

---

# Optimal and Efficient Dynamic Regret Algorithms for Non-Stationary Dueling Bandits

---

Aadirupa Saha <sup>\*1</sup> Shubham Gupta <sup>\*2</sup>

## Abstract

We study the problem of *dynamic regret minimization* in  $K$ -armed Dueling Bandits under non-stationary or time-varying preferences. This is an online learning setup where the agent chooses a pair of items at each round and observes only a relative binary ‘win-loss’ feedback for this pair sampled from an underlying preference matrix at that round. We first study the problem of static-regret minimization for adversarial preference sequences and design an efficient algorithm with  $\tilde{O}(\sqrt{KT})$  regret bound. We next use similar algorithmic ideas to propose an efficient and provably optimal algorithm for dynamic-regret minimization under two notions of non-stationarities. In particular, we show  $\tilde{O}(\sqrt{SKT})$  and  $\tilde{O}(V_T^{1/3}K^{1/3}T^{2/3})$  dynamic-regret guarantees, respectively, with  $S$  being the total number of ‘effective-switches’ in the underlying preference relations and  $V_T$  being a measure of ‘continuous-variation’ non-stationarity. These rates are provably optimal as justified with matching lower bound guarantees. Moreover, our proposed algorithms are flexible as they can be easily ‘blackboxed’ to yield dynamic regret guarantees for other notions of dueling bandits regret, including condorcet regret, best-response bounds, and Borda regret. The complexity of these problems have not been studied prior to this work despite the practicality of non-stationary environments. Extensive simulations corroborate our results.

## 1. Introduction

The problem of *Dueling-Bandits* has gained much attention in the machine learning community (Yue et al., 2012; Zoghi

---

<sup>\*</sup>Equal contribution <sup>1</sup>Microsoft Research, New York City, United States. <sup>2</sup>IBM Research, Orsay, France.. Correspondence to: Aadirupa Saha <aasa@microsoft.com>, Shubham Gupta <shubham.gupta1@ibm.com>.

et al., 2014b; 2015). This is an online learning framework that generalizes the standard multi-armed bandit (MAB) setting (Auer et al., 2002a; Agrawal & Goyal, 2012; Lattimore & Szepesvári, 2018; Audibert & Bubeck, 2010; Kalyanakrishnan et al., 2012) by querying relative preference feedback of actively chosen item pairs instead of an absolute feedback for a single item. More formally, in dueling bandits, the learning proceeds in rounds: At each round, the learner selects a pair of arms and observes stochastic preference feedback of the winner of a comparison (duel) between the selected arms. The objective of the learner is to minimize the regret with respect to a (or a set of) ‘best’ arm(s) in hindsight. Towards this, several algorithms have been proposed (Ailon et al., 2014; Zoghi et al., 2014a; Komiyama et al., 2015; Gajane et al., 2015). Due to the inherent exploration-vs-exploitation tradeoff of the learning framework and several advantages of preference feedback (Bengs et al., 2021; Yue & Joachims, 2009), many real-world applications can be modeled as dueling bandits, including movie recommendations, retail management, search engine optimization, and job scheduling, etc.

While most existing works have studied the dueling bandit problem in a stochastic setting where the underlying preferences are assumed to be fixed over time, it is often unjustified in reality. Specifically, depending on the time, season, demographics, occasions, etc., the underlying preferences may often change over time. E.g., a recommendation system must consider that a user’s preferred list of items (videos) depends on factors such as the time of the day and seasonal needs. Similarly, the relevance of a search query changes with the location and market influence. It is therefore more reasonable to assume adversarial feedback models for dueling bandits where, unlike the stochastic setting, the underlying preference matrices can vary over time. This naturally leads to the following question.

**Q.1:** *Can we design an efficient algorithm that is competitive under adversarial changes in the preference model?*

Unfortunately, except for some handful of attempts (Ailon et al., 2014; Gajane et al., 2015), which are restricted to a very special class of linear-score based preference structures (see details in Sec. 1.1), none of the existing dueling bandits algorithms can guarantee an order optimal  $O(\sqrt{T})$

regret rate against any fixed benchmark (competitor) for any arbitrary adversarial sequence of preference matrices.

More importantly, drawing inspiration from the necessity of analyzing non-stationary preference models, a more ambitious and practical regret objective is to evaluate the performance of learning algorithms in terms of how quickly they adapt to a new environment, i.e., where they are actually compared against a dynamic benchmark based on the underlying preferences instead of a single fixed benchmark considered above (see Q.1). This problem is often studied as dynamic regret minimization in standard multiarmed bandits (MAB) literature (Besbes et al., 2014; 2015; Auer et al., 2002b; Luo et al., 2018; Chen et al., 2019; Wei & Luo, 2021; Cheung et al., 2019a;b), where the learner gets to observe the absolute reward upon pulling an arm. The next, and perhaps a daunting, question to ask is

**Q.2:** *How can we design efficient and optimal dynamic regret algorithms for the dueling bandits setup?*

Surprisingly, despite dueling bandits being a significantly well studied extension of MAB, existing literature still lacks efficient algorithms for dynamic regret minimization in this framework. To the best of our knowledge, *we are the first to extend the stochastic  $K$ -armed dueling bandit problem to non-stationary preferences* (a.k.a. NST-DB).

---

<b>Parameters.</b> Item set: $[K]$
<b>For</b> $t = 1, 2, \dots, T$ , the learner:
Chooses $(k_{+1,t}, k_{-1,t}) \in [K] \times [K]$
Observes $o_t := \mathbb{I}\{k_{+1,t} > k_{-1,t}\} \sim \text{Ber}(P_t(k_{+1,t}, k_{-1,t}))$
Incurs $r_t := \frac{1}{2} (P_t(i_t^*, k_{+1,t}) + P_t(i_t^*, k_{-1,t}) - 1)$ ; $i_t^* \in [K]$

---

Figure 1: Setting of Non-Stationary DB (NST-DB)

**Problem definition (informal).** Let  $[K]$  be the item (action) set and  $\mathbf{P}_t \in [0, 1]^{K \times K}$  denote the underlying preference matrix at time  $t$ . The learner is allowed to choose two actions  $\{k_{+1,t}, k_{-1,t}\} \subseteq [K]$  at time  $t$ , upon which it observes an outcome  $o_t \sim \text{Ber}(P_t(k_{+1,t}, k_{-1,t}))$  and incurs a regret  $r_t := \frac{1}{2} (P_t(i_t^*, k_{+1,t}) + P_t(i_t^*, k_{-1,t}) - 1)$  w.r.t. any ‘good arm’  $i_t^* \in [K]$  (see Fig. 1). The goal is to provide optimal dynamic-regret algorithms under different ‘measures of non-stationarity’ (see details in Sec. 2).

**Contributions.** Our specific contributions are:

- As motivated in Q.1, we first study the ‘static dueling bandit regret’ (see Sec. 2 for definition) w.r.t. any fixed arm  $i^* \in [K]$  (or more generally any distribution  $\pi^* \in \Delta_K$ ) for adversarially chosen preferences. Our algorithm (Alg. 1) guarantees an  $\tilde{O}(\sqrt{KT})$  high probability regret bound (Thm. 3.3, Rem. 3.4, Sec. 3). This is one of our primary contributions that shows how techniques from standard EXP3 algorithm (Auer et al.,

2002b) can be borrowed to establish high-probability  $\tilde{O}(\sqrt{KT})$  regret bounds for dueling bandits against any fixed benchmark. Sec. 3.2 details the main ideas behind the analysis. The result is not only useful on its own but also serves as the basic building block for our subsequent non-stationary algorithms.

- Towards addressing the NST-DB problem (see Q.2), our second (and main) contribution lies in analyzing *Non-Stationary Dueling Bandits*. We design an algorithm (Alg. 2) with high probability dynamic-regret guarantees for two measures of non-stationarities: ‘Switching-Variation’  $S$  and ‘Continuous-Variation’  $V_T$  (Definitions in Sec. 2). More precisely, the upper bound on the dynamic regret of our proposed algorithm are shown to be  $O(\sqrt{SKT})$  and  $O(V_T^{1/3} K^{1/3} T^{2/3})$ , respectively (see Thm. 4.1, 4.4, Sec. 4).
- The regret optimality of our algorithms is justified by proving matching lower bound guarantees for dynamic-regret under both the notions of non-stationarities (Thm. 5.1, 5.2, Sec. 5).
- We further show the versatility of the proposed algorithm, which can be modularized and applied as a blackbox recipe to more general forms of dynamic regrets, including one that we analyzed above as well as Borda score based dynamic regret (Sec. 6).
- Finally, our theoretical analysis is corroborated with detailed empirical evaluations to compare the performance of the proposed techniques over state-of-the-art dueling bandit algorithms (Sec. 7).

## 1.1. Related Works

The problem of regret minimization in stochastic multi-armed bandits (MAB) is very well-studied in the online learning literature (Auer et al., 2002b; Lattimore & Szepesvári, 2018; Audibert & Bubeck, 2010) where the learner pulls a single arm per round and sees a noisy sample of the absolute reward feedback for this arm. A well-motivated generalization of MAB framework is extending it to non-stationary environments where the underlying reward model changes with time (Besbes et al., 2015; Auer et al., 2002b) and the goal of the learner is to minimize the *dynamic regret* w.r.t. the ‘best-performing’ benchmark (arm) at each round. This problem is well-studied in the online learning community for the standard  $K$ -item MAB (Besbes et al., 2014; 2015; Auer et al., 2002b), and the non-stationary contextual bandits (Chen et al., 2019; Wei & Luo, 2021) under different non-stationarities (Luo et al., 2018).

On the other hand, the relative feedback variants of the stochastic MAB problem have seen a widespread resurgence in the form of Dueling Bandits over the last two decades. Several algorithms have been proposed to address regret minimization in dueling bandits for different notions of ‘best performing’ arms depending on the underlying pref-

erence models, e.g. Condorcet, Borda, and Copeland winner (Yue et al., 2012; Szörényi et al., 2015; Zoghi et al., 2014b; 2015; Wu & Liu, 2016; Dudík et al., 2015; Saha & Gopalan, 2019a; 2020). Recent works have also extended the pairwise preference feedback model to include subset-wise preferences (Sui et al., 2017; Brost et al., 2016; Saha & Gopalan, 2018; 2019b; Ren et al., 2018).

Ailon et al. (2014) first studied the dueling bandit problem for adversarial setup and introduced the famous sparring EXP3 idea (without regret guarantees) which was later used in many follow-up works for proving regret guarantees on dueling bandits in adversarial environments (Gajane et al., 2015; Saha et al., 2021; Sui et al., 2017). However, unlike us, *they do not consider the dynamic regret objective* w.r.t. a time-varying benchmark  $i_t^*$ . The settings in Ailon et al. (2014) and (Gajane et al., 2015) are restricted to only utility based preferences where each item has a scalar utility at each time. This entails a complete ordering between the items at each step, which only covers a small subclass of the general  $[K] \times [K]$  preference matrices. Additionally, Gajane et al. (2015) assume that the feedback includes not only the winner but also the difference in the utilities between the winning and losing item, which is more similar to MAB feedback and far from our 0/1 one bit preference feedback. Saha et al. (2021) did consider the dueling bandit setup for general adversarial preferences, but they measure (static) regret in terms of *Borda-scores* for which the fundamental regret performance limit is shown to be  $\Omega(K^{1/3}T^{2/3})$ . This measure of regret is very different from our preference-based regret objective (see Sec. 2). Dudík et al. (2015) adapted sparring-EXP3 to the contextual dueling bandit setup where the learner competes against the best policy from a given policy class  $\Pi$ . Their proposed sparring-EXP4 algorithm is shown to give a  $\tilde{O}(\sqrt{KT} \log |\Pi|)$  rate, however the algorithm is computationally intractable for arbitrary context sets. *To the best of our knowledge, we are the first to design efficient and optimal dynamic regret algorithms for ‘non-stationary dueling bandits’.*

## 2. Problem Formulation

**Notations.** Let  $[n] := \{1, 2, \dots, n\}$  for any  $n \in \mathbb{N}$ . Given a set  $\mathcal{A}$ , for any  $a, b \in \mathcal{A}$ , we denote by  $a \succ b$  the event that  $a$  is preferred over  $b$ .  $\Delta_K := \{\mathbf{p} \in [0, 1]^K \mid \sum_{i=1}^K p_i = 1, p_i \geq 0, \forall i \in [K]\}$  denotes the  $K$ -simplex.

**Setting.** We assume a decision space of  $K$  arms denoted by  $\mathcal{A} := [K]$ . At each round  $t$ , the task of the learner is to select a pair of actions  $(k_{+1,t}, k_{-1,t}) \in [K] \times [K]$ , upon which a preference feedback  $o_t \sim \text{Ber}(P_t(k_{+1,t}, k_{-1,t}))$  is revealed to the learner according to the underlying preference matrix  $\mathbf{P}_t \in [0, 1]^{K \times K}$  (chosen adversarially), such that the probability of  $k_{+1,t}$  being preferred over  $k_{-1,t}$  is given by  $\mathbb{P}(o_t = 1) := \mathbb{P}(k_{+1,t} \succ k_{-1,t}) = P_t(k_{+1,t}, k_{-1,t})$ ,

and hence  $\mathbb{P}(o_t = 0) := \mathbb{P}(k_{-1,t} \succ k_{+1,t}) = 1 - P_t(k_{+1,t}, k_{-1,t})$ .

**Measures of Non-Stationarities.** We measure the dynamic regret w.r.t. the following notions of non-stationarity:

1. **Switching-Variation**  $S := \sum_{t=2}^T \mathbf{1}[\mathbf{P}_t \neq \mathbf{P}_{t-1}]$ .
2. **Continuous-Variation:**

$$V_T := \sum_{t=2}^T \max_{(a,b) \in [K] \times [K]} |P_t(a, b) - P_{t-1}(a, b)|.$$

**Objective (Static and Dynamic Regret).** Assuming the learner selects the duel  $(k_{+1,t}, k_{-1,t})$  at round  $t$ , one can measure its performance w.r.t. a single fixed arm  $i^* \in [K]$ <sup>1</sup> in hindsight by calculating the Static Regret:

$$\text{SR}_T := \max_{i^* \in [K]} \sum_{t=1}^T \frac{1}{2} (P_t(i^*, k_{+1,t}) + P_t(i^*, k_{-1,t}) - 1). \quad (1)$$

We further study a *stronger notion of dynamic regret* with respect to a time varying benchmark sequence  $j^T = (j_1, j_2, \dots, j_T)$ , defined as:

$$\text{DR}_T(j^T) := \sum_{t=1}^T \frac{1}{2} (P_t(j_t, k_{+1,t}) + P_t(j_t, k_{-1,t}) - 1). \quad (2)$$

Let  $r_t(j) = \frac{1}{2}(P_t(j, k_{+1,t}) + P_t(j, k_{-1,t}) - 1)$ . The learner competes against the ‘best-response’ of  $(k_{+1,t}, k_{-1,t})$  at time  $t$  if  $j_t = j_t^* = \arg\max_{j \in [K]} r_t(j)$ . On the other hand, it is easy to see that if we set  $j_t^* = j^* = \arg\max_{j \in [K]} \sum_{t=1}^T r_t(j)$ , we have  $\text{DR}_T(j^T) = \text{SR}_T$ . We now proceed to describe our algorithms for Static and Dynamic regret minimization.

