

Improved High-Probability Regret for Adversarial Bandits with Time-Varying Feedback Graphs

Haipeng Luo*

University of Southern California

HAIPENGL@USC.EDU

Hanghang Tong*

University of Illinois at Urbana-Champaign

HTONG@ILLINOIS.EDU

Mengxiao Zhang*

University of Southern California

MENGXIAO.ZHANG@USC.EDU

Yuheng Zhang*

University of Illinois at Urbana-Champaign

YUHENGZ2@ILLINOIS.EDU

Editors: Shipra Agrawal and Francesco Orabona

Abstract

We study high-probability regret bounds for adversarial K -armed bandits with time-varying feedback graphs over T rounds. For general strongly observable graphs, we develop an algorithm that achieves the optimal regret $\tilde{O}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ with high probability, where α_t is the independence number of the feedback graph at round t . Compared to the best existing result (Neu, 2015) which only considers graphs with self-loops for all nodes, our result not only holds more generally, but importantly also removes any $\text{poly}(K)$ dependence that can be prohibitively large for applications such as contextual bandits. Furthermore, we also develop the first algorithm that achieves the optimal high-probability regret bound for weakly observable graphs, which even improves the best expected regret bound of (Alon et al., 2015b) by removing the $\mathcal{O}(\sqrt{KT})$ term with a refined analysis. Our algorithms are based on the online mirror descent framework, but importantly with an innovative combination of several techniques. Notably, while earlier works use optimistic biased loss estimators for achieving high-probability bounds, we find it important to use a pessimistic one for nodes without self-loop in a strongly observable graph.

Keywords: multi-armed bandits, bandits with feedback graph, high-probability regret bounds

1. Introduction

In this work, we study adversarial multi-armed bandits (MAB) with directed feedback graphs, which is a generalization of the expert problem (Freund and Schapire, 1997) and the standard MAB problem (Auer et al., 2002). The interaction between the learner and the environment lasts for T rounds. In each round, the learner needs to choose one of K actions while simultaneously an adversary decides the loss for each action. After that, the learner suffers the loss of the chosen action, and her observation is determined based on a directed graph with the K actions as nodes. Specifically, she observes the loss of every action to which the chosen action is connected. When the graph only contains self-loops, this recovers the standard MAB problem, and when the graph is a complete graph, this recovers the expert problem. By allowing arbitrary feedback graphs, however, this model captures many other interesting problems; see (Mannor and Shamir, 2011) for example.

* Authors are listed in alphabetical order.

Table 1: Summary of our results and comparisons with prior work. T is the number of rounds. K is the number of actions. α_t and d_t are respectively the independence number and the weak domination number of feedback graph G_t at round t . The results of (Alon et al., 2015b; Neu, 2015) are for a fixed feedback graph G (so $G_t = G$, $\alpha_t = \alpha$, and $d_t = d$ for all t). Our high-probability regret bound for weakly observable graphs omits some lower-order terms; see Theorem 4 for the complete form.

Graph Type		Expected Regret	High-probability Regret	
		(Alon et al., 2015b)	(Neu, 2015)	Our work
Strongly Observable	General	$\tilde{O}(\sqrt{\alpha T})$	N/A	$\tilde{O}\left(\sqrt{\sum_{t=1}^T \alpha_t} + \max_{t \in [T]} \alpha_t\right)$
	Self-aware		$\tilde{O}(\sqrt{\alpha T} + K)$	
Weakly Observable		$\tilde{O}(d^{1/3}T^{2/3} + \sqrt{KT})$	N/A	$\tilde{O}((\sum_{t=1}^T d_t)^{1/3}T^{1/3} + \frac{1}{T} \sum_{t=1}^T d_t)$

Alon et al. (2015b) characterized the minimax expected regret bound for this problem with a fixed feedback graph G . Specifically, for a strongly observable graph (see Section 2 for all formal definitions), their algorithm achieves $\tilde{O}(\sqrt{\alpha T})$ expected regret where α is the independence number of G , while for a weakly observable graph, they achieve $\tilde{O}(d^{1/3}T^{2/3})$ expected regret (ignoring a $\mathcal{O}(\sqrt{KT})$ term), where d is the weak domination number of G . Both are shown to be near-optimal.

Despite these near-optimal expected regret guarantees, it is known that these algorithm exhibit a huge variance and can in fact suffer $\Theta(T)$ regret with a constant probability (see (Lattimore and Szepesvári, 2020)), which is clearly undesirable in practice. To mitigate this issue, Alon et al. (2017) designed an algorithm called ELP.P, which ensures $\tilde{O}(\sqrt{mT} + m^2T^{1/4})$ regret with high probability for self-aware graphs (a special case of strongly observable graphs in which every node has a self-loop), where m is the size of the maximal acyclic graph in G and can be much larger than α . On the other hand, Neu (2015) developed the EXP3-IX algorithm which uses implicit exploration in the loss estimator construction and achieves $\tilde{O}(\sqrt{\alpha T} + K)$ high-probability regret bound also for self-aware graphs. While the bound is almost optimal, the additional K term could be prohibitively large for applications such as contextual bandits where K is the number of policies (usually considered as exponentially large).

In this work, we significantly improve these results and extend them to more general graphs. For full generality, we also consider a sequence of time-varying feedback graphs G_1, \dots, G_T , each of which can be chosen adaptively by the environment based on the learner's previous actions. We denote the independence number of G_t by α_t and its weakly domination number by d_t . Our main contributions are (see also Table 1):

- In Section 3, we start with a refined analysis showing that EXP3-IX of (Neu, 2015) in fact achieves $\tilde{O}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ high-probability regret bound for self-aware graphs, removing the $\tilde{O}(K)$ dependence of (Neu, 2015). We then extend the same bound to the more general strongly observable graphs via a new algorithm that, on top of the implicit exploration technique of EXP3-IX, further injects certain positive bias to the loss estimator of an action that has no self-loop but is selected with more than $1/2$ probability, making it a pessimistic estimator.
- In Section 4, we propose an algorithm with high-probability regret $\tilde{O}((\sum_{t=1}^T d_t)^{1/3}T^{1/3} + \frac{1}{T} \sum_{t=1}^T d_t)$ for weakly observable graphs (ignoring some lower-order terms). To the best of our knowledge,

this is the first algorithm with (near-optimal) high-probability regret guarantees for such graphs. Moreover, our bound even improves the expected regret bound of (Alon et al., 2015b) by removing the $\tilde{O}(\sqrt{KT})$ term.

We remark that for simplicity we prove our results by assuming the knowledge of $\alpha_1, \dots, \alpha_T$ or d_1, \dots, d_T to tune the parameters, but this can be easily relaxed using the standard doubling trick, making our algorithms completely parameter-free.

Techniques. Our algorithms are based on the well-known Online Mirror Descent (OMD) framework with the entropy regularizer. However, several crucial techniques are needed to achieve our results, including implicit exploration, explicit uniform exploration, injected positive bias, and a loss shifting trick. Among them, using positive bias and thus a pessimistic loss estimator is especially notable since most earlier works use optimistic underestimators for achieving high-probability regret bounds. The combination of these techniques also requires non-trivial analysis.

Related Works. Since Mannor and Shamir (2011) initiated the study of online learning with feedback graphs, many follow-up works consider different variants of the problem, including stochastic feedback graphs (Caron et al., 2012; Baccapatnam et al., 2018; Marinov et al., 2022), minimax regret bounds for different feedback graph types (Alon et al., 2015b; Chen et al., 2021), small-loss bounds (Lykouris et al., 2018; Lee et al., 2020b), best-of-both-world algorithms (Erez and Koren, 2021; Ito et al., 2022), and uninformed time-varying feedback graphs (Cohen et al., 2016).

This work focuses on achieving high-probability regret bounds, which is relatively less studied in the bandit literature but as mentioned extremely important due to the potentially large variance of the regret. As far as we know, to achieve high-probability regret bounds for adversarial bandit problems, there are three categories of methods as discussed below.

The first method is to inject a negative bias to the loss estimators, making them optimistic and trading unbiasedness for lower variance. Examples include the very first work in this line for standard MAB (Auer et al., 2002), linear bandits (Bartlett et al., 2008; Abernethy and Rakhlin, 2009; Zimmert and Lattimore, 2022), and bandits with self-aware feedback graphs (Alon et al., 2017).

The second method is the so-called implicit exploration approach (Kocák et al., 2014) (which also leads to optimistic estimators). Neu (2015) used this method to achieve $\tilde{O}(\sqrt{KT})$ regret for MAB and $\tilde{O}(\sqrt{\alpha T} + K)$ regret for bandits with a fixed self-aware feedback graph, improving over the results by (Alon et al., 2017). Lykouris et al. (2018) also used implicit exploration and achieved high-probability first-order regret bound for bandits with self-aware undirected feedback graphs. However, their regret bounds are either suboptimal in T or in terms of the clique partition number of the graph (which can be much larger than the independence number).

The third method is to use OMD with a self-concordant barrier and an increasing learning rate scheduling, proposed by Lee et al. (2020a). They used this method to achieve high-probability data-dependent regret bounds for MAB, linear bandits, and episodic Markov decision processes. However, using a self-concordant barrier regularizer generally leads to $\tilde{O}(\sqrt{KT})$ regret in bandits with strongly observable feedback graphs and $\tilde{O}(K^{1/3}T^{2/3})$ regret in bandits with weakly observable feedback graphs, making it suboptimal compared to the minimax regret bound $\tilde{O}(\sqrt{\alpha T})$ and $\tilde{O}(d^{1/3}T^{2/3})$ respectively.

All our algorithms adopt the implicit exploration technique for nodes with self-loop. For strongly observable graphs, we find it necessary to further adopt the injected bias idea for nodes without self-loop, but contrary to prior works, our bias is positive, which makes the loss overestimated and intuitively prevents the algorithm from picking such nodes too often without seeing their actual loss.

2. Problem Setup and Notations

Throughout the paper, we denote $\{1, 2, \dots, N\}$ by $[N]$ for some positive integer N . At each round $t \in [T]$, the learner selects one of the K available actions $i_t \in [K]$, while the adversary decides a loss vector $\ell_t \in [0, 1]^K$ with $\ell_{t,i}$ being the loss for action i , and a directed feedback graph $G_t = ([K], E_t)$ where $E_t \subseteq [K] \times [K]$. The adversary can be adaptive and chooses ℓ_t and G_t based on the learner's previous actions i_1, \dots, i_{t-1} in an arbitrary way. At the end of round t , the learner observes some information about ℓ_t according to the feedback graph G_t . Specifically, she observes the loss of action j for all j such that $i_t \in N_t^{\text{in}}(j)$, where $N_t^{\text{in}}(j) = \{i \in [K] : (i, j) \in E_t\}$ is the set of nodes that can observe node j . The standard measure of the learner's performance is the regret, defined as the difference between the total loss of the learner and that of the best fixed action in hindsight

$$\text{Reg} \triangleq \sum_{t=1}^T \ell_{t,i_t} - \sum_{t=1}^T \ell_{t,i^*},$$

where $i^* = \operatorname{argmin}_{i \in [K]} \sum_{t=1}^T \ell_{t,i}$. In this work, we focus on designing algorithms with *high-probability* regret guarantees.

We refer the reader to (Alon et al., 2015b) for the many examples of such a general model, and only point out that the contextual bandit problem (Langford and Zhang, 2007) is indeed a special case where each node corresponds to a policy and each G_t is the union of several cliques. Each such clique consists of all policies that make the same decision for the current context at round t . In this case, K , the number of policies, is usually considered as exponentially large, and only $\text{polylog}(K)$ dependence on the regret is acceptable. This justifies the significance of our results that indeed remove $\text{poly}(K)$ dependence from existing regret bounds.

Strongly/Weakly Observable Graphs. For a directed graph $G = ([K], E)$, 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 its nodes are observable. An observable graph is *strongly observable* if all nodes are strongly observable, and *weakly observable* otherwise. *Self-aware* graphs are a special type of strongly observable graphs where $i \in N^{\text{in}}(i)$ for all $i \in [K]$.

