



HAL
open science

Versatile Dueling Bandits: Best-of-both-World Analyses for Online Learning from Preferences

Aadirupa Saha, Pierre Gaillard

► **To cite this version:**

Aadirupa Saha, Pierre Gaillard. Versatile Dueling Bandits: Best-of-both-World Analyses for Online Learning from Preferences. ICML 2022 - 39th International Conference on Machine Learning, Jul 2022, Baltimore (MA), United States. pp.1-25. hal-03922380

HAL Id: hal-03922380

<https://inria.hal.science/hal-03922380>

Submitted on 6 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Versatile Dueling Bandits: Best-of-both-World Analyses for Online Learning from Preferences

Aadirupa Saha* Pierre Gaillard †

Abstract

We study the problem of K -armed dueling bandit for both stochastic and adversarial environments, where the goal of the learner is to aggregate information through relative preferences of pair of decisions points queried in an online sequential manner. We first propose a novel reduction from any (general) dueling bandits to multi-armed bandits and despite the simplicity, it allows us to improve many existing results in dueling bandits. In particular, *we give the first best-of-both world result for the dueling bandits regret minimization problem*—a unified framework that is guaranteed to perform optimally for both stochastic and adversarial preferences simultaneously. Moreover, our algorithm is also the first to achieve an optimal $O(\sum_{i=1}^K \frac{\log T}{\Delta_i})$ regret bound against the Condorcet-winner benchmark, which scales optimally both in terms of the arm-size K and the instance-specific suboptimality gaps $\{\Delta_i\}_{i=1}^K$. This resolves the long standing problem of designing an instancewise gap-dependent order optimal regret algorithm for dueling bandits (with matching lower bounds up to small constant factors). We further justify the robustness of our proposed algorithm by proving its optimal regret rate under adversarially corrupted preferences—this outperforms the existing state-of-the-art corrupted dueling results by a large margin. In summary, we believe our reduction idea will find a broader scope in solving a diverse class of dueling bandits setting, which are otherwise studied separately from multi-armed bandits with often more complex solutions and worse guarantees. The efficacy of our proposed algorithms is empirically corroborated against the existing dueling bandit methods.

1 Introduction

Studies have shown that it is often easier, faster and less expensive to collect feedback on a relative scale rather than asking ratings on an absolute scale. E.g., to understand the liking for a given pair of items, say (A,B), it is easier for the users to answer preference-based queries like: “Do you prefer Item A over B?”, rather than their absolute counterparts: “How

*Microsoft Research, NYC. aadirupa.saha@microsoft.com.

†Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France. pierre.gaillard@inria.fr

much do you score items A and B in a scale of [0-10]?" . From a system designer’s point of view, exploiting such user preference information could greatly aid in improving systems performances, especially when data can be collected on a relative scale and online fashion; such as recommendation systems, crowd-sourcing platforms, training bots, multi-player games, search-engine optimization, online retail, just to name a few. In many real world problems, especially where human preferences are elicited in an online fashion, e.g., design of surveys and expert reviews, assortment selection, search engine optimization, recommender systems, ranking in multiplayer games, etc, or even more general reinforcement learning problems where rewards shaping is often a challenging problem (e.g. if multi-objective rewards etc.), and instead, a preference feedback is much easier to elicit.

Due to the widespread applicability and ease of data collection with relative feedback, learning from preferences has gained much popularity in the machine learning community and widely studied as the problem of *Dueling-Bandits* (DB) over last decade Yue et al. [2012], Ailon et al. [2014], Zoghi et al. [2014a,b, 2015a], which is an online learning framework that generalizes the standard multiarmed bandit (MAB) Auer et al. [2002] setting for identifying a set of ‘good’ arms from a fixed decision-space (set of items) by querying preference feedback of actively chosen item-pairs.

Dueling Bandits Problem (DB) More formally, in classical dueling bandits with K arms, the learning proceeds in rounds, where at each time step $t \in \{1, 2, \dots, T\}$, the learner selects a pair of arms $(k_{+1,t}, k_{-1,t})$ and receives the winner of the duel in terms of a binary preference feedback $o_t(k_{+1,t}, k_{-1,t}) \sim \text{Ber}(P_t(k_{+1,t}, k_{-1,t}))$, sampled according to an underlying preference matrix $\mathbf{P}_t \in [0, 1]^{K \times K}$ (chosen adversarially in the most general setup). The objective of the learner is to minimize the regret with respect to a (or set of) ‘best’ arm(s) in hindsight.

