

The Real Price of Bandit Information in Multiclass Classification

Liad Erez

Tel Aviv University, Tel Aviv, Israel.

LIADEREZ@MAIL.TAU.AC.IL

Alon Cohen

Tel Aviv University, Tel Aviv, Israel and Google Research Tel Aviv, Israel.

ALONCO@TAUEX.TAU.AC.IL

Tomer Koren

Tel Aviv University, Tel Aviv, Israel and Google Research Tel Aviv, Israel.

TKOREN@TAUEX.TAU.AC.IL

Yishay Mansour

Tel Aviv University, Tel Aviv, Israel and Google Research Tel Aviv, Israel.

MANSOUR.YISHAY@GMAIL.COM

Shay Moran

Technion, Haifa, Israel and Google Research Tel Aviv, Israel.

SHAYMORANI@GMAIL.COM

Editors: Shipra Agrawal and Aaron Roth

Abstract

We revisit the classical problem of multiclass classification with bandit feedback (Kakade, Shalev-Shwartz, and Tewari, 2008), where each input classifies to one of K possible labels and feedback is restricted to whether the predicted label is correct or not. Our primary inquiry is with regard to the dependency on the number of labels K , and whether T -step regret bounds in this setting can be improved beyond the \sqrt{KT} dependence exhibited by existing algorithms. Our main contribution is in showing that the minimax regret of bandit multiclass is in fact more nuanced, and is of the form $\Theta(\min\{|\mathcal{H}| + \sqrt{T}, \sqrt{KT \log|\mathcal{H}|\})$, where \mathcal{H} is the underlying (finite) hypothesis class. In particular, we present a new bandit classification algorithm that guarantees regret $\tilde{O}(|\mathcal{H}| + \sqrt{T})$, improving over classical algorithms for moderately-sized hypothesis classes, and give a matching lower bound establishing tightness of the upper bounds (up to log-factors) in all parameter regimes.

1. Introduction

Online multiclass classification is an important and fundamental learning problem, which has practical relevance in a variety of applications. This is a sequential problem, where in any given round $t = 1, 2, \dots, T$ the learner receives an example x_t and is tasked with predicting a label y'_t from a set of K possible labels, after which the true label y_t is revealed to the learner who suffers a loss if $y'_t \neq y_t$. The learner's performance is measured with respect to an underlying *hypothesis class* \mathcal{H} of mappings from examples to labels, with the objective of minimizing the *regret*, that is, the total loss of the learner compared to that of the best hypothesis in \mathcal{H} .

In *bandit multiclass classification* (Kakade et al., 2008), rather than observing the true label after each prediction (a setting which is often referred to as *full information*), the learner only observes whether the prediction was correct or not (this is reminiscent of *bandit information* in the context of online learning). As a practical example, consider a classification task over the ImageNet dataset, where a learner is tasked with classifying images into one of $K > 10000$ possible labels. After the learner makes a prediction, the image and the prediction are shown to a human rater that is asked whether or not the prediction is correct. If we consider the rater as a part of the environment, the learner essentially faces a bandit multiclass classification problem. Notably, while a given image

may have multiple correct labels, this number is usually small (say, at most 10) and does not scale with K .

The bandit multiclass classification problem has been extensively studied since the early work of Kakade et al. 2008 (e.g., Daniely et al., 2011; Daniely and Helbertal, 2013; Long, 2020; Raman et al., 2023), where the primary focus has been on characterizing the price of bandit information as compared to the standard full-information setting. Several of these works placed a particular focus on characterizing the properties of \mathcal{H} under which regret minimization is at all possible (i.e., render the problem learnable), and have introduced refinements of the best known regret bound of $O(\sqrt{KT \log|\mathcal{H}|})$ (that can be extracted from Auer et al., 2002) in terms of various structural properties of the class \mathcal{H} ; for instance, Daniely and Helbertal (2013) proved a bound which scales with the Littlestone dimension of \mathcal{H} rather than $\log|\mathcal{H}|$, while Raman et al. (2023) replaced the dependence on K by a quantity called the “Bandit Littlestone dimension” of \mathcal{H} , which, for general classes can be as large as $|\mathcal{H}|$.

Despite this abundance of previous work on the fundamental bandit multiclass problem, it remains unclear what is the correct dependence of the regret on the number of labels, K , even for finite hypothesis classes. Drawing analogy to the vast literature on bandit problems, one could conjecture that the right dependence should be roughly \sqrt{KT} , as is the case in stochastic and non-stochastic K -armed bandit problems and their contextual counterparts (Auer et al., 2002; Lattimore and Szepesvari, 2020), which is indeed the case in all existing upper bounds for bandit multiclass. However, a close inspection of the available lower bounds for bandit multiclass classification (Daniely and Helbertal, 2013), reveals that they can only rule out bounds better than \sqrt{KT} in the *multi-label setting*, where each example may be labeled with a multitude of labels rather than just one or a few (the lower bounds require that $\Theta(K)$ correct labels are possible). When considering the canonical single-label setting exclusively, which has a significant “sparsity” structure compared to a general bandit problem, the only available lower bound becomes the $\tilde{\Omega}(\sqrt{T})$ rate of the full-information case.

Thus, the following question remains unresolved: *what is the real price of bandit information in (standard, single-label) multiclass classification?*¹ In this work, we provide a complete answer to this question by fully characterizing the minimax regret in single-label bandit multiclass classification over finite hypothesis classes (up to logarithmic factors). Concretely, we show that the minimax regret for this problem is of the form

$$\tilde{\Theta}(\min(|\mathcal{H}| + \sqrt{T}, \sqrt{KT \log|\mathcal{H}|})).$$

In particular, this bound improves upon known algorithms (such as EXP4, Auer et al., 2002) for hypothesis class of moderate size, that is, $|\mathcal{H}| \lesssim \sqrt{KT}$. Somewhat surprisingly, in the latter case there is therefore no price for the bandit information: namely, the (leading T -dependent term in the) minimax rate become identical, up to logarithmic factors, to the rate in the full-information case of the problem. On the flip side, for larger classes, our main result shows that the price of bandit information is fully pronounced (i.e., a \sqrt{KT} dependence is unavoidable) despite the inherent “sparsity” structure of the single-label setting.

1. Henceforth, we focus on the classical single-label setup where each example has a single, unique correct label, but we note that all of our results are equally valid in the case where each example may have at most a constant number of correct labels.

1.1. Summary of contributions

In some more detail, our main results in this paper are the following:

- We describe and analyze an algorithm which is an instantiation of the follow-the-regularized-leader framework with a regularization that combines negative entropy and log-barrier components (see Section 3). Our analysis shows that this algorithm achieves the expected regret bound stated above in general adversarial environments. In fact, our result holds in a more general s -sparse contextual bandits setup (see Section 3 for a formal description of the problem setup) over a finite policy class Π of mappings from contexts to actions, in which our algorithm obtains a regret bound of

$$\tilde{O}(\min\{|\Pi| + \sqrt{sT}, \sqrt{KT \log|\Pi|}\}).$$

In the case of bandit multiclass, the sparsity is $s = 1$ which yields the upper bound stated above. (The same result holds true if the maximal number of correct labels is any other constant that does not scale with K .)

- We prove a lower bound establishing the fact that the regret guarantees provided by our algorithm are essentially optimal (see Section 4). Our construction is of a stochastic i.i.d. bandit multiclass classification instance, implying that this guarantee cannot be improved even when the environment is stochastic rather than adversarial. Specifically, we show that for any bandit multiclass classification algorithm there exists a stochastic instance on which it must incur an expected regret of at least

$$\tilde{\Omega}(\min\{|\mathcal{H}| + \sqrt{T}, \sqrt{KT}\}).$$

This bound implies that our new bandit classification algorithm is tight for moderately-sized hypothesis classes, while for larger classes the $\tilde{O}(\sqrt{KT})$ bound attained by EXP4 cannot be improved, despite the inherent sparsity structure in the problem.

1.2. Overview of main ideas and techniques

One of the main challenges in obtaining improved regret bounds for bandit multiclass classification comes in taking advantage of structural properties which are absent in more general bandit scenarios, where regret bounds of the form \sqrt{KT} are optimal. One basic observation is that bandit multiclass classification is essentially a special case of *contextual bandits* (Auer et al., 2002; Beygelzimer et al., 2011; Agarwal et al., 2014) with additional important structural properties of the loss functions. Namely, the loss functions in bandit multiclass classification are given by the *zero-one loss*, that is, a zero loss for the correct prediction and one otherwise. In the setting where each example has a unique correct label, the losses exhibit a certain sparsity property. We are thus motivated to investigate whether it is possible to take advantage of the sparse structure of the loss functions in this setting in order to obtain guarantees that improve upon those given for contextual bandits in general.

This type of sparsity in multi-armed bandits has been a topic of interest in several previous works (Bubeck et al., 2018; Audibert et al., 2009), where it has been shown that the optimal regret bound of $\tilde{O}(\sqrt{KT})$ can be improved to $\tilde{O}(K + \sqrt{sT})$ if the loss vectors are s -sparse, that is, every loss vector is bounded in ℓ_2 -norm by s . These results motivated the investigation of whether or not

such improvements are possible in the more general bandit multiclass classification setup. That is, whether or not we can obtain regret bounds of the form $\tilde{O}(K + \sqrt{T \log |\mathcal{H}|})$, since in this case the loss vectors are 1-sparse. Perhaps surprisingly, it turns out that achieving regret bounds of the form $K + \sqrt{sT}$ is not possible in general in bandit multiclass classification. However, by applying an instantiation of *follow-the-regularized-leader* (FTRL) with a log-barrier component in the contextual bandit setting we can establish a regret bound of $\tilde{O}(|\Pi| + \sqrt{sT})$ when the losses are s -sparse, which in turn implies a bound of $\tilde{O}(|\mathcal{H}| + \sqrt{T})$ for bandit multiclass classification. This turns out to be optimal, up to logarithmic factors, when the hypothesis class is not too large, that is, $|\mathcal{H}| \lesssim \sqrt{KT}$.