Independent Set and Weakly Dominating Set. An independence set of a directed graph is a subset of nodes in which no two distinct nodes are connected. The size of the largest independence set in G_t , called the *independence number* of G_t , is denoted by α_t . For a weakly observable graph G , a weakly dominating set is a subset \mathcal{D} of nodes such that for any node j in G without self-loop, there exists $i \in \mathcal{D}$ such that i is connected to j . The size of the smallest weakly dominating set of G_t , called the *weak domination number* of G_t , is denoted by d_t .¹

Informed/Uninformed Setting. Under the *informed setting*, the feedback graph G_t is shown to the learner at the beginning of round t before she selects i_t . In other words, the learner's decision at round t can be dependent on G_t . In contrast, the *uninformed setting* is a harder setting, in which the learner observes G_t only at the end of round t after she selects i_t . For strongly observable graphs, we study the harder uninformed setting, while for weakly observable graphs, in light of the $\Omega(K^{1/3}T^{2/3})$ regret lower bound of (Alon et al., 2015a, Theorem 9), we only study the informed setting.

1. We follow the definition in (Ito et al., 2022), which differs by at most 1 compared to that in (Alon et al., 2015b).

Other Notations. Define $S_t \triangleq \{i \in [K] : i \in N_t^{\text{in}}(i)\}$ as the set of nodes with self-loop in G_t . For a differentiable convex function ψ defined on a convex set Ω , we denote the induced Bregman divergence by $D_\psi(x, y) = \psi(x) - \psi(y) - \langle \nabla \psi(y), x - y \rangle$ for any two points $x, y \in \Omega$. For notational convenience, for two vectors $x, y \in \mathbb{R}^K$ and an arbitrary index set $U \subseteq [K]$, we define $\langle x, y \rangle_U \triangleq \sum_{i \in U} x_i y_i$ to be the partial inner product with respect to the coordinates in U . We denote the $(K - 1)$ -dimensional simplex by Δ_K , the all-one vector in \mathbb{R}^K by $\mathbf{1}$, and the i -th standard basis vector in \mathbb{R}^K by e_i . We use the $\tilde{\mathcal{O}}(\cdot)$ notation to hide factors that are logarithmic in K and T .²

3. Optimal High-Probability Regret for Strongly Observable Graphs

In this section, we consider the uninformed setting with strongly observable graphs, that is, each G_t is strongly observable and revealed to the learner after she selects i_t at round t . We propose an algorithm which achieves $\tilde{\mathcal{O}}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ high-probability regret bound. As mentioned, this result improves over those from (Neu, 2015; Alon et al., 2017) in two aspects: first, they only consider self-aware graphs;³ second, our bound enjoys the optimal $(\sum_{t=1}^T \alpha_t)^{1/2}$ dependence with no $\text{poly}(K)$ dependence at all.

To present our algorithm, which is built on top of the EXP3-IX algorithm of (Neu, 2015), we start by reviewing how EXP3-IX works and how it achieves $\tilde{\mathcal{O}}((\sum_{t=1}^T \alpha_t)^{1/2} + K)$ high-probability regret bound for self-aware graphs. At each round t , after picking the action i_t randomly according to $p_t \in \Delta_K$ and observing the loss $\ell_{t,j}$ for all j such that $i_t \in N_t^{\text{in}}(j)$, EXP3-IX constructs the underestimator $\hat{\ell}_t$ for ℓ_t , such that $\hat{\ell}_{t,i} = \frac{\ell_{t,i}}{W_{t,i} + \gamma} \cdot \mathbf{1}\{i_t \in N_t^{\text{in}}(i)\}$ where $W_{t,i} = \sum_{j \in N_t^{\text{in}}(i)} p_{t,j}$ is the probability of observing $\ell_{t,i}$ and $\gamma > 0$ is a bias parameter. Then, the strategy at round $t + 1$ is computed via the standard multiplicative weight update (equivalent to OMD with entropy regularizer): $p_{t+1,i} \propto p_{t,i} \exp(-\eta \hat{\ell}_{t,i})$ for all $i \in [K]$ where $\eta > 0$ is the learning rate.

Following standard analysis of OMD, we know that for any $j \in [K]$,

$$\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle \leq \underbrace{\frac{\log K}{\eta}}_{\text{BIAS-TERM}} + \underbrace{\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \hat{\ell}_{t,i}^2}_{\text{STABILITY-TERM}}. \quad (1)$$

To derive the high-probability regret bound from Eq. (1), Neu (2015) first shows that with probability at least $1 - \delta$, the following two inequalities hold due to the underestimation:

$$\sum_{t=1}^T (\hat{\ell}_{t,j} - \ell_{t,j}) \leq \frac{\log(2K/\delta)}{2\gamma}, \quad \forall j \in [K] \quad (2)$$

$$\sum_{t=1}^T \sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma} (\hat{\ell}_{t,i} - \ell_{t,i}) \leq \frac{K \log(2K/\delta)}{2\gamma}. \quad (3)$$

2. In the text, $\mathcal{O}(\cdot)$ and $\tilde{\mathcal{O}}(\cdot)$ often further hide lower-order terms (in terms of T dependence) and $\text{poly}(\log(1/\delta))$ factors for simplicity. However, in all formal theorem/lemma statements, we use $\mathcal{O}(\cdot)$ to hide universal constants only and $\tilde{\mathcal{O}}(\cdot)$ to also hide factors logarithmic in K and T .

3. Although (Neu, 2015) only considers a fixed feedback graph (i.e. $G_t = G$ for all $t \in [T]$), their result can be directly generalized to time-varying feedback graphs. On the other hand, we point out that (Alon et al., 2017) only considers the easier informed setting.

Define $Q_t \triangleq \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma}$, which is simply $\sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma}$ for self-aware graphs. Using [Eq. \(3\)](#), the stability term can be upper bounded as follows:

$$\text{STABILITY-TERM} \leq \eta \sum_{t=1}^T \sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} \leq \eta \sum_{t=1}^T Q_t + \widetilde{\mathcal{O}}\left(\frac{K\eta}{\gamma}\right). \quad (4)$$

To connect the true regret $\sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,i^*})$ with $\sum_{t=1}^T \langle p_t - e_{i^*}, \widehat{\ell}_t \rangle$, direct calculation shows:

$$\begin{aligned} \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,i^*}) &= \sum_{t=1}^T \langle p_t - e_{i^*}, \widehat{\ell}_t \rangle + \sum_{t=1}^T (\ell_{t,i_t} - \langle p_t, \ell_t \rangle) + \sum_{t=1}^T (\widehat{\ell}_{t,i^*} - \ell_{t,i^*}) \\ &\quad + \sum_{t=1}^T \sum_{i=1}^K (W_{t,i} - \mathbf{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} + \sum_{t=1}^T \sum_{i=1}^K \frac{\gamma p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}. \end{aligned} \quad (5)$$

In this last expression (summation of five terms), the first term is bounded using [Eq. \(1\)](#) and [Eq. \(4\)](#); the second term is upper bounded by $\widetilde{\mathcal{O}}(\sqrt{T})$ via standard Azuma's inequality; the third term is bounded by $\widetilde{\mathcal{O}}(1/\gamma)$ according to [Eq. \(2\)](#); the fourth term is a summation over a martingale sequence and can be bounded by $\widetilde{\mathcal{O}}\left(\sqrt{\sum_{t=1}^T Q_t} + K\right)$ with high probability via Freedman's inequality; and the last term can be bounded by $\gamma \sum_{t=1}^T Q_t$. Combining all the bounds above, we obtain that with high probability, the regret is bounded as follows:

$$\sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,i^*}) \leq \widetilde{\mathcal{O}}\left(\frac{1}{\eta} + \frac{\eta K}{\gamma} + \frac{1}{\gamma} + \sqrt{\sum_{t=1}^T Q_t} + K + (\eta + \gamma) \sum_{t=1}^T Q_t\right).$$

Finally, using the fact that $Q_t = \widetilde{\mathcal{O}}(\alpha_t)$ (Lemma 1 of [\(Kocák et al., 2014\)](#), included as [Lemma 11](#) in this work) and choosing γ and η optimally gives $\widetilde{\mathcal{O}}((\sum_{t=1}^T \alpha_t)^{1/2} + K)$ high-probability bound.

Improvement from $\widetilde{\mathcal{O}}(K)$ to $\widetilde{\mathcal{O}}(\max_{t \in [T]} \alpha_t)$. We now show that with a refined analysis, the undesirable $\widetilde{\mathcal{O}}(K)$ dependence can be improved to $\widetilde{\mathcal{O}}(\max_{t \in [T]} \alpha_t)$ (still for self-aware graphs using the same EXP3-IX algorithm). From the previous analysis sketch of [\(Neu, 2015\)](#), we can see that the $\widetilde{\mathcal{O}}(K)$ dependency comes from two terms: STABILITY-TERM and the fourth term in [Eq. \(5\)](#). The upper bound of STABILITY-TERM is derived by using [Eq. \(3\)](#) and the fourth term in [Eq. \(5\)](#) is bounded via Freedman's inequality. We show that both of these two bounds are in fact loose and can be improved by using a strengthened Freedman's inequality (Lemma 9 of [\(Zimmert and Lattimore, 2022\)](#), included as [Lemma 13](#) in the appendix). Specifically, we prove the following lemma to bound these two terms. Note that this lemma is not restricted to self-aware graphs, and we will use it later for both general strongly observable graphs and weakly observable graphs.

Lemma 1 *For all t and $i \in S_t$, let $\widehat{\ell}_{t,i}$ be the underestimator $\frac{\ell_{t,i}}{W_{t,i} + \gamma} \cdot \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}$ with $\gamma \leq \frac{1}{2}$. Then, with probability at least $1 - \delta$, the following two inequalities hold:*

$$\sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} (\widehat{\ell}_{t,i} - \ell_{t,i}) \leq \mathcal{O}\left(\sum_{t=1}^T \frac{Q_t^2}{\gamma U_T} + U_T \log\left(\frac{KT}{\delta\gamma}\right)\right), \quad (6)$$

Algorithm 1 Algorithm for Strongly Observable Graphs

Input: Parameter $\gamma, \beta, \eta, \mathcal{T} = \emptyset$.

Define: Regularizer $\psi(p) = \frac{1}{\eta} \sum_{i=1}^K p_i \log p_i$.

Initialize: p_1 is such that $p_{1,i} = \frac{1}{K}$ for all $i \in [K]$.

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

- 1 Calculate $\tilde{p}_t = (1 - \eta)p_t + \frac{\eta}{K} \mathbf{1}$.
 - 2 Sample action i_t from \tilde{p}_t .
 - 3 Receive the feedback graph G_t and the feedback $\ell_{t,j}$ for all j such that $i_t \in N_t^{\text{in}}(j)$.
 - 4 Construct estimator $\hat{\ell}_t \in \mathbb{R}^K$ such that $\hat{\ell}_{t,i} = \frac{\ell_{t,i} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}}{W_{t,i} + \gamma \mathbb{1}\{i \in S_t\}}$ where $W_{t,i} = \sum_{j \in N_t^{\text{in}}(i)} \tilde{p}_{t,j}$.
 - 5 If there exists a node $j_t \in \bar{S}_t$ with $\tilde{p}_{t,j_t} > \frac{1}{2}$ (at most one such j_t exists), set $\mathcal{T} \leftarrow \{j_t\} \cup \mathcal{T}$.
 - 6 Construct bias $b_t \in \mathbb{R}^K$ such that $b_{t,i} = \frac{\beta}{W_{t,i}} \mathbb{1}\{t \in \mathcal{T}, i = j_t\}$.
 - 7 Compute $p_{t+1} = \operatorname{argmin}_{p \in \Delta_K} \{\langle p, \hat{\ell}_t + b_t \rangle + D_\psi(p, p_t)\}$.
-

$$\sum_{t=1}^T \sum_{i \in S_t} (W_{t,i} - \mathbf{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} \leq \mathcal{O} \left(\sqrt{\sum_{t=1}^T Q_t \iota_1} + \max_{t \in [T]} Q_t \iota_1 \right), \quad (7)$$

where $Q_t = \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma}$, $U_t = \max \left\{ 1, \frac{2 \max_{t \in [T]} Q_t}{\gamma} \right\}$, and $\iota_1 = \log \left(\frac{2 \max_{t \in [T]} Q_t + 2\sqrt{\sum_{t=1}^T Q_t}}{\delta} \right)$.