Related Works Over the years, the problem of Dueling Bandits has been studied with various objectives and generalizations. This includes analyzing the learning rate under various preference structures, such as total ordering, transitivity, stochastic triangle inequality Falahatgar et al. [2017], Yue and Joachims [2011], utility based preference structure Ailon et al. [2014], Szorenyi et al. [2015], Saha and Gopalan [2018a], Chen et al. [2018], or under any general pairwise preference matrices Dudík et al. [2015], Jamieson et al. [2015], Komiyama et al. [2016]. Consequently, depending on the underlying preference structure, the learner’s performance has been evaluated w.r.t. different benchmarks including among other promising generalizations. The problem of stochastic dueling bandits has been studied for both PAC Falahatgar et al. [2018], Szorenyi et al. [2015], Sui et al. [2018] as well as regret minimization setting Zoghi et al. [2014a], Yue and Joachims [2009], Chen and Frazier [2017], Zoghi et al. [2015a] under several notions of benchmarks including best arm identification Saha and Gopalan [2019b], Busa-Fekete et al. [2014], Falahatgar et al. [2017], top-set detection Busa-Fekete et al. [2013], Saha and Gopalan [2019a], ranking Ren et al. [2018], Saha and Gopalan [2018b], amongst many. Some recent works have also looked into the problem for adversarial sequence of preference matrices Gupta and Saha [2021], Gajane et al. [2015], Saha et al. [2021], robustness to corruptions Agarwal et al. [2021], or extending dueling bandits to potentially infinite arm sets Saha [2021], Kumagai [2017] and contextual scenarios Dudík et al. [2015],

Saha and Krishnamurthy [2021]. Another interesting line of work along dueling bandits is to study the implications for its subsetwise generalization Ren et al. [2018], Sui et al. [2018], Brost et al. [2016], Chen et al. [2018], also studied as battling bandits Saha and Gopalan [2018a, 2019a], Bengs et al. [2021].

Despite widespread surge of interest along this line of research and multiple attempts there are some fundamental long standing open questions in dueling bandits which are (surprisingly!) yet unresolved.

Unresolved Question #1 One of the longest and most widely studied objective in stochastic dueling bandit is to measure regret w.r.t. the Condorcet winner (CW) arm: Given a preference matrix P , an arm $k^{(cw)} \in [K]$ is termed as the CW of P if $P(k^{(cw)}, k) > 0.5 \forall k \in [K] \setminus \{k\}$ Zoghi et al. [2014a]. Assuming P contains a CW, there have been several attempts to design an optimal regret dueling bandit algorithm against the CW arm $k^{(cw)}$ (see Eq. (2) for details) Zoghi et al. [2014a], Wu and Liu [2016], Komiyama et al. [2015], Bengs et al. [2021]. Without loss of generality, assuming $k^{(cw)} = 1$ and by denoting $\Delta_i = P(1, i) - 0.5$ to be the suboptimality gap of item i w.r.t. the CW, it is well known that the dueling bandit regret lower bound (w.r.t. the CW arm) is $\Omega\left(\sum_{i=2}^K \frac{\log T}{\Delta_i}\right)$ Yue et al. [2012], Komiyama et al. [2015]. However, despite several attempts, it is still unknown how to design an order optimal dueling bandit algorithm for the CW regret. Existing upper-bounds suffer all suboptimal Δ_{\min}^{-2} and/or K^2 dependencies.

Notably, under more restricted structures, e.g. total ordering Yue and Joachims [2011], or utility based preferences Szorenyi et al. [2015], Saha and Gopalan [2020], or even special preference structures where the suboptimality gaps of all items ($\Delta_i, i \in [K] \setminus \{1\}$) are equal, the problem is easier to solve and tight regret guarantees are available with matching upper and lower bound analysis. However, for the case of any general preference matrix with CW, none of the existing attempts were able to close this regret analysis gap successfully Zoghi et al. [2014a], Wu and Liu [2016], Chen and Frazier [2017], Komiyama et al. [2015], Saha and Gaillard [2021]. Subsequently, the natural questions to ask are:

(1). *Is the lower bound tight?* (2). *How to close the gap between the upper and lower bound for CW regret?*

Unresolved Question #2 Till date, all the proposed algorithms of dueling bandits need to know underlying preference structure ahead of time in order to yield optimal regret bounds. In fact, different algorithms have been proposed based on the nature/structures of the underlying preference matrices, e.g. Yue and Joachims [2011] for preferences with total orderings in terms of (relaxed) stochastic transitivity and stochastic triange inequality; Ailon et al. [2014], Gajane et al. [2015] for linear-utility based preferences, Szorenyi et al. [2015] for BTL models, Zoghi