The incorporation of log-barrier regularization in order to stabilize the FTRL iterates turns out to be crucial for our approach. The reason stems from the fact that when working with a shifted version of the zero-one loss functions (which is required for sparsity to hold), the loss values are necessarily non-positive, which is notoriously more challenging compared to nonnegative losses (Kwon and Perchet, 2016). Together with the fact that the learner observes bandit feedback, the standard importance-weighted loss estimators become both negative and extremely large in magnitude, and as a result, existing algorithms for contextual bandits such as EXP4 (Auer et al., 2002) become unstable. The addition of a log-barrier component to the regularization aims to resolve this very problem, albeit with the penalty of incurring an additional additive term of $|\mathcal{H}|$ in the regret bound.

1.3. Open problems and future work

In this work, we addressed the fundamental question of characterizing the optimal regret bound in multiclass classification with bandit feedback for finite classes. Specifically, we provided a tight bound that is expressed in terms of the class size $|\mathcal{H}|$ and the number of labels K . This achievement, however, can be seen as an initial step within a broader research context. Below, we detail several open problems and natural avenues for future research.

1. **Structured (possibly infinite) hypothesis classes.** A logical extension of our work is to refine these bounds by including dependencies on more refined class properties. This methodology has a solid foundation in learning theory; a classical example is the role of the VC dimension (Blumer et al., 1989) in PAC learning, which refines the dependence on $|\mathcal{H}|$ and yields optimal bounds in terms of sample complexity that are applicable also to infinite classes.

In fact, our lower bound in terms of $|\mathcal{H}|$ reveals a natural candidate for such a combinatorial parameter. For an integer κ , define the κ -list star number of a hypothesis class \mathcal{H} as the maximal s for which there exist s examples x_1, \dots, x_s and a pivot hypothesis $h_0 \in \mathcal{H}$ such that, for each x_i , there are κ hypotheses $h_{i,1}, \dots, h_{i,\kappa}$ that coincide with h_0 on every $x_j \neq x_i$ and diverge on x_i , with all labels $h_0(x_i), h_{i,1}(x_i), \dots, h_{i,\kappa}(x_i)$ being distinct. When $\kappa = 1$ this parameter specializes to the star-number which characterizes the optimal query complexity in active learning (Hanneke and Yang, 2015). An adaptation of our lower bound for finite classes shows an $\tilde{\Omega}(\kappa s + \sqrt{T \log K})$ lower bound for classes with κ -list star number s , thus positioning the κ -list star number as a natural barrier. It remains an open question determine whether it is the only barrier and in particular whether there exist a corresponding upper bound which replaces $|\mathcal{H}|$ with the κ -list star number, offering tighter bounds on the regret which apply also to infinite classes.

2. **Computationally efficient algorithms.** Several previous works on contextual bandits (Langford and Zhang, 2007; Dudik et al., 2011; Agarwal et al., 2014) focus on the stochastic i.i.d. setting and aim to achieve low regret with efficient algorithms based on various types of optimization oracles. An interesting question for future research is whether or not we can utilize the sparse nature of bandit multiclass classification in the stochastic setting and obtain more efficient algorithms that guarantee optimal regret in this setting. In the simpler setting of stochastic K -armed bandits, variants of optimistic algorithms such as UCB (Audibert et al., 2009) that are adaptive to the “sparsity” of the rewards have been established, but it is not yet clear to us if and how such techniques could be extended to the contextual setup.
3. **Tight sample complexity bounds.** Another possible learning objective in a stochastic multiclass classification setup is a PAC objective. In this setting, the learner does not incur any loss during the learning process, but is tasked with outputting a hypothesis which is nearly optimal with a low *sample complexity*. A straightforward application of an online-to-batch reduction to our online algorithm gives a sample complexity guarantee of $\tilde{O}(N/\varepsilon + 1/\varepsilon^2)$ in the PAC setting. Moreover, the construction of our lower bound can be used to show a sample complexity lower bound of $\Omega(K/\varepsilon)$, which does not preclude the possibility of having an algorithm in this framework that guarantees finding an ε -optimal hypothesis using $\tilde{O}(K/\varepsilon + 1/\varepsilon^2)$ samples (suppressing dependence on $\log|\mathcal{H}|$) with bandit feedback, which would yield a significant improvement over the bound obtained by the reduction from the online setting when the number of labels is large. Such a result would show in particular that there is a fundamental gap between the PAC and online objectives in bandit multiclass classification.

1.4. Additional related work

Bandit multiclass classification. In the realizable setting, the optimal mistake bound under bandit feedback is $\tilde{O}(K \log |\mathcal{H}|)$ in the worst case. In fact, even more refined bounds exist where $\log |\mathcal{H}|$ is replaced by the *Littlestone dimension* of \mathcal{H} , which equals the optimal mistake bound in the full information setting (Auer and Long, 1999; Daniely and Helbertal, 2013; Long, 2020; Geneson, 2021). In the agnostic setting, the best known general upper bound is $\tilde{O}(\sqrt{KT \log |\mathcal{H}|})$; this bound has been further refined by substituting $\log |\mathcal{H}|$ with the Littlestone dimension of \mathcal{H} (Daniely and Helbertal, 2013), and by replacing K with the effective number of labels, defined as $\sup_{x \in \mathcal{X}} |\{h(x) : h \in \mathcal{H}\}|$ (Raman et al., 2023). These aforementioned studies introduce additional refinements in terms of the *Bandit Littlestone Dimension*, which characterizes the optimal mistake bound in the realizable setting for deterministic learners (Daniely et al., 2011). However, to the best of our knowledge, our work is the first to offer a general improvement in terms of K in front of the dominant \sqrt{T} term in the regret.

Contextual bandits. The bandit version of online multiclass classification is closely related to the contextual bandit problem that has been studied extensively in the online learning literature. The EXP4 algorithm by Auer et al. (2002) and other variants (McMahan and Streeter, 2009; Beygelzimer et al., 2011) obtain the optimal instance-independent regret bound of $O(\sqrt{KT \log |\Pi|})$ in an adversarial environment, however they are not computationally efficient if the underlying policy class is exponentially large in the natural problem parameters (e.g., the dimension of the instance space). Another line of work (Chu et al., 2011; Filippi et al., 2010; Li et al., 2017; Foster et al.,

2018; Foster and Rakhlin, 2020; Foster et al., 2020b) investigates contextual bandit problems from a different point of view, often referred to as a *realizability-based* approach. In this approach, instead of an underlying policy class of mappings from contexts to actions, there is instead a function class $\mathcal{F} \subseteq \{\mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}\}$ of mappings from context-action pairs to rewards, and it is assumed that there exists $f^* \in \mathcal{F}$ which realizes the true expected reward of the actions given the context, an assumption referred to as realizability. In this point of view, the learner’s goal is to compete against the policy π^* which selects actions according to f^* via $\pi^*(x) = \arg \max_{a \in \mathcal{A}} f^*(x, a)$. In contrast, we do not assume realizability in any form.

Log-barrier regularization. Enhancing the stability of Follow the Regularized Leader (FTRL) and Online Mirror Descent (OMD) online algorithms using log-barrier regularization has been used in various online learning scenarios. In multi-armed bandits, Wei and Luo (2018) used a log-barrier regularizer with an optimistic OMD update in order to obtain variation-dependent regret bounds, and similar techniques (Anagnostides et al., 2022, 2023) have been used in order to obtain improved regret in general sum games. Other works (Jin and Luo, 2020; Jin et al., 2021; Erez and Koren, 2021) took advantage on the iterate stability provided by the log barrier regularization in order to obtain best-of-both-worlds guarantees in RL and bandit scenarios. The use of log-barrier in order to take advantage of sparsity of the loss vectors in K -armed bandits can be seen in Bubeck et al. (2018), and in contextual bandits models a variant of log-barrier has been used in Foster et al. (2020a) in order to obtain regret bounds for infinite action sets. For the upper bound we present in Section 3 we utilize an approach most similar to that of Bubeck et al. (2018) in the sense the role of the log-barrier is to make use of action-level sparsity properties of the loss vectors.

2. Problem setup

Bandit multiclass classification. We consider a learning setup involving classification of objects from a set of examples \mathcal{X} with one of K possible labels from the set $\mathcal{Y} \triangleq \{1, \dots, K\}$. An instance of bandit multiclass classification is specified by a hypothesis class $\mathcal{H} \subseteq \{\mathcal{X} \rightarrow \mathcal{Y}\}$ of labeling functions. Our focus in this paper is on finite classes, and we denote by $N \triangleq |\mathcal{H}|$ the number of hypotheses in the class.

A learning algorithm operates in the bandit multiclass classification setting according to the following protocol, over prediction rounds $t = 1, 2, \dots$:

- (i) The environment generates an example-label pair $(x_t, y_t) \in \mathcal{X} \times \mathcal{Y}$ and x_t is presented to the algorithm;
- (ii) The algorithm classifies x_t into one of the K labels by choosing $y'_t \in \mathcal{Y}$;
- (iii) The algorithm incurs loss $\ell(y'_t; y_t) \triangleq \mathbb{1}[y'_t \neq y_t]$, where $\mathbb{1}[\cdot]$ is the indicator function;
- (iv) The algorithm observes the loss value $\ell(y'_t, y_t)$ as feedback, but *not* the true label y_t .

Note that the algorithm only observes bandit feedback at each round, namely, only whether its prediction is correct or not, rather than full feedback as in standard classification settings where the algorithm typically observes the true label y_t . We also remark the the setting describe above is that

of single-label multiclass, as opposed to the multi-label variant studied in some of the previous work (e.g., Daniely and Helbertal, 2013).²

Learning objective. The goal of the online classification algorithm is to minimize its *expected regret* over T rounds compared to any hypothesis from the class \mathcal{H} , defined by³

$$\mathcal{R}_T(\mathcal{H}) \triangleq \sup_{h^* \in \mathcal{H}} \sum_{t=1}^T \mathbb{E}[\ell(y'_t; y_t) - \ell(h^*(x_t); y_t)], \quad (1)$$

where expectations are taken with respect to any randomization present in the environment and any internal randomization in the algorithm.

Types of environments. We distinguish between two different types of environments with respect to how the example-label pairs $(x_1, y_1), \dots, (x_T, y_T)$ are generated.