The full proof is deferred to [Appendix A.1](#). As $Q_t/\gamma \leq U_T = \Theta(\max_{t \in [T]} Q_t)/\gamma$ for all $t \in [T]$ and $Q_t \leq \tilde{\mathcal{O}}(\alpha_t)$, [Lemma 1](#) shows that STABILITY-TERM is bounded by $\tilde{\mathcal{O}}(\eta \sum_{t=1}^T Q_t + \eta U_T) = \tilde{\mathcal{O}}(\eta \sum_{t=1}^T \alpha_t + \eta \max_{t \in [T]} \alpha_t/\gamma)$, which only has logarithmic dependence on K , unlike the $\tilde{\mathcal{O}}(\eta K/\gamma)$ bound of [Eq. \(3\)](#)! For the fourth term in [Eq. \(5\)](#), [Lemma 1](#) shows that it is bounded by $\tilde{\mathcal{O}}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$, which again has no $\text{poly}(K)$ dependence. Combining [Lemma 1](#) with the rest of the analysis outlined earlier, we know that EXP3-IX in fact achieves $\tilde{\mathcal{O}}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ high-probability regret for self-aware graphs, formally stated in the following theorem. The full proof is deferred to [Appendix A.1](#).

Theorem 2 EXP3-IX with the optimal choice of $\eta > 0$ and $\gamma > 0$ guarantees that with probability at least $1 - \delta$, $\text{Reg} = \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \alpha_t \log \frac{1}{\delta}} + \max_{t \in [T]} \alpha_t \log \frac{1}{\delta} \right)$.

Generalization to Strongly Observable Graphs. Next, we show how to deal with general strongly observable graphs. The pseudocode of our proposed algorithm is shown in [Algorithm 1](#). Compared to EXP3-IX, there are three main differences. First, we have an additional η amount of uniform exploration over all actions ([Line 1](#)). Second, while keeping the same loss estimator construction for node $i \in S_t$ at each round t , for $i \in \bar{S}_t$ (nodes without self-loop), we construct a standard unbiased estimator ([Line 4](#)). Third, if there exists $j_t \in \bar{S}_t$ such that the probability of choosing action j_t is larger than $\frac{1}{2}$, then we further add positive bias $\beta/W_{t,j_t}$ to the loss estimator $\hat{\ell}_{t,j_t}$ (encoded via the b_t vector; see [Line 6](#) and [Line 7](#)), making it a pessimistic over-estimator. Intuitively, the reason of doing so is that we should avoid picking actions without self-loop too often even if past data suggest that it is a good action, since the only way to observe its loss and see if it stays good is by selecting

other actions. A carefully chosen positive bias injected to the loss estimator of such actions would exactly allow us to achieve this goal. In what follows, by outlining the analysis of our algorithm, we further explain why we make each of these three modifications from a technical perspective.

First, we note that with a nonempty \bar{S}_t , the key fact used earlier $\sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma} = \tilde{\mathcal{O}}(\alpha_t)$ is no longer true, making the fourth and the fifth term in Eq. (5) larger than desired if we still do implicit exploration for all nodes. Therefore, for nodes in \bar{S}_t , we go back to the original inverse importance weighted *unbiased* estimators (Line 4), and decompose the regret against any fixed action $j \in [K]$ differently into the following six terms:

$$\begin{aligned}
 \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) &\leq \underbrace{\left(\sum_{t=1}^T \ell_{t,i_t} - \sum_{t=1}^T \langle \tilde{p}_t, \ell_t \rangle \right)}_{\text{Term (1)}} + \underbrace{\left(\sum_{t=1}^T \langle \tilde{p}_t - p_t, \ell_t \rangle \right)}_{\text{Term (2)}} \\
 &\quad + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \ell_t \rangle_{\bar{S}_t} - \sum_{t=1}^T \mathbb{E}_t \left[\langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t} \right] \right)}_{\text{Term (3)}} \\
 &\quad + \underbrace{\left(\sum_{t=1}^T \mathbb{E}_t \left[\langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t} \right] - \sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t} \right)}_{\text{Term (4)}} \\
 &\quad + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \ell_t - \hat{\ell}_t \rangle_{S_t} \right)}_{\text{Term (5)}} + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle \right)}_{\text{Term (6)}}. \tag{8}
 \end{aligned}$$

Term (1) can be bounded again by $\tilde{\mathcal{O}}(\sqrt{T})$ via standard Azuma's inequality. Term (2) can be bounded by ηT due to the $\mathcal{O}(\eta)$ amount of uniform exploration. Term (3) is simply 0 as $\hat{\ell}_{t,i}$ is unbiased for $i \in \bar{S}_t$. Term (5) can similarly be written as $\sum_{t=1}^T (\hat{\ell}_{t,j} - \ell_{t,j}) \mathbb{1}\{j \in S_t\} + \sum_{t=1}^T \sum_{i \in S_t} (W_{t,i} - \mathbf{1}\{i \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} + \sum_{t=1}^T \sum_{i \in S_t} \frac{\gamma p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}$ since the loss estimator construction for $i \in S_t$ is the same as the one in EXP3-IX. Following how we handled the last three terms in Eq. (5), we have with high probability,

$$\text{Term (5)} \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \alpha_t} + \max_{t \in [T]} \alpha_t + \gamma \sum_{t=1}^T \alpha_t + \frac{1}{\gamma} \right). \tag{9}$$

The formal statement and the proof are deferred to Lemma 8 in Appendix A.2.

The key challenge lies in controlling Term (4) and Term (6). To this end, let us first consider the variance of the $\langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t}$. Let $\mathcal{T} = \{t : \exists j \in \bar{S}_t, \tilde{p}_{t,j_t} > \frac{1}{2}\}$ be the final value of \mathcal{T} in Algorithm 1. If $t \notin \mathcal{T}$, then $\hat{\ell}_{t,i} \leq \frac{1}{1 - \tilde{p}_{t,i}} \leq 2$ for all $i \in \bar{S}_t$ and the variance of $\langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t}$ is a constant; otherwise, we know that there is at most one node $j_t \in \bar{S}_t$ such that $p_{t,j_t} > \frac{1}{2}$. Direct calculation shows that the variance of $\langle p_t - e_j, \hat{\ell}_t \rangle_{\bar{S}_t}$ is bounded by $\tilde{\mathcal{O}}(\frac{1}{W_{t,j_t}} \cdot \mathbb{1}\{j \neq j_t\})$. With the help of the Freedman's inequality and the fact $W_{t,j_t} = \Omega(\eta)$ thanks to the uniform exploration

(Line 1), we prove in Lemma 7 of Appendix A.2 that Term (4) can be bounded as follows:

$$\text{Term (4)} \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \left(1 + \frac{\mathbb{1}\{j \neq j_t, t \in \mathcal{T}\}}{W_{t,j_t}} \right)} + \frac{1}{\eta} \right). \quad (10)$$

Handling this potentially large deviation is exactly the reason we inject a positive bias to the loss estimator (Line 6). Specifically, when $t \in \mathcal{T}$, we add $b_{t,j_t} = \frac{\beta}{W_{t,j_t}}$ to the loss estimator $\hat{\ell}_{t,j_t}$ for some parameter $\beta > 0$. With the help of this positive bias, we can decompose Term (6) as follows:

$$\text{Term (6)} = \sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t + b_t \rangle - \sum_{t=1}^T \langle p_t - e_j, b_t \rangle. \quad (11)$$

Direct calculation shows that the second negative term is of order $-\Theta(\sum_{t \in \mathcal{T}} \frac{\beta}{W_{t,j_t}} \cdot \mathbb{1}\{j \neq j_t\}) + \Theta(\sum_{t \in \mathcal{T}} \beta \cdot \mathbb{1}\{j = j_t\})$, large enough to cancel the large deviation in Eq. (10). Specifically, using AM-GM inequality, we obtain

$$\text{Term (4)} - \sum_{t=1}^T \langle p_t - e_{i^*}, b_t \rangle \leq \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \sqrt{T} + \frac{1}{\beta} + \beta T \right). \quad (12)$$

The final step is to handle the first term in Eq. (11). Similar to Eq. (1), standard analysis of online mirror descent shows that

$$\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t + b_t \rangle \leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \left(\hat{\ell}_{t,i} + b_{t,i} \right)^2, \quad (13)$$

However, unlike the case for self-aware graphs, when there exist nodes without a self-loop, the second term may be prohibitively large when $t \in \mathcal{T}$. Inspired by (Alon et al., 2015b), we address this issue with a *loss shifting trick*. Specifically, the following refined version of Eq. (13) holds:

$$\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t + b_t \rangle \leq \frac{\log K}{\eta} + 2\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \left(\hat{\ell}_{t,i} + b_{t,i} - z_t \right)^2, \quad (14)$$

for any $z_t \leq \frac{3}{\eta}$, $t \in [T]$. We choose $z_t = 0$ when $t \notin \mathcal{T}$ and $z_t = \hat{\ell}_{t,j_t} + b_{t,j_t}$ when $t \in \mathcal{T}$, which satisfies the condition $z_t \leq \frac{3}{\eta}$ again thanks to the $\mathcal{O}(\eta)$ amount of uniform exploration over all nodes (Line 1). With such a loss shift z_t , continuing with Eq. (14) it can be shown that:

$$\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t + b_t \rangle \leq \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \eta \sum_{t \notin \mathcal{T}} \sum_{i \in [K]} p_{t,i} \hat{\ell}_{t,i}^2 + \eta \sum_{t \in \mathcal{T}} \frac{\mathbb{1}\{i_t \neq j_t\}}{W_{t,j_t}} + \beta^2 T \right). \quad (15)$$

Note that for $t \in \mathcal{T}$, $\hat{\ell}_{t,i} \leq 2$ for all $i \in \bar{S}_t$. Therefore, the second term in Eq. (15) can be bounded by $\tilde{\mathcal{O}}(\eta T + \eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \hat{\ell}_{t,i}) \leq \tilde{\mathcal{O}}(\eta T + \eta \sum_{t=1}^T \alpha_t + \eta/\gamma \cdot \max_{t \in [T]} \alpha_t)$ where the inequality is by using Lemma 1. The third term can be bounded by $\tilde{\mathcal{O}}(\sqrt{\eta T})$ with high probability by using Freedman's inequality. Together with Eq. (12), Eq. (9), the bounds for Term (1), Term (2), and Term (3), and the optimal choice of the parameters η, β , and γ , we arrive at the following main theorem for general strongly observable graphs (see Appendix A.2 for the full proof).

Algorithm 2 Algorithm for Weakly Observable Graphs

Input: Parameters $\eta, \gamma, \varepsilon$.

Define: Regularizer $\psi(p) = \frac{1}{\eta} \sum_{i=1}^K p_i \log p_i$.

Initialize: $p_{1,i} = \frac{1}{K}$ for all $i \in [K]$.