## 3. DEX3 . P and Static Regret

This section describes our base algorithm DEX3 . P and a high-probability *Static Regret* upper bound on its regret with respect to any fixed benchmark  $i^* \in [K]$  (see  $\text{SR}_T$ , Sec. 2).

### 3.1. DEX3 . P Algorithm

Algorithm 1 presents the pseudocode for DEX3 . P. It pits two players against each other, each parallelly running its own copy of the standard EXP3 algorithm (Auer et al., 2002b). This idea was introduced as the sparring technique

<sup>1</sup>Note that this is equivalent to maximizing the expected regret w.r.t. any fixed distribution  $\pi^* \in \Delta_K$ , i.e. when  $i^* \sim \pi^*$ . This is because the regret objective is linear in the entries of  $\pi^*$ , so the maximizer  $\pi^*$  is always one hot.

**Algorithm 1** DEX3.P

---

```

1: Input:  $\eta > 0, \gamma, \beta \in (0, 1)$ 
2: Set  $W_{+1,1}(k) = W_{-1,1}(k) = 1$  for all  $k \in [K]$ 
3: for  $t = 1, \dots, T$  do
4:   for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
5:      $p_{i,t}(k) = (1 - \gamma) \frac{W_{i,t}(k)}{\sum_{j=1}^K W_{i,t}(j)} + \frac{\gamma}{K}$ 
6:   end for
7:   For  $i \in \{+1, -1\}$ , sample  $k_{i,t}$  from the distribution
   ( $p_{i,t}(1), \dots, p_{i,t}(K)$ )
8:   Observe preference feedback  $o_t(k_{+1,t}, k_{-1,t})$ 
9:   Compute  $\hat{g}_{i,t}(k)$  for  $i \in \{+1, -1\}$  and  $k \in [K]$ 
   using (3)
10:  for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
11:     $W_{i,t+1}(k) = W_{i,t}(k)e^{\eta \hat{g}_{i,t}(k)}$ 
12:  end for
13: end for

```

---

by Ailon et al. (2014), but without any regret guarantees (see Sec. 1.1). We refer to these players as the *row player* and the *column player*, and denote them by +1 and -1, respectively.

As in EXP3, these players use weights,  $\{W_{+1,t}(k)\}_{k=1}^K$  and  $\{W_{-1,t}(k)\}_{k=1}^K$ , to compute distributions over  $K$  arms,  $\mathbf{p}_{+1,t} = (p_{+1,t}(1), \dots, p_{+1,t}(K))$  and  $\mathbf{p}_{-1,t} = (p_{-1,t}(1), \dots, p_{-1,t}(K))$  (Line 5), from which their actions are sampled (Line 7). We use  $k_{+1,t}$  and  $k_{-1,t}$  to denote the arms chosen by the row and column player at time  $t$ . The players then update their weights using the observed preference feedback  $o_t(k_{+1,t}, k_{-1,t})$  that identifies the winner of a duel between the chosen arms.

The row player considers  $r_{+1,t} = o_t(k_{+1,t}, k_{-1,t})$  as its reward at time  $t$  whereas the column player uses  $r_{-1,t} = 1 - o_t(k_{+1,t}, k_{-1,t})$ . The weight update for each player is similar to EXP3 (Line 11) and uses the following expression as an importance-weighted estimate of the rewards. For  $i \in \{+1, -1\}$ ,

$$\hat{g}_{i,t}(k) = \begin{cases} (r_{i,t} + \beta)/p_{i,t}(k) & \text{if } k = k_{i,t} \\ \beta/p_{i,t}(k) & \text{otherwise.} \end{cases} \quad (3)$$

The quantity  $\beta \in (0, 1)$  in the expression above prevents  $\hat{g}_{i,t}(k)$  from becoming too large when  $p_{i,t}(k)$ 's are small. This is crucial for obtaining a high-probability regret guarantee (see proof of Lemma 3.1 in Appendix A.1).

As both players try to maximize their cumulative rewards, they start putting a higher weight on arms that are more likely to win against the choice made by the opponent. An equilibrium is reached when both players choose the optimal arm. The quantity  $\gamma \in (0, 1)$  in Line 5 ensures a minimum exploration of all arms and encourages the players to reach the optimal equilibrium strategy. Next, we establish a high-

probability static regret guarantee for DEX3.P.

### 3.2. Analysis: Static Regret Guarantee

Given the actions  $\{k_{+1,t}, k_{-1,t}\}_{t=1}^T$  taken by the players, the total regret with respect to fixed action  $j \in [K]$  is defined as:

$$\text{SR}_T(j) = \sum_{t=1}^T \frac{P_t(j, k_{+1,t}) + P_t(j, k_{-1,t}) - 1}{2}.$$

The *central idea* in our analysis is to decompose the total regret into the regret incurred by the row ( $\text{SR}_T^{+1}(j)$ ) and the column ( $\text{SR}_T^{-1}(j)$ ) player, defined as

$$\begin{aligned} \text{SR}_T^{+1}(j) &= \sum_{t=1}^T P_t(j, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t}) \\ \text{SR}_T^{-1}(j) &= \sum_{t=1}^T P_t(j, k_{+1,t}) - P_t(k_{-1,t}, k_{+1,t}). \end{aligned}$$

It is easy to see that  $\text{SR}_T(j) = \frac{1}{2} (\text{SR}_T^{+1}(j) + \text{SR}_T^{-1}(j))$ . The next two lemmas establish high-probability upper bounds on  $\text{SR}_T^{+1} := \max_{j \in [K]} \text{SR}_T^{+1}(j)$  and  $\text{SR}_T^{-1} := \max_{j \in [K]} \text{SR}_T^{-1}(j)$ .

**Lemma 3.1.** *Let  $\eta = \frac{1}{2} \sqrt{\frac{\ln K}{KT}}$ ,  $\beta = \sqrt{\frac{\ln K}{KT}}$ , and  $\gamma = 2\eta K \geq (1 + \beta)\eta K$ . With probability at least  $1 - \delta$ ,*

$$\text{SR}_T^{+1} := \max_{j \in [K]} \text{SR}_T^{+1}(j) = O(\sqrt{KT} \ln(K/\delta)).$$

**Lemma 3.2.** *Let  $\eta$ ,  $\beta$ , and  $\gamma$  be set as in Lemma 3.1. With probability at least  $1 - \delta$ ,*

$$\text{SR}_T^{-1} := \max_{j \in [K]} \text{SR}_T^{-1}(j) = O(\sqrt{KT} \ln(K/\delta))$$

The following regret bound on  $\max_{j \in [K]} \text{SR}_T(j)$  is now a straightforward consequence of the two lemmas above.

**Theorem 3.3** (Static-Regret of DEX3.P). *Set  $\eta$ ,  $\beta$ ,  $\gamma$  as in Lem. 3.1. Then with probability at least  $1 - 2\delta$*

$$\max_{j \in [K]} \text{SR}_T(j)_{[\text{DEX3.P}]} = O\left(\sqrt{KT} \ln\left(\frac{K}{\delta}\right)\right)$$

Lemmas 3.1 and 3.2 have been proved in Appendix A. *From the perspective of each agent, the actions of other agents, and hence their own rewards, are adaptively chosen by an adversary.* With this insight, one can borrow techniques from the analysis of the standard EXP3 algorithm to establish the high-probability regret bound in Thm. 4.1.

*Remark 3.4.* The regret guarantee above has been written in terms of any fixed (pure) action  $j \in [K]$  in hindsight.

**Algorithm 2** DEX3 . S


---

```

1: Input:  $\eta > 0, \alpha, \gamma, \beta \in (0, 1)$ 
2: Set  $W_{+1,1}(k) = W_{-1,1}(k) = 1$  for all  $k \in [K]$ 
3: for  $t = 1, \dots, T$  do
4:   for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
5:      $p_{i,t}(k) = (1 - \gamma) \frac{W_{i,t}(k)}{\sum_{j=1}^K W_{i,t}(j)} + \frac{\gamma}{K}$ 
6:   end for
7:   For  $i \in \{+1, -1\}$ , sample  $k_{i,t}$  from the distribution
   ( $p_{i,t}(1), \dots, p_{i,t}(K)$ )
8:   Observe preference feedback  $o_t(k_{+1,t}, k_{-1,t})$ 
9:   Compute  $\hat{g}_{i,t}(k)$  for  $i \in \{+1, -1\}$  and  $k \in [K]$ 
   using (3)
10:  for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
11:     $W_{i,t+1}(k) = W_{i,t}(k)e^{\eta \hat{g}_{i,t}(k)} + e\alpha \sum_{j=1}^K W_{i,t}(j)$ 
12:  end for
13: end for

```

---

However, the adversary can also play any fixed distribution over actions  $\pi^* \in \Delta_K$ , and the same regret guarantee will be valid against  $\pi^*$ . Consequently, this implies that our algorithm achieves an  $\tilde{O}(\sqrt{KT})$  regret w.r.t. the Nash-Equilibrium (a.k.a the Von-Neumann winner<sup>2</sup>) of the aggregated matrix  $\frac{1}{T} \sum_{t=1}^T \mathbf{P}_t$ , or even against the best response  $\pi^* := \arg \max_{\pi \in \Delta_K} \text{SR}_T(\pi)$  [DEX3 . P].

With the base algorithm in place, we now modify DEX3 . P to obtain DEX3 . S and derive high-probability dynamic regret guarantees for this algorithm.

#### 4. DEX3 . S and Dynamic Regret

Competing against a fixed action in hindsight is an interesting but modest goal. In this section, we consider the dynamic regret objective  $\text{DR}_T(j^T)$  defined with respect to any sequence of actions  $j^T = (j_1, \dots, j_T)$  (see Section 2). It is clear that  $\text{DR}_T(j^T)$  can be  $O(T)$  in the worst case unless one imposes further assumptions. Sections 4.1 and 4.2 respectively assume the ‘Switching-Variation’ and ‘Continuous-Variation’ model of non-stationarity (see Section 2) to derive sub-linear dynamic regret guarantees.

To achieve these bounds, we propose DEX3 . S, a variant of DEX3 . P. This algorithm uses the same procedure as DEX3 . P, except at Line 11, where it instead uses the following weight update rule:

$$W_{i,t+1}(k) = W_{i,t}(k)e^{\eta \hat{g}_{i,t}(k)} + e\alpha W_{i,t}.$$

Here,  $\alpha \in (0, 1)$  is a user-specified hyper-parameter and  $W_{i,t} := \sum_{j=1}^K W_{i,t}(j)$ . Algorithm 2 presents this procedure.

<sup>2</sup>Given a preference matrix  $\mathbf{P}$ , a Von-Neumann winner of  $\mathbf{P} \in [0, 1]^{K \times K}$  is defined to be a probability distribution  $\pi \in \Delta_K$ , such that  $\sum_{i=1}^n \pi(i)P(i, b) \geq 1/2, \forall b \in [K]$ . (see Eqn. 1, (Dudik et al., 2015))

ture. We next move to the analysis of DEX3 . S under the two non-stationarity models.

#### 4.1. Dynamic Regret Analysis of DEX3 . S under Switching-Variation S

As before, we write  $\text{DR}_T(j^T)$  as  $\text{DR}_T(j^T) = \frac{1}{2}(\text{DR}_T^{+1}(j^T) + \text{DR}_T^{-1}(j^T))$ , where  $\text{DR}_T^{+1}(j^T)$  and  $\text{DR}_T^{-1}(j^T)$  are the dynamic regret incurred by the row and column player with respect to the sequence  $j^T$ :

$$\begin{aligned} \text{DR}_T^{+1}(j^T) &= \sum_{t=1}^T P_t(j_t, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t}) \\ \text{DR}_T^{-1}(j^T) &= \sum_{t=1}^T P_t(j_t, k_{+1,t}) - P_t(k_{-1,t}, k_{+1,t}). \end{aligned}$$

We assume that the sequence  $j^T$  has at most  $S$  switches, i.e.,  $1 + |\{1 \leq \ell < T : j_\ell \neq j_{\ell+1}\}| \leq S$ . The result below bounds  $\text{DR}_T(j^T)$  under two conditions, first where  $S$  is known to the algorithm apriori and second where it is unknown.

**Theorem 4.1** (DEX3 . S regret under Switching-Variation). *Let  $j^T = (j_1, \dots, j_T)$  be any arbitrary sequence of actions with  $S$  switches. Then the following bounds hold for DEX3 . S with probability at least  $1 - 2\delta$  when  $\alpha = \frac{1}{T}$  and  $\gamma = 2\eta K$ :*

(1)  $\text{DR}_T(j^T)_{[\text{DEX3 . S}]} = O\left(\sqrt{SKT} \ln \frac{KT}{\delta}\right)$ , and

(2)  $\text{DR}_T(j^T)_{[\text{DEX3 . S}]} = O\left(S\sqrt{KT} \ln \frac{KT}{\delta}\right)$ , by respectively setting  $\beta = \eta = \sqrt{\frac{S}{KT}}$ , and  $\beta = \eta = \frac{1}{\sqrt{KT}}$ .

In the first case above,  $S$  is known apriori to the algorithm and is used for setting  $\beta$  and  $\eta$ . If  $S$  is unknown, the regret grows linearly with  $S$ .

*Remark 4.2.* Theorem 4.1 is valid for any sequence  $j^T$  with  $S$  switches. In particular,  $j^T$  can be constructed as follows. **(i)** Divide the time horizon  $T$  into  $S$  sub-intervals  $[T_1, \dots, T_2), [T_2, \dots, T_3), \dots, [T_S, \dots, T_{S+1})$ , where  $T_1 = 1$  and  $T_{S+1} = T + 1$ . **(ii)** Assume  $P_t(i, j) = P^s(i, j)$  for all  $t \in [T_s, T_{s+1})$ ,  $s \in [S]$ , and  $i, j \in [K]$ . **(iii)** Set  $j_{T_s} = j_{T_s+1} = \dots = j_{T_{s+1}-1} = j^s$ , where

$$j^s = \arg \max_{j \in [K]} \sum_{t=T_s}^{T_{s+1}-1} \frac{P_s(j, k_{+1,t}) + P_s(j, k_{-1,t}) - 1}{2}.$$

In other words, if the winner probabilities remain constant within sub-intervals  $[T_s, T_{s+1})$ , then  $j_{T_s}, \dots, j_{T_{s+1}-1}$  can be set to the best action within this sub-interval in hindsight for each  $s \in [S]$ . Theorem 4.1 will be valid for such a sequence  $j^T$ . Therefore, if preferences  $\mathbf{P}_1, \dots, \mathbf{P}_T$  have

Switching-Variation  $S$ , then, with probability at least  $1 - 2\delta$ ,

$$\text{DR}_T(j^1, \dots, j^S)[\text{DEX3.S}] = \begin{cases} O\left(\sqrt{SKT} \ln \frac{KT}{\delta}\right), \\ O\left(S\sqrt{KT} \ln \frac{KT}{\delta}\right), \end{cases}$$

for known and unknown  $S$  respectively.