- In the *stochastic setting*, it is assumed that there exists a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ such that each example-label pair (x_t, y_t) is sampled i.i.d. from \mathcal{D} .
- In the *adversarial setting*, we assume that the example-label pairs are generated by a (possibly adaptive) adversary, which chooses (x_t, y_t) based on the entire history up to round t .

3. Main algorithm and upper bounds

In this section, we describe and analyze an online algorithm which obtains the optimal regret (up to logarithmic factors) for bandit multiclass classification with finite hypothesis classes. In fact, we design the algorithm in a more general setting of contextual bandits with adversarial sparse losses; the algorithm thus applies to the adversarial variant of bandit multiclass classification and thereby also to its stochastic variant. Below, we first detail a reduction from bandit multiclass classification to sparse contextual bandits, and later describe an algorithm for the latter problem and its regret analysis.

3.1. Reduction to Sparse Contextual Bandits

Our first step is in reducing the bandit multiclass problem into an instance of a Sparse Contextual Bandits problem, defined as follows:

- The setup involves prediction over instance space \mathcal{X} of *contexts*, a finite set of *actions* \mathcal{A} with $|\mathcal{A}| = K$ and a finite *policy class* $\Pi \subseteq \{\mathcal{X} \rightarrow \mathcal{A}\}$ of size $|\Pi| = N$.
- At each round t , the environment generates a context x_t and a corresponding loss vector $\ell_t \in \mathbb{R}^K$ that assigns losses to actions. The learning algorithm is given the context x_t , chooses an action $a_t \in \mathcal{A}$ (possibly at random) and suffers loss $\ell_t(a_t)$, which is then observed by the algorithm.

2. This does not imply that labels are necessarily consistent; e.g., in a stochastic setup the label y might not be deterministic given the example x . Both the upper and lower bounds we will derive will be agnostic to this matter (that is, upper bounds will apply in the more general setting, while lower bounds will hold in the specific setting of deterministic labels).

3. The quantity defined in Eq. (1) is often referred to as *pseudo-regret* in the online learning literature.

- The goal of the algorithm is to minimize the expected regret, given by

$$\mathcal{R}_T(\Pi) \triangleq \inf_{\pi^* \in \Pi} \left\{ \mathbb{E} \left[\sum_{t=1}^T \ell_t(a_t) \right] - \sum_{t=1}^T \ell_t(\pi^*(x_t)) \right\},$$

- The instance is said to be s -sparse (with respect to the L_2 -norm), if the loss vectors satisfy $\|\ell_t\|_2^2 \leq s$ for all t .⁴

It is straightforward to frame bandit multiclass classification as a 1-sparse contextual bandits problem, albeit with *negative losses*. Instances are treated as contexts and labels as possible actions; for each incoming instance-label pair (x_t, y_t) at round t , the loss vector at round t is set to $\ell_t \in \{-1, 0\}^K$ such that $\ell_t(a) = -\mathbb{1}(a = y_t)$ for all $a \in [K]$. This particular assignment of losses renders the problem 1-sparse, since $\|\ell_t\|_2^2 = 1$ for all t . Importantly, minimizing regret in the contextual problem is equivalent to minimizing regret in the original multiclass setting, since our losses are such that $\ell_t(a) = -\mathbb{1}(a = y_t) = \mathbb{1}(a \neq y_t) - 1$ and the latter is the zero-one classification loss at step t shifted by 1 (note that shifting the losses of all actions by the same constant does not affect the regret).

3.2. Algorithm for Sparse Contextual Bandits

The reduction described above is designed so as to make the contextual online problem 1-sparse, at the cost of arriving at a problem with negative losses. In the literature on bandit problems, negative losses (or equivalently, nonnegative rewards) are notorious to be significantly more challenging technically compared to nonnegative losses (see, e.g., the discussion in [Kwon and Perchet, 2016](#)). In particular, the standard EXP4 algorithm for the contextual setting fails with negative losses since the exponential weights updates it employs become highly unstable in this case due to the loss-estimates that become negative and prohibitively large in absolute value.⁵ Drawing inspiration from related (non-contextual) bandit problems, our approach to address this is through the use of an additional log-barrier regularization that promotes stability regardless of the sign of the losses (e.g., [Wei and Luo, 2018](#); [Bubeck et al., 2018](#); [Jin and Luo, 2020](#); [Jin et al., 2021](#)).

Our algorithm for the Sparse Contextual Bandits problem is detailed in [Algorithm 1](#). The algorithm performs Follow-the-Regularized-Leader (FTRL) updates over the ε -shrunk N -dimensional probability simplex Δ_N^ε , defined by

$$\Delta_N^\varepsilon \triangleq \{p \in \Delta_N \mid p_i \geq \varepsilon \quad \forall i \in [N]\}. \quad (3)$$

for a parameter $0 \leq \varepsilon < 1/K$, where Δ_N is the probability simplex in \mathbb{R}^N containing probability vectors over the policy class $\Pi = \{\pi_1, \dots, \pi_N\}$. The FTRL regularizer $R_{\eta, \nu}(\cdot)$ in [Eq. \(2\)](#) is parameterized by $\eta, \nu > 0$ is defined by

$$R_{\eta, \nu}(p) = H_\eta(p) + \psi_\nu(p); \quad \text{where} \quad H_\eta(p) = \frac{1}{\eta} \sum_{i=1}^N p_i \log p_i, \quad \psi_\nu(p) = -\frac{1}{\nu} \sum_{i=1}^N \log p_i. \quad (4)$$

4. We refer to this as sparsity since the condition $\|\ell_t\|_2^2 \leq s$ is satisfied if ℓ_t is an s -sparse binary vector (with entries in $\{0, 1\}$ and at most s non-zero entries), but it allows for more general loss vectors that are not sparse per-se.

5. One common fix for using EXP4 with nonnegative rewards, that already appears in the original paper of [Auer et al. \(2002\)](#), is to implement mixing with a uniform distribution that prevents the loss estimates from becoming too negative. However, this mixing hinders the algorithm from leveraging sparsity in the losses, and the regret accumulated just due to the mixing is of order \sqrt{KT} ; see also the discussion in [Bubeck et al. \(2018\)](#).

Algorithm 1 Sparse Contextual Bandits

- 1: **Input:** policy class $\Pi = \{\pi_1, \dots, \pi_N\}$, parameters η, ν, ε
- 2: Initialize $p_{1,i} = 1/N$ for all $i \in [N]$
- 3: **for** rounds $t = 1, 2, 3, \dots, T$ **do**
- 4: Sample a policy $\pi_{i_t} \sim p_t$
- 5: Receive context x_t and choose action $a_t = \pi_{i_t}(x_t)$
- 6: For all $i \in [N]$ compute the importance-weighted loss estimate for policy π_i :

$$\hat{c}_{t,i} \triangleq \frac{\ell_t(a_t) \mathbb{1}[\pi_i(x_t) = a_t]}{\sum_{j=1}^N p_{t,j} \mathbb{1}[\pi_j(x_t) = a_t]}$$

- 7: Update p_t via:

$$p_{t+1} = \arg \min_{p \in \Delta_N^\varepsilon} \left\{ p \cdot \sum_{s=1}^t \hat{c}_s + R_{\eta,\nu}(p) \right\} \quad (2)$$

- 8: **end for**
-

Here $H_\eta(\cdot)$ is the *negative entropy* regularization and $\psi_\nu(\cdot)$ is the *log-barrier* regularization. We remark that allowing $\nu = \infty$ amounts to $R_{\eta,\nu}(p) = H_\eta(p)$, in which case the algorithm reduces to a version of the known EXP4 algorithm with an appropriate choice of η .

The main result of this section is given by the following theorem which provides a regret bound for Algorithm 1 in the general s -sparse contextual bandit setup, under an appropriate choice of parameters:

Theorem 1 *Let $\Pi \subseteq \{\mathcal{X} \rightarrow \mathcal{A}\}$ be a finite policy class of size N where $|\mathcal{A}| = K$, and let $T \geq 1$. Then for any s -sparse contextual bandit instance over Π , the expected regret of Algorithm 1 with $\eta = \sqrt{\log(N)/sT}$, $\nu = 1/16$ and $\varepsilon = 1/NT$ is at most*

$$O\left(N \log(NT) + \sqrt{sT \log N}\right).$$

Note that for the special case of bandit multiclass classification, setting $s = 1$ provides a regret bound of $\tilde{O}(N + \sqrt{T})$. In the regime where $N \ll \sqrt{KT}$, the bound given in Theorem 1 improves upon the guarantee of $O(\sqrt{KT \log N})$ given by the EXP4 algorithm, which does not take advantage of sparsity. Thus, we obtain the following immediate corollary which provides a regret upper bound for bandit multiclass classification:

Corollary 2 *Let $\mathcal{H} \subseteq \{\mathcal{X} \rightarrow \mathcal{Y}\}$ be a finite hypothesis class with $|\mathcal{H}| = N$ and $|\mathcal{Y}| = K \leq N$, and let $T \geq 1$. Then there exists an algorithm which, for any bandit multiclass classification instance over \mathcal{H} , guarantees an expected regret bound of at most*

$$O\left(\min\left\{\sqrt{KT \log N}, N \log(NT) + \sqrt{T \log N}\right\}\right).$$

3.3. Regret analysis

We conclude this section by sketching the proof of Theorem 1.

Proof (sketch) It suffices to bound the regret of the algorithm compared to any fixed policy in the truncated simplex, provided that ε is chosen sufficiently small ($\varepsilon = 1/NT$ is a valid choice). With that in mind, fix $p^\star \in \Delta_N^\varepsilon$. Using the fact that, at any round t , the importance-weighted loss estimators \hat{c}_t are conditionally unbiased given the randomness in previous rounds (that is, $\mathbb{E}_t[\hat{c}_t] = c_t$, where $c_t \in [-1, 0]^N$ is the loss vector induced over policies, with $c_{t,i} = \ell_t(\pi_i(x_t))$ for all i), together with a standard regret bound for FTRL (see Lemma 4 in Appendix A), we can obtain the following regret bound compared to p^\star :

$$\sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \frac{\eta}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N z_{t,i} \hat{c}_{t,i}^2 \right],$$

where z_t is a point on the line segment connecting p_t and p_{t+1} (crucially, it is a random variable that depends on the randomness at round t , and thus is *not* conditionally independent of \hat{c}_t given the history). Here, the first term comes from a bound over the magnitude of the log-barrier regularizer inside the truncated simplex Δ_N^ε , and the second term is an upper bound on the entropy component over Δ^N .