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

- 1 Receive the feedback graph G_t and (approximately) find a smallest weakly dominating set D_t .
 - 2 Let $\tilde{p}_t = (1 - \varepsilon|D_t|)p_t + \varepsilon \cdot \mathbf{1}_{D_t}$ ($\mathbf{1}_{D_t}$ is a vector with 1 for coordinates in D_t and 0 otherwise).
 - 3 Sample action $i_t \sim \tilde{p}_t$.
 - 4 Receive feedback $\ell_{t,j}$ for all j such that $i_t \in N_t^{\text{in}}(j)$.
 - 5 Construct estimator $\hat{\ell}_t \in \mathbb{R}^K$ such that $\hat{\ell}_{t,i} = \frac{\ell_{t,i} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}}{W_{t,i} + \gamma \mathbb{1}\{i \in S_t\}}$ where $W_{t,i} = \sum_{j \in N_t^{\text{in}}(i)} \tilde{p}_{t,j}$.
 - 6 Compute $p_{t+1} = \operatorname{argmin}_{p \in \Delta_K} \{ \langle p, \hat{\ell}_t \rangle + D_\psi(p, p_t) \}$.
-

Theorem 3 *Algorithm 1* with parameter $\gamma = \beta = \eta = \min \left\{ \frac{1}{\sqrt{\sum_{t=1}^T \alpha_t \log(1/\delta)}}, \frac{1}{2} \right\}$ guarantees that with probability at least $1 - 6\delta$, the regret is bounded as follows:

$$\text{Reg} \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \alpha_t \log \frac{1}{\delta}} + \max_{t \in [T]} \alpha_t \log \frac{1}{\delta} \right).$$

To the best of our knowledge, this is the first optimal high-probability regret bound for general strongly observable graphs, importantly without any $\text{poly}(K)$ dependence. While in this theorem we assume the knowledge of α_t for all $t \in [T]$ to tune the parameters η, β and γ , a standard doubling trick can be applied to remove this requirement and make [Algorithm 1](#) completely parameter-free.⁴

4. High-Probability Regret for Weakly Observable Graphs

In this section, we study the setting where the feedback graph G_t is weakly observable for all $t \in [T]$. Under the uninformed setting, ([Alon et al., 2015a](#), Theorem 9) proves that the lower bound of expected regret is $\Omega(K^{1/3}T^{2/3})$. To get rid of the $\text{poly}(K)$ dependence, we thus consider the informed setting, in which G_t is revealed to the learner before she selects i_t . We propose a simple algorithm to achieve $\tilde{\mathcal{O}}(T^{1/3}(\sum_{t=1}^T d_t)^{1/3} + \frac{1}{T} \sum_{t=1}^T d_t)$ high-probability regret bound.

Our algorithm is summarized in [Algorithm 2](#), which is a combination of EXP3.G ([Alon et al., 2015b](#)) and EXP3-IX. Following EXP3.G, we add uniform exploration over a smallest weakly dominating set ([Line 2](#)).⁵ In this way, each weakly observable node has at least ε probability to be observed, which is essential to control the variance of the loss estimators. Similar to [Algorithm 1](#), we apply implicit exploration for nodes with self-loops when constructing their loss estimators ([Line 5](#)). Different from [Algorithm 1](#), we do not need to inject bias any more. This is because with the combination of uniform exploration and implicit exploration, our algorithm already achieves $\tilde{\mathcal{O}}(T^{2/3})$ bound which is optimal for weakly observable graphs. Our main result is summarized below (see [Appendix B](#) for the proof).

4. This can be achieved efficiently by applying a standard doubling trick on the quantity $B_t = \frac{1}{\sqrt{\sum_{\tau=1}^t Q_\tau}}$, $t \in [T]$.

5. While finding it exactly is computational hard, it suffices to find an approximate one with size $\mathcal{O}(d_t \log K)$, which can be done in polynomial time.

Theorem 4 *Algorithm 2 with parameter $\varepsilon = \min\{\frac{1}{2}, T^{1/3}(\sum_{t=1}^T d_t)^{-2/3} \log(1/\delta)^{1/3}\}$, $\gamma = \sqrt{\frac{\log(1/\delta)}{\sum_{t=1}^T \tilde{\alpha}_t}}$, $\eta = \min\{T^{-1/3}(\sum_{t=1}^T d_t)^{-1/3} \log(1/\delta)^{-1/3}, \gamma\}$ ensures with probability at least $1 - 4\delta$:*

$$\text{Reg} \leq \tilde{\mathcal{O}} \left(T^{\frac{1}{3}} \left(\sum_{t=1}^T d_t \right)^{\frac{1}{3}} \log^{\frac{1}{3}}(1/\delta) + \frac{1}{T} \sum_{t=1}^T d_t \log \frac{1}{\delta} + \sqrt{\sum_{t=1}^T \tilde{\alpha}_t \log \frac{1}{\delta} + \max_{t \in [T]} \tilde{\alpha}_t \log \frac{1}{\delta}} \right),$$

where $\tilde{\alpha}_t$ is the independence number of the subgraph induced by nodes with self-loops in G_t .

When $G_t = G$ for all t , our bound becomes $\tilde{\mathcal{O}}(d^{1/3}T^{2/3} + \sqrt{\alpha T} + \alpha + d)$, where α is the independence number of the subgraph of G induced by its nodes with self-loops and d is the weak domination number of G . This even improves over the $\tilde{\mathcal{O}}(d^{1/3}T^{2/3} + \sqrt{KT})$ expected regret bound of (Alon et al., 2015b), removing any poly(K) dependence.

To prove Theorem 4, similar to the analysis in Section 3, we first decompose the regret against any fixed action $j \in [K]$ as follows:

$$\underbrace{\sum_{t=1}^T (\ell_{t,i_t} - \langle \tilde{p}_t, \ell_t \rangle)}_{\text{Term (a)}} + \underbrace{\sum_{t=1}^T (\tilde{p}_t - p_t, \ell_t)}_{\text{Term (b)}} + \underbrace{\sum_{t=1}^T \langle p_t - e_j, \ell_t - \hat{\ell}_t \rangle}_{\text{Term (c)}} + \underbrace{\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle}_{\text{Term (d)}}. \quad (16)$$

Term (a) is of order $\tilde{\mathcal{O}}(\sqrt{T})$ via Azuma's inequality. By the definition of \tilde{p}_t , Term (b) is of order $\mathcal{O}(\varepsilon \sum_{t=1}^T d_t)$. To bound Term (c), we state the following lemma, which controls the deviation between real losses and loss estimators. The proof starts by considering nodes in S_t and \bar{S}_t separately, followed by standard concentration inequalities; see Appendix B for details.

Lemma 5 *Algorithm 2 guarantees the following with probability at least $1 - \delta$*

$$\sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \tilde{\alpha}_t \log \frac{1}{\delta}} + \gamma \sum_{t=1}^T \tilde{\alpha}_t + \max_{t \in [T]} \tilde{\alpha}_t \log \frac{1}{\delta} + \sqrt{\frac{T}{\varepsilon} \log \frac{1}{\delta}} \right).$$

Furthermore, with probability at least $1 - \delta$, for any $i \in [K]$, the following inequality holds:

$$\sum_{t=1}^T (\hat{\ell}_{t,i} - \ell_{t,i}) \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{T}{\varepsilon} \log \frac{1}{\delta}} + \frac{1}{\varepsilon} \log \frac{1}{\delta} + \frac{1}{\gamma} \log \frac{1}{\delta} \right).$$

Next, we prove the following lemma bounding Term (d) (see Appendix B again for the full proof).

Lemma 6 *Algorithm 2 guarantees that with probability at least $1 - \delta$*

$$\text{Term (d)} \leq \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \eta \sum_{t=1}^T \tilde{\alpha}_t + \frac{\eta}{\gamma} \max_t \tilde{\alpha}_t \log \frac{1}{\delta} + \eta \sqrt{\frac{T}{\varepsilon^3} \log \frac{1}{\delta}} + \frac{\eta T}{\varepsilon} + \frac{\eta}{\varepsilon^2} \log \frac{1}{\delta} \right).$$

Proof sketch. First, we apply standard OMD analysis (Bubeck et al., 2012) and obtain

$$\text{Term (d)} \leq \frac{\log K}{\eta} + 2\eta \sum_{t=1}^T \sum_{i \in S_t} p_{t,i} \widehat{\ell}_{t,i}^2 + 2\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2.$$

We can bound the second term by $\tilde{O}\left(\eta \sum_{t=1}^T \tilde{\alpha}_t + \frac{\eta}{\gamma} \max_t \tilde{\alpha}_t \log(1/\delta)\right)$ using Eq. (6) in Lemma 1.

For the third term, based on the definition of $\widehat{\ell}_{t,i}$ for $i \in \bar{S}_t$, we decompose it as follows

$$\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2 \leq \eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} (\widehat{\ell}_{t,i} - \ell_{t,i}) + \eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} \ell_{t,i}.$$

We bound the first term by $\tilde{O}(\eta\sqrt{T/\varepsilon^3} + \eta/\varepsilon^2)$ using Freedman's inequality. With the help of uniform exploration, we know that $W_{t,i} \geq \varepsilon$ and thus the second term is bounded by $\frac{\eta T}{\varepsilon}$. ■

With the help of Lemma 5 and Lemma 6, we are ready to prove Theorem 4.

Proof [Proof of Theorem 4] Putting results from Eq. (16), Lemma 5, and Lemma 6 together, our regret bound becomes

$$\begin{aligned} \text{Reg} \leq & \tilde{O}\left(\varepsilon \sum_{t=1}^T d_t + \sqrt{\sum_{t=1}^T \tilde{\alpha}_t \log \frac{1}{\delta}} + \gamma \sum_{t=1}^T \tilde{\alpha}_t + \max_t \tilde{\alpha}_t \log \frac{1}{\delta} + \frac{1}{\gamma} \log \frac{1}{\delta} + \sqrt{\frac{T}{\varepsilon} \log \frac{1}{\delta}}\right) \\ & + \tilde{O}\left(\frac{1}{\eta} + \eta \sum_{t=1}^T \tilde{\alpha}_t + \frac{\eta}{\gamma} \max_t \tilde{\alpha}_t \log \frac{1}{\delta} + \eta \sqrt{\frac{T}{\varepsilon^3} \log \frac{1}{\delta}} + \frac{\eta T}{\varepsilon} + \frac{\eta}{\varepsilon^2} \log \frac{1}{\delta}\right). \end{aligned}$$

By picking ε , η , and γ as stated in Theorem 4, we achieve that with probability at least $1 - \delta$,

$$\text{Reg} \leq \tilde{O}\left(T^{\frac{1}{3}} \left(\sum_{t=1}^T d_t\right)^{\frac{1}{3}} \log^{\frac{1}{3}} \frac{1}{\delta} + \frac{1}{T} \sum_{t=1}^T d_t \log \frac{1}{\delta} + \sqrt{\sum_{t=1}^T \tilde{\alpha}_t \log \frac{1}{\delta}} + \max_t \tilde{\alpha}_t \log \frac{1}{\delta}\right). \quad \blacksquare$$

Again, we can apply the standard doubling trick to tune η , γ , and ε adaptively without requiring the knowledge of d_t and $\tilde{\alpha}_t$ for $t \in [T]$ ahead of time.

5. Conclusions and Open Problems

In this work, we design algorithms that achieve near-optimal high-probability regret bounds for adversarial MAB with time-varying feedback graphs for both the strongly observable case and the weakly observable case. We achieve $\tilde{O}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ regret for strongly observable graphs, improving and extending the results of (Neu, 2015), which only considers self-aware graphs and suffers an $\tilde{O}(K)$ term. In addition, we derive the first high-probability regret bound for weakly observable graph setting, which also depends on K only logarithmically and is order optimal.

One open problem is whether one can achieve high-probability data-dependent regret bounds for this problem, such as the so-called small-loss bounds which scales with the loss of the best action. Lee et al. (2020a) achieved expected regret bound $\tilde{O}(\sqrt{\kappa L_\star})$ for a fixed graph where κ is the clique partition number and L_\star is the loss of the best action. Achieving the same bound with high-probability under an adaptive adversary appears to require new ideas.

References

- Jacob Abernethy and Alexander Rakhlin. Beating the adaptive bandit with high probability. In *2009 Information Theory and Applications Workshop*, pages 280–289. IEEE, 2009.
- Noga Alon, Nicolò Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. *arXiv preprint arXiv:1502.07617*, 2015a.
- Noga Alon, Nicolo Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *Conference on Learning Theory*, pages 23–35. PMLR, 2015b.
- Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 46(6):1785–1826, 2017.
- 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.
- Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Conference on Learning Theory*, pages 335–342. Omnipress, 2008.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26. JMLR Workshop and Conference Proceedings, 2011.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Swapna Buccapatnam, Fang Liu, Atilla Eryilmaz, and Ness B Shroff. Reward maximization under uncertainty: Leveraging side-observations on networks. *Journal of Machine Learning Research*, 18:1–34, 2018.
- Stéphane Caron, Branislav Kveton, Marc Lelarge, and Smriti Bhagat. Leveraging side observations in stochastic bandits. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 142–151, 2012.
- Houshuang Chen, Shuai Li, Chihao Zhang, et al. Understanding bandits with graph feedback. *Advances in Neural Information Processing Systems*, 34:24659–24669, 2021.
- Alon Cohen, Tamir Hazan, and Tomer Koren. Online learning with feedback graphs without the graphs. In *International Conference on Machine Learning*, pages 811–819. PMLR, 2016.
- Liad Erez and Tomer Koren. Towards best-of-all-worlds online learning with feedback graphs. *Advances in Neural Information Processing Systems*, 34:28511–28521, 2021.
- 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*, 55(1):119–139, 1997.