*Remark 4.3.* As in  $\text{DEX3.P}$ , a similar regret guarantee as Theorem 4.1 would still hold for  $\text{DEX3.S}$  even when the adversary plays a fixed distribution over actions in each of the sub-intervals  $[T_s, T_{s+1})$  instead of playing a single fixed action, as explained in Rem. 3.4.

The proof of Theorem 4.1 uses the same regret decomposition strategy as was used in Theorem 3.3. Lemmas B.1 and B.2 in Appendix B.1 establish high-probability upper bounds on  $\text{DR}_T^+(j^T)$  and  $\text{DR}_T^-(j^T)$ . Theorem 4.1 then follows by combining the two lemmas.

#### 4.2. Dynamic Regret Analysis of $\text{DEX3.S}$ under Continuous-Variation $V_T$

One can decompose the dynamic regret along the same lines as in Section 4.1 to establish a dynamic regret guarantee for  $\text{DEX3.S}$  under the Continuous-Variation model of non-stationarity (Sec. 2).

**Theorem 4.4** ( $\text{DEX3.S}$  regret under Continuous-Variation). *Consider any preference sequence  $\mathbf{P}_1, \dots, \mathbf{P}_T$  with Continuous-Variation  $V_T$ . Then, setting  $\alpha, \gamma$  same as defined in Thm. 4.1 and  $\beta = \eta = \frac{V_T^{1/3}}{4K^{2/3}T^{1/3}}$ , for any time-variant benchmark  $j^T = (j_1, \dots, j_T)$ , with probability at least  $1 - \delta$ , the dynamic regret bound of  $\text{DEX3.S}$  satisfies  $\text{DR}_T(j^T)[\text{DEX3.S}] = O\left((V_T^{1/3}K^{1/3}T^{2/3} + 4K^{2/3}T^{1/3}V_T^{-1/3}) \ln \frac{KT}{\delta}\right)$ .*

The proof of Theorem 4.4 relies on the observation that for any sub-interval  $\mathcal{I} \subseteq [T]$  with limited Continuous-Variation budget, the performance difference between any fixed benchmark  $j^* \in [K]$  vs a time-varying benchmark (say a sequence of arms  $\{j_\tau^*\}_{\tau \in \mathcal{I}}$ ) is bounded. Hence, one can divide the full time horizon  $[T]$  into small sub-intervals  $\mathcal{I}_1, \dots, \mathcal{I}_S$ , each of length say  $\Delta$  (a tuning parameter to be chosen appropriately), thus  $S = \lceil T/\Delta \rceil$ , and apply the results from Thm. 4.1 to bound the learner's regret within each sub-interval against a fixed benchmark. The bound in Thm. 4.4 now follows by combining both these factors and by properly adjusting  $\Delta$  along with the algorithm parameters  $\alpha, \beta, \eta$ , and  $\gamma$ . The complete proof is given in Appendix B.2.

*Remark 4.5.* Theorems 4.1 and 4.4, together with Theorems 5.1 and 5.2 respectively, immediately justify the optimality of  $\text{DEX3.S}$  and tightness of our regret analysis in terms of Switching and Continuous Variation.

*Remark 4.6.* The regret guarantees in Theorems 3.3, 4.1, and 4.4 hold with high probability. The corresponding expected regret guarantees of the same order can be derived by setting  $\delta = 1/T$ . For example, against any sequence  $j^T$  with Switching-Variation  $S$ , Thm. 4.1 gives  $\text{SR}_T[\text{DEX3.P}] = O(\sqrt{KT} \ln KT)$  with probability  $1 - \delta$ , and  $\text{SR}_T[\text{DEX3.P}] \leq T$  with probability  $\delta$ . Thus choosing  $\delta = \frac{1}{T}$  one may conclude that  $\mathbb{E}[\text{SR}_T[\text{DEX3.P}]] = O(\sqrt{KT} \ln KT)$ .

## 5. Lower Bound

In this section, we derive a lower bound on the expected dynamic regret of any algorithm for the  $\text{NST-DB}$  problem. The high-level idea is to construct a class of problem instances  $\mathcal{V}$  and argue that there exists an instance  $\nu \in \mathcal{V}$  and a sequence  $j^T = (j_1, \dots, j_T)$  for which  $\mathbb{E}[\text{DR}_T(j^T)]$  is sufficiently large (see Theorems 5.1 and 5.2) for any algorithm.  $\mathcal{V}$  is defined such that for all instances  $\nu \in \mathcal{V}$ :

1. The time horizon  $[T]$  is divided into  $S$  sub-intervals  $[T_1, \dots, T_2), \dots, [T_S, \dots, T_{S+1})$ , where  $T_1 = 1$  and  $T_{S+1} = T + 1$ . For all  $s \in [S]$  and  $i, j \in [K]$ , the winner probabilities satisfy

$$t_1, t_2 \in [T_s, T_{s+1}) \Rightarrow P_{t_1}(i, j) = P_{t_2}(i, j).$$

2. For all  $s \in [S]$ , an arm  $j^s \neq 1$  is chosen uniformly at random as the *condorcet winner*<sup>3</sup>, i.e.,  $P_t(j^s, j) > 0.5$  for all  $j \neq j^s$  and for all  $t \in [T_s, T_{s+1})$ .
3. With  $\epsilon > 0$  as a small constant, for all  $s \in [S]$ ,  $t \in [T_s, T_{s+1})$ , and  $i < j$ , the winner probabilities are:

$$P_t(i, j) = \begin{cases} \frac{1}{2} + \epsilon & \text{if } i = j^s \vee (i = 1 \wedge j \neq j^s), \\ \frac{1}{2} - \epsilon & \text{if } j \in \{1, j^s\}, \\ \frac{1}{2} & \text{otherwise.} \end{cases} \quad (4)$$

In other words, an instance  $\nu \in \mathcal{V}$  is created by choosing a condorcet winner  $j^s \in [K] \setminus \{1\}$  for each sub-interval  $[T_s, T_{s+1})$ . In sub-interval  $[T_s, T_{s+1})$ , arm  $j^s$  beats every arm with probability  $0.5 + \epsilon$ , arm 1 beats every arm except  $j^s$  with probability  $0.5 + \epsilon$ , and all other arms beat each other with equal probability.

For an instance  $\nu \in \mathcal{V}$ , define the sequence of actions  $j^T(\nu)$  such that  $j_t(\nu) = j^s$  for all  $t \in [T_s, T_{s+1})$ . Here,  $j_t(\nu)$  is the  $t^{\text{th}}$  element of  $j^T(\nu)$  and  $j^s$  is the condorcet winner in  $\nu$  during the sub-interval  $[T_s, T_{s+1})$ . The results below establish lower bounds on  $\mathbb{E}[\text{DR}_T(j^T(\nu))]$  under Switching-Variation and Continuous-Variation model of non-stationarity.

<sup>3</sup>Any preference matrix  $\mathbf{P} \in [0, 1]^{K \times K}$  is said to have a condorcet winner  $i^* \in [K]$ , if there exists an arm  $i^* \in [K]$  s.t.  $P(i^*, j) > 0.5$  for all  $j \in [K] \setminus \{i^*\}$  (Bengs et al., 2021).

**Algorithm 3** Borda-DEX3.S

---

```

1: Input:  $\eta > 0, \alpha, \gamma, \beta \in (0, 1)$ 
2: Set  $W_{+1,1}(k) = W_{-1,1}(k) = 1$  for all  $k \in [K]$ 
3: for  $t = 1, \dots, T$  do
4:   for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
5:      $p_{i,t}(k) = (1 - \gamma) \frac{W_{i,t}(k)}{\sum_{j=1}^K W_{i,t}(j)} + \frac{\gamma}{K}$ 
6:   end for
7:   For  $i \in \{+1, -1\}$ , sample  $k_{i,t}$  from the distribution
   ( $p_{i,t}(1), \dots, p_{i,t}(K)$ )
8:   Observe preference feedback  $o_t(k_{+1,t}, k_{-1,t})$ 
9:   Compute  $s'_{i,t}(k)$  for  $i \in \{+1, -1\}$  and  $k \in [K]$  as:
       
$$s'_{i,t}(k) = \frac{\mathbf{1}(k_{i,t} = k)}{K p_{i,t}(k)} \sum_{j \in [K]} \frac{\mathbf{1}(k_{-i,t} = j) o_t(k, j)}{p_{-i,t}(j)} + \frac{\beta}{p_{i,t}(k)}$$

10:  for  $k \in [K]$  and  $i \in \{+1, -1\}$  do
11:     $W_{i,t+1}(k) = W_{i,t}(k) e^{\eta s'_{i,t}(k)} + e\alpha \sum_{j=1}^K W_{i,t}(j)$ 
12:  end for
13: end for

```

---

**Theorem 5.1** (Dynamic Regret Lower Bound for Switching-Variation). *Assume that  $K \geq 3$  and  $\epsilon \in (0, 1/4)$  in (4). With  $\mathcal{V}$  and  $j^T(\nu)$  defined above, there exists an instance  $\nu \in \mathcal{V}$  such that:*

$$\mathbb{E} [DR_T(j^T(\nu))] = \Omega(\sqrt{SKT}),$$

for any non-stationary dueling bandits algorithm.

Given a change budget  $V_T$  and the number of switches  $S$ ,  $\epsilon$  can be set to ensure that the total change in  $\mathbf{P}_t$ 's across  $T$  steps is at most  $V_T$  (see Appendix C.2). Using this value of  $\epsilon$  yields the following result:

**Theorem 5.2** (Dynamic Regret Lower Bound for Continuous-Variation). *Assume that  $K \geq 3$ . With  $\mathcal{V}$  and  $j^T(\nu)$  defined above, there exists an instance  $\nu \in \mathcal{V}$  such that:*

$$\mathbb{E} [DR_T(j^T(\nu))] = \Omega((KV_T)^{1/3} T^{2/3}),$$

for any non-stationary dueling bandits algorithm.

The proofs of Theorems 5.1 and 5.2 are given in Appendix C. The key idea is to establish a lower bound on the cumulative regret within each sub-interval  $[T_s, T_{s+1})$ . This can be done by upper-bounding the number of times arms 1 and  $j^s$  (these arms are more likely to win against others) are pulled in expectation within each sub-interval  $[T_s, T_{s+1})$ .

## 6. Dynamic Regret with Borda Scores

In this section, we briefly switch to a different regret measure that uses the so-called Borda scores. The *Borda score* of an arm  $i \in [K]$  with respect to the preference matrix  $\mathbf{P}_t$  at time  $t$  measures the average probability of arm  $i$  winning

against a randomly chosen opponent  $j \in [K] \setminus \{i\}$  at time  $t$ :

$$b_t(i) := \frac{1}{K-1} \sum_{j \neq i} P_t(i, j).$$

While the ideas presented here apply equally well to the static Borda regret objective (regret with respect to the arm with highest cumulative Borda score) that was already addressed in (Saha et al., 2021), we only focus on the more challenging dynamic Borda regret objective:

$$\text{DBR}_T(j^T) := \sum_{t=1}^T \frac{2b_t(j_t) - b_t(k_{+1,t}) - b_t(k_{-1,t})}{2}, \quad (5)$$

where, as before,  $j^T = (j_1, \dots, j_T)$  is an arbitrary sequence of arms, and  $k_{+1,t}$  and  $k_{-1,t}$ , respectively, are the arms chosen by the row and column player at time  $t$ .

Algorithm 3 describes a variant of DEX3.S that uses a different arm rewards estimate instead of  $\hat{g}_{i,t}(k)$ , but is otherwise same as DEX3.S. We call this algorithm Borda-DEX3.S. The reward estimate in Borda-DEX3.S was also used in Saha et al. (2021) where the authors studied only the static Borda regret objective. The result below establishes an upper bound on the more challenging dynamic regret  $\text{DBR}_T(j^T)$  of Alg. 3 under the switching variation model of non-stationarity. The proof uses the same regret-decomposition idea as Theorem 4.1, but differs in technical details to accommodate the change in the reward estimate. See Appendix D for details.

**Theorem 6.1** (Borda-DEX3.S regret under Switching-Variation). *Let  $j^T = (j_1, \dots, j_T)$  be any sequence of actions with  $S$  switches. Then, with probability at least  $1 - \delta$ ,  $\text{DBR}_T(j^T)[\text{Borda-DEX3.S}] = \tilde{O}(S^{1/6} K^{-1/3} T^{5/6} + S^{1/2} K^{1/3} T^{2/3})$ , when  $\eta = \left(\frac{S \ln K}{T \sqrt{2K}}\right)^{2/3}$ ,  $\beta = \frac{S^{1/3} \sqrt{\ln(2K/\delta)}}{(2\eta)^{1/4} K^{3/4} \sqrt{T}}$ ,  $\gamma = \sqrt{2\eta K}$ , and  $\alpha = \frac{1}{KT}$ .*

It is worth noting that, as in Sec. 4.2, one can also extend Thm. 6.1 to analyze the continuous variation dynamic Borda regret of Alg. 3 using similar arguments as in the proof of Thm. 4.4. We skip this analysis to avoid redundancy.

## 7. Experiments

We now present numerical results to complement our theoretical findings and compare with baselines. Section 7.1 studies static regret for DEX3.P against the best fixed action in hindsight. Section 7.2 studies dynamic regret for DEX3.S under Switching-Variation and Continuous-Variation models of non-stationarity.

**Baseline algorithms:** We compare our algorithms against two baselines: RAND and REX3. RAND independently samples arms  $k_{+1,t}$  and  $k_{-1,t}$  from a uniform distribution over




Figure 2: Comparison between DEX3.P, RAND, and REX3 in terms of static regret with respect to the best fixed action in hindsight. The shaded region represents standard deviation across 10 independent runs.

$K$  arms at each time  $t$ . REX3 (Gajane et al., 2015) uses a time-varying distribution  $\mathbf{p}_t \in \Delta_K$  and independently samples  $k_{+1,t}, k_{-1,t} \sim \mathbf{p}_t$ .