The next step in the analysis aims at removing z_t from the above bound using stability properties induced by the added log-barrier regularization. In a nutshell, using the fact that the algorithm minimizes at step t a convex objective regularized by the log-barrier ψ_ν , we can show that the iterates it produces are stable in the sense that $\|p_{t+1} - p_t\|_{\nabla^2 \psi_\nu(p_t)}^2 \leq 1/\nu$,⁶ for a suitable choice of the parameter ν (we show that setting $\nu = 1/16$ is sufficient). Since $\nabla^2 \psi_\nu(p) = \nu^{-1} \text{diag}(1/p_1^2, \dots, 1/p_N^2)$, this implies that

$$\sum_{i=1}^N \frac{(p_{t+1,i} - p_{t,i})^2}{\nu p_{t,i}^2} \leq \frac{1}{\nu} \quad \implies \quad \forall i \in [N] : \left(\frac{p_{t+1,i}}{p_{t,i}} - 1 \right)^2 \leq 1.$$

Thus, we are guaranteed that $p_{t+1,i}/p_{t,i} \leq 2$, which is the content of Lemma 5 in Appendix A. This also implies that $z_{t,i} \leq 2p_{t,i}$ for all t and i , which can be used to eliminate $z_{t,i}$ from the regret bound above, resulting in

$$\sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \eta \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N p_{t,i} \hat{c}_{t,i}^2 \right].$$

Finally, we bound the remaining second-order term in the right-hand side. For every round t and action a denote by $q_{t,a} \triangleq \sum_{i=1}^N p_{t,i} \mathbb{1}[\pi_i(x_t) = a_t]$ the probability that Algorithm 1 takes action a at round t . We then use the definition of the loss estimators together with the s -sparsity assumption, as follows:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N p_{t,i} \hat{c}_{t,i}^2 \right] &= \sum_{t=1}^T \sum_{i=1}^N \mathbb{E} \left[p_{t,i} \frac{\ell_{t,a_t}^2 \mathbb{1}[\pi_i(x_t) = a_t]}{q_{t,a_t}^2} \right] \\ &= \sum_{t=1}^T \mathbb{E} \left[\frac{\ell_{t,a_t}^2}{q_{t,a_t}^2} \sum_{i=1}^N p_{t,i} \mathbb{1}[\pi_i(x_t) = a_t] \right] \end{aligned}$$

6. The left-hand side is the ‘‘local norm’’ at p_t with respect to ψ_ν ; here we use the notation $\|x\|_A = \sqrt{x^T A x}$.

$$\begin{aligned}
 &= \sum_{t=1}^T \mathbb{E} \left[\frac{\ell_{t,a_t}^2}{q_{t,a_t}} \right] \\
 &= \sum_{t=1}^T \mathbb{E} \left[\sum_{a:q_{t,a}>0} q_{t,a} \frac{\ell_{t,a}^2}{q_{t,a}} \right] \\
 &\leq \sum_{t=1}^T \mathbb{E} [\|\ell_t\|_2^2] \leq sT.
 \end{aligned}$$

To summarize, we have the following regret bound compared to p^* :

$$\sum_{t=1}^T \mathbb{E} [c_t \cdot (p_t - p^*)] \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \eta sT,$$

where ν is set to $1/16$. Plugging in the values for ε, ν, η given in the statement of the theorem, and noting that the left-hand side is equal to the expected regret of the algorithm up to $O(\varepsilon NT)$, we conclude the proof. \blacksquare

4. Lower bound

In this section we establish a regret lower bound for bandit multiclass classification which proves that upper bound given in corollary 2 is tight, up to logarithmic factors. The lower bound is stated formally in the following theorem:

Theorem 3 *For any (possibly randomized) bandit multiclass online algorithm and for all integers $K, N, T \geq 1$, there exists a stochastic bandit multiclass instance with $K + 1$ labels over a hypothesis class \mathcal{H} of size N , where the algorithm must incur an expected regret of at least*

$$\tilde{\Omega}(\min\{N + \sqrt{T}, \sqrt{KT}\}).$$

We note that the traditional $\Omega(\sqrt{KT})$ lower bounds for K -armed bandits (Auer et al., 2002; Lattimore and Szepesvari, 2020; Slivkins, 2020), that only needs a single example/context and $N = K$ hypotheses, do not translate to our setting due to our “sparsity” constraint; namely, the sum of rewards of all labels for a given instance must be one (rather than $\Theta(K)$ as in the standard lower bound constructions). Alternatively, it is straightforward to meet the sparsity constraint by allowing N to be exponential in K , in which case the minimax rates become trivial (in this case $\log N = \Theta(K)$). The challenge in proving the lower bound is in striking a balance between these two extremes, taking advantage of multiple examples without exploding the number of hypotheses.

Let us give here a sketch of the proof, deferring the formal details to Appendix B.

Proof of Theorem 3 (sketch) Consider a bandit multiclass problem where there are a finite number of C possible examples and $K + 1$ labels. The distribution over examples is uniform, and the conditional distribution of the label given an example is designed as follows. One of the labels, $y = 0$, is a “default label” whose conditional probability is $1/3$, regardless of the example x . In addition, there is one “hidden” example-label pair, (x^*, y^*) , such that the conditional probability of the label y^* given the example x^* is $2/3$. All other example-label pairs have zero probability of appearing.

The hypothesis class is the set of functions that label all instances with the default label, except for one example that may be labeled with any label $y \neq 0$; there are precisely $N = CK$ such functions. Notice that the optimal hypothesis (which is in this class) is the one that labels x^* with the label y^* , and any other example $x \neq x^*$ with the default label $y = 0$.

For the learning algorithm that tries to compete with the optimal hypothesis, there are intuitively two different strategies to choose from. The first is to opt out of identifying the “hidden” pair (x^*, y^*) and simply choose the default label $y = 0$ for all examples, receiving a reward of $1/3$ in expectation. This strategy incurs regret compared to optimal whenever the example x^* appears (with prob. $1/C$), in which case it receives expected reward $1/3$ whereas optimal receives expected reward $2/3$; the overall regret here is therefore $\Omega(T/C)$ in expectation. The second strategy is to explore and try to identify (x^*, y^*) . Once this hidden pair is found the algorithm stops incurring regret, however the number of exploration rounds required for doing so is at least $\Omega(CK)$, and on each of them the algorithm suffers constant regret (since its reward is zero while optimal receives $1/3$); the overall regret for this strategy is therefore $\Omega(CK)$ in expectation.

Balancing these two extreme choices for the algorithm, we set $C = \Theta(\sqrt{T/K})$ in which case the algorithm suffers $\Omega(\sqrt{KT})$ regret either way. Note that crucially, the hypothesis class in this construction is therefore of size $N = CK = \Theta(\sqrt{KT})$, and in particular, polynomial in K and T . In the regime where N is smaller than $\Theta(\sqrt{KT})$, we can adjust the parameters such that the lower bound becomes $\tilde{\Omega}(N + \sqrt{T})$.

There is one issue with the construction above we neglected thus far: the conditional label distributions for examples $x \neq x^*$ do not sum to one, or in other words, it allowed for the case that some (in fact, most) examples are not labeled with any of the labels, whereas in the multiclass setting there should always be a single correct label for every example. Our formal construction and proof address this issue by allowing a small probability for all labels $y \neq 0$ given an example $x \neq x^*$, such that the sum of these probabilities is $2/3$ which therefore makes the conditional label distribution sum up to one. The analysis becomes more challenging due to this modification, since there is now a small “information leakage” about (x^*, y^*) whenever the algorithm sees an instance $x \neq x^*$ and chooses any label $y \neq 0$. Our formal argument in this case shows that this leakage is not too harmful.

Another, more minor nuisance is that the lower bound we described does not directly apply to the natural case where the true labels are determined deterministically given the examples (our construction forms a case where there is a distribution over labels for each possible example). Nevertheless, we can easily adapt our construction to use a deterministic mapping from example to label by duplicating each example many times and setting the single label of each such copy according to the label probabilities before duplication (while keeping the hypothesis class agnostic to the duplication, and thus still of size CK). ■

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreements No. 101078075; 882396). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This project has also received

funding from the Israel Science Foundation (ISF, grant numbers 2549/19; 2250/22), the Yandex Initiative for Machine Learning at Tel Aviv University, the Tel Aviv University Center for AI and Data Science (TAD), the Len Blavatnik and the Blavatnik Family foundation, and from the Adelis Foundation.

Shay Moran is a Robert J. Shillman Fellow; supported by ISF grant 1225/20, by BSF grant 2018385, by an Azrieli Faculty Fellowship, by Israel PBC-VATAT, by the Technion Center for Machine Learning and Intelligent Systems (MLIS), and by the European Union (ERC, GENERALIZATION, 101039692). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646. PMLR, 2014.
- Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. Uncoupled learning dynamics with $o(\log t)$ swap regret in multiplayer games. *Advances in Neural Information Processing Systems*, 35:3292–3304, 2022.
- Ioannis Anagnostides, Gabriele Farina, and Tuomas Sandholm. Near-optimal ϕ -regret learning in extensive-form games. In *International Conference on Machine Learning*, pages 814–839. PMLR, 2023.
- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- Peter Auer and Philip M. Long. Structural results about on-line learning models with and without queries. *Mach. Learn.*, 36(3):147–181, 1999. doi: 10.1023/A:1007614417594.
- 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.
- 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.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989.
- Sébastien Bubeck, Michael Cohen, and Yuanzhi Li. Sparsity, variance and curvature in multi-armed bandits. In *Algorithmic Learning Theory*, pages 111–127. PMLR, 2018.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.

