1}^T \sum_{j \in S \setminus \mathcal{D}} p_{t,j} \hat{\ell}_{t,j}^2 + \delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \hat{\ell}_{t,j} + \frac{1}{T} \sum_{t=1}^T \sum_{j \in S \setminus \mathcal{D}} \hat{\ell}_{t,j} \\
 & \leq \frac{2s \ln T}{\eta} + 2\eta \sum_{t=1}^T \hat{\ell}_{t,i} + 2\delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \hat{\ell}_{t,j} + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j}.
 \end{aligned}$$

The second inequality is because $p_{t,j} \hat{\ell}_{t,j} \leq 1$ for $j \in S$, $\hat{\ell}_{t,j} \leq \frac{1}{\delta}$ for $j \in \bar{S}$, $2\bar{\eta} \leq \delta$ and $2\eta \leq 1$. The reason that $\hat{\ell}_{t,j} \leq \frac{1}{\delta}$ holds for $j \in \bar{S}$ is that if i is weakly observable, this trivially holds according to the definition of Ω . Otherwise, j can be observed by all the other nodes, which include at least one weakly observable node. This shows that $\hat{\ell}_{t,j} \leq \frac{1}{\delta}$.

On the other hand, when comparing with $i \in \bar{S}$, we have

$$\begin{aligned}
 & \sum_{j \in S} h\left(\frac{u_j}{p_{1,j}}\right) + \sum_{j \in \bar{S}} \left(u_j \ln \frac{u_j}{p_{1,j}}\right) \\
 &= \sum_{j \in S \setminus \mathcal{D}} \left(\frac{2s}{T} - 1 - \ln \frac{2s}{T}\right) + \sum_{j \in S \cap \mathcal{D}} (2s\delta - 1 - \ln 2s\delta) + (|\bar{S} \cap \mathcal{D}| - 1) (\delta \ln(2\bar{s}\delta)) + u_i \ln(2\bar{s}u_i) \\
 & \leq s \ln \frac{T}{2s} + \ln(2\bar{s}) \leq s \ln T + \ln(2\bar{s})
 \end{aligned}$$

because $T \geq 2K \geq 2s$, $\frac{1}{T} \leq \delta \leq \frac{1}{4s}$ and $u_i \leq 1$. Therefore,

$$\sum_{t=1}^T \langle p_t - e_i, \hat{\ell}_t \rangle \leq \frac{s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + 2\eta \sum_{t=1}^T \hat{\ell}_{t,i}^2 + \delta \sum_{t=1}^T \left(2\eta \sum_{j \in S \cap \mathcal{D}} p_{t,j} \hat{\ell}_{t,j}^2 + 2\bar{\eta} \sum_{j \in \bar{S} \cap \mathcal{D}} \hat{\ell}_{t,j}^2\right)$$

Algorithm 7 Adaptive Version for [Algorithm 2](#) for General Weakly Observable Graphs

Input: Feedback graph G and parameter δ .

Initialize: p_1 is such that $p_{1,i} = \frac{1}{2s}$ for $i \in S$ and $p_{1,i} = \frac{1}{2\bar{s}}$ for $i \in \bar{S}$.

```

1 for  $\lambda = 1, 2, \dots$  do
2    $p_t = p_1, \eta = \delta^{\frac{1}{2\gamma}}, \bar{\eta} = \delta^{\frac{1}{2} + \frac{1}{2\gamma}}, T_\lambda = t - 1.$ 
3   Define decision set  $\Omega = \{p \in \Delta(K) : p_i \geq \delta, i \in \mathcal{D}\}.$ 
4   Define hybrid regularizer  $\psi(p) = \frac{1}{\eta} \sum_{i \in S} \ln \frac{1}{p_i} + \frac{1}{\bar{\eta}} \sum_{i \in \bar{S}} p_i \ln p_i.$ 
5   while  $t \leq T$  do
6     Execute Line 1 to Line 4 of Algorithm 2.
7     if  $\delta^{-\frac{1}{\gamma}} \leq \sum_{\tau=T_\lambda+1}^t \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$  then
8        $\delta \leftarrow \frac{\delta}{2}, t \leftarrow t + 1.$ 
9       Break.
10     $t \leftarrow t + 1.$ 
    
```