The values of parameters  $\alpha, \beta, \eta$ , and  $\gamma$  for DEX3.P and DEX3.S were set in accordance with Theorems 3.3 and 4.1 (or 4.4 as appropriate from the context), respectively. REX3 uses a single parameter  $\gamma$  and we set its value to  $\gamma = \sqrt{\frac{2K \ln K}{eT}}$ , as recommended in Gajane et al. (2015).

### 7.1. DEX3.P: Static Regret

Recall that  $P_t(i, j)$  is the probability with which arm  $i$  beats arm  $j$  at time  $t$ . We simulate an environment where these values follow a Gaussian random walk. That is, for every  $t \in [T]$  and  $i < j$ ,

$$P_{t+1}(i, j) = P_t(i, j) + \epsilon_t(i, j),$$

where  $\epsilon_t(i, j) \sim \mathcal{N}(0, 0.002)^4$ . We ensure that  $P_t(i, i) = 0.5$ ,  $P_t(j, i) = 1 - P_t(i, j)$ , and clip the entries of  $\mathbf{P}_t$  in the range  $[0, 1]$  so that they represent valid probabilities. The initial values  $P_1(i, j) \sim \text{Uniform}(0, 1)$ .

Let  $j^* = \operatorname{argmax}_{j \in [K]} \sum_{t=1}^T (P_t(j, k_{+1,t}) + P_t(j, k_{-1,t}) - 1)/2$  be the best fixed action in hindsight. Figure 2a sets  $K = 10$  and plots the static regret  $\text{SR}_T$  (with respect to fixed arm  $j^*$ ) incurred by DEX3.P, RAND, and REX3 for different values of  $T$  in the range  $10^4$  and  $10^6$ . The dotted line in the plot represents  $O(\sqrt{T})$  growth. It can be seen that  $\text{SR}_T$  grows as  $O(\sqrt{T})$  for DEX3.P and REX3. Moreover, DEX3.P outperforms REX3 when  $T$  is large. Note that  $\text{SR}_T$  can indeed be negative as  $\mathbf{P}_t$  changes over time and no fixed benchmark (arm) could be competitive against the learner’s dueling sequence  $\{(k_{+1,t}, k_{-1,t})\}_{t \in [T]}$ .

Figure 2b similarly fixes  $T = 10^6$  and plots  $\text{SR}_T$  for different values of  $K$ . The dotted line in this plot represents

<sup>4</sup> $\mathcal{N}(\mu, \sigma)$  denotes a Gaussian distribution with mean  $\mu \in \mathbb{R}$  and standard deviation  $\sigma > 0$ .




Figure 3: Dynamic regret under Switching-Variation non-stationarity with respect to sequence  $j^T$  defined in Remark 4.2. The shaded region represents standard deviation across 10 independent runs.

$O(\sqrt{K})$  growth and shows that  $\text{SR}_T$  grows as  $O(\sqrt{K})$  for DEX3.P and REX3. As  $K$  increases, RAND’s regret slightly decreases as it becomes harder to have a fixed action that performs well across time. Moreover, as expected, the gap between the performance of DEX3.P (and REX3) and RAND decreases as  $K$  becomes large (recall that  $T$  is fixed).

### 7.2. DEX3.S: Dynamic Regret

Next, we turn to DEX3.S and focus on dynamic regret under the two notions of non-stationarity (see Sec. 2).

**Switching-Variation non-stationarity:** For these experiments, we use a similar Gaussian random walk model for the probabilities  $P_t(i, j)$  as in Section 7.1, but change the values only after every  $\Delta = T/(S - 1)$  steps. That is,

$$P_{t+1}(i, j) = \begin{cases} P_t(i, j) + \epsilon_t(i, j) & \text{if } \exists c \in \mathbb{N} : t = \Delta c \\ P_t(i, j) & \text{otherwise.} \end{cases}$$

Here,  $\mathbb{N}$  is the set of all natural numbers. We sample  $\epsilon_t(i, j) \sim \mathcal{N}(0, 0.05)$  where we use a higher standard deviation as compared to Section 7.1 as the changes in  $P_t(i, j)$  happen less frequently. As before, we set  $P_t(i, i) = 0.5$ ,  $P_t(j, i) = 1 - P_t(i, j)$ , and clip all values in the range  $[0, 1]$ .

The dynamic regret is measured with respect to the sequence  $j^T$  defined in Remark 4.2. This sequence considers the best fixed action in hindsight for each of the sub-intervals of length  $\Delta$  in which the probability values remain unchanged. See Remark 4.2 for details.

Figure 3a fixes  $K = S = 10$  and plots  $\text{DR}_T(j^T)$  as a function of  $T$ . It can be verified that the growth of  $\text{DR}_T(j^T)$  is  $O(\sqrt{T})$  for DEX3.S, as established in Theorem 4.1. Moreover, DEX3.S slightly outperforms both DEX3.P and REX3 for large  $T$ . Figure 3b fixes  $S = 10$ ,  $T = 10^6$ , and plots  $\text{DR}_T(j^T)$  as a function of  $K$ , confirming that  $\text{DR}_T(j^T)$  indeed grows as  $O(\sqrt{K})$ .



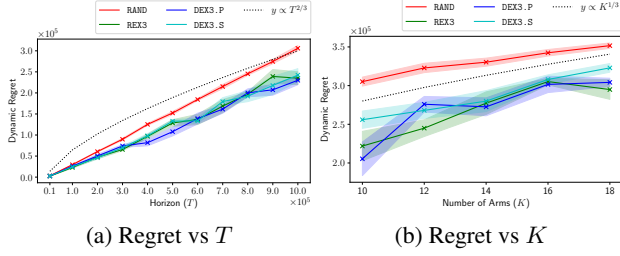


Figure 4: Dynamic regret under Continuous-Variation non-stationarity w.r.t. sequence  $j^T$  (Eqn. (6)). The shared region represents standard deviation across 10 independent runs.

**Continuous-Variation non-stationarity:** Recall that  $\Delta_T$  denotes a  $T$ -simplex and let  $\mu \in \Delta_T$ . Given Continuous-Variation  $V_T$ , we construct a sequence of probability values  $P_t(i, j)$  such that for all  $t \in [T]$

$$\max_{i, j \in [K]} |P_{t+1}(i, j) - P_t(i, j)| = \mu_t V_T.$$

Here  $\mu_t$  denotes the  $t^{\text{th}}$  element of  $\mu$ . As before, we sample  $P_1(i, j) \sim \text{Uniform}(0, 1)$  and ensure that  $P_t(i, i) = 0.5$ ,  $P_t(j, i) = 1 - P_t(i, j)$ , and  $P_t(i, j) \in [0, 1]$  for all  $t$ . For the plots in Figure 4, we sample  $\mu \sim \text{Dirichlet}(1, \dots, 1)$ . The dynamic regret is measured with respect to the best action at each time step, i.e.,  $j^T = (j_1, \dots, j_T)$  is such that

$$j_t = \operatorname{argmax}_{j \in [K]} \frac{1}{2} (P_t(j, k_{+1, t}) + P_t(j, k_{-1, t}) - 1). \quad (6)$$

Figures 4a and 4b show that  $\text{DR}_T(j^T)$  grows as  $O(T^{2/3})$  and  $O(K^{1/3})$ , respectively, as established in Theorem 4.4. Figure 4a uses  $K = V_T = 10$  and Figure 4b uses  $V_T = 10$  and  $T = 10^6$ . Finally, we would like to emphasize that while the dynamic regret for REX3 is close to DEX3.S in our experiments, theoretically REX3 only comes with a static regret guarantee, whereas we have worst-case dynamic regret guarantees for DEX3.S.

## 8. Discussions

We analyzed the complexity of regret minimization in non-stationary dueling bandits. We first proposed an  $\tilde{O}(\sqrt{KT})$  adversarial dueling bandit algorithm which is shown to be provably competitive against any fixed benchmark (arm-distribution). This algorithm was further extended to yield  $\tilde{O}(\sqrt{SKT})$  and  $\tilde{O}(V_T^{1/3} K^{1/3} T^{2/3})$  dynamic-regret algorithms for Switching-Variation ( $S$ ) and Continuous-Variation ( $V_T$ ) non-stationarities. The optimality and effectiveness of our proposed techniques were justified through matching lower bound analysis and extensive empirical evaluations. It was also demonstrated that the ideas in the proposed algorithm apply to a different, Borda score based,

regret objective. To the best of our knowledge, we are the first to analyze dynamic regret guarantees for dueling bandits in non-stationary environments.

**Future Works.** The potential future scope of this work is enormous, this being a nearly unexplored direction in preference bandits literature. An obvious next step is to formulate and study the  $\text{NSt-DB}$  problem complexity for other meaningful definitions of non-stationarities. Extending the results to contextual scenarios and subsetwise preference feedback setup would also be of much practical importance. Another useful generalization could be to make the algorithms delay tolerant where the preference feedback are revealed only after a certain unknown delay (Grover et al., 2018; Segal et al., 2021) or analyzing the dynamic regret with paid observations (Seldin et al., 2014).

## Acknowledgement

We thank the anonymous reviewers for their insightful suggestions to improve the paper.

## References

- Agrawal, S. and Goyal, N. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pp. 39–1, 2012.
- Ailon, N., Karnin, Z. S., and Joachims, T. Reducing dueling bandits to cardinal bandits. In *ICML*, volume 32, pp. 856–864, 2014.
- Audibert, J.-Y. and Bubeck, S. Best arm identification in multi-armed bandits. In *COLT-23th Conference on Learning Theory-2010*, pp. 13–p, 2010.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002a.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The nonstochastic multiarmed bandit problem. *SIAM Journal of Computing*, 32(1):48–77, 2002b.
- Bengs, V., Busa-Fekete, R., El Mesaoudi-Paul, A., and Hüllermeier, E. Preference-based online learning with dueling bandits: A survey. *J. Mach. Learn. Res.*, 22:7–1, 2021.
- Besbes, O., Gur, Y., and Zeevi, A. Stochastic multi-armed-bandit problem with non-stationary rewards. *Advances in Neural Information Processing Systems*, 27:199–207, 2014.
- Besbes, O., Gur, Y., and Zeevi, A. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.

- Brost, B., Seldin, Y., Cox, I. J., and Lioma, C. Multi-dueling bandits and their application to online ranker evaluation. *CoRR*, abs/1608.06253, 2016.
- Chen, Y., Lee, C.-W., Luo, H., and Wei, C.-Y. A new algorithm for non-stationary contextual bandits: Efficient, optimal, and parameter-free. In *Proceedings of the 32nd Conference on Learning Theory*, 99:1–30, 2019.
- Cheung, W. C., Simchi-Levi, D., and Zhu, R. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1079–1087. PMLR, 2019a.
- Cheung, W. C., Simchi-Levi, D., and Zhu, R. Non-stationary reinforcement learning: The blessing of (more) optimism. Available at SSRN 3397818, 2019b.
- Dudík, M., Hofmann, K., Schapire, R. E., Slivkins, A., and Zoghi, M. Contextual dueling bandits. In *Conference on Learning Theory*, pp. 563–587, 2015.
- Dudik, M., Hofmann, K., Schapire, R. E., Slivkins, A., and Zoghi, M. Contextual dueling bandits. *Conference on Learning Theory*, pp. 563–587, 2015.
- Gajane, P., Urvoy, T., and Clérot, F. A relative exponential weighing algorithm for adversarial utility-based dueling bandits. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 218–227, 2015.
- Grover, A., Markov, T., Attia, P., Jin, N., Perkins, N., Cheong, B., Chen, M., Yang, Z., Harris, S., Chueh, W., et al. Best arm identification in multi-armed bandits with delayed feedback. In *International Conference on Artificial Intelligence and Statistics*, pp. 833–842. PMLR, 2018.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. PAC subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pp. 655–662, 2012.
- Komiyama, J., Honda, J., Kashima, H., and Nakagawa, H. Regret lower bound and optimal algorithm in dueling bandit problem. In *COLT*, pp. 1141–1154, 2015.
- Lattimore, T. and Szepesvári, C. Bandit algorithms. *preprint*, 2018.
- Luo, H., Wei, C.-Y., Agarwal, A., and Langford, J. Efficient contextual bandits in non-stationary worlds. In *Proceedings of the 31st Conference On Learning Theory*, 75:1739–1776, 2018.
- Ren, W., Liu, J., and Shroff, N. B. PAC ranking from pairwise and listwise queries: Lower bounds and upper bounds. *arXiv preprint arXiv:1806.02970*, 2018.
- Saha, A. and Gopalan, A. Active ranking with subset-wise preferences. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- Saha, A. and Gopalan, A. Combinatorial bandits with relative feedback. In *Advances in Neural Information Processing Systems*, 2019a.
- Saha, A. and Gopalan, A. PAC Battling Bandits in the Plackett-Luce Model. In *Algorithmic Learning Theory*, pp. 700–737, 2019b.
- Saha, A. and Gopalan, A. Best-item learning in random utility models with subset choices. In *International Conference on Artificial Intelligence and Statistics*, pp. 4281–4291. PMLR, 2020.
- Saha, A., Koren, T., and Mansour, Y. Adversarial dueling bandits. In *International Conference on Machine Learning*, pp. 9235–9244. PMLR, 2021.
- Segal, S., Koren, T., Mansour, Y., et al. Stochastic multi-armed bandits with unrestricted delay distributions. In *International Conference on Machine Learning*, 2021.
- Seldin, Y., Bartlett, P., Crammer, K., and Abbasi-Yadkori, Y. Prediction with limited advice and multiarmed bandits with paid observations. In *International Conference on Machine Learning*, pp. 280–287. PMLR, 2014.
- Sui, Y., Zhuang, V., Burdick, J., and Yue, Y. Multi-dueling bandits with dependent arms. In *Conference on Uncertainty in Artificial Intelligence*, UAI’17, 2017.
- Szörényi, B., Busa-Fekete, R., Paul, A., and Hüllermeier, E. Online rank elicitation for plackett-luce: A dueling bandits approach. In *Advances in Neural Information Processing Systems*, pp. 604–612, 2015.
- Wei, C.-Y. and Luo, H. Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. In *Proceedings of the 32nd International Conference on Learning Theory*, 2021.
- Wu, H. and Liu, X. Double Thompson sampling for dueling bandits. In *Advances in Neural Information Processing Systems*, pp. 649–657, 2016.
- Yue, Y. and Joachims, T. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 1201–1208. ACM, 2009.
- Yue, Y., Broder, J., Kleinberg, R., and Joachims, T. The  $k$ -armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.

- Zoghi, M., Whiteson, S., Munos, R., Rijke, M. d., et al. Relative upper confidence bound for the  $k$ -armed dueling bandit problem. In *JMLR Workshop and Conference Proceedings*, number 32, pp. 10–18. JMLR, 2014a.
- Zoghi, M., Whiteson, S. A., De Rijke, M., and Munos, R. Relative confidence sampling for efficient on-line ranker evaluation. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pp. 73–82. ACM, 2014b.
- Zoghi, M., Karnin, Z. S., Whiteson, S., and De Rijke, M. Copeland dueling bandits. In *Advances in Neural Information Processing Systems*, pp. 307–315, 2015.