- Thomas M Cover and Joy A Thomas. Network information theory. *Elements of information theory*, pages 374–458, 1991.
- Amit Daniely and Tom Helbertal. The price of bandit information in multiclass online classification. In *Conference on Learning Theory*, pages 93–104. PMLR, 2013.
- Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz. Multiclass learnability and the erm principle. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 207–232. JMLR Workshop and Conference Proceedings, 2011.
- Miroslav Dudík, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. Efficient optimal learning for contextual bandits. *arXiv preprint arXiv:1106.2369*, 2011.
- Liad Erez and Tomer Koren. Best-of-all-worlds bounds for online learning with feedback graphs. In *NeurIPS*, pages 28511–28521, 2021.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. *Advances in neural information processing systems*, 23, 2010.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Dylan Foster, Alekh Agarwal, Miroslav Dudík, Haipeng Luo, and Robert Schapire. Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR, 2018.
- Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in contextual bandits. *Advances in Neural Information Processing Systems*, 33:11478–11489, 2020a.
- Dylan J Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020b.
- Jesse Geneson. A note on the price of bandit feedback for mistake-bounded online learning. *Theoretical Computer Science*, 874:42–45, 2021.
- Steve Hanneke and Liu Yang. Minimax analysis of active learning. *J. Mach. Learn. Res.*, 16(1): 3487–3602, 2015.
- Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in Neural Information Processing Systems*, 33, 2020.
- Tiancheng Jin, Longbo Huang, and Haipeng Luo. The best of both worlds: stochastic and adversarial episodic mdps with unknown transition. *Advances in Neural Information Processing Systems*, 34:20491–20502, 2021.

- Sham M Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. Efficient bandit algorithms for online multiclass prediction. In *Proceedings of the 25th international conference on Machine learning*, pages 440–447, 2008.
- Joon Kwon and Vianney Perchet. Gains and losses are fundamentally different in regret minimization: The sparse case. *The Journal of Machine Learning Research*, 17(1):8106–8137, 2016.
- 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 Szepesvari. *Bandit Algorithms*. Cambridge University Press, 2020.
- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080. PMLR, 2017.
- Philip M. Long. New bounds on the price of bandit feedback for mistake-bounded online multiclass learning. *Theor. Comput. Sci.*, 808:159–163, 2020. doi: 10.1016/J.TCS.2019.11.017.
- H Brendan McMahan and Matthew Streeter. Tighter bounds for multi-armed bandits with expert advice. In *COLT*, 2009.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Ananth Raman, Vinod Raman, Unique Subedi, and Ambuj Tewari. Multiclass online learnability under bandit feedback. *arXiv preprint arXiv:2308.04620*, 2023.
- Aleksandrs Slivkins. Introduction to multi-armed bandits. *SIGecom Exch.*, 18(1):28–30, 2020.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Conference On Learning Theory*, pages 1263–1291. PMLR, 2018.

Appendix A. Proof of Section 3

Theorem 1 mainly follows from following lemma which provides a regret bound for FTRL with log-barrier regularization:

Lemma 4 *Suppose we run Algorithm 1 on arbitrary loss vectors $g_t \in \mathbb{R}^N$ where the parameters ν, η satisfy $0 < \nu \leq 1$ and $\eta > 0$. Then, for all $\varepsilon > 0$ and all $p^\star \in \Delta_N^\varepsilon$,*

$$\sum_{t=1}^T g_t \cdot (p_t - p^\star) \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \frac{\eta}{2} \sum_{t=1}^T \sum_{i=1}^N z_{t,i} g_{t,i}^2,$$

where $z_t \in [p_t, p_{t+1}]$ is some point on the line segment connecting p_t and p_{t+1} .

We also rely on the following lemma, which provides a multiplicative stability property of the FTRL iterates which follows from the log-barrier regularization. For a proof, see Section A.1.

Lemma 5 *If $\nu \leq \frac{1}{16}$, then for every round t and every $i \in [N]$ it holds that $p_{t+1,i} \leq \frac{1}{8\nu} p_{t,i}$.*

We now have what we need in order to complete the proof of Theorem 1.

Proof [Proof of Theorem 1] We prove a regret bound of $O(N \log N + \sqrt{S \log N})$ against any mixture of policies $p^\star \in \Delta_N^\varepsilon$. This will imply a regret bound of the same magnitude against the best fixed policy for the following reason: Denote by $\pi^\star = \pi_{i^\star} = \arg \max_{\pi \in \Pi} \sum_{t=1}^T \mathbb{E}[r(\pi(x_t) | x_t)]$ the benchmark policy. The regret of ALG is then given by

$$\mathcal{R}_T(\Pi) = \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - \mathbf{e}_{i^\star})],$$

where $c_{t,i} = \ell_{t,\pi_t(x_t)} \in [-1, 0]$ are the policy cost vectors, and \mathbf{e}_i denotes the i 'th standard basis vector in \mathbb{R}^N . Define $p^\star \in \Delta_N$ by

$$p^\star = \varepsilon \sum_{i \neq i^\star} \mathbf{e}_i + (1 - (N-1)\varepsilon) \mathbf{e}_{i^\star}.$$

It is straightforward to see that $p^\star \in \Delta_N^\varepsilon$, and using Hölder's inequality it holds that

$$\begin{aligned} \mathcal{R}_T(\Pi) &= \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] + \sum_{t=1}^T \mathbb{E}[c_t \cdot (p^\star - \mathbf{e}_{i^\star})] \\ &\leq \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] + T \cdot \|p^\star - \mathbf{e}_{i^\star}\|_1 \\ &= \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] + 2\varepsilon(N-1)T \\ &\leq \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] + 2. \end{aligned}$$

Thus it suffices to bound the regret of Algorithm 1 compared to any fixed $p^\star \in \Delta_N^\varepsilon$. With that in mind, fix $p^\star \in \Delta_N^\varepsilon$. Using the fact that the importance-weighted loss estimators are conditionally unbiased together with Lemma 4, we obtain the following regret bound compared to p^\star :

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] &= \sum_{t=1}^T \mathbb{E}[\mathbb{E}_t[\hat{c}_t] \cdot (p_t - p^\star)] \\ &= \sum_{t=1}^T \mathbb{E}[\hat{c}_t \cdot (p_t - p^\star)] \\ &\leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \frac{\eta}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N z_{t,i} \hat{c}_{t,i}^2 \right], \end{aligned}$$

where z_t is a point on the line segment connecting p_t and p_{t+1} . Using Lemma 5, for every $i \in [N]$ it holds that $z_{t,i} \leq 2p_{t,i}$, and we thus obtain

$$\sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \eta \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N p_{t,i} \hat{c}_{t,i}^2 \right].$$

Next, for every round t and action a denote by $q_{t,a} \triangleq \sum_{i=1}^N p_{t,i} \mathbb{1}[\pi_i(x_t) = a_t]$ the probability that Algorithm 1 performs action a in round t . To bound the last term in the above equation, we use the definition of the loss estimators together with the ℓ_2 -sparseness assumption as follows:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N p_{t,i} \hat{c}_{t,i}^2 \right] &= \sum_{t=1}^T \sum_{i=1}^N \mathbb{E} \left[p_{t,i} \frac{\ell_{t,a_t}^2 \mathbb{1}[\pi_i(x_t) = a_t]}{q_{t,a_t}^2} \right] \\ &= \sum_{t=1}^T \mathbb{E} \left[\frac{\ell_{t,a_t}^2}{q_{t,a_t}^2} \sum_{i=1}^N p_{t,i} \mathbb{1}[\pi_i(x_t) = a_t] \right] \\ &= \sum_{t=1}^T \mathbb{E} \left[\frac{\ell_{t,a_t}^2}{q_{t,a_t}} \right] \\ &= \sum_{t=1}^T \mathbb{E} \left[\sum_{a: q_{t,a} > 0} q_{t,a} \frac{\ell_{t,a}^2}{q_{t,a}} \right] \\ &\leq \sum_{t=1}^T \mathbb{E} [\|\ell_t\|_2^2] \leq sT. \end{aligned}$$

To summarize, we have the following regret bound compared to p^\star :

$$\sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] \leq \frac{N}{\nu} \log \frac{1}{\varepsilon} + \frac{1}{\eta} \log N + \eta sT,$$

and plugging in the values for ε, ν, η given in the statement of the theorem we obtain

$$\sum_{t=1}^T \mathbb{E}[c_t \cdot (p_t - p^\star)] \leq 16N \log(NT) + 2\sqrt{sT \log N},$$

which concludes the proof. ■

A.1. Proof of Lemma 4

For completeness, we provide a proof of the general regret bound for FTRL given in Lemma 4. We consider here a general FTRL framework, in which an online algorithm generates predictions $w_1, w_2, \dots, w_T \in \mathcal{W}$, given a sequence of arbitrary loss vectors g_1, g_2, \dots, g_T and a convex regularization function R , via the update rule:

$$w_t = \arg \min_{w \in \mathcal{W}} \left\{ w \cdot \sum_{s=1}^{t-1} g_s + R(w) \right\}.$$

We first prove the following general first-order regret bound for FTRL, whose proof can be found in the literature (see, e.g., [Hazan et al., 2016](#); [Orabona, 2019](#); [Lattimore and Szepesvari, 2020](#)).

Theorem 6 *There exists a sequence of points $z_t \in [w_t, w_{t+1}]$ such that, for all $w^\star \in \mathcal{W}$,*

$$\sum_{t=1}^T g_t \cdot (w_t - w^\star) \leq R(w^\star) - R(w_1) + \frac{1}{2} \sum_{t=1}^T (\|g_t\|_{z_t}^*)^2.$$

Here $\|w\|_{z_t} = \sqrt{w^\top \nabla^2 R(z_t) w}$ is the local norm induced by R at z_t , and $\|\cdot\|_{z_t}^*$ is its dual.