$$\begin{aligned}
 & + \frac{2\eta}{T} \sum_{t=1}^T \sum_{j \in S \setminus \mathcal{D}} p_{t,j} \hat{\ell}_{t,j}^2 + \delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \hat{\ell}_{t,j} + \frac{1}{T} \sum_{t=1}^T \sum_{j \in S \setminus \mathcal{D}} \hat{\ell}_{t,j} \\
 & \leq \frac{s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + \frac{2\eta}{\delta} \sum_{t=1}^T \hat{\ell}_{t,i} + 2\delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \hat{\ell}_{t,j} + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \hat{\ell}_{t,j},
 \end{aligned}$$

where the second inequality holds by the same reasons for the case $i \in S$. Taking expectation over both sides, we have for $i \in S$

$$\begin{aligned}
 \text{Reg}_i & \leq \frac{2s \ln T}{\eta} + 2\eta \sum_{t=1}^T \mathbb{E} [\hat{\ell}_{t,i}] + 2\delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \mathbb{E} [\hat{\ell}_{t,j}] + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \mathbb{E} [\hat{\ell}_{t,j}] \\
 & \leq \frac{2s \ln T}{\eta} + 2\eta L_i + 2\delta d L_{\mathcal{D}} + 2s.
 \end{aligned}$$

Similarly for $i \in \bar{S}$, we have

$$\begin{aligned}
 \text{Reg}_i & \leq \frac{s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + \frac{2\bar{\eta}}{\delta} \mathbb{E} \left[\sum_{t=1}^T \hat{\ell}_{t,i} \right] + 2\delta \sum_{t=1}^T \sum_{j \in \mathcal{D}} \mathbb{E} [\hat{\ell}_{t,j}] + \frac{2}{T} \sum_{t=1}^T \sum_{j \in S} \mathbb{E} [\hat{\ell}_{t,j}] \\
 & \leq \frac{s \ln T}{\eta} + \frac{\ln(2\bar{s})}{\bar{\eta}} + \frac{2\bar{\eta} L_i}{\delta} + 2\delta d L_{\mathcal{D}} + 2s.
 \end{aligned}$$

This finishes the proof. ■

E.5. Adaptive Version of [Algorithm 2](#) for General Weakly Observable Graphs

In this section, we introduce the adaptive version of [Algorithm 2](#) for general weakly observable graphs. For conciseness, we ignore the dependence on s and d and prove that we can achieve the same result of [Theorem 8](#). We also assume that the weakly dominating set \mathcal{D} contains at least one

node in S if S is not empty. Otherwise, we can add an arbitrary node in S to \mathcal{D} and the new set is still a weakly dominant set (just with size increased by 1). [Algorithm 7](#) shows the full pseudocode. Unlike the directed complete bipartite graphs case, here we do not need the clipping technique. Formally, we have the following theorem.

Theorem 18 [Algorithm 7](#) with $\delta = \min\{\frac{1}{125}, \frac{1}{4d}, \frac{1}{4s}\}$ guarantees:

$$\text{Reg} = \begin{cases} \tilde{\mathcal{O}}\left(L_{\mathcal{D}}^{1-\gamma}\right), & \text{if } i^* \in S. \\ \tilde{\mathcal{O}}\left(L_{\mathcal{D}}^{(1+\gamma)/2}\right), & \text{if } i^* \in \bar{S}. \end{cases}$$

Proof We call the time steps between two resets an epoch (indexed by λ) and let δ_λ , η_λ , and $\bar{\eta}_\lambda$ be the value of δ , η , and $\bar{\eta}$ during epoch λ so that $\delta_\lambda = 2^{1-\lambda}\delta_1$, $\eta_\lambda = \delta_\lambda^{1/2\gamma}$, and $\bar{\eta}_\lambda = \delta_\lambda^{(1+\gamma)/2\gamma}$. Also let λ^* be the index of the last epoch. According to the analysis in [Theorem 8](#), for fixed η , $\bar{\eta}$ and δ , we have