## Supplementary Material: Optimal and Efficient Dynamic Regret Algorithms for Non-Stationary Dueling Bandits

### A. Missing details from the Static Regret Analysis of DEX3 . P (Section 4)

The proof of Lemmas 3.1 and 3.2 use similar arguments. Therefore, we only present the proof of Lemma 3.1 in Section A.1. Let us begin with a technical result.

**Lemma A.1.** *For any  $\delta > 0$ ,  $\beta \in (0, 1)$ , with probability at least  $1 - \delta/2$ ,*

$$\forall k \in [K], \quad \sum_{t=1}^T \hat{g}_{+1,t}(k) \geq \sum_{t=1}^T P_t(k, k_{-1,t}) - \frac{\ln(2K/\delta)}{\beta}.$$

*Proof.* Let  $\mathcal{F}_{t-1}^{+1} = \sigma(\{\mathbf{p}_{i,s}, k_{i,s}, o_s\}_{s < t, i \in \{+1, -1\}} \cup \{\mathbf{p}_{+1,t}, \mathbf{p}_{-1,t}, k_{-1,t}\})$  be the  $\sigma$ -algebra generated by using all information till time  $t$ , except the choice of the row player's action  $k_{+1,t}$  and the feedback from the environment  $o_t(k_{+1,t}, k_{-1,t})$  at time  $t$ . Note that

$$\mathbb{E}[\hat{g}_{+1,t}(k) | \mathcal{F}_{t-1}^{+1}] = P_t(k, k_{-1,t}) + \frac{\beta}{p_{+1,t}(k)}. \quad (7)$$

By Markov's inequality,

$$\begin{aligned} \mathbb{P}\left(e^{-\beta \sum_{t=1}^T (\hat{g}_{+1,t}(k) - P_t(k, k_{-1,t}))} \leq \frac{2}{\delta}\right) &\geq 1 - \mathbb{P}\left(e^{-\beta \sum_{t=1}^T (\hat{g}_{+1,t}(k) - P_t(k, k_{-1,t}))} \geq \frac{2}{\delta}\right) \\ &\geq 1 - \frac{\delta}{2} \mathbb{E}\left[\exp\left(\beta \sum_{t=1}^T (P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))\right)\right]. \end{aligned}$$

We only need to show that  $\mathbb{E}\left[\exp\left(\beta \sum_{t=1}^T (P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))\right)\right] \leq 1$  to finish the proof. Consider the quantity  $\mathbb{E}\left[\exp(\beta(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))) | \mathcal{F}_{t-1}^{+1}\right]$ .

$$\begin{aligned} \mathbb{E}\left[e^{\beta(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))} | \mathcal{F}_{t-1}^{+1}\right] &= \mathbb{E}\left[e^{\beta(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k) + \beta/p_{+1,t}(k))} \Big| \mathcal{F}_{t-1}^{+1}\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\ &\stackrel{(a)}{\leq} \mathbb{E}\left[1 + \beta\left(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k) + \frac{\beta}{p_{+1,t}(k)}\right) + \beta^2\left(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k) + \frac{\beta}{p_{+1,t}(k)}\right)^2 \Big| \mathcal{F}_{t-1}^{+1}\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\ &\stackrel{(b)}{=} \left[1 + \beta^2 \mathbb{E}\left[\left(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k) + \frac{\beta}{p_{+1,t}(k)}\right)^2 \Big| \mathcal{F}_{t-1}^{+1}\right]\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\ &\stackrel{(c)}{=} \left[1 + \beta^2 \text{Var}(\hat{g}_{+1,t}(k) | \mathcal{F}_{t-1}^{+1})\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\ &= \left[1 + \beta^2 \text{Var}\left(\frac{o_t(k, k_{-1,t}) \mathbf{1}\{k_{+1,t} = k\}}{p_{+1,t}(k)} \Big| \mathcal{F}_{t-1}^{+1}\right)\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\ &\leq \left[1 + \beta^2 \mathbb{E}\left[\left(\frac{o_t(k, k_{-1,t}) \mathbf{1}\{k_{+1,t} = k\}}{p_{+1,t}(k)}\right)^2 \Big| \mathcal{F}_{t-1}^{+1}\right]\right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \end{aligned}$$

$$\begin{aligned}
 &= \left[ 1 + \frac{\beta^2}{p_{+1,t}^2(k)} \mathbb{E} [o_t(k, k_{-1,t}) \mathbf{1}\{k_{+1,t} = k\} | \mathcal{F}_{t-1}^{+1}] \right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\
 &\stackrel{(d)}{\leq} \left[ 1 + \frac{\beta^2}{p_{+1,t}(k)} \right] e^{-\frac{\beta^2}{p_{+1,t}(k)}} \\
 &\stackrel{(e)}{\leq} 1.
 \end{aligned}$$

Here, (a) follows as  $\beta(P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k) + \beta/p_{+1,t}(k)) \leq 1$  because  $\hat{g}_{+1,t}(k) \geq \beta/p_{+1,t}(k)$  by (3). Thus, we can use the identity  $e^x \leq 1 + x + x^2$  for  $x \leq 1$ . (b) and (c) follow from (7), (d) uses the fact that  $o_t(i, j) \leq 1$ , and (e) follows as  $1 + x \leq e^x$ . Now, by induction,

$$\begin{aligned}
 \mathbb{E} \left[ e^{\beta \sum_{t=1}^T (P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))} \right] &= \mathbb{E} \left[ \mathbb{E} \left[ e^{\beta(P_T(k, k_{-1,T}) - \hat{g}_{+1,T}(k))} \middle| \mathcal{F}_{T-1}^{+1} \right] e^{\beta \sum_{t=1}^{T-1} (P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))} \right] \\
 &\leq \mathbb{E} \left[ e^{\beta \sum_{t=1}^{T-1} (P_t(k, k_{-1,t}) - \hat{g}_{+1,t}(k))} \right] \leq \dots \leq 1.
 \end{aligned}$$

Setting  $\delta = \delta/K$  and taking a union bound over all  $k \in [K]$  finishes the proof.  $\square$

### A.1. Proof of Lemma 3.1

**Lemma 3.1.** Let  $\eta = \frac{1}{2} \sqrt{\frac{\ln K}{KT}}$ ,  $\beta = \sqrt{\frac{\ln K}{KT}}$ , and  $\gamma = 2\eta K \geq (1 + \beta)\eta K$ . With probability at least  $1 - \delta$ ,

$$SR_T^{+1} := \max_{j \in [K]} SR_T^{+1}(j) = O(\sqrt{KT} \ln(K/\delta)).$$

*Proof.* Recall from Algorithm 1 that  $W_{+1,t}(k) = \exp(\eta \sum_{s=1}^{t-1} \hat{g}_{+1,s}(k))$  and define  $W_{+1,t} = \sum_{k=1}^K W_{+1,t}(k)$ . Then,

$$\begin{aligned}
 W_{+1,t+1} &= \sum_{k=1}^K W_{+1,t+1}(k) \\
 &= \sum_{k=1}^K W_{+1,t}(k) \exp(\eta \hat{g}_{+1,t}(k)) \\
 &= W_{+1,t} \sum_{k=1}^K \frac{W_{+1,t}(k)}{W_{+1,t}} \exp(\eta \hat{g}_{+1,t}(k)) \\
 &\stackrel{(a)}{=} W_{+1,t} \sum_{k=1}^K \frac{p_{+1,t}(k) - \gamma/K}{1 - \gamma} \exp(\eta \hat{g}_{+1,t}(k)) \\
 &\stackrel{(b)}{\leq} W_{+1,t} \sum_{k=1}^K \frac{p_{+1,t}(k) - \gamma/K}{1 - \gamma} (1 + \eta \hat{g}_{+1,t}(k) + \eta^2 \hat{g}_{+1,t}^2(k)) \\
 &\leq W_{+1,t} \left( 1 + \frac{\eta}{1 - \gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta^2}{1 - \gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k) \right) \\
 &\stackrel{(c)}{\leq} W_{+1,t} \exp \left( \frac{\eta}{1 - \gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta^2}{1 - \gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k) \right)
 \end{aligned}$$

Here, (a) follows from the definition of  $p_{+1,t}(k)$  in Algorithm 1. Note that, by assumption on  $\eta$ ,  $\beta$ ,  $\gamma$ , and  $K$

$$\eta \hat{g}_{+1,t}(k) \leq \eta \frac{1 + \beta}{p_{+1,t}(k)} \leq \eta K \frac{1 + \beta}{\gamma} \leq 1.$$

Thus, (b) follows from the inequality  $e^x \leq 1 + x + x^2$  if  $x \leq 1$ . (c) follows from the inequality  $1 + x \leq e^x$ . We now have,

$$\ln \frac{W_{+1,t+1}}{W_{+1,t}} \leq \frac{\eta}{1-\gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta^2}{1-\gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k).$$

Summing over  $t = 1, 2, \dots, T$ , we get

$$\ln \frac{W_{+1,T+1}}{W_{+1,1}} \leq \frac{\eta}{1-\gamma} \sum_{t=1}^T \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta^2}{1-\gamma} \sum_{t=1}^T \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k). \quad (8)$$

Note that,  $W_{+1,1} = K$  (see Algorithm 1). For any action  $j \in [K]$ ,

$$\ln \frac{W_{+1,T+1}}{W_{+1,1}} \geq \ln W_{+1,T+1}(j) - \ln K = \eta \sum_{t=1}^T \hat{g}_{+1,t}(j) - \ln K \quad (9)$$

Combining (8) and (9) gives

$$\sum_{t=1}^T \hat{g}_{+1,t}(j) \leq \frac{1}{1-\gamma} \sum_{t=1}^T \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta}{1-\gamma} \sum_{t=1}^T \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k) + \frac{\ln K}{\eta}.$$

Now consider the sum  $\sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k)$ .

$$\sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) = \sum_{k=1}^K p_{+1,t}(k) \frac{o_t(k, k_{-1,t}) \mathbf{1}\{k_{+1,t} = k\} + \beta}{p_{+1,t}(k)} = o_t(k_{+1,t}, k_{-1,t}) + K\beta. \quad (10)$$

As  $p_{+1,t}(k) \hat{g}_{+1,t}(k) \leq 1 + \beta$ , we get

$$\begin{aligned} (1-\gamma) \sum_{t=1}^T \hat{g}_{+1,t}(j) &\leq \beta KT + \sum_{t=1}^T o_t(k_{+1,t}, k_{-1,t}) + \eta(1+\beta) \sum_{t=1}^T \sum_{k=1}^K \hat{g}_{+1,t}(k) + \frac{\ln K}{\eta} \\ &\leq \beta KT + \sum_{t=1}^T o_t(k_{+1,t}, k_{-1,t}) + \eta(1+\beta) K \max_{k \in [K]} \sum_{t=1}^T \hat{g}_{+1,t}(k) + \frac{\ln K}{\eta} \\ &\leq \beta KT + \sum_{t=1}^T o_t(k_{+1,t}, k_{-1,t}) + \gamma \max_{k \in [K]} \sum_{t=1}^T \hat{g}_{+1,t}(k) + \frac{\ln K}{\eta} \end{aligned}$$

The last inequality uses the fact that  $\gamma \geq \eta(1+\beta)K$  by assumption. As the inequality holds for any arbitrary  $j \in [K]$ , we can take a max over  $j \in [K]$  on the left hand side. Simplifying this further yields:

$$(1-2\gamma) \max_{j \in [K]} \sum_{t=1}^T \hat{g}_{+1,t}(j) \leq \beta KT + \sum_{t=1}^T o_t(k_{+1,t}, k_{-1,t}) + \frac{\ln K}{\eta}.$$

By Azuma-Hoeffding's inequality, with probability at least  $1 - \delta/2$ ,

$$\sum_{t=1}^T o_t(k_{+1,t}, k_{-1,t}) \leq \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) + 2\sqrt{2T \ln(2/\delta)}.$$

Combining this with Lemma A.1, we get with probability at least  $1 - \delta$ ,

$$(1-2\gamma) \max_{j \in [K]} \left[ \sum_{t=1}^T P_t(j, k_{-1,t}) - \frac{\ln(2K/\delta)}{\beta} \right] \leq \beta KT + \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) + 2\sqrt{2T \ln(2/\delta)} + \frac{\ln K}{\eta}.$$

Let  $\gamma = 2\eta K \geq \eta(1+\beta)K$ . Then,

$$\max_{j \in [K]} \sum_{t=1}^T P_t(j, k_{-1,t}) - \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) \leq \beta KT + 2\sqrt{2T \ln(2/\delta)} + \frac{\ln K}{\eta} + \frac{\ln(2K/\delta)}{\beta} + 4\eta KT.$$

Setting  $\eta = \frac{1}{2} \sqrt{\frac{\ln K}{KT}}$  and  $\beta = \sqrt{\frac{\ln K}{KT}}$  finishes the proof.  $\square$

## B. Missing details from Dynamic Regret Analysis of algs (Section 4)

### B.1. Missing Details from Dynamic Regret Analysis in terms of Switching Variation (Section 4.1)

Lemmas B.1 and B.2 bound  $\text{DR}_T^+(j^T)$  and  $\text{DR}_T^-(j^T)$ , respectively.