Proof Denote $\Phi_t(w) = w \cdot \sum_{s=1}^{t-1} g_s + R(w)$, so that $w_t = \arg \min_{w \in \mathcal{W}} \Phi_t(w)$. We first write

$$\begin{aligned} \sum_{t=1}^T g_t \cdot w_{t+1} &= \sum_{t=1}^T (\Phi_{t+1}(w_{t+1}) - \Phi_t(w_{t+1})) \\ &= \Phi_{T+1}(w_{T+1}) - \Phi_1(w_1) + \sum_{t=1}^T (\Phi_t(w_t) - \Phi_t(w_{t+1})). \end{aligned}$$

Since w_t is the minimizer of Φ_t over \mathcal{W} , first-order optimality conditions imply

$$\begin{aligned} \Phi_t(w_t) - \Phi_t(w_{t+1}) &= -\nabla \Phi_t(w_t) \cdot (w_{t+1} - w_t) - D_{\Phi_t}(w_{t+1}, w_t) \\ &\leq -D_{\Phi_t}(w_{t+1}, w_t) \\ &= -D_R(w_{t+1}, w_t), \end{aligned}$$

where $D_{\Phi}(w, w') \triangleq \Phi(w) - \Phi(w') - \nabla \Phi(w') \cdot (w - w')$ denotes the Bregman divergence with respect to Φ , and we have used the fact that the Bregman divergence is invariant to linear terms. On the other hand, since w_{T+1} is the minimizer of Φ_{T+1} , we have that

$$\sum_{t=1}^T g_t \cdot w^\star = \Phi_{T+1}(w^\star) - R(w^\star) \geq \Phi_{T+1}(w_{T+1}) - R(w^\star).$$

Combining inequalities and observing that $\Phi_1(w_1) = R(w_1)$, we obtain

$$\sum_{t=1}^T g_t \cdot (w_{t+1} - w^\star) \leq R(w^\star) - R(w_1) - \sum_{t=1}^T D_R(w_{t+1}, w_t).$$

On the other hand, a Taylor expansion of $R(\cdot)$ around w_t with an explicit second-order remainder term implies that, for some intermediate point $z_t \in [w_t, w_{t+1}]$, it holds that

$$D_R(w_{t+1}, w_t) = \frac{1}{2}(w_{t+1} - w_t)^\top \nabla^2 R(z_t) (w_{t+1} - w_t) = \frac{1}{2} \|w_{t+1} - w_t\|_{z_t}^2.$$

An application of Holder's inequality then gives

$$\begin{aligned} g_t \cdot (w_t - w_{t+1}) &\leq \|g_t\|_{z_t}^* \|w_t - w_{t+1}\|_{z_t} \\ &\leq \frac{1}{2} (\|g_t\|_{z_t}^*)^2 + \frac{1}{2} \|w_t - w_{t+1}\|_{z_t}^2 \\ &= \frac{1}{2} (\|g_t\|_{z_t}^*)^2 + D_R(w_{t+1}, w_t). \end{aligned}$$

The proof is finalized by summing over $t = 1, \dots, T$ and adding to the inequality above. \blacksquare

We can now prove Lemma 4.

Proof Using Theorem 6, the fact that $\psi_\nu(\cdot)$ is nonnegative and $H_\eta(\cdot)$ is non-positive, we obtain that for any $p^\star \in \Delta_N^\varepsilon$:

$$\sum_{t=1}^T g_t \cdot (p_t - p^\star) \leq \psi_\nu(p^\star) - H_\eta(p_1) + \frac{1}{2} \sum_{t=1}^T (\|g_t\|_{z_t}^*)^2.$$

The proof is now concluded once we make use of the fact that $\psi_\nu(\cdot) \leq \frac{N}{\nu} \log \frac{1}{\varepsilon}$ over Δ_N^ε , $H_\eta(p_1) = -\frac{1}{\eta} \log N$ and the fact that $\nabla^2 R_{\eta,\nu}(\cdot) \geq \nabla^2 H_\eta(\cdot)$. \blacksquare

A.2. Proof of Lemma 5

Finally, we provide a proof of Lemma 5 that establishes a crucial stability property of the FTRL iterates when employing log-barrier regularization. This lemma is analogous, e.g., to Lemma 12 of Jin and Luo (2020) and the proof is along similar lines. Throughout the proof, we suppress the subscripts η, ν of $R_{\eta,\nu}(\cdot)$, $H_\eta(\cdot)$ and $\psi_\nu(\cdot)$ as they are clear from the context.

Proof [Proof of Lemma 5] For any t denote

$$F_t(p) = p \cdot \sum_{s=1}^{t-1} \hat{c}_s + R(p),$$

i.e., F_t is the potential function minimized by Algorithm 1 at round t . Note that

$$\nabla^2 \psi(p) = \frac{1}{\nu} \text{diag}(p_1^2, \dots, p_N^2) \quad \forall p \in \Delta_N^\varepsilon.$$

Thus for all $p, p', p'' \in \Delta_N^\varepsilon$ it holds that

$$\|p' - p''\|_{\nabla^2 \psi(p)}^2 = \frac{1}{\nu} \sum_{i=1}^N \frac{(p'_i - p''_i)^2}{p_i^2}.$$

Denote $\alpha = \frac{1}{8\nu}$. To complete the proof, it suffices to show that $\|p_{t+1} - p_t\|_{\nabla^2 \psi(p_t)}^2 \leq \frac{(\alpha-1)^2}{\nu}$. It then suffices to show that for any $p' \in \Delta_N^\varepsilon$ with $\|p' - p_t\|_{\nabla^2 \psi(p_t)}^2 = (\alpha-1)^2/\nu$ it holds that $F_{t+1}(p') \geq$

$F_{t+1}(p_t)$, the reason being that the latter implies that p_{t+1} , which minimizes the convex function F_{t+1} , must be contained in the convex set $\{p \in \Delta_N^\varepsilon : \|p - p_t\|_{\nabla^2 \psi(p_t)}^2 \leq (\alpha - 1)^2/\nu\}$. With that in mind, fix $p' \in \Delta_N^\varepsilon$ with $\|p' - p_t\|_{\nabla^2 \psi(p_t)}^2 = (\alpha - 1)^2/\nu$. We lower bound $F_{t+1}(p')$ as follows:

$$\begin{aligned} F_{t+1}(p') &= F_{t+1}(p_t) + \nabla F_{t+1}(p_t) \cdot (p' - p_t) + \frac{1}{2} \|p' - p_t\|_{\nabla^2 R(\tilde{p})}^2 \\ &= F_{t+1}(p_t) + \nabla F_t(p_t) \cdot (p' - p_t) + \hat{c}_t \cdot (p' - p_t) + \frac{1}{2} \|p' - p_t\|_{\nabla^2 R(\tilde{p})}^2 \\ &\geq F_{t+1}(p_t) + \hat{c}_t \cdot (p' - p_t) + \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\tilde{p})}^2, \end{aligned}$$

where the first equality is a Taylor expansion of F_{t+1} about p_t , with \tilde{p} being a point on the line segment connecting p_t and p' , and the last inequality follows from a first-order optimality condition as p_t is the minimizer of F_t , and the fact that $\nabla^2 R(\tilde{p}) \geq \nabla^2 \psi(\tilde{p})$. Note that since $\|p' - p_t\|_{\nabla^2 \psi(p_t)}^2 = (\alpha - 1)^2/\nu$, using the same argument as in the start of the proof we can conclude that $p'_{t,i} \leq \alpha p_{t,i}$ for all $i \in [N]$. Since \tilde{p} lies between p_t and p' we also conclude that $\tilde{p}_i \leq \alpha p_{t,i}$ for all $i \in [N]$. Thus we can lower bound the final term as follows:

$$\begin{aligned} \frac{1}{2} \|p' - p_t\|_{\nabla^2 \psi(\tilde{p})}^2 &= \frac{1}{2\nu} \sum_{i=1}^N \frac{(p'_i - p_{t,i})^2}{\tilde{p}_i^2} \\ &\geq \frac{1}{2\nu\alpha^2} \sum_{i=1}^N \frac{(p'_i - p_{t,i})^2}{p_{t,i}^2} \\ &= \frac{1}{2\alpha^2} \|p' - p_t\|_{\nabla^2 \psi(p_t)}^2 \\ &= \frac{1}{2\nu} \left(\frac{\alpha - 1}{\alpha} \right)^2. \end{aligned}$$

To conclude the proof, we need to show that $\hat{c}_t \cdot (p' - p_t) \geq -\frac{1}{2\nu} \left(\frac{\alpha - 1}{\alpha} \right)^2$. For $a \in \mathcal{A}$ denote by $q_{t,a}$ the probability that Algorithm 1 picks action a at round t . We then have,

$$\begin{aligned} \hat{c}_t \cdot (p' - p_t) &= \frac{\ell_t(a_t)}{q_{t,a_t}} \sum_{i=1}^N (p'_{t,i} - p_{t,i}) \mathbb{1}[\pi_i(x_t) = a_t] \\ &\geq \frac{\ell_t(a_t)}{q_{t,a_t}} \sum_{i=1}^N p'_{t,i} \mathbb{1}[\pi_i(x_t) = a_t] \\ &\geq \alpha \frac{\ell_t(a_t)}{q_{t,a_t}} \sum_{i=1}^N p_{t,i} \mathbb{1}[\pi_i(x_t) = a_t] \\ &= \alpha \ell_t(a_t) \geq -\alpha, \end{aligned}$$

where the first inequality follows from the fact that the losses are non-positive, the second inequality by the fact that $p'_{t,i} \leq \alpha p_{t,i}$ and the last inequality by the fact that $\|\ell_t\|_\infty \leq 1$. The proof is concluded once we show that $\alpha \leq \frac{1}{2\nu} \left(\frac{\alpha - 1}{\alpha} \right)^2$ which clearly holds if $\alpha - 1 \geq \alpha/2$, i.e. $\alpha \geq 2$. \blacksquare

Appendix B. Proofs of Section 4

B.1. Construction of the hard instances

Consider instances with label set $\mathcal{Y} = \{0, 1, \dots, K\}$ and a finite set of examples $\mathcal{X} = \{1, 2, \dots, C\}$, where the first label $y = 0$ is a “default label” and has an expected reward of $\frac{1}{3}$ across all the instances we construct, independently of the example. The instances we consider are each labeled with a specific example-label pair (x, y) with $y \neq 0$, and are denoted by $\mathcal{I}_{x,y}$. The underlying policy class \mathcal{H} will be of size $CK + 1$, defined via

$$\mathcal{H} = \{h_0\} \cup \{h_{x,y} \mid x \in \mathcal{X}, y \in \mathcal{Y} \setminus \{0\}\},$$

where h_0 is the hypothesis which always predicts the default label $y = 0$, and $h_{x,y}$ are given by

$$h_{x,y}(x') = \begin{cases} y & x' = x, \\ 0 & x' \neq x. \end{cases}$$

Definition of $\mathcal{I}_{x,y}$. At every round t , An example x_t is sampled uniformly at random from \mathcal{X} . The reward vectors $r_t(\cdot \mid x_t)$ are sampled via

$$\begin{aligned} \text{For } x_t \neq x: \quad r_t(\cdot \mid x_t) &= \begin{cases} e_0 & \text{w.p. } \frac{1}{3}; \\ e_{y'}, \quad y' \neq 0 & \text{w.p. } \frac{2}{3K}, \end{cases} \\ \text{For } x_t = x: \quad r_t(\cdot \mid x_t) &= \begin{cases} e_0 & \text{w.p. } \frac{1}{3}; \\ e_y & \text{w.p. } \frac{2}{3} - \frac{K-1}{K^2}; \\ e_{y'}, \quad y' \notin \{0, y\} & \text{w.p. } \frac{1}{K^2}, \end{cases} \end{aligned}$$

where e_j denotes the j 'th standard basis vector in \mathbb{R}^{K+1} . We also define an additional instance \mathcal{I}_0 in which the reward vector is sampled independently of the example via

$$r_t(\cdot \mid x_t) = \begin{cases} e_0 & \text{w.p. } \frac{1}{3}; \\ e_y, \quad y \neq 0 & \text{w.p. } \frac{2}{3K}. \end{cases}$$

Our aim is to show that any online algorithm Alg must incur the desired regret lower bound on one of the $CK + 1$ instances constructed above.