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \leq \begin{cases} \tilde{\mathcal{O}}\left(\frac{1}{\eta} + \eta \sum_{t=1}^T \hat{\ell}_{t,i^*} + \delta \sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i}\right), & \text{if } i^* \in S. \\ \tilde{\mathcal{O}}\left(\frac{1}{\eta} + \frac{1}{\bar{\eta}} + \frac{\bar{\eta}}{\delta} \sum_{t=1}^T \hat{\ell}_{t,i^*} + \delta \sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i}\right), & \text{if } i^* \in \bar{S}. \end{cases}$$

Here, $\tilde{\mathcal{O}}(\cdot)$ suppresses all the constant factors, log factors and the dependence on s and d . Now we consider the regret of epoch λ . We will prove that

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \begin{cases} \tilde{\mathcal{O}}\left(\delta_\lambda^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}}\left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right), & \text{if } i^* \in \bar{S}. \end{cases} \quad (26)$$

If $i^* \in S$, we have from the proof of [Theorem 8](#)

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \tilde{\mathcal{O}}\left(\frac{1}{\eta_\lambda} + \eta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right).$$

Compare the terms $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*}$ and $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$. If $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$, then we have

$$\begin{aligned} & \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ & \leq \tilde{\mathcal{O}}\left(\frac{1}{\eta_\lambda} + \left(\delta_\lambda^{\frac{1}{2\gamma}} + \delta_\lambda\right) \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right) \quad (\eta_\lambda = \delta_\lambda^{\frac{1}{2\gamma}}) \\ & \leq \tilde{\mathcal{O}}\left(\delta_\lambda^{-\frac{1}{2\gamma}} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right) \quad (\delta_\lambda \leq 1 \text{ and } \gamma \in [1/3, 1/2]) \\ & \leq \tilde{\mathcal{O}}\left(\delta_\lambda^{-\frac{1}{2\gamma}} + \delta_\lambda^{1-\frac{1}{\gamma}} + \delta_\lambda \sum_{i \in \mathcal{D}} \hat{\ell}_{T_{\lambda+1},i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i}\right) \quad (\delta_\lambda^{-1/\gamma} > \sum_{\tau=T_\lambda+1}^{T_{\lambda+1}-1} \sum_{i \in \mathcal{D}} \hat{\ell}_{\tau,i}) \end{aligned}$$

$$\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right).$$

The last inequality is because for $i \in \mathcal{D}$, if i has a self-loop, then $\hat{\ell}_{t,i} \leq \frac{1}{p_{t,i}} \leq \frac{1}{\delta_\lambda}$; if i does not have a self-loop, it can be observed either by all the other nodes or at least one node in \mathcal{D} , which also means that $\hat{\ell}_{t,i} \leq \frac{1}{\delta_\lambda}$. Therefore, $\delta_\lambda \sum_{i \in \mathcal{D}} \hat{\ell}_{T_\lambda+1,i} \leq d$. We also use $\delta_\lambda^{-1/2\gamma} \leq \delta_\lambda^{1-1/\gamma}$ as $\delta_\lambda \leq 1$ and $\gamma \in [1/3, 1/2]$.

If $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \geq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$, then let $i_0 \in \mathcal{D}$ be the node with a self-loop in \mathcal{D} and we have $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \geq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0}$. Therefore,

$$\begin{aligned} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} &\leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0} \\ &\leq \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \eta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2\gamma}} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right). \end{aligned} \quad (27)$$

The third inequality is because $i_0 \in \mathcal{D}$ and also $\eta_\lambda \leq \delta_\lambda$, and the last inequality is by the same reason as in the previous case.

Next we consider the case $i^* \in \bar{S}$. We have again from the proof of [Theorem 8](#):

$$\begin{aligned} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\ \leq \tilde{\mathcal{O}} \left(\frac{1}{\eta_\lambda} + \frac{1}{\bar{\eta}_\lambda} + \frac{\bar{\eta}_\lambda}{\delta_\lambda} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\ = \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right), \end{aligned}$$

where the last step is because $\eta_\lambda = \delta_\lambda^{1/2\gamma} \geq \delta_\lambda^{(\gamma+1)/2\gamma} = \bar{\eta}_\lambda$. Consider the following two cases. If $S = \emptyset$, then \mathcal{D} contains at least one node in \bar{S} , which means that $\min_{i \in \bar{S}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$. Therefore, following similar steps as in previous cases, we have