**Lemma B.1.** *Let  $j^T = (j_1, \dots, j_T)$  be an arbitrary sequence of actions with  $S$  switches. The following bounds hold with probability at least  $1 - \delta$  when  $\alpha = \frac{1}{T}$  and  $\gamma = 2\eta K$ :*

$$\begin{aligned} \text{DR}_T^+(j^T) &= O\left(S\sqrt{KT} \ln \frac{KT}{\delta}\right) \text{ if } \beta = \eta = \frac{1}{\sqrt{KT}} \\ \text{DR}_T^+(j^T) &= O\left(\sqrt{SKT} \ln \frac{KT}{\delta}\right) \text{ if } \beta = \eta = \sqrt{\frac{S}{KT}}. \end{aligned}$$

*Proof.* Assuming that  $\eta K \frac{1+\beta}{\gamma} \leq 1$ , a similar calculation as in the proof of Lemma 3.1 gives

$$\ln \frac{W_{+1,t+1}}{W_{+1,t}} \leq \frac{\eta}{1-\gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}(k) + \frac{\eta^2}{1-\gamma} \sum_{k=1}^K p_{+1,t}(k) \hat{g}_{+1,t}^2(k) + e\alpha.$$

Using (10) and noting that  $p_{+1,t}(k) \hat{g}_{+1,t}(k) \leq 1 + \beta$ , we get

$$\ln \frac{W_{+1,t+1}}{W_{+1,t}} \leq \frac{\eta}{1-\gamma} (\alpha_t(k_{+1,t}, k_{-1,t}) + \beta K) + (1 + \beta) \frac{\eta^2}{1-\gamma} \sum_{k=1}^K \hat{g}_{+1,t}(k) + e\alpha.$$

Recall that  $S = 1 + |\{1 \leq \ell < T : j_\ell \neq j_{\ell+1}\}|$ . Divide the time horizon  $T$  into sub-intervals

$$[T_1, \dots, T_2), [T_2, \dots, T_3), \dots, [T_S, \dots, T_{S+1}),$$

where  $T_1 = 1$  and  $T_{S+1} = T + 1$  and  $j_{T_s} = j_{T_{s+1}} = \dots = j_{T_{s+1}-1}$  for all sub-intervals  $s = 1, \dots, S$ . Further, define  $\Delta^s = T_{s+1} - T_s$ . Summing over  $t \in [T_s, T_{s+1})$  results in

$$\ln \frac{W_{+1,T_{s+1}}}{W_{+1,T_s}} \leq \frac{\eta}{1-\gamma} \beta K \Delta^s + \frac{\eta}{1-\gamma} \sum_{t=T_s}^{T_{s+1}-1} \alpha_t(k_{+1,t}, k_{-1,t}) + (1 + \beta) \frac{\eta^2}{1-\gamma} \sum_{t=T_s}^{T_{s+1}-1} \sum_{k=1}^K \hat{g}_{+1,t}(k) + e\alpha \Delta^s.$$

For any action  $j^s \in [K]$ ,

$$\begin{aligned} W_{+1,T_{s+1}}(j^s) &\geq W_{+1,T_s}(j^s) \exp\left(\eta \sum_{t=T_s+1}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s)\right) \\ &\geq e\alpha W_{+1,T_s} \exp\left(\eta \sum_{t=T_s+1}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s)\right) \\ &\geq \alpha W_{+1,T_s} \exp\left(\eta \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s)\right), \end{aligned}$$

where the last step uses  $\eta \hat{g}_{+1,t}(j^s) \leq \eta \frac{1+\beta}{p_{+1,t}(j^s)} \leq \eta K \frac{1+\beta}{\gamma} \leq 1$ . We have

$$\ln \frac{W_{+1,T_{s+1}}}{W_{+1,T_s}} \geq \ln \frac{W_{+1,T_{s+1}}(j^s)}{W_{+1,T_s}} \geq \ln \alpha + \eta \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s).$$

Combining the upper-bound and lower-bound on  $\ln \frac{W_{+1,T_{s+1}}}{W_{+1,T_s}}$  results in

$$\sum_{t=T_s}^{T_{s+1}-1} \alpha_t(k_{+1,t}, k_{-1,t}) \geq (1-\gamma) \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s) - \eta(1+\beta) \sum_{t=T_s}^{T_{s+1}-1} \sum_{k=1}^K \hat{g}_{+1,t}(k) - \left(\frac{e\alpha}{\eta} - \beta K\right) \Delta^s - \frac{1}{\eta} \ln 1/\alpha.$$

Using Azuma-Hoeffding inequality, with probability at least  $1 - \delta/2$ ,

$$\begin{aligned}
 \sum_{t=T_s}^{T_{s+1}-1} P_t(k_{+1,t}, k_{-1,t}) + 2\sqrt{2\Delta^s \ln(2/\delta)} &\geq \sum_{t=T_s}^{T_{s+1}-1} o_t(k_{+1,t}, k_{-1,t}) \\
 &\geq (1-\gamma) \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s) - \eta(1+\beta) \sum_{t=T_s}^{T_{s+1}-1} \sum_{k=1}^K \hat{g}_{+1,t}(k) - \left(\frac{e\alpha}{\eta} - \beta K\right) \Delta^s - \frac{1}{\eta} \ln 1/\alpha \\
 &\geq (1-\gamma) \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(j^s) - \gamma \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(k) - \left(\frac{e\alpha}{\eta} - \beta K\right) \Delta^s - \frac{1}{\eta} \ln 1/\alpha \\
 &\geq (1-2\gamma) \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(k) - \left(\frac{e\alpha}{\eta} - \beta K\right) \Delta^s - \frac{1}{\eta} \ln 1/\alpha.
 \end{aligned}$$

The last inequality holds as the expression above it is valid for all choices of  $j^s \in [K]$ , and hence we can take a maximum over all such choices. A calculation similar to the one in the proof of Lemma A.1 implies that  $\sum_{t=T_s}^{T_{s+1}-1} \hat{g}_{+1,t}(k) \geq \sum_{t=T_s}^{T_{s+1}-1} P_t(k, k_{-1,t}) - \frac{\ln(2K/\delta)}{\beta}$  for all  $k \in [K]$  with probability at least  $1 - \delta/2$ . Therefore, with probability at least  $1 - \delta$ ,

$$\sum_{t=T_s}^{T_{s+1}-1} P_t(k_{+1,t}, k_{-1,t}) + 2\sqrt{2\Delta^s \ln(2/\delta)} \geq (1-2\gamma) \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} P_t(k, k_{-1,t}) - \left(\frac{e\alpha}{\eta} - \beta K\right) \Delta^s - \frac{1}{\eta} \ln \frac{1}{\alpha} - \frac{\ln(2K/\delta)}{\beta}. \quad (11)$$

Setting  $\gamma = 2\eta K \geq \eta(1+\beta)K$ , rearranging, and summing over  $s = 1, \dots, S$ , we get

$$\begin{aligned}
 \sum_{s=1}^S \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} P_t(k, k_{-1,t}) - \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) &\leq \sum_{s=1}^S 2\sqrt{2\Delta^s \ln(2/\delta)} + \left(\frac{e\alpha}{\eta} - \beta K\right) T + \frac{S}{\eta} \ln \frac{1}{\alpha} \\
 &\quad + \frac{S}{\beta} \ln \frac{2K}{\delta} + 4\eta KT.
 \end{aligned}$$

In particular, we can use the sequence  $j^T$  on the left side above as this sequence uses a constant action within intervals  $[T_s, T_{s+1})$ .

$$\sum_{t=1}^T P_t(j_t, k_{-1,t}) - \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) \leq \sum_{s=1}^S 2\sqrt{2\Delta^s \ln(2/\delta)} + \left(\frac{e\alpha}{\eta} - \beta K\right) T + \frac{S}{\eta} \ln \frac{1}{\alpha} + \frac{S}{\beta} \ln \frac{2K}{\delta} + 4\eta KT$$

Using  $\alpha = \frac{1}{T}$  and  $\eta = \beta = \sqrt{\frac{1}{KT}}$  results in:

$$\begin{aligned}
 \sum_{t=1}^T P_t(j_t, k_{-1,t}) - \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) &\leq 2\sqrt{2ST \ln(2/\delta)} + (e-1)\sqrt{KT} + S\sqrt{KT} \ln \frac{2KT}{\delta} + 4\sqrt{KT} \\
 &= O\left(S\sqrt{KT} \ln \frac{KT}{\delta}\right).
 \end{aligned}$$

On the other hand, if  $\alpha = \frac{1}{T}$  and  $\beta = \eta = \sqrt{\frac{S}{KT}}$  (i.e.,  $S$  is used for setting  $\beta$  and  $\eta$ ), we get

$$\begin{aligned}
 \sum_{t=1}^T P_t(j_t, k_{-1,t}) - \sum_{t=1}^T P_t(k_{+1,t}, k_{-1,t}) &\leq 2\sqrt{2ST \ln(2/\delta)} + e\sqrt{\frac{KT}{S}} - \sqrt{SKT} + \sqrt{SKT} \ln \frac{2KT}{\delta} + 4\sqrt{SKT} \\
 &= O\left(\sqrt{SKT} \ln \frac{KT}{\delta}\right),
 \end{aligned}$$

which concludes the claim.  $\square$



**Lemma B.2.** Assume the same setup as Lemma B.1. With probability at least  $1 - \delta$ , when  $\alpha = \frac{1}{T}$  and  $\gamma = 2\eta K$ :

$$\begin{aligned} DR_T^{-1}(j^T) &= O\left(S\sqrt{KT} \ln \frac{KT}{\delta}\right) \text{ if } \beta = \eta = \frac{1}{\sqrt{KT}} \\ DR_T^{-1}(j^T) &= O\left(\sqrt{SKT} \ln \frac{KT}{\delta}\right) \text{ if } \beta = \eta = \sqrt{\frac{S}{KT}}. \end{aligned}$$

*Proof.* The proof uses similar arguments as Lemma B.1 and has been omitted.  $\square$

## B.2. Missing Details of Dynamic Regret Analysis in Terms of Continuous Variation (Section 4.2)

**Theorem 4.4** (DEX3.S regret under Continuous-Variation). Consider any preference sequence  $\mathbf{P}_1, \dots, \mathbf{P}_T$  with Continuous-Variation  $V_T$ . Then, setting  $\alpha, \gamma$  same as defined in Thm. 4.1 and  $\beta = \eta = \frac{V_T^{1/3}}{4K^{2/3}T^{1/3}}$ , for any time-variant benchmark  $j^T = (j_1, \dots, j_T)$ , with probability at least  $1 - \delta$ , the dynamic regret bound of DEX3.S satisfies  $DR_T(j^T)_{[DEX3.S]} = O\left((V_T^{1/3}K^{1/3}T^{2/3} + 4K^{2/3}T^{1/3}V_T^{-1/3}) \ln \frac{KT}{\delta}\right)$ .

*Proof.* Lem. B.3 builds the key result which shows for any arbitrary sequence of preference matrices  $\mathbf{P}_1, \dots, \mathbf{P}_T$  with Continuous-Variation  $V_T$ , the dynamic regret is upper bounded by  $\tilde{O}(V_T^{1/3}K^{1/3}T^{2/3})$ . The proof now follows recalling that by construction  $DR_T(j^T) := DR_T^{+1}(j^T) + DR_T^{-1}(j^T)$  and  $DR_T(DEX3.S) = \max_{j^T \in [K]^T} DR_T(j^T)$ .  $\square$

**Lemma B.3.** Let  $j^T = (j_1, \dots, j_T)$  be an arbitrary sequence of  $T$  actions and  $\mathbf{P}_1, \dots, \mathbf{P}_T$  be the underlying sequence of preference matrices with Continuous-Variation  $V_T$ . Then setting  $\alpha = \frac{1}{T}$ ,  $\gamma = 2\eta K$ ,  $\beta = \eta = \frac{V_T^{1/3}}{4K^{2/3}T^{1/3}}$ , with probability at least  $1 - \delta$ :

$$\begin{aligned} DR_T^{+1}(j^T) &= O\left((V_T^{1/3}K^{1/3}T^{2/3} + 4K^{2/3}T^{1/3}V_T^{-1/3}) \ln \frac{KT}{\delta}\right) \\ DR_T^{-1}(j^T) &= O\left((V_T^{1/3}K^{1/3}T^{2/3} + 4K^{2/3}T^{1/3}V_T^{-1/3}) \ln \frac{KT}{\delta}\right). \end{aligned}$$

*Proof.* Drawing ideas from the proof analyses of (Besbes et al., 2014), let us first divide the time interval  $[T]$  into sub-intervals  $T_1, \dots, T_S$ , each of size at most  $\Delta$  (hence  $S := \lceil T/\Delta \rceil$ ), such that

$$T_j := \{t \mid (j-1)\Delta + 1 \leq t \leq \min\{j\Delta, T\}\}.$$

Recall at each time  $t$  we may break the instantaneous regret w.r.t. any arm  $j_t \in [K]$  as,  $r_t(j_t) := [P_t(j_t, k_{-1,t}) + P_t(j_t, k_{+1,t})] = r_t^{+1}(j_t) + r_t^{-1}(j_t)$ .

$$\begin{aligned} r_t^{+1}(j_t) &= [P_t(j_t, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t})] \\ r_t^{-1}(j_t) &= [P_t(j_t, k_{+1,t}) - P_t(k_{-1,t}, k_{+1,t})] \end{aligned}$$

For any phase  $i \in [\frac{T}{\Delta}]$ , note we can break the total regret of the left-arm at sub-interval  $i$  w.r.t any arm-sequence  $\{j_t\}_{t \in \mathcal{T}_i}$  as:

$$\begin{aligned} R_{\mathcal{T}_i}^{+1}(\{j_t\}_{t \in \mathcal{T}_i}) &:= \sum_{t \in \mathcal{T}_i} [P_t(j_t, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t})] \\ &= \sum_{t \in \mathcal{T}_i} [P_t(j_t, k_{-1,t}) - P_t(i^*, k_{-1,t}) + P_t(i^*, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t})] \\ &= \sum_{t \in \mathcal{T}_i} [P_t(j_t, k_{-1,t}) - P_t(i^*, k_{-1,t})] + \sum_{t \in \mathcal{T}_i} [P_t(i^*, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t})] \\ &\leq 2V_{\mathcal{T}_i}\Delta + \sum_{t \in \mathcal{T}_i} [P_t(i^*, k_{-1,t}) - P_t(k_{+1,t}, k_{-1,t})] \end{aligned}$$

$$\leq 2V_{\mathcal{T}_i}\Delta + \left[ 2\sqrt{2|\mathcal{T}_i|\ln(2/\delta)} + \left( \frac{e\alpha}{\eta} - \beta K \right) |\mathcal{T}_i| + \frac{1}{\eta} \ln \frac{1}{\alpha} + \frac{1}{\beta} \ln \frac{2K}{\delta} + 4\eta K |\mathcal{T}_i| \right],$$

here  $i^*$  is any representative arm of sub-interval  $i$ . Then summing over all sub-intervals  $\mathcal{T}_1, \dots, \mathcal{T}_S$ , the total regret of the left-arm w.r.t any arm-sequence  $j^T$  is:

$$\begin{aligned} \text{DR}_T^{+1}(j^T) &:= \sum_{s=1}^S R_{\mathcal{T}_i}^{+1}(\{j_t\}_{t \in \mathcal{T}_i}) \\ &\leq \sum_{s=1}^S \left[ 2V_{\mathcal{T}_i}\Delta + [2\sqrt{2|\mathcal{T}_i|\ln(2/\delta)} + \left( \frac{e\alpha}{\eta} - \beta K \right) |\mathcal{T}_i| + \frac{1}{\eta} \ln \frac{1}{\alpha} + \frac{1}{\beta} \ln \frac{2K}{\delta} + 4\eta K |\mathcal{T}_i|] \right] \\ &= 2V_T\Delta + 2\sqrt{2T\ln(2/\delta)} + \left( \frac{e\alpha T}{\eta} - \beta K T \right) + \frac{T}{\eta\Delta} \ln \frac{1}{\alpha} + \frac{1}{\eta} \ln \frac{1}{\alpha} + \frac{T}{\beta\Delta} \ln \frac{2K}{\delta} + \frac{1}{\beta} \ln \frac{2K}{\delta} + 8\eta K T \\ &\stackrel{(a)}{=} 2V_T\Delta + 2\sqrt{2T\ln(2/\delta)} + \left( \frac{e}{\eta} - \sqrt{KT} \right) + \frac{T}{\eta\Delta} \ln T + \frac{1}{\eta} \ln T + \frac{T}{\beta\Delta} \ln \frac{2K}{\delta} + \frac{1}{\beta} \ln \frac{2K}{\delta} + 8\eta K T \\ &\stackrel{(b)}{=} 6V_T^{1/3} K^{1/3} T^{2/3} + 2\sqrt{2} V_T^{1/3} K^{1/3} T^{2/3} (\ln T + \ln(2K/\delta)) + 4K^{2/3} T^{1/3} V_T^{-1/3} (\ln T + \ln(2K/\delta)) \\ &= O\left( (V_T^{1/3} K^{1/3} T^{2/3} + 4K^{2/3} T^{1/3} V_T^{-1/3}) (\ln T + \ln(2K/\delta)) \right), \end{aligned}$$

where (a) holds by setting  $\alpha = 1/T$ , (b) holds by setting  $\Delta = \frac{T^{2/3} K^{1/3} \sqrt{2}}{V_T^{2/3}}$ , and  $\beta = \eta = \frac{V_T^{1/3}}{4K^{2/3} T^{1/3}}$ .