- Shinji Ito, Taira Tsuchiya, and Junya Honda. Nearly optimal best-of-both-worlds algorithms for online learning with feedback graphs. *Advances in Neural Information Processing Systems*, 2022.
- Tomás Kocák, Gergely Neu, Michal Valko, and Rémi Munos. Efficient learning by implicit exploration in bandit problems with side observations. *Advances in Neural Information Processing Systems*, 27, 2014.
- John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. *Advances in Neural Information Processing Systems*, 20(1):96–1, 2007.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang. Bias no more: high-probability data-dependent regret bounds for adversarial bandits and mdps. *Advances in Neural Information Processing Systems*, 33:15522–15533, 2020a.
- Chung-Wei Lee, Haipeng Luo, and Mengxiao Zhang. A closer look at small-loss bounds for bandits with graph feedback. In *Conference on Learning Theory*, pages 2516–2564. PMLR, 2020b.
- Thodoris Lykouris, Karthik Sridharan, and Éva Tardos. Small-loss bounds for online learning with partial information. In *Conference on Learning Theory*, pages 979–986. PMLR, 2018.
- Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. *Advances in Neural Information Processing Systems*, 24, 2011.
- Teodor V Marinov, Mehryar Mohri, and Julian Zimmert. Stochastic online learning with feedback graphs: Finite-time and asymptotic optimality. *Advances in Neural Information Processing Systems*, 2022.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Advances in Neural Information Processing Systems*, 28, 2015.
- Julian Zimmert and Tor Lattimore. Return of the bias: Almost minimax optimal high probability bounds for adversarial linear bandits. In *Conference on Learning Theory*, pages 3285–3312. PMLR, 2022.

Appendix A. Omitted Details in [Section 3](#)

A.1. Proof of [Theorem 2](#)

In this section, we prove [Theorem 2](#), which shows that the regret of the EXP3-IX algorithm ([Neu, 2015](#)) does not necessarily has linear dependence on the number of actions K (that appears in the original analysis), but is instead $\tilde{O}((\sum_{t=1}^T \alpha_t)^{1/2} + \max_{t \in [T]} \alpha_t)$ with high probability.

First, we prove [Lemma 1](#), which shows a tighter concentration between $\hat{\ell}_t$ and ℓ_t and is crucial to the improvement from $\tilde{O}(K)$ to $\tilde{O}(\max_{t \in [T]} \alpha_t)$.

Proof [Proof of [Lemma 1](#)] We first prove [Eq. \(6\)](#). Let $X_{t,1} = \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} (\widehat{\ell}_{t,i} - \ell_{t,i})$ and $Q_t = \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma}$. According to the definition of $\widehat{\ell}_{t,i}$ and the fact that $\ell_t \in [0, 1]^K$, we know that

$$|X_{t,1}| \leq \sum_{i \in S_t} \frac{p_{t,i}}{(W_{t,i} + \gamma)} \left(\frac{1}{\gamma} + 1 \right) \leq \frac{2Q_t}{\gamma},$$

where we use the fact that $\gamma \leq \frac{1}{2}$. Next, consider the term $\mathbb{E}_t[X_{t,1}^2]$.

$$\begin{aligned} \mathbb{E}_t[X_{t,1}^2] &\leq \mathbb{E}_t \left[\left(\sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} \right)^2 \right] \\ &\leq \mathbb{E}_t \left[\left(\sum_{i \in S_t} \frac{p_{t,i}}{(W_{t,i} + \gamma)^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\} \right) \left(\sum_{j \in S_t} \frac{p_{t,j}}{(W_{t,j} + \gamma)^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(j)\} \right) \right] \\ &\leq \mathbb{E}_t \left[\left(\sum_{i \in S_t} \frac{p_{t,i}}{(W_{t,i} + \gamma)^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\} \right) \left(\sum_{j \in S_t} \frac{p_{t,j}}{(W_{t,j} + \gamma)^2} \right) \right] \\ &\leq \frac{Q_t}{\gamma} \mathbb{E}_t \left[\sum_{i \in S_t} \frac{p_{t,i}}{(W_{t,i} + \gamma)^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\} \right] \leq \frac{Q_t^2}{\gamma}. \end{aligned}$$

Note that $Q_t \leq K$. Therefore, $X_{t,1} \leq \frac{2K}{\gamma}$ and $\mathbb{E}_t[X_{t,1}^2] \leq \frac{K^2}{\gamma}$. Then, using [Lemma 13](#), we know that with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{t=1}^T X_{t,1} &\leq 3 \sqrt{\sum_{t=1}^T \frac{Q_t^2}{\gamma} \log \left(\frac{2K}{\delta} \sqrt{\frac{T}{\gamma}} \right)} + 2 \max\{1, \max_{t \in [T]} X_{t,1}\} \log \left(\frac{2K}{\delta} \sqrt{\frac{T}{\gamma}} \right) \\ &\leq \mathcal{O} \left(\sum_{t=1}^T \frac{Q_t^2}{\gamma U_T} + U_T \log \left(\frac{KT}{\delta \gamma} \right) \right), \end{aligned}$$

where $U_T = \max\{1, \max_{t \in [T]} X_{t,1}\}$ and the last inequality is because of AM-GM inequality.

Next, we prove [Eq. \(7\)](#). Let $X_{t,2} = \sum_{i \in S_t} (W_{t,i} - \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}$. Direct calculation shows that $|X_{t,2}| \leq 2Q_t$. Consider its conditional variance:

$$\begin{aligned} \mathbb{E}_t[X_{t,2}^2] &\leq \mathbb{E}_t \left[\left(\sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\} \right)^2 \right] \\ &= \sum_{i \in S_t} p_{t,i} \sum_{j \in S_t} \frac{p_{t,j}}{W_{t,j} + \gamma} \leq Q_t. \end{aligned}$$

Define $\ell_1 = \log \left(\frac{2 \max_t Q_t + 2 \sqrt{\sum_{t=1}^T Q_t}}{\delta} \right)$. Applying [Lemma 13](#), we can obtain that with probability at least $1 - \delta$,

$$\sum_{t=1}^T X_{t,2} \leq \mathcal{O} \left(\sqrt{\sum_{t=1}^T Q_t \ell_1} + \max_{t \in [T]} Q_t \ell_1 \right).$$

■

Next, we are ready to prove [Theorem 2](#).

Proof [Proof of [Theorem 2](#)] According to [Eq. \(5\)](#), for an arbitrary comparator $j \in [K]$, we decompose the overall regret as follows:

$$\begin{aligned} \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) &= \sum_{t=1}^T \langle p_t - e_j, \widehat{\ell}_t \rangle + \sum_{t=1}^T (\ell_{t,i_t} - \langle p_t, \ell_t \rangle) + \sum_{t=1}^T (\widehat{\ell}_{t,j} - \ell_{t,j}) \\ &\quad + \sum_{t=1}^T \sum_{i=1}^K (W_{t,i} - \mathbf{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} + \sum_{t=1}^T \sum_{i=1}^K \frac{\gamma p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}. \end{aligned} \quad (17)$$

According to the standard analysis of EXP3 ([Bubeck et al., 2012](#)), the first term of [Eq. \(17\)](#) can be bounded as follows:

$$\begin{aligned} &\sum_{t=1}^T \langle p_t - e_j, \widehat{\ell}_t \rangle \\ &\leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \widehat{\ell}_{t,i}^2 \\ &\leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} \\ &\leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K \frac{p_{t,i}}{W_{t,i} + \gamma} \ell_{t,i} + \mathcal{O} \left(\eta \sum_{t=1}^T \frac{Q_t^2}{\max_{\tau \in [T]} Q_\tau} + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \left(\frac{KT}{\delta \gamma} \right) \right), \end{aligned}$$

where the last inequality holds with probability at least $1 - \delta$ according to [Lemma 1](#).

According to standard Hoeffding-Azuma inequality, we know that with probability at least $1 - \delta$, the second term of [Eq. \(17\)](#) is bounded as follows:

$$\sum_{t=1}^T (\ell_{t,i_t} - \langle p_t, \ell_t \rangle) \leq \mathcal{O} \left(\sqrt{T \log \frac{1}{\delta}} \right).$$

Based on Corollary 1 in ([Neu, 2015](#)), with probability at least $1 - \delta$, the third term is bounded as follows: for all $j \in [K]$,

$$\sum_{t=1}^T (\widehat{\ell}_{t,j} - \ell_{t,j}) \leq \frac{\log(K/\delta)}{2\gamma}.$$

The fourth term of [Eq. \(17\)](#) can be bounded by using [Lemma 1](#) and the final term of [Eq. \(17\)](#) is bounded by $\mathcal{O}(\gamma \sum_{t=1}^T Q_t)$ (recall that $Q_t = \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma}$). Combining all the terms, we know that with probability at least $1 - 3\delta$,

$$\sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) \leq \widetilde{\mathcal{O}} \left(\frac{1}{\eta} + \frac{\log(1/\delta)}{\gamma} + (\eta + \gamma) \sum_{t=1}^T Q_t + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \frac{1}{\delta} \right)$$

$$+ \tilde{\mathcal{O}} \left(\sqrt{T \log \frac{1}{\delta}} + \sqrt{\sum_{t=1}^T \log \frac{1}{\delta}} + \max_{t \in [T]} Q_t \log \frac{1}{\delta} \right).$$

According to [Lemma 11](#), we know that $Q_t = \tilde{\mathcal{O}}(\alpha_t)$. Finally, choosing $\eta = \gamma = \sqrt{\frac{\log(1/\delta)}{\sum_{t=1}^T \alpha_t}}$ and picking $\delta' = \frac{\delta}{3}$ finishes the proof. \blacksquare

A.2. Proof of [Theorem 3](#)

In this section, we prove our main result [Theorem 3](#) in the strongly observable setting. To prove [Theorem 3](#), according to [Eq. \(8\)](#), we can decompose the overall regret with respect to any $j \in [K]$ as follows

$$\begin{aligned} \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) &\leq \underbrace{\left(\sum_{t=1}^T \ell_{t,i_t} - \sum_{t=1}^T \langle \tilde{p}_t, \ell_t \rangle \right)}_{\text{Term (1)}} + \underbrace{\left(\sum_{t=1}^T \langle \tilde{p}_t - p_t, \ell_t \rangle \right)}_{\text{Term (2)}} \\ &\quad + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \ell_t \rangle_{\tilde{S}_t} - \sum_{t=1}^T \mathbb{E}_t \left[\langle p_t - e_j, \hat{\ell}_t \rangle_{\tilde{S}_t} \right] \right)}_{\text{Term (3)}} \\ &\quad + \underbrace{\left(\sum_{t=1}^T \mathbb{E}_t \left[\langle p_t - e_j, \hat{\ell}_t \rangle_{\tilde{S}_t} \right] - \sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle_{\tilde{S}_t} \right)}_{\text{Term (4)}} \\ &\quad + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \ell_t - \hat{\ell}_t \rangle_{S_t} \right)}_{\text{Term (5)}} + \underbrace{\left(\sum_{t=1}^T \langle p_t - e_j, \hat{\ell}_t \rangle \right)}_{\text{Term (6)}}. \end{aligned} \quad (18)$$

With the help of Hoeffding-Azuma's inequality, we know that with probability at least $1 - \delta$, $\text{Term (1)} \leq \mathcal{O}(\sqrt{T \log(1/\delta)})$. $\text{Term (2)} \leq \mathcal{O}(\eta T)$ because of the definition of \tilde{p}_t and p_t . $\text{Term (3)} = 0$ as $\hat{\ell}_{t,i}$ is an unbiased estimator of $\ell_{t,i}$ for $i \in \tilde{S}_t$. In the next three sections, we bound Term (4) , Term (5) and Term (6) respectively.