B.2. Proof of Theorem 3

We turn to proving Theorem 3 based on the construction detailed above. First, let us establish some additional notation. Given a deterministic algorithm Alg we denote by $\mathbb{P}_{x,y}[\cdot]$ the probability distribution over length- T sequences of example-reward pairs induced by the instance $\mathcal{I}_{x,y}$ and the decisions of Alg , and similarly $\mathbb{P}_0[\cdot]$ for \mathcal{I}_0 . We also denote by $\mathbb{E}_{x,y}[\cdot]$ and $\mathbb{E}_0[\cdot]$ expectations taken with respect to $\mathbb{P}_{x,y}[\cdot]$ and $\mathbb{P}_0[\cdot]$ respectively. Let the random variable $N_{x,y}$ denote the number of times the example x was sampled and Alg predicted the label y . Additionally, denote the history up to round t by

$$\tau^t := (x_1, r_1, \dots, x_t, r_t),$$

where we define $\tau^0 := \emptyset$. We emphasize that the algorithm's predictions y_t are not included in the definition of a history, since y_t is deterministic given τ^{t-1} and x_t .

The first step in the proof of Theorem 3 is in establishing the following key lemma, whose proof can be found in Section B.3:

Lemma 7 *Fix a deterministic algorithm Alg. For any $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $y \neq 0$ it holds that*

$$\|\mathbb{P}_{x,y} - \mathbb{P}_0\|_1 \leq 2\sqrt{\mathbb{E}_0[N_{x,y}] + \frac{\log K}{K} \sum_{y' \notin \{0,y\}} \mathbb{E}_0[N_{x,y'}]},$$

where $\|\cdot\|_1$ denotes the total variation distance between distributions.

Informally, Lemma 7 states that for distinguishing between the instances $\mathcal{I}_{x,y}$ and \mathcal{I}_0 (equivalently, between the distributions $\mathbb{P}_{x,y}$ and \mathbb{P}_0), any algorithm would either need a constant number of samples from the reward at (x, y) , or roughly K samples of the rewards at (x, y') for $y' \neq y$ (in both cases, in expectation over \mathbb{P}_0). As our analysis will show, in each of these cases (and in any combination of the two), the regret incurred while collecting samples is significant.

Proof [Proof of Theorem 3] In the case when $N \leq \sqrt{T}$, it is straightforward to construct a 2-armed bandit instance in which Alg must incur $\Omega(\sqrt{T})$ regret. We therefore focus on the case when $N \geq \sqrt{T}$. We begin by making the observation that by Yao's principle, it suffices to consider deterministic algorithms since the instances we defined do not depend on the algorithm's decisions. With that in mind, fix a deterministic algorithm Alg and denote by G_{Alg} its total reward. Given an instance $\mathcal{I}_{x,y}$, denote by $G_{x,y}^*$ the total reward of the optimal policy in the instance $\mathcal{I}_{x,y}$. Also denote by G^* the total reward of the policy that always predicts $y = 0$, i.e., the optimal policy in the instance \mathcal{I}_0 . We separate the analysis into two cases according to the behavior of Alg. The first case is relatively straightforward, where we assume that $\mathbb{E}_0[N_0] \leq T - T/C$. In this case, we claim that the regret that Alg suffers in the instance \mathcal{I}_0 is at least $T/6C$. Indeed, we have:

$$\begin{aligned} \mathbb{E}_0[\mathcal{R}_T(\mathcal{H})] &= \mathbb{E}_0[G^* - G_{Alg}] \\ &= \frac{T}{3} - \left(\frac{1}{3} \cdot \mathbb{E}_0[N_0] + \frac{2}{3K}(T - \mathbb{E}_0[N_0]) \right) \\ &= \frac{T}{3} - \left(\frac{1}{3} - \frac{2}{3K} \right) \mathbb{E}_0[N_0] - \frac{2T}{3K} \\ &\geq \frac{T}{3} - \left(\frac{1}{3} - \frac{2}{3K} \right) \left(T - \frac{T}{C} \right) - \frac{2T}{3K} \\ &= \frac{T}{C} \left(\frac{1}{3} - \frac{2}{3K} \right) \\ &\geq \frac{T}{6C}, \end{aligned}$$

where the last inequality holds for $K \geq 4$. Setting $C = \min\{N/K, \sqrt{T/K}\}$ we obtain the desired regret lower bound for Alg in the instance \mathcal{I}_0 :

$$\mathbb{E}_0[\mathcal{R}_T(\mathcal{H})] \geq \frac{1}{6} \min\{N, \sqrt{KT}\}.$$

Assume now that $\mathbb{E}_0[N_0] \geq T - T/C$. In this case, we show that there exists $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $y \neq 0$ such that Alg suffers sufficiently large regret on $\mathcal{I}_{x,y}$. Denote by $G_{x,y}^*$ the total reward of $h_{x,y}$ which is the optimal hypothesis in the instance $\mathcal{I}_{x,y}$, and the total reward of Alg by G_{Alg} . Note that $h_{x,y}$ always predicts the default label $y' = 0$ as long as $x_t \neq x$, and otherwise predicts $y' = y$. Therefore:

$$\mathbb{E}_{x,y}[G_{x,y}^*] = \left(\frac{2}{3} - \frac{K-1}{K^2}\right)\frac{T}{C} + \left(1 - \frac{1}{C}\right)\frac{T}{3} \geq \frac{1}{2}\frac{T}{C} + \left(1 - \frac{1}{C}\right)\frac{T}{3} = \frac{T}{3} + \frac{T}{6C}, \quad (5)$$

where the inequality uses the fact that $K \geq 5$. The expected reward of the algorithm Alg in $\mathcal{I}_{x,y}$ is bounded by

$$\begin{aligned} & \mathbb{E}_{x,y}[G_{Alg}] \\ &= \frac{1}{3}\mathbb{E}_{x,y}[N_0] + \left(\frac{2}{3} - \frac{K-1}{K^2}\right)\mathbb{E}_{x,y}[N_{x,y}] + \frac{1}{K^2} \sum_{y' \neq 0} \mathbb{E}_{x,y}[N_{x,y'}] + \frac{2}{3K} \sum_{x' \neq x} \sum_{y' \neq 0} \mathbb{E}_{x,y}[N_{x',y'}] \\ &\leq \frac{1}{3}\mathbb{E}_{x,y}[N_0] + \frac{2}{3}\mathbb{E}_{x,y}[N_{x,y}] + \frac{2}{3K}(T - \mathbb{E}_{x,y}[N_0] - \mathbb{E}_{x,y}[N_{x,y}]) \\ &= \frac{2T}{3K} + \left(\frac{1}{3} - \frac{2}{3K}\right)\mathbb{E}_{x,y}[N_0] + \left(\frac{2}{3} - \frac{2}{3K}\right)\mathbb{E}_{x,y}[N_{x,y}], \end{aligned} \quad (6)$$

where in the inequality we have used the fact that $N_0 + N_{x,y} + \sum_{y' \neq 0} N_{x,y'} + \sum_{x' \neq x} \sum_{y' \neq 0} N_{x',y'} = T$. Denote by $\bar{\mathbb{P}}[\cdot]$ the distribution induced by sampling an instance among the instances $\mathcal{I}_{x,y}$ uniformly at random, i.e., $\bar{\mathbb{P}}[\cdot] = \frac{1}{CK} \sum_{x,y} \mathbb{P}_{x,y}[\cdot]$ (where the sum over y doesn't include the default label), and let $\bar{\mathbb{E}}[\cdot]$ denote the expectation with respect to this distribution. Averaging the above over the pairs (x, y) and using the fact that N_0 is bounded by T , the total reward of Alg on the average instance is bounded by

$$\begin{aligned} \bar{\mathbb{E}}[G_{Alg}] &\leq \frac{2T}{3K} + \left(\frac{1}{3} - \frac{2}{3K}\right)\frac{1}{CK} \sum_{x,y} \mathbb{E}_{x,y}[N_0] + \frac{2}{3CK} \sum_{x,y} \mathbb{E}_{x,y}[N_{x,y}] \\ &\leq \frac{2T}{3K} + \left(\frac{1}{3} - \frac{2}{3K}\right)T + \frac{2}{3CK} \sum_{x,y} \mathbb{E}_{x,y}[N_{x,y}] \\ &= \frac{T}{3} + \frac{2}{3CK} \sum_{x,y} \mathbb{E}_{x,y}[N_{x,y}]. \end{aligned}$$

We now define the random variable $\tilde{N}_{x,y}$ by

$$\tilde{N}_{x,y} = \sum_t \mathbb{E}_{x,y}[\mathbf{1}(y_t = y) \mid x_t = x, \tau^{t-1}],$$