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \min_{i \in \bar{S}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i}$$

$$\begin{aligned}
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \min_{i \in S} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \sum_{i \in \mathcal{D}} \hat{\ell}_{T_{\lambda+1},i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}}' \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right).
 \end{aligned}$$

If $S \neq \emptyset$, then \mathcal{D} contains at least one node in S . Let this node be i_0 . Now we compare $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*}$ with $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0}$. If $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0}$, then we have

$$\begin{aligned}
 &\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \\
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0} + \delta_\lambda \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \delta_\lambda^{-\frac{1}{2}+\frac{1}{2\gamma}} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_\lambda^{-\frac{1}{2}-\frac{1}{2\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right),
 \end{aligned}$$

where the second inequality is because $\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}$. Otherwise, according to Eq. (27), we have

$$\sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i^*} \leq \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \langle p_t, \hat{\ell}_t \rangle - \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \hat{\ell}_{t,i_0} \leq \tilde{\mathcal{O}} \left(\delta_\lambda^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right).$$

By combining all the above cases, we finish the proof of Eq. (26). Summing up the regret from $\lambda = 1, 2, \dots, \lambda^*$, we have for $i^* \in S$:

$$\begin{aligned}
 \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \hat{\ell}_{t,i^*} &\leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}} \left(\delta_\lambda^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=T_\lambda+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_{\lambda^*}^{1-\frac{1}{\gamma}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_1^{1-\frac{1}{\gamma}} \cdot 2^{\lambda^* \left(-1+\frac{1}{\gamma}\right)} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right).
 \end{aligned}$$

and for $i^* \in \bar{S}$:

$$\begin{aligned}
 \sum_{t=1}^T \langle p_t, \hat{\ell}_t \rangle - \sum_{t=1}^T \hat{\ell}_{t,i^*} &\leq \sum_{\lambda=1}^{\lambda^*} \tilde{\mathcal{O}} \left(\delta_{\lambda}^{-\frac{1}{2}-\frac{1}{2\gamma}} + \frac{1}{T} \sum_{t=T_{\lambda}+1}^{T_{\lambda+1}} \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_{\lambda^*}^{-\frac{1}{2}-\frac{1}{2\gamma}} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\delta_1^{-\frac{1}{2}-\frac{1}{2\gamma}} \cdot 2^{\lambda^* \left(\frac{1}{2} + \frac{1}{2\gamma} \right)} + \frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} \right).
 \end{aligned}$$

Below we upper bound λ^* . If $\lambda^* \geq 2$, consider the last round of epoch $\lambda^* - 1$. According to the update rule, we have

$$\left(\delta_1 \cdot 2^{-\lambda^*+1} \right)^{-\frac{1}{\gamma}} = \delta_{\lambda^*-1}^{-\frac{1}{\gamma}} \leq \sum_{t=T_{\lambda^*-1}+1}^{T_{\lambda^*}} \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} \leq \sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i}.$$

Plugging the above result to both cases, we have

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \hat{\ell}_t \rangle \leq \begin{cases} \tilde{\mathcal{O}} \left(\frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} + \left(\sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} \right)^{1-\gamma} \right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}} \left(\frac{1}{T} \sum_{t=1}^T \sum_{i \in S} \hat{\ell}_{t,i} + \left(\sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} \right)^{\frac{(1+\gamma)}{2}} \right), & \text{if } i^* \in \bar{S}. \end{cases}$$

Combining the case $\lambda^* = 1$, taking expectation over both sides and using Jensen's inequality, we have:

$$\text{Reg} \leq \begin{cases} \tilde{\mathcal{O}} \left(\mathbb{E} \left[\left(\sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} \right)^{1-\gamma} \right] \right) \leq \tilde{\mathcal{O}} \left(L_{\mathcal{D}}^{1-\gamma} \right), & \text{if } i^* \in S, \\ \tilde{\mathcal{O}} \left(\mathbb{E} \left[\left(\sum_{t=1}^T \sum_{i \in \mathcal{D}} \hat{\ell}_{t,i} \right)^{\frac{(1+\gamma)}{2}} \right] \right) \leq \tilde{\mathcal{O}} \left(L_{\mathcal{D}}^{\frac{1+\gamma}{2}} \right), & \text{if } i^* \in \bar{S}. \end{cases}$$

This completes the proof. ■