Bounding Term (4) . Using Freedman's inequality, we prove the following lemma:

Lemma 7 *With probability at least $1 - \delta$,*

$$\text{Term (4)} \leq \left(2 + \frac{4}{\eta} \right) \log \frac{1}{\delta} + 2 \sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}}.$$

Proof Let $Y_t = \mathbb{E}_t \left[\left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t} \right] - \left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t}$. Then,

$$|Y_t| \leq \mathbb{E}_t \left[\left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t} \right] + 2 \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{\frac{K-1}{K}\eta} \leq 2 + \frac{4}{\eta}.$$

If $t \notin \mathcal{T}$, then we know that $W_{t,i} \geq 1/2$ for all $i \in \bar{S}_t$ and

$$\mathbb{E}_t[Y_t^2] \leq \mathbb{E}_t \left[\left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t}^2 \right] \leq 4.$$

If $t \in \mathcal{T}$, then we know that $W_{t,i} \geq 1/2$ for all $i \in \bar{S}_t$ except for $i = j_t$. When $j \neq j_t$, we can bound $\mathbb{E}_t[Y_t^2]$ as follows:

$$\begin{aligned} \mathbb{E}_t[Y_t^2] &\leq \mathbb{E}_t \left[\left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t}^2 \right] \\ &\leq 2\mathbb{E}_t \left[\left\langle p_t, \widehat{\ell}_t \right\rangle_{\bar{S}_t}^2 \right] + 2\mathbb{E}_t \left[\left\langle e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t}^2 \right] \\ &\leq 2\mathbb{E}_t \left[\sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2 \right] + \frac{2}{W_{t,j_t}} \leq 2\mathbb{E}_t \left[\sum_{i \in \bar{S}_t} \frac{p_{t,i} \widehat{\ell}_{t,i}}{W_{t,i}} \right] + \frac{2}{W_{t,j_t}} \leq 4 + \frac{4}{W_{t,j_t}}. \end{aligned}$$

If $j = j_t$, we know that $\left\langle p_t - e_j, \widehat{\ell}_t \right\rangle_{\bar{S}_t} = \sum_{i \in \bar{S}_t, i \neq j_t} p_{t,i} \widehat{\ell}_{t,i} + (1 - p_{t,j_t}) \cdot \frac{1}{W_{t,j_t}} \leq 2 + \frac{1}{1-\eta} \leq 4$ as $\eta \leq \frac{1}{2}$. Then we know that $\mathbb{E}_t[Y_t^2] \leq 16$.

Therefore, according to Freedman's inequality [Lemma 12](#), we know that with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{t=1}^T Y_t &\leq \min_{\lambda \in [0, 1/(2+4/\eta)]} \left\{ \frac{\log(1/\delta)}{\lambda} + \lambda \sum_{t=1}^T \mathbb{E}_t[Y_t^2] \right\} \\ &\leq \min_{\lambda \in [0, 1/(2+4/\eta)]} \left\{ \frac{\log(1/\delta)}{\lambda} + \lambda \left(16T + 4 \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \right\} \\ &\leq \left(2 + \frac{4}{\eta} \right) \log \frac{1}{\delta} + 2 \sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}}. \end{aligned}$$

■

Bounding Term (5). The following lemma gives a bound on Term (5). The proving technique is similar to the one that we use to bound the last three terms in [Eq. \(17\)](#).

Lemma 8 *With probability at least $1 - 2\delta$,*

$$\text{Term (5)} \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \max_{t \in [T]} Q_t \log \frac{1}{\delta} + \gamma \sum_{t=1}^T Q_t + \frac{1}{\gamma} \log \frac{1}{\delta} \right).$$

Proof We bound $\sum_{t=1}^T \langle p_t, \ell_t - \widehat{\ell}_t \rangle_{S_t}$ and $\sum_{t=1}^T \langle e_j, \ell_t - \widehat{\ell}_t \rangle_{S_t}$ separately. Note that $\widehat{\ell}_{t,i}$ is an under-biased estimator of $\ell_{t,i}$ for $i \in S_t$. Direct calculation shows that

$$\begin{aligned} \sum_{t=1}^T \langle p_t, \ell_t - \widehat{\ell}_t \rangle_{S_t} &= \sum_{t=1}^T \sum_{i \in S_t} p_{t,i} (\ell_{t,i} - \widehat{\ell}_{t,i}) \\ &= \sum_{t=1}^T \sum_{i \in S_t} (W_{t,i} - \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} + \sum_{t=1}^T \sum_{i \in S_t} \gamma \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}. \end{aligned} \quad (19)$$

Therefore, according to [Lemma 1](#), with probability at least $1 - \delta$,

$$\sum_{t=1}^T \langle p_t, \ell_t - \widehat{\ell}_t \rangle_{S_t} \leq \widetilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \max_{t \in [T]} Q_t \log \frac{1}{\delta} + \gamma \sum_{t=1}^T Q_t \right).$$

Next, consider the term $-\sum_{t=1}^T \langle e_j, \ell_t - \widehat{\ell}_t \rangle_{S_t} = \sum_{t=1}^T (\widehat{\ell}_{t,j} - \ell_{t,j}) \cdot \mathbb{1}\{j \in S_t\}$. Similar to the proof of [Corollary 1](#) in [\(Neu, 2015\)](#), define $\bar{\ell}_{t,i} = \frac{\ell_{t,i}}{W_{t,i}} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}$ and we know that for any $i \in [K]$,

$$\begin{aligned} \widehat{\ell}_{t,i} \mathbb{1}\{i \in S_t\} &\leq \frac{\ell_{t,i}}{W_{t,i} + \gamma} \mathbb{1}\{i_t \in N_t^{\text{in}}(i), i \in S_t\} \leq \frac{\ell_{t,i}}{W_{t,i} + \gamma \ell_{t,i}} \mathbb{1}\{i_t \in N_t^{\text{in}}(i), i \in S_t\} \\ &\leq \frac{1}{2\gamma} \frac{2\gamma \ell_{t,i} / W_{t,i}}{1 + \gamma \ell_{t,i} / W_{t,i}} \mathbb{1}\{i_t \in N_t^{\text{in}}(i), i \in S_t\} \\ &\leq \frac{1}{2\gamma} \log(1 + 2\gamma \bar{\ell}_{t,i}) \mathbb{1}\{i \in S_t\} \\ &= \frac{1}{2\gamma} \log(1 + 2\gamma \bar{\ell}_{t,i} \mathbb{1}\{i \in S_t\}). \end{aligned}$$

Therefore, we know that

$$\mathbb{E}_t \left[\exp \left(2\gamma \bar{\ell}_{t,i} \mathbb{1}\{i \in S_t\} \right) \right] \leq \mathbb{E}_t \left[1 + 2\gamma \bar{\ell}_{t,i} \mathbb{1}\{i \in S_t\} \right] = 1 + 2\gamma \ell_{t,i} \mathbb{1}\{i \in S_t\} \leq \exp(2\gamma \ell_{t,i} \mathbb{1}\{i \in S_t\}).$$

Define $Z_t = \exp \left(2\gamma \mathbb{1}\{i \in S_t\} (\widehat{\ell}_{t,i} - \ell_{t,i}) \right)$ and according to previous analysis, we know that Z_t is a super-martingale and by Markov inequality, we obtain that

$$\Pr \left[\sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in S_t\} > \varepsilon \right] = \Pr \left[\exp \left(2\gamma \sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in S_t\} \right) > \exp(2\gamma \varepsilon) \right] \leq \exp(-2\gamma \varepsilon).$$

Taking a union bound over $i \in [K]$, we know that with probability at least $1 - \delta$, for all $i \in [K]$,

$$\sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in S_t\} \leq \frac{\log(K/\delta)}{2\gamma}. \quad (20)$$

Combining both parts gives the bound for Term (5): with probability at least $1 - 2\delta$,

$$\text{Term (5)} \leq \widetilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \max_{t \in [T]} Q_t \log \frac{1}{\delta} + \gamma \sum_{t=1}^T Q_t + \frac{1}{\gamma} \log \frac{1}{\delta} \right). \quad (21)$$

■

Bounding Term (6). For completeness, before bounding Term (6), we show the following OMD analysis lemma.

Lemma 9 *Suppose that $p' = \operatorname{argmin}_{p \in \Delta_K} \{\langle p, \ell \rangle + D_\psi(p, p_t)\}$ with $\psi(p) = \frac{1}{\eta} \sum_{i=1}^K p_i \log p_i$. If $\eta \ell_i \geq -3$ for all $i \in [K]$, then for any $u \in \Delta_K$, the following inequality hold:*

$$\langle p - u, \ell \rangle \leq D_\psi(u, p) - D_\psi(u, p') + 2\eta \sum_{i=1}^K p_i \ell_i^2.$$

Proof Let $q_i = p_i \exp(-\eta \ell_i)$ and direct calculation shows that $p' = \operatorname{argmin}_{p \in \Delta_K} D_\psi(p, q)$ and for all $u \in \Delta_K$,

$$\begin{aligned} \langle p - u, \ell \rangle &= D_\psi(u, p) - D_\psi(u, q) + D_\psi(p, q) \\ &\leq D_\psi(u, p) - D_\psi(u, p') + D_\psi(p, q), \end{aligned}$$

where the second step uses the generalized Pythagorean theorem. On the other hand, using the inequality $\exp(-x) \leq 1 - x + 2x^2$ for any $x \geq -3$, we know that

$$\begin{aligned} D_\psi(p, q) &= \frac{1}{\eta} \sum_{i=1}^K \left(p_i \log \frac{p_i}{q_i} + q_i - p_i \right) \\ &\leq \frac{1}{\eta} \sum_{i=1}^K p_i (\exp(-\eta \ell_i) - 1 + \eta \ell_i) \leq 2\eta \sum_{i=1}^K p_i \ell_i^2, \end{aligned}$$

where the inequality is because $\eta \ell_i \geq -3$. ■

Now we are ready to bound Term (6) as follows.

Lemma 10 *With probability at least $1 - 2\delta$,*

$$\begin{aligned} \text{Term (6)} &\leq \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \eta T + \eta \sum_{t=1}^T Q_t + \log \frac{1}{\delta} + 48\sqrt{\eta T} + \beta^2 T + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \frac{1}{\delta} \right) \\ &\quad - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}\} + \beta \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}}. \end{aligned}$$

Proof Recall that according to the definition in [Algorithm 1](#), $\mathcal{T} \triangleq \{t \mid \text{there exists } j_t \in \bar{S}_t \text{ and } p_{t,j_t} > 1/2\}$. To apply [Lemma 9](#), we first need to verify the scale of $\widehat{\ell}_t + b_t - z_t$ where $z_t = \widehat{\ell}_{t,j_t} + b_{t,j_t}$ if $t \in \mathcal{T}$. If $t \notin \mathcal{T}$, then we know that for all $i \in [K]$, $\eta(\widehat{\ell}_{t,i} + b_{t,i}) = \eta \widehat{\ell}_{t,i} \geq 0$. If $t \in \mathcal{T}$, note that with an η amount of uniform exploration,

$$\eta z_t = \eta \left(\widehat{\ell}_{t,j_t} + b_{t,j_t} \right) \leq \eta(1 + \beta) \cdot \frac{1}{\frac{K-1}{K}\eta} \leq 2(1 + \beta) \leq 3,$$

where the second inequality is because $K \geq 2$ and the last inequality is because $\beta \leq \frac{1}{2}$. Therefore, we know that $\eta(\widehat{\ell}_{t,i} + b_{t,i} - z_t) \geq -3$ for all $i \in [K]$.