Following a similar analysis also yields

$$\text{DR}_T^{-1}(j^T) = O\left( (V_T^{1/3} K^{1/3} T^{2/3} + 4K^{2/3} T^{1/3} V_T^{-1/3}) (\ln T + \ln(2K/\delta)) \right),$$

combining which with the earlier concludes the proof.  $\square$

## C. Missing details from the Lower Bound (Section 5)

### C.1. Proof of Theorem 5.1

In what follows, we consider only deterministic algorithms. An execution of a randomized algorithm corresponds to a particular choice from the set of all deterministic algorithms. As the changes in the environment are not affected by the algorithm's decisions, it is enough to show the lower bound for all deterministic algorithms (Auer et al., 2002b).

Let  $H_t = \{(k_{+1,t'}, k_{-1,t'}, o_{t'})\}_{t' \leq t}$  be the history of observations upto time  $t$ , where  $o_{t'} := o_{t'}(k_{+1,t'}, k_{-1,t'})$  is the feedback returned by the environment at time  $t'$ . We use  $P_j^s$  to denote the probability distribution over histories induced by assuming that arm  $j \in [K] \setminus \{1\}$  is the condorcet winner in the sub-interval  $[T_s, T_{s+1})$ . Similarly,  $P_1^s$  is the probability distribution assuming that arm 1 is the condorcet winner in the sub-interval  $[T_s, T_{s+1})$  with probability of winning  $0.5 + \epsilon$  against all arms and  $P_t(i, j) = 0.5$  for all  $i, j \in [K] \setminus \{1\}$ . The corresponding expectations are denoted by  $E_j^s[\cdot]$  and  $E_1^s[\cdot]$ , respectively.

The following technical lemma assumes a simpler case where  $P_t(i, j)$  remains constant for all  $t \in [T]$ . One can think of this as the case when  $S = 1$ . We omit the superscript in  $P_j^s, P_1^s, E_j^s[\cdot], E_1^s[\cdot]$ , and  $j^s$  in the lemma below to simplify the notation.

**Lemma C.1.** *Let  $\mathcal{H}_T$  be the set of all possible histories  $H_T$  and let  $f : \mathcal{H}_T \rightarrow [0, M]$  be a measurable function that maps a history  $H_T$  to a number in the interval  $[0, M]$ . Then,*

$$E_j[f(H_T)] \leq E_1[f(H_T)] + M \sqrt{\epsilon \ln \left( \frac{1+2\epsilon}{1-2\epsilon} \right) E_1[N_{1j} + N_j]},$$

where,  $N_{1j}$  is the number of times the algorithm chooses to duel between arms 1 and  $j$  and  $N_j$  is the number of times arm  $j$  is compared with an arm other than itself and arm 1.

*Proof.* For any two probability distributions  $p$  and  $q$ , let  $\text{TV}(p, q)$  and  $\text{KL}(p, q)$ , respectively, denote the total variation distance and KL-divergence between  $p$  and  $q$ . Note that,

$$\begin{aligned}
 \mathbb{E}_j [f(H_T)] - \mathbb{E}_1 [f(H_T)] &= \sum_{H_T \in \mathcal{H}_T} f(H_T) (P_j(H_T) - P_1(H_T)) \\
 &\leq \sum_{H_T: P_j(H_T) \geq P_1(H_T)} f(H_T) (P_j(H_T) - P_1(H_T)) \\
 &\leq M \sum_{H_T: P_j(H_T) \geq P_1(H_T)} (P_j(H_T) - P_1(H_T)) \\
 &\stackrel{(a)}{=} M \text{TV}(P_j, P_1) \\
 &\stackrel{(b)}{\leq} M \sqrt{\frac{\ln 2}{2} \text{KL}(P_1, P_j)}. \tag{12}
 \end{aligned}$$

Here, (a) follows from the definition of the total-variation distance and (b) follows from Pinsker's inequality. Let  $h_t = (k_{+1,t}, k_{-1,t}, o_t)$  be the  $t^{\text{th}}$  element of  $H_T$ . Using the chain-rule for KL-divergence, we get

$$\begin{aligned}
 \text{KL}(P_1, P_j) &= \sum_{t=1}^T \text{KL}(P_1(h_t | H_{t-1}), P_j(h_t | H_{t-1})) \\
 &= \sum_{t=1}^T P_1((k_{+1,t} = k_{-1,t}) \vee (k_{+1,t} \notin \{1, j\} \wedge k_{-1,t} \notin \{1, j\})) \text{KL}\left(\frac{1}{2}, \frac{1}{2}\right) + \\
 &\quad P_1((k_{+1,t} = 1 \wedge k_{-1,t} = j) \vee (k_{+1,t} = j \wedge k_{-1,t} = 1)) 2\epsilon \log_2\left(\frac{1+2\epsilon}{1-2\epsilon}\right) - \\
 &\quad P_1((k_{+1,t} = j \wedge k_{-1,t} \notin \{1, j\}) \vee (k_{+1,t} \notin \{1, j\} \wedge k_{-1,t} = j)) \frac{1}{2} \log_2(1-4\epsilon^2) \\
 &= 2\mathbb{E}_1[N_{1j}] \epsilon \log_2\left(\frac{1+2\epsilon}{1-2\epsilon}\right) - \frac{1}{2} \mathbb{E}_1[N_j] \log_2(1-4\epsilon^2) \\
 &\leq 2\epsilon \log_2\left(\frac{1+2\epsilon}{1-2\epsilon}\right) \mathbb{E}_1[N_{1j} + N_j].
 \end{aligned}$$

The last line follows from the fact that  $-\frac{1}{2} \log_2(1-4x^2) \leq 2x \log_2 \frac{1+2x}{1-2x}$  for all  $x \in (0, 0.5)$ . We have abused the notation and used  $\text{KL}(\alpha, \beta)$  to represent  $\text{KL}(p, q)$  where  $p = \text{Bernoulli}(\alpha)$  and  $q = \text{Bernoulli}(\beta)$  for  $\alpha, \beta \in [0, 1]$ . Combining this inequality with (12) yields the desired result.  $\square$

Theorem 5.1 has been reproduced verbatim from Section 5 below. See Section 5 for the definition of  $\mathcal{V}$  and  $j^T(\nu)$ .

**Theorem 5.1** (Dynamic Regret Lower Bound for Switching-Variation). *Assume that  $K \geq 3$  and  $\epsilon \in (0, 1/4)$  in (4). With  $\mathcal{V}$  and  $j^T(\nu)$  defined above, there exists an instance  $\nu \in \mathcal{V}$  such that:*

$$\mathbb{E} [DR_T(j^T(\nu))] = \Omega(\sqrt{SKT}),$$

for any non-stationary dueling bandits algorithm.

*Proof.* Let  $t \in [T_s, T_{s+1})$  and define the following events at time  $t$ :

1.  $\mathcal{E}_a^t = \{(k_{+1,t} = 1 \wedge k_{-1,t} = j^s) \vee (k_{+1,t} = j^s \wedge k_{-1,t} = 1)\}$
2.  $\mathcal{E}_b^t = \{(k_{+1,t} = j^s \wedge k_{-1,t} \notin \{1, j^s\}) \vee (k_{+1,t} \notin \{1, j^s\} \wedge k_{-1,t} = j^s)\}$
3.  $\mathcal{E}_c^t = \{k_{+1,t} \neq j^s \wedge k_{-1,t} \neq j^s\}$
4.  $\mathcal{E}_d^t = \{k_{+1,t} = j^s \wedge k_{-1,t} = j^s\}$

For  $t \in [T_s, T_{s+1})$ , let  $r_t = (P_t(j^s, k_{+1,t}) + P_t(j^s, k_{-1,t}) - 1)/2$ . Now consider the expected regret at time  $t \in [T_s, T_{s+1})$  under the distribution  $\mathbb{P}_{j^s}^s$ .

$$\begin{aligned} \mathbb{E}_{j^s}^s [r_t] &= \frac{\epsilon}{2} \mathbb{P}_{j^s}^s(\mathcal{E}_a^t) + \frac{\epsilon}{2} \mathbb{P}_{j^s}^s(\mathcal{E}_b^t) + \epsilon \mathbb{P}_{j^s}^s(\mathcal{E}_c^t) \\ &= \frac{\epsilon}{2} \mathbb{P}_{j^s}^s(\mathcal{E}_a^t) + \frac{\epsilon}{2} \mathbb{P}_{j^s}^s(\mathcal{E}_b^t) + \epsilon (1 - \mathbb{P}_{j^s}^s(\mathcal{E}_a^t) - \mathbb{P}_{j^s}^s(\mathcal{E}_b^t) - \mathbb{P}_{j^s}^s(\mathcal{E}_d^t)) \\ &= \epsilon - \epsilon \mathbb{P}_{j^s}^s(\mathcal{E}_d^t) - \frac{\epsilon}{2} (\mathbb{P}_{j^s}^s(\mathcal{E}_a^t) + \mathbb{P}_{j^s}^s(\mathcal{E}_b^t)). \end{aligned}$$

Define  $N_{1j^s}^s = \sum_{t=T_s}^{T_{s+1}-1} \mathbf{1}\{\mathcal{E}_a^t\}$ ,  $N_{j^s}^s = \sum_{t=T_s}^{T_{s+1}-1} \mathbf{1}\{\mathcal{E}_b^t\}$ ,  $\bar{N}_{j^s}^s = \sum_{t=T_s}^{T_{s+1}-1} \mathbf{1}\{\mathcal{E}_d^t\}$ , and  $\Delta^s = T_{s+1} - T_s$ . We have,

$$\begin{aligned} \mathbb{E}_{j^s}^s \left[ \sum_{t=T_s}^{T_{s+1}-1} r_t \right] &= \epsilon \Delta^s - \epsilon \mathbb{E}_{j^s}^s [\bar{N}_{j^s}^s] - \frac{\epsilon}{2} (\mathbb{E}_{j^s}^s [N_{1j^s}^s] + \mathbb{E}_{j^s}^s [N_{j^s}^s]) \\ &= \epsilon \Delta^s - \frac{\epsilon}{2} \mathbb{E}_{j^s}^s [2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s]. \end{aligned}$$

We bound the expectation on the right hand side above using Lemma C.1. As the learner only uses the history  $H_t$  to choose the arms at time  $t$ ,  $\bar{N}_{j^s}^s$ ,  $N_{1j^s}^s$ ,  $N_{j^s}^s$ , and hence  $2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s$  is a measurable function of the history. Applying Lemma C.1 and noting that the maximum value of  $2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s$  is  $2\Delta^s$  results in

$$\mathbb{E}_{j^s}^s [2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s] \leq \mathbb{E}_1^s [2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s] + 2\Delta^s \sqrt{\epsilon \ln \left( \frac{1+2\epsilon}{1-2\epsilon} \right) \mathbb{E}_1 [N_{1j}^s + N_j^s]}.$$

Therefore, using Cauchy-Schwarz inequality,

$$\begin{aligned} \sum_{j^s=2}^K \mathbb{E}_{j^s}^s [2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s] &\leq \mathbb{E}_1^s \left[ \sum_{j^s=2}^K 2\bar{N}_{j^s}^s + N_{1j^s}^s + N_{j^s}^s \right] + 2\Delta^s \sqrt{\epsilon \ln \left( \frac{1+2\epsilon}{1-2\epsilon} \right) (K-1) \sum_{j^s=2}^K \mathbb{E}_1 [N_{1j}^s + N_j^s]} \\ &\leq 2\Delta^s + 2\Delta^s \sqrt{2\Delta^s (K-1) \epsilon \ln \left( \frac{1+2\epsilon}{1-2\epsilon} \right)}. \end{aligned}$$

This results in the following bound:

$$\begin{aligned} \sum_{j^s=2}^K \mathbb{E}_{j^s}^s \left[ \sum_{t=T_s}^{T_{s+1}-1} r_t \right] &\geq \epsilon \Delta^s \left( K-1-1 - \sqrt{2\Delta^s (K-1) \epsilon \ln \left( \frac{1+2\epsilon}{1-2\epsilon} \right)} \right) \\ &\geq \epsilon \Delta^s \left( K-1-1 - 8\epsilon \sqrt{\Delta^s (K-1) \ln(4/3)} \right). \end{aligned}$$

The last line uses the inequality  $2x \ln \frac{1+2x}{1-2x} \leq 32x^2 \ln(4/3)$  for all  $x \in [0, 1/4]$ .