(this is a random variable that depends on the randomness in the histories τ^{t-1}). Now, observe that $N_{x,y}$ and $\tilde{N}_{x,y}$ are related as follows:

$$\begin{aligned} \mathbb{E}_{x,y}[N_{x,y}] &= \sum_t \mathbb{P}_{x,y}[x_t = x, y_t = y] \\ &= \sum_t \mathbb{P}_{x,y}[x_t = x] \mathbb{P}_{x,y}[y_t = y \mid x_t = x] \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{C} \sum_t \mathbb{P}_{x,y}[y_t = y \mid x_t = x] \\
 &= \frac{1}{C} \sum_t \mathbb{E}_{x,y}[\mathbb{E}_{x,y}[\mathbf{1}(y_t = y) \mid x_t = x, \tau^{t-1}]] \\
 &= \frac{1}{C} \mathbb{E}_{x,y}[\tilde{N}_{x,y}].
 \end{aligned}$$

Therefore the reward of Alg on the average instance is bounded by

$$\bar{\mathbb{E}}[G_A] \leq \frac{T}{3} + \frac{2}{3C^2K} \sum_{x,y} \mathbb{E}_{x,y}[\tilde{N}_{x,y}].$$

Using the fact that $\tilde{N}_{x,y}$ are (deterministically) bounded by T , it is straightforward to show that

$$\mathbb{E}_{x,y}[\tilde{N}_{x,y}] - \mathbb{E}_0[\tilde{N}_{x,y}] \leq \frac{T}{2} \|\mathbb{P}_{x,y} - \mathbb{P}_0\|_1,$$

We can use this fact together with Lemma 7 as follows:

$$\begin{aligned}
 \bar{\mathbb{E}}[G_{Alg}] &\leq \frac{T}{3} + \frac{2}{3C^2K} \sum_{x,y} \mathbb{E}_0[\tilde{N}_{x,y}] + \frac{1}{3} \frac{T}{C^2K} \sum_{x,y} \|\mathbb{P}_{x,y} - \mathbb{P}_0\|_1 \\
 &\leq \frac{T}{3} + \frac{2T}{3CK} + \frac{2}{3} \frac{T}{C^2K} \sum_{x,y} \sqrt{\log K \cdot \mathbb{E}_0[N_{x,y}] + \frac{1}{K} \sum_{y' \neq y} \mathbb{E}_0[N_{x,y'}]} \\
 &\leq \frac{T}{3} + \frac{2T}{3CK} + \frac{2T\sqrt{\log K}}{3C^2K} \sum_{x,y} \sqrt{\mathbb{E}_0[N_{x,y}] + \frac{1}{K} \sum_{y' \neq y} \mathbb{E}_0[N_{x,y'}]}
 \end{aligned}$$

where in the second line we have also used the fact that $\sum_j \tilde{N}_{x,j} \leq T$. Using the Cauchy-Schwarz inequality, we further bound the reward of Alg by

$$\begin{aligned}
 \bar{\mathbb{E}}[G_{Alg}] &\leq \frac{T}{3} + \frac{2T}{3CK} + \frac{2T\sqrt{\log K}}{3C^{3/2}\sqrt{K}} \sqrt{\sum_{x,y} \mathbb{E}_0[N_{x,y}] + \frac{1}{K} \sum_y \sum_x \sum_{y' \neq y} \mathbb{E}_0[N_{x,y'}]} \\
 &\leq \frac{T}{3} + \frac{2T}{3CK} + \frac{4T\sqrt{\log K}}{3C^{3/2}\sqrt{K}} \sqrt{T - \mathbb{E}_0[N_0]},
 \end{aligned}$$

where in the second line we used the fact that the second summand in the square root is an average of terms bounded by $T - \mathbb{E}_0[N_0]$ each. We now use our assumption that $\mathbb{E}_0[N_0] \geq T - T/C$ to obtain

$$\bar{\mathbb{E}}[G_{Alg}] \leq \frac{T}{3} + \frac{2T}{3CK} + \frac{4T^{3/2}\sqrt{\log K}}{3C^2\sqrt{K}} \leq \frac{T}{3} + \frac{T}{9C} + \frac{4T^{3/2}\sqrt{\log K}}{3C^2\sqrt{K}},$$

where we have used the assumption that $K \geq 6$. Using the probabilistic method argument we conclude that there exists some instance labeled by a example-label pair (x, y) , $\mathcal{I}_{x,y}$, in which the total reward of Alg is bounded by

$$\mathbb{E}_{x,y}[G_{Alg}] \leq \frac{T}{3} + \frac{T}{9C} + \frac{4T^{3/2}\sqrt{\log K}}{3C^2\sqrt{K}}.$$

Putting this together with Eq. (5), we obtain the following regret lower bound for A :

$$\mathbb{E}_{x,y}[G_{x,y}^* - G_A] \geq \frac{T}{18C} - \frac{4T^{3/2}\sqrt{\log K}}{3C^2\sqrt{K}}.$$

Setting $C = 100 \min\{(N/K)\sqrt{\log K}, \sqrt{(T/K)\log K}\}$, we obtain the desired regret lower bound:

$$\mathbb{E}_0[\mathcal{R}_T(\mathcal{H})] \geq \frac{10^{-4}}{\sqrt{\log K}} \min\{N, \sqrt{KT}\}.$$

■

B.3. Proof of Lemma 7

The proof uses the following notation. Given two distributions P and Q over a common discrete domain \mathcal{W} , the *KL divergence* between P and Q is given by

$$\text{KL}(P, Q) = \sum_{w \in \mathcal{W}} P(w) \log \frac{P(w)}{Q(w)} = \mathbb{E}_{w \sim P} \left[\log \frac{P(w)}{Q(w)} \right],$$

and given a random variable Z with distribution \mathcal{D}_Z , the *conditional KL divergence* between P and Q conditioned on Z is given by

$$\text{KL}(P(\cdot | Z), Q(\cdot | Z)) = \mathbb{E}_{z \sim \mathcal{D}_Z} [\text{KL}(P(\cdot | Z = z), Q(\cdot | Z = z))].$$

Proof [Proof of Lemma 7] Using Pinsker's inequality, the squared total variation distance is bounded by the KL divergence:

$$\|\mathbb{P}_{x,y} - \mathbb{P}_0\|_1^2 \leq 2 \cdot \text{KL}(\mathbb{P}_0, \mathbb{P}_{x,y}),$$

so it suffices to provide an upper bound on the latter. In what follows, given $p, q \in (0, 1)$ we denote by $\text{KL}(p, q)$ the KL-divergence of two independent Bernoulli variables with parameters p and q respectively. Using the chain rule for relative entropy (e.g., Theorem 2.5.3 of [Cover and Thomas, 1991](#)), we have that

$$\begin{aligned} \text{KL}(\mathbb{P}_0, \mathbb{P}_{x,y}) &= \sum_t \text{KL}(\mathbb{P}_0[r_t | \tau^{t-1}, x_t], \mathbb{P}_{x,y}[r_t | \tau^{t-1}, x_t]) \\ &= \sum_t \mathbb{P}_0[x_t = x] \text{KL}(\mathbb{P}_0[r_t | \tau^{t-1}, x_t = x], \mathbb{P}_{x,y}[r_t | \tau^{t-1}, x_t = x]), \end{aligned}$$

where in the second line we used the fact that the KL divergence terms are zero unless $x_t = x$. To further simplify this expression, we note that the prediction of A at round t , y_t , is a deterministic function of τ^{t-1} and x_t , and also the distribution of the reward r_t depends only on x_t and y_t , not on the entire history before that. Therefore,

$$\text{KL}(\mathbb{P}_0, \mathbb{P}_{x,y}) = \sum_t \mathbb{P}_0[x_t = x] \text{KL}(\mathbb{P}_0[r_t | x_t = x, y_t], \mathbb{P}_{x,y}[r_t | x_t = x, y_t])$$

$$\begin{aligned}
 &= \sum_t \sum_{y'} \mathbb{P}_0[x_t = x, y_t = y'] \text{KL}(\mathbb{P}_0[r_t | x_t = x, y_t = y'], \mathbb{P}_{x,y}[r_t | x_t = x, y_t = y']) \\
 &= \sum_t \mathbb{P}_0[x_t = x, y_t = y] \text{KL}\left(\frac{2}{3K}, \frac{2}{3} - \frac{K-1}{K^2}\right) \\
 &\quad + \sum_t \sum_{y' \notin \{0,y\}} \mathbb{P}_0[x_t = x, y_t = y'] \text{KL}\left(\frac{2}{3K}, \frac{1}{K^2}\right), \tag{7}
 \end{aligned}$$

where in the second line we used the fact that the reward distribution whenever $y_t = 0$ does not depend on the instance. We now make use of the following upper bounds on the KL divergence between Bernoulli random variables:

$$\text{KL}\left(\frac{2}{3K}, \frac{2}{3} - \frac{K-1}{K^2}\right) = \frac{2}{3K} \log\left(\frac{\frac{2}{3K}}{\frac{2}{3} - \frac{K-1}{K^2}}\right) + \left(1 - \frac{2}{3K}\right) \log\left(\frac{1 - \frac{2}{3K}}{\frac{1}{3} + \frac{K-1}{K^2}}\right) \leq \log 3 \leq 2,$$

and

$$\text{KL}\left(\frac{2}{3K}, \frac{1}{K^2}\right) = \frac{2}{3K} \log\left(\frac{2K}{3}\right) + \left(1 - \frac{2}{3K}\right) \log\left(\frac{1 - \frac{2}{3K}}{1 - \frac{1}{K^2}}\right) \leq \frac{4}{3K} \log K \leq 2 \frac{\log K}{K}.$$

Plugging these two bounds into Eq. (7), we have

$$\begin{aligned}
 \text{KL}(\mathbb{P}_0, \mathbb{P}_{x,y}) &\leq 2 \sum_t \mathbb{P}_0[x_t = x, y_t = y] + 2 \frac{\log K}{K} \sum_{y' \notin \{0,y\}} \sum_t \mathbb{P}_0[x_t = x, y_t = y'] \\
 &= 2 \cdot \mathbb{E}_0[N_{x,y}] + 2 \frac{\log K}{K} \sum_{y' \notin \{0,y\}} \mathbb{E}_0[N_{x,y'}],
 \end{aligned}$$

concluding the proof. ■