Therefore, applying [Lemma 9](#) with $p = p_t$ and $p' = p_{t+1}$ and taking summation over $t \in [T]$, we know that for any $u \in \Delta_K$,

$$\begin{aligned}
 & \sum_{t=1}^T \langle p_t - u, \widehat{\ell}_t + b_t \rangle \\
 & \leq \sum_{t=1}^T (D_\psi(u, p_t) - D_\psi(u, p_{t+1})) + 2\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \widehat{\ell}_{t,i}^2 \cdot \mathbb{1}\{t \notin \mathcal{T}\} \\
 & \quad + 2\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} (\widehat{\ell}_{t,i} + b_{t,i} - z_t)^2 \cdot \mathbb{1}\{t \in \mathcal{T}\} \\
 & \leq \frac{D_\psi(u, p_1)}{\eta} + 2\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \widehat{\ell}_{t,i}^2 \mathbb{1}\{t \notin \mathcal{T}\} + 6\eta \sum_{t=1}^T \sum_{i \neq j_t} p_{t,i} (\widehat{\ell}_{t,i}^2 + b_{t,i}^2 + z_t^2) \mathbb{1}\{t \in \mathcal{T}\} \\
 & = \frac{D_\psi(u, p_1)}{\eta} + 2\eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \widehat{\ell}_{t,i}^2 \mathbb{1}\{t \notin \mathcal{T}\} + 6\eta \sum_{t=1}^T \sum_{i \neq j_t} p_{t,i} (\widehat{\ell}_{t,i}^2 + z_t^2) \mathbb{1}\{t \in \mathcal{T}\}.
 \end{aligned}$$

For $t \notin \mathcal{T}$, we know that $\widehat{\ell}_{t,i} \leq 2$ for $i \in \bar{S}_t$ and

$$\begin{aligned}
 \sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2 & \leq 4 \sum_{i \notin S_t} p_{t,i} \leq 4, \\
 \sum_{i \in S_t} p_{t,i} \widehat{\ell}_{t,i}^2 & \leq \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i}.
 \end{aligned}$$

For $t \in \mathcal{T}$, we know that $\widehat{\ell}_{t,i} \leq 2$ for all $i \in \bar{S}_t \setminus \{j_t\}$ and

$$\begin{aligned}
 \sum_{i \in S_t} p_{t,i} \widehat{\ell}_{t,i}^2 & \leq \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i}, \\
 \sum_{i \in \bar{S}_t, i \neq j_t} p_{t,i} \widehat{\ell}_{t,i}^2 & \leq 4 \sum_{i \in \bar{S}_t, i \neq j_t} p_{t,i} \leq 4, \\
 \sum_{i \neq j_t} p_{t,i} z_t^2 & \leq W_{t,j_t} \cdot (\widehat{\ell}_{t,j_t} + b_{t,j_t})^2 \leq 2W_{t,j_t} \widehat{\ell}_{t,j_t}^2 + 2\frac{\beta^2}{W_{t,j_t}} \leq 2\widehat{\ell}_{t,j_t} + \frac{4\beta^2}{\eta},
 \end{aligned}$$

where the last inequality uses the fact that $W_{t,j_t} \geq \frac{K-1}{K}\eta \geq \frac{1}{2}\eta$. For any $j \in [K]$, let $u = e_j \in \Delta_K$. Combining all the above inequalities, we can obtain that

$$\begin{aligned}
 & \sum_{t=1}^T \langle p_t - e_j, \widehat{\ell}_t + b_t \rangle \\
 & \leq \frac{D_\psi(e_j, p_1)}{\eta} + 24\eta T + 6\eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} + 12\eta \sum_{t=1}^T \widehat{\ell}_{t,j_t} \mathbb{1}\{t \in \mathcal{T}\} + 24\beta^2 \sum_{t=1}^T \mathbb{1}\{t \in \mathcal{T}\} \\
 & \leq \frac{D_\psi(e_j, p_1)}{\eta} + 24\eta T + 6\eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} + 12\eta \sum_{t=1}^T \widehat{\ell}_{t,j_t} \mathbb{1}\{t \in \mathcal{T}\} + 24\beta^2 T. \tag{22}
 \end{aligned}$$

We first bound the term $\sum_{t=1}^T \widehat{\ell}_{t,j_t} \mathbb{1}\{t \in \mathcal{T}\}$. Let $Z_t = \widehat{\ell}_{t,j_t} \mathbb{1}\{t \in \mathcal{T}\} - \ell_{t,j_t} \mathbb{1}\{t \in \mathcal{T}\}$. We know that $\mathbb{E}_t[Z_t] = 0$ and $|Z_t| \leq \frac{1}{\frac{K-1}{K}\eta} \leq \frac{2}{\eta}$. In addition,

$$\mathbb{E}_t[Z_t^2] \leq \mathbb{E}_t \left[\frac{1}{W_{t,j_t}^2} \cdot \mathbb{1}\{i_t \neq j_t\} \right] \cdot \mathbb{1}\{t \in \mathcal{T}\} = \frac{\mathbb{1}\{t \in \mathcal{T}\}}{W_{t,j_t}}.$$

Therefore, by Freedman's inequality ([Lemma 12](#)), we can obtain that with probability at least $1 - \delta$,

$$\sum_{t=1}^T Z_t \leq \min_{\lambda \in [0, \frac{2}{\eta}]} \left\{ \frac{\log(1/\delta)}{\lambda} + \lambda \sum_{t=1}^T \mathbb{E}_t[Z_t^2] \right\} \leq \frac{2 \log(1/\delta)}{\eta} + 2 \sqrt{\sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}\}}{W_{t,j_t}}} \leq \frac{2 \log(1/\delta)}{\eta} + 4 \sqrt{\frac{T}{\eta}}.$$

Combining with [Eq. \(22\)](#), we know that with probability at least $1 - 2\delta$

$$\begin{aligned} & \sum_{t=1}^T \langle p_t - e_j, \widehat{\ell}_t \rangle \\ & \leq \frac{D_\psi(e_j, p_1)}{\eta} + 24\eta T + 6\eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} + 24 \log(1/\delta) + 48\sqrt{\eta T} + 24\beta^2 T - \sum_{t=1}^T \langle p_t - e_j, b_t \rangle \\ & \leq \frac{\log K}{\eta} + 24\eta T + 6\eta \sum_{t=1}^T Q_t + 24 \log \frac{1}{\delta} + 48\sqrt{\eta T} + 24\beta^2 T \\ & \quad + \widetilde{\mathcal{O}} \left(\eta \sum_{t=1}^T Q_t + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \frac{1}{\delta} \right) - \sum_{t=1}^T \langle p_t - e_j, b_t \rangle \\ & \leq \widetilde{\mathcal{O}} \left(\frac{1}{\eta} + \eta T + \eta \sum_{t=1}^T Q_t + \log \frac{1}{\delta} + \sqrt{\eta T} + \beta^2 T + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \frac{1}{\delta} \right) \\ & \quad - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}\} + \beta \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}}, \end{aligned}$$

where the second inequality is because of [Lemma 1](#) and the choice of $p_1 = \frac{1}{K} \cdot \mathbf{1}$. \blacksquare

With [Lemma 7](#), [Lemma 8](#) and [Lemma 10](#) on hand, we are ready to prove [Theorem 3](#).

Proof [Proof of [Theorem 3](#)] According to the regret decomposition [Eq. \(18\)](#), [Lemma 7](#), [Lemma 8](#) and [Lemma 10](#) and the bounds on Term (1), Term (2) and Term (3), we know that with probability at least $1 - 6\delta$, for any $j \in [K]$,

$$\begin{aligned} & \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) \\ & \leq \widetilde{\mathcal{O}} \left(\sqrt{T \log \frac{1}{\delta}} \right) + \mathcal{O}(\eta T) + \left(2 + \frac{4}{\eta} \right) \log \frac{1}{\delta} + 2 \sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}} \\ & \quad + \widetilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \max_{t \in [T]} Q_t \log \frac{1}{\delta} + \gamma \sum_{t=1}^T Q_t + \frac{\log(1/\delta)}{\gamma} \right) \end{aligned}$$

$$\begin{aligned}
 & + \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \eta T + \eta \sum_{t=1}^T Q_t + \log \frac{1}{\delta} + 48\sqrt{\eta T} + \beta^2 T + \frac{\eta \max_{t \in [T]} Q_t}{\gamma} \log \frac{1}{\delta} \right) \\
 & - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}\} + \beta \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}} \\
 \leq & \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \frac{\log(1/\delta)}{\gamma} + (\eta + \gamma) \sum_{t=1}^T Q_t + \sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \sqrt{\eta T} + \beta^2 T + \left(\frac{\eta}{\gamma} + 1 \right) \max_{t \in [T]} Q_t \log \frac{1}{\delta} \right) \\
 & + 2\sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}} - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}\} + \beta \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}}.
 \end{aligned} \tag{23}$$

Consider the last three terms:

$$\begin{aligned}
 & 2\sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}} - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}\} + \beta \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}} \\
 & = 2\sqrt{\left(4T + \sum_{t=1}^T \frac{\mathbb{1}\{t \in \mathcal{T}, j \neq j_t\}}{W_{t,j_t}} \right) \log \frac{1}{\delta}} - \beta \sum_{t=1}^T \frac{p_{t,j_t}}{W_{t,j_t}} \mathbb{1}\{t \in \mathcal{T}, j \neq j_t\} \\
 & \quad + \beta \sum_{t=1}^T \frac{(1 - p_{t,j_t}) \mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{W_{t,j_t}} \\
 & \leq \mathcal{O} \left(\sqrt{T \log \frac{1}{\delta}} \right) + \frac{1}{\beta} + \beta \sum_{t=1}^T \frac{\sum_{i \neq j_t} p_{t,i} \mathbb{1}\{t \in \mathcal{T}, j = j_t\}}{(1 - \eta) \sum_{i \neq j_t} p_{t,i}} \\
 & \leq \mathcal{O} \left(\sqrt{T \log \frac{1}{\delta}} + \frac{1}{\beta} + \beta T \right).
 \end{aligned}$$

where the first inequality uses the AM-GM inequality and the second inequality uses the fact that $\eta \leq \frac{1}{2}$. Combining with Eq. (23), we obtain

$$\begin{aligned}
 & \sum_{t=1}^T (\ell_{t,i_t} - \ell_{t,j}) \\
 \leq & \tilde{\mathcal{O}} \left(\frac{1}{\eta} + \frac{\log(1/\delta)}{\gamma} + (\eta + \gamma) \sum_{t=1}^T Q_t + \sqrt{\sum_{t=1}^T Q_t \log \frac{1}{\delta}} + \sqrt{\eta T} + \beta^2 T + \left(\frac{\eta}{\gamma} + 1 \right) \max_{t \in [T]} Q_t \log \frac{1}{\delta} \right) \\
 & + \mathcal{O} \left(\sqrt{T \log \frac{1}{\delta}} + \frac{1}{\beta} + \beta T \right).
 \end{aligned}$$

Using Lemma 11, we know that

$$Q_t \leq 2\alpha_t \log \left(1 + \frac{\lceil K^2/\gamma \rceil + K}{\alpha_t} \right) + 2 \leq 4\alpha_t \log \left(1 + \frac{\lceil K^2/\gamma \rceil + K}{\alpha_t} \right) = \tilde{\mathcal{O}}(\alpha_t).$$

Picking $\eta = \beta = \gamma = 1/\sqrt{\sum_{t=1}^T \alpha_t \log(1/\delta)}$, we achieve that with probability at least $1 - 6\delta$,

$$\text{Reg}_T \leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \alpha_t \log \frac{1}{\delta}} + \max_{t \in [T]} \alpha_t \log \frac{1}{\delta} \right).$$

This finishes our proof. \blacksquare

Appendix B. Proofs for Section 4

In this section, we prove [Lemma 5](#) and [Lemma 6](#). The key of the proof is to use a careful analysis of Freedman's inequality with the help of uniform exploration and implicit exploration.