Recall that  $\text{DR}_T(j^T(\nu)) = \sum_{s=1}^S \sum_{t=T_s}^{T_{s+1}-1} r_t$ . The expected value of  $\text{DR}_T(j^T(\nu))$  under a randomly sampled environment  $\nu$  from  $\mathcal{V}$  is given by:

$$\begin{aligned} \mathbb{E} [\text{DR}_T(j^T(\nu))] &= \sum_{s=1}^S \frac{1}{K-1} \sum_{j^s=2}^K \mathbb{E}_{j^s}^s \left[ \sum_{t=T_s}^{T_{s+1}-1} r_t \right] \\ &\geq \sum_{s=1}^S \left( \epsilon \Delta^s - \frac{\epsilon}{K-1} \Delta^s - \frac{8\epsilon^2}{K-1} \Delta^s \sqrt{\Delta^s (K-1) \ln(4/3)} \right) \\ &\geq \sum_{s=1}^S \left( \epsilon \Delta^s - \frac{\epsilon}{K-1} \Delta^s - \frac{8\epsilon^2}{K-1} \Delta^s \sqrt{\frac{T(K-1)}{S} \ln(4/3)} \right) \\ &= \epsilon T - \epsilon \frac{T}{K-1} - 8\epsilon^2 \sqrt{\frac{T(K-1)}{S} \ln(4/3)} \frac{T}{K-1} \end{aligned}$$

$$\begin{aligned}
 &= \epsilon \frac{K-2}{K-1} T - 8\epsilon^2 \sqrt{\frac{T(K-1)}{S} \ln(4/3)} \frac{T}{K-1} \\
 &\stackrel{(a)}{\geq} \epsilon \frac{T}{2} - 8\epsilon^2 \sqrt{\frac{T(K-1)}{S} \ln(4/3)} \frac{T}{K-1}.
 \end{aligned}$$

Here, (a) follows by assuming that  $K \geq 3$ . Choosing the value of  $\epsilon = \frac{1}{2} \sqrt{\frac{S(K-1)}{T \ln(4/3)}}$  to maximize the expression on the right hand side results in

$$\begin{aligned}
 \mathbb{E} [\text{DR}_T(j^T(\nu))] &\geq \frac{1}{4\sqrt{\ln(4/3)}} \sqrt{S(K-1)T} - \frac{2}{\sqrt{\ln(4/3)}} \sqrt{S(K-1)T} \\
 &= \Omega(\sqrt{SKT}).
 \end{aligned}$$

As the expected value of  $\text{DR}_T(j^T(\nu))$  under a random choice of environment  $\nu \in \mathcal{V}$  is  $\Omega(\sqrt{SKT})$ , there is at least one environment in  $\mathcal{V}$  for which  $\mathbb{E} [\text{DR}_T(j^T(\nu))] = \Omega(\sqrt{SKT})$ .  $\square$

## C.2. Proof of Theorem 5.2

As before, the theorem below has been reproduced verbatim from Section 5. Please refer to Section 5 for the definition of  $\mathcal{V}$  and  $j^T(\nu)$ .

**Theorem 5.2** (Dynamic Regret Lower Bound for Continuous-Variation). *Assume that  $K \geq 3$ . With  $\mathcal{V}$  and  $j^T(\nu)$  defined above, there exists an instance  $\nu \in \mathcal{V}$  such that:*

$$\mathbb{E} [\text{DR}_T(j^T(\nu))] = \Omega((KV_T)^{1/3} T^{2/3}),$$

for any non-stationary dueling bandits algorithm.

*Proof.* Let us define  $\mathcal{V}' \subseteq \mathcal{V}$  such that for all  $\nu \in \mathcal{V}'$ ,  $\Delta^1 = \Delta^2 = \dots = \Delta^{S-1} = \Delta$  and  $\Delta^S < \Delta$ , where recall that  $\Delta^s = T_{s+1} - T_s$ . As  $\max_{i,j \in [K]} |P_{t+1}(i,j) - P_t(i,j)| \leq 2\epsilon$  for all  $t > 0$ , for any environment  $\nu \in \mathcal{V}'$ ,

$$\sum_{t=1}^{T-1} \max_{i,j \in [K]} |P_{t+1}(i,j) - P_t(i,j)| = \sum_{s=2}^S \max_{i,j \in [K]} |P_{T_s}(i,j) - P_{T_{s-1}}(i,j)| \leq 2\epsilon \left\lceil \frac{T}{\Delta} \right\rceil \leq V_T.$$

The last inequality above assumes that  $\epsilon \leq \frac{V_T}{2\lceil T/\Delta \rceil}$ .

Proceeding as in the proof of Theorem 5.1, and noting that  $S = \lceil \frac{T}{\Delta} \rceil$  we get:

$$\mathbb{E} [\text{DR}_T(j^T(\nu))] \geq \epsilon T - \frac{\epsilon}{K-1} T - \frac{8\epsilon^2}{K-1} T \sqrt{\frac{T(K-1)}{\lceil T/\Delta \rceil} \ln(4/3)},$$

where the expectation this time includes a random choice of environment from  $\mathcal{V}'$ . Further assuming that  $\epsilon = \min \left\{ \frac{V_T}{2\lceil T/\Delta \rceil}, \frac{1}{16\sqrt{\Delta K \ln(4/3)}} \right\}$  results in

$$\begin{aligned}
 \mathbb{E} [\text{DR}_T(j^T(\nu))] &\geq \epsilon \frac{K-2}{K-1} T - \frac{8\epsilon^2}{K-1} T \sqrt{\frac{T(K-1)}{\lceil T/\Delta \rceil} \ln(4/3)} \\
 &\geq \frac{\epsilon}{2} T - \frac{8\epsilon^2}{2} T \sqrt{\Delta(K-1) \ln(4/3)} \\
 &\geq \epsilon T \left( \frac{1}{2} - 4\epsilon \sqrt{\Delta K \ln(4/3)} \right) \geq \frac{T}{64\sqrt{\Delta K \ln(4/3)}}.
 \end{aligned}$$

Setting  $\Delta = \frac{1}{K^{5/3}} \left( \frac{T}{V_T} \right)^{2/3}$  yields the desired result.  $\square$

## D. Missing details from Dynamic Regret Analysis under Borda Scores (Section 6)

Algorithm 3 details the procedure for  $\text{Borda-DEX3.S}$  and defines the reward estimate  $s'_{i,t}(k)$  for  $i \in \{+1, -1\}$  and  $k \in [K]$  in L9. Before we prove the high-probability dynamic regret guarantee from Theorem 6.1, we reproduce Lemma 10 from Saha et al. (2021), which serves the same purpose as Lemma A.1 for the Borda score estimate  $s'_{i,t}(k)$  defined in Algorithm 3.

**Lemma D.1.** For any  $\delta > 0$ ,  $\beta, \gamma \in (0, 1)$ ,  $i \in \{+1, -1\}$ , with probability at least  $1 - \delta/2$ ,

$$\forall k \in [K], \quad \sum_{t=1}^T s'_{i,t}(k) \geq \sum_{t=1}^T s_t(k) - \frac{1}{\gamma\beta} \ln \left( \frac{2K}{\delta} \right),$$

where  $s_t(k) = \frac{1}{K} \sum_{j \in [K]} P_t(k, j)$  is the shifted Borda score of arm  $k$  at time  $t$ .

The basic regret decomposition idea that was used in Appendix B to bound the regret of  $\text{DEX3.S}$  can be used for  $\text{Borda-DEX3.S}$  as well. We define  $\text{DBR}_T^+(j^T)$  and  $\text{DBR}_T^-(j^T)$  as:

$$\begin{aligned} \text{DBR}_T^+(j^T) &= \sum_{t=1}^T b_t(j_t) - b_t(k_{+1,t}) = \frac{K}{K-1} \sum_{t=1}^T s_t(j_t) - s_t(k_{+1,t}) \\ \text{DBR}_T^-(j^T) &= \sum_{t=1}^T b_t(j_t) - b_t(k_{-1,t}) = \frac{K}{K-1} \sum_{t=1}^T s_t(j_t) - s_t(k_{-1,t}). \end{aligned}$$

As before,  $\text{DBR}_T(j^T) = \frac{1}{2} (\text{DBR}_T^+(j^T) + \text{DBR}_T^-(j^T))$ . We only show the bound on  $\text{DBR}_T^+(j^T)$  below. The bound on  $\text{DBR}_T^-(j^T)$ , and hence on  $\text{DBR}_T(j^T)$ , follows along the lines of Appendix B. We begin with a simple lemma.

**Lemma D.2.** With probability at least  $1 - \delta/2$ ,

$$\sum_{t=1}^T \sum_{k=1}^K p_{+1,t}(k) s'_{+1,t}(k) \leq \sum_{t=1}^T (s_t(k_{+1,t}) + \beta K) + \frac{\gamma+1}{\gamma} \sqrt{2T \ln(2/\delta)}.$$

*Proof.* Let  $X_n = \sum_{t=1}^n \left[ \sum_{k=1}^K p_{+1,t}(k) s'_{+1,t}(k) - (s_t(k_{+1,t}) + \beta K) \right]$ . One can easily show that the sequence  $X_1, X_2, \dots$  is a martingale. Using Azuma's inequality,

$$\mathbb{P}(X_T \geq \epsilon) \leq \exp \left( -\frac{\epsilon^2}{2 \sum_{t=1}^T c_t^2} \right), \quad (13)$$

where  $c_t = b_t - a_t$  where  $a_t \leq X_t - X_{t-1} \leq b_t$  almost surely. We next derive such a  $c_t$ .

$$\begin{aligned} X_t - X_{t-1} &= \sum_{k=1}^K p_{+1,t}(k) s'_{+1,t}(k) - (s_t(k_{+1,t}) + \beta K) \\ &= \sum_{k=1}^K \left[ \frac{\mathbb{I}\{k_{+1,t} = k\}}{K} \sum_{j=1}^K \frac{\mathbb{I}\{k_{-1,t} = j\} o_t(k, j)}{p_{-1,t}(j)} + \beta \right] - \frac{1}{K} \sum_{j=1}^K P_t(k_{+1,t}, j) - \beta K \\ &\leq \frac{1}{\gamma} + \beta K - \beta K = \frac{1}{\gamma} = b_t. \end{aligned}$$

Similarly,

$$X_t - X_{t-1} = \sum_{k=1}^K p_{+1,t}(k) s'_{+1,t}(k) - (s_t(k_{+1,t}) + \beta K) \geq \beta K - 1 - \beta K = -1 = a_t.$$

Therefore,  $c_t = b_t - a_t = \frac{\gamma+1}{\gamma}$ . Using this value of  $c_t$  in (13) and setting  $\exp \left( -\frac{\epsilon^2}{2 \sum_{t=1}^T c_t^2} \right) \leq \frac{\delta}{2}$  gives  $\epsilon = \frac{\gamma+1}{\gamma} \sqrt{2T \ln \frac{2}{\delta}}$  finishing the proof.  $\square$

We now show a high probability regret bound on  $\text{DBR}_T^+(j^T)$ .

**Lemma D.3.** *Let  $j^T = (j_1, \dots, j_T)$  be an arbitrary sequence of  $T$  actions with  $S$  switches. The following bound holds with probability at least  $1 - \delta$  when  $\eta = \left(\frac{S \ln K}{T \sqrt{2K}}\right)^{2/3}$ ,  $\beta = \frac{S^{1/3} \sqrt{\ln(2K/\delta)}}{(2\eta)^{1/4} K^{3/4} \sqrt{T}}$ ,  $\gamma = \sqrt{2\eta K} \geq \sqrt{\eta K(1 + \beta)}$ , and  $\alpha = \frac{1}{KT}$ :*

$$\text{DBR}_T^+(j^T) = \tilde{O}(S^{1/6} K^{-1/3} T^{5/6} + S^{1/2} K^{1/3} T^{2/3}).$$

*Proof.* Divide the time horizon into sub-intervals  $[T_1, T_2), [T_2, T_3), \dots, [T_S, T_{S+1})$  as in Appendix B. Using Lemma D.2 and the condition  $\frac{\eta K}{\gamma^2}(1 + \gamma\beta) \leq 1$ , and performing a similar calculation as in Appendix B yields for any action  $j^s$ , with probability at least  $1 - \delta/2$ ,

$$\begin{aligned} -\ln \frac{1}{\alpha} + \eta \sum_{t=T_s}^{T_{s+1}-1} s'_{+1,t}(j^s) &\leq \ln \frac{W_{+1,T_{s+1}}}{W_{+1,T_s}} \leq \frac{\eta}{1-\gamma} \beta K \Delta^s + \frac{\eta}{1-\gamma} \sum_{t=T_s}^{T_{s+1}-1} s_t(k_{+1,t}) + \frac{\eta(1+\gamma)}{\gamma(1-\gamma)} \sqrt{2\Delta^s \ln \frac{2}{\delta}} \\ &\quad \frac{\eta^2(1+\gamma\beta)}{\gamma(1-\gamma)} \sum_{t=T_s}^{T_{s+1}-1} \sum_{k=1}^K s'_{+1,t}(k) + e\alpha K \Delta^s, \end{aligned}$$

where  $\Delta^s = T_{s+1} - T_s$  is the length of the  $s^{\text{th}}$  sub-interval. Simplifying as in Appendix B, we get with probability at least  $1 - \delta/2$ ,

$$\sum_{t=T_s}^{T_{s+1}-1} s_t(k_{+1,t}) + \frac{1+\gamma}{\gamma} \sqrt{2\Delta^s \ln(2/\delta)} \geq (1-2\gamma) \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} s'_{+1,t}(k) - \left(\frac{e\alpha}{\eta} + \beta\right) K \Delta^s - \frac{1}{\eta} \ln \frac{1}{\alpha}.$$

Using Lemma D.1, we get with probability at least  $1 - \delta$

$$\begin{aligned} \sum_{t=T_s}^{T_{s+1}-1} s_t(k_{+1,t}) + \frac{1+\gamma}{\gamma} \sqrt{2\Delta^s \ln(2/\delta)} &\geq (1-2\gamma) \max_{k \in [K]} \sum_{t=T_s}^{T_{s+1}-1} s_t(k) - \left(\frac{e\alpha}{\eta} + \beta\right) K \Delta^s - \frac{1}{\eta} \ln \frac{1}{\alpha} - \frac{1}{\gamma\beta} \ln \frac{2K}{\delta} \\ &\geq (1-2\gamma) \sum_{t=T_s}^{T_{s+1}-1} s_t(j^s) - \left(\frac{e\alpha}{\eta} + \beta\right) K \Delta^s - \frac{1}{\eta} \ln \frac{1}{\alpha} - \frac{1}{\gamma\beta} \ln \frac{2K}{\delta}. \end{aligned}$$

Here, recall that  $j^s$  is the constant arm in the sequence  $j^T = (j_1, j_2, \dots, j_T)$  in the sub-interval  $[T_s, T_{s+1})$ . Noting that  $s_t(k) \leq 1$  for any arm  $k$ , summing over  $s = 1, \dots, S$ , and rearranging, we get,

$$\sum_{t=1}^T (s_t(j_t) - s_t(k_{+1,t})) \leq 2\gamma T + \frac{1+\gamma}{\gamma} \sqrt{2ST \ln(2/\delta)} + \left(\frac{e\alpha}{\eta} + \beta\right) KT + \frac{S}{\eta} \ln \frac{1}{\alpha} + \frac{S}{\gamma\beta} \ln \frac{2K}{\delta}.$$

Setting  $\eta = \left(\frac{S \ln K}{T \sqrt{2K}}\right)^{2/3}$ ,  $\beta = \frac{S^{1/3} \sqrt{\ln(2K/\delta)}}{(2\eta)^{1/4} K^{3/4} \sqrt{T}}$ ,  $\gamma = \sqrt{2\eta K} \geq \sqrt{\eta K(1 + \beta)}$ , and  $\alpha = \frac{1}{KT}$  finishes the proof.  $\square$