Proof of [Lemma 5](#). Recall that $\langle p, \ell \rangle_S \triangleq \sum_{i \in S} p_i \ell_i$ for any $S \subseteq [K]$. Therefore, we decompose the target $\sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle$ as follows

$$\sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle = \sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle_{S_t} + \sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle_{\bar{S}_t}.$$

Bounding $\sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle_{S_t}$: We proceed as follows

$$\sum_{t=1}^T \sum_{i \in S_t} p_{t,i} (\ell_{t,i} - \hat{\ell}_{t,i}) = \sum_{t=1}^T \sum_{i \in S_t} \gamma \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} + \sum_{t=1}^T \sum_{i \in S_t} (W_{t,i} - \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma}.$$

Recall that $Q_t = \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma}$. As $\ell_t \in [0, 1]^K$, it is clear that $\sum_{i \in S_t} \gamma \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} \leq \gamma Q_t$. To bound the second term, according to [Lemma 1](#), let $\iota_1 = \log \left(\frac{2 \max_t Q_t + 2\sqrt{\sum_{t=1}^T Q_t}}{\delta'} \right)$, we know that with probability at least $1 - \delta'$,

$$\sum_{t=1}^T \sum_{i \in S_t} (W_{t,i} - \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\}) \frac{p_{t,i} \ell_{t,i}}{W_{t,i} + \gamma} \leq \mathcal{O} \left(\sqrt{\sum_{t=1}^T Q_t \iota_1} + \max_{t \in [T]} Q_t \iota_1 \right).$$

Consider the subgraph \tilde{G}_t of $G_t = ([K], E_t)$ where $\tilde{G}_t = (S_t, \tilde{E}_t)$ and $\tilde{E}_t \subseteq E_t$ is the set of edges with respect to the nodes in S_t . Applying [Lemma 11](#) on the subgraph \tilde{G}_t , we know that

$$Q_t \leq \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{p_{t,i} + \sum_{j \in N_t^{\text{in}}(i)} p_{t,j} + \gamma} \leq \tilde{\mathcal{O}}(\alpha_t). \quad (24)$$

Combining all the above equations, we know that with probability at least $1 - \delta'$,

$$\begin{aligned} \sum_{t=1}^T \langle p_t, \ell_t - \hat{\ell}_t \rangle &\leq \mathcal{O} \left(\gamma \sum_{t=1}^T Q_t + \sqrt{\sum_{t=1}^T Q_t \iota_1} + \max_{t \in [T]} Q_t \iota_1 \right) \\ &\leq \tilde{\mathcal{O}} \left(\gamma \sum_{t=1}^T \tilde{\alpha}_t + \sqrt{\sum_{t=1}^T \tilde{\alpha}_t \iota_1} + \max_{t \in [T]} \tilde{\alpha}_t \iota_1 \right). \end{aligned} \quad (25)$$

Bounding $\sum_{t=1}^T \langle p_t, \ell_t - \widehat{\ell}_t \rangle_{\bar{S}_t}$: Since the loss estimators for nodes without a self-loop are unbiased, we directly apply [Lemma 12](#) to bound $\sum_{t=1}^T \langle p_t, \ell_t - \widehat{\ell}_t \rangle_{\bar{S}_t}$. Note that

$$\begin{aligned} \sum_{i \in \bar{S}_t} p_{t,i} (\ell_{t,i} - \widehat{\ell}_{t,i}) &\leq \sum_{i \in \bar{S}_t} p_{t,i} \ell_{t,i} \leq 1 \\ \mathbb{E}_t \left[\left(\sum_{i \in \bar{S}_t} p_{t,i} (\ell_{t,i} - \widehat{\ell}_{t,i}) \right)^2 \right] &\leq \mathbb{E}_t \left[\left(\sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i} \right)^2 \right] \\ &\leq \frac{1}{\varepsilon} \mathbb{E}_t \left[\sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i} \right] \\ &\leq \frac{1}{\varepsilon}, \end{aligned}$$

where the second inequality is because $\widehat{\ell}_{t,i} \leq \frac{1}{\varepsilon}$ for all $i \in \bar{S}_t$. Therefore, using [Lemma 12](#), we obtain that with probability at least $1 - \delta'$

$$\sum_{t=1}^T \sum_{i \in \bar{S}_t} p_{t,i} (\ell_{t,i} - \widehat{\ell}_{t,i}) \leq 2\sqrt{\frac{T}{\varepsilon} \log(1/\delta')} + \log(1/\delta'). \quad (26)$$

Let $\delta' = \frac{\delta}{2}$. Combining [Eq. \(25\)](#), [Eq. \(26\)](#), we prove the result of first part.

For the second part, we consider the cases where $i \in S_t$ and $i \in \bar{S}_t$ separately.

$$\sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) = \sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in S_t\} + \sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in \bar{S}_t\}.$$

The analysis for the first term is the same as [Eq. \(20\)](#) and we can obtain that with probability at least $1 - \delta'$, for all $i \in [K]$,

$$\sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in S_t\} \leq \frac{\log(K/\delta')}{2\gamma}. \quad (27)$$

For the second term, note that $(\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in \bar{S}_t\} \leq \frac{1}{\varepsilon}$ as $\widehat{\ell}_{t,i} \leq \frac{1}{\varepsilon}$ for all $i \in \bar{S}_t$. In addition, the conditional variance is bounded as follows

$$\mathbb{E}_t \left[\left((\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in \bar{S}_t\} \right)^2 \right] \leq \mathbb{E}_t \left[\widehat{\ell}_{t,i}^2 \mathbb{1}\{i \in \bar{S}_t\} \right] \leq \frac{1}{\varepsilon}.$$

Using [Lemma 12](#) and an union bound over $[K]$, for all $i \in [K]$, we have that with probability at least $1 - \delta'$

$$\sum_{t=1}^T (\widehat{\ell}_{t,i} - \ell_{t,i}) \mathbb{1}\{i \in \bar{S}_t\} \leq \sqrt{\frac{T}{\varepsilon} \log \frac{K}{\delta'}} + \frac{1}{\varepsilon} \log \frac{K}{\delta'}. \quad (28)$$

Combining [Eq. \(27\)](#) and [Eq. \(28\)](#) and picking $\delta' = \frac{\delta}{2}$, we prove the second part. ■

Next, we prove [Lemma 6](#), which bounds the estimated regret Term (d) in [Eq. \(16\)](#).

Proof of Lemma 6. We apply standard OMD analysis ([Bubeck et al., 2012](#)) and obtain that

$$\begin{aligned} \text{Term (d)} &\leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^K p_{t,i} \widehat{\ell}_{t,i}^2 \\ &\leq \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{i \in S_t} p_{t,i} \widehat{\ell}_{t,i}^2 + \eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2. \end{aligned}$$

For the second term, using [Lemma 1](#), we know that with probability at least $1 - \delta'$

$$\begin{aligned} \eta \sum_{t=1}^T \sum_{i \in S_t} p_{t,i} \widehat{\ell}_{t,i}^2 &\leq \eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \widehat{\ell}_{t,i} \\ &\leq \eta \sum_{t=1}^T \sum_{i \in S_t} \frac{p_{t,i}}{W_{t,i} + \gamma} \ell_{t,i} + \widetilde{\mathcal{O}} \left(\eta \sum_{t=1}^T Q_t + \max_{t \in [T]} \frac{Q_t}{\gamma} \log \frac{1}{\delta'} \right) \\ &\leq \widetilde{\mathcal{O}} \left(\eta \sum_{t=1}^T \widetilde{\alpha}_t + \frac{\eta}{\gamma} \max_t \widetilde{\alpha}_t \log \frac{1}{\delta'} \right). \end{aligned} \quad (29)$$

For the third term, we decompose it as

$$\begin{aligned} \eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} p_{t,i} \widehat{\ell}_{t,i}^2 &\leq \eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} \widehat{\ell}_{t,i} \\ &\leq \underbrace{\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} (\widehat{\ell}_{t,i} - \ell_{t,i})}_{\text{Term (i)}} + \underbrace{\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} \ell_{t,i}}_{\text{Term (ii)}}. \end{aligned}$$

To bound Term (i), note that with uniform exploration on the dominating set, $\sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} (\widehat{\ell}_{t,i} - \ell_{t,i}) \leq \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} \widehat{\ell}_{t,i} \leq \frac{1}{\varepsilon^2}$. Next, we consider the conditional variance:

$$\mathbb{E}_t \left[\left(\sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} (\widehat{\ell}_{t,i} - \ell_{t,i}) \right)^2 \right] \leq \mathbb{E}_t \left[\sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(i)\} \sum_{j \in \bar{S}_t} \frac{p_{t,j}}{W_{t,j}^2} \mathbb{1}\{i_t \in N_t^{\text{in}}(j)\} \right] \leq \frac{1}{\varepsilon^3}.$$

Using Freedman's inequality [Lemma 12](#), we know that with probability at least $1 - \delta'$,

$$\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} (\widehat{\ell}_{t,i} - \ell_{t,i}) \leq \eta \sqrt{\frac{T}{\varepsilon^3} \log \frac{1}{\delta'}} + \frac{\eta}{\varepsilon^2} \log \frac{1}{\delta'}. \quad (30)$$

For Term (ii), we directly bound it by noticing that $W_{t,i} \geq \varepsilon$ for all $i \in \bar{S}_t$

$$\eta \sum_{t=1}^T \sum_{i \in \bar{S}_t} \frac{p_{t,i}}{W_{t,i}} \ell_{t,i} \leq \frac{\eta T}{\varepsilon}. \quad (31)$$

Combining [Eq. \(29\)](#), [Eq. \(30\)](#), [Eq. \(31\)](#) and picking $\delta' = \frac{\delta}{2}$, we finish the proof. \blacksquare

Appendix C. Auxiliary Lemmas

In this section, we show several auxiliary lemmas that are useful in the analysis.

Lemma 11 (Lemma 1 in (Kocák et al., 2014)) *Let $G = (V, E)$ be a directed graph with $|V| = K$, in which each node $i \in V$ is assigned a positive weight w_i . Assume that $\sum_{i \in V} w_i \leq 1$, then*

$$\sum_{i \in V} \frac{w_i}{w_i + \sum_{j \in N^{\text{in}}(i)} w_j + \gamma} \leq 2\alpha \log \left(1 + \frac{\lceil K^2/\gamma \rceil + K}{\alpha} \right) + 2,$$

where α is the independence number of G .

Lemma 12 (Freedman's inequality, Theorem 1 (Beygelzimer et al., 2011)) *Let X_1, X_2, \dots, X_T be a martingale difference sequence with respect to a filtration $F_1 \subseteq F_2 \subseteq \dots \subseteq F_T$ such that $\mathbb{E}[X_t | F_t] = 0$. Assume for all t , $X_t \leq R$. Let $V = \sum_{t=1}^T \mathbb{E}[X_t^2 | F_t]$. Then for any $\delta > 0$, with probability at least $1 - \delta$, we have the following guarantee:*

$$\begin{aligned} \sum_{t=1}^T X_t &\leq \inf_{\lambda \in \left[0, \frac{\sqrt{e-2}}{R\sqrt{\log(1/\delta)}}\right]} \left\{ \sqrt{(e-2) \log(1/\delta)} \left(\lambda V + \frac{1}{\lambda} \right) \right\} \\ &= \inf_{\lambda' \in \left[0, \frac{1}{R}\right]} \left\{ \frac{\log(1/\delta)}{\lambda'} + (e-2)\lambda' V \right\}. \end{aligned}$$

Lemma 13 (Strengthened Freedman's inequality, Theorem 9 (Zimmert and Lattimore, 2022))

Let X_1, X_2, \dots, X_T be a martingale difference sequence with respect to a filtration $F_1 \subseteq F_2 \subseteq \dots \subseteq F_T$ such that $\mathbb{E}[X_t | F_t] = 0$ and assume $\mathbb{E}[|X_t| | F_t] < \infty$ a.s. Then with probability at least $1 - \delta$

$$\sum_{t=1}^T X_t \leq 3\sqrt{V_T \log \left(\frac{2 \max\{U_T, \sqrt{V_T}\}}{\delta} \right)} + 2U_T \log \left(\frac{2 \max\{U_T, \sqrt{V_T}\}}{\delta} \right),$$

where $V_T = \sum_{t=1}^T \mathbb{E}[X_t^2 | F_t]$, $U_T = \max\{1, \max_{s \in [T]} X_s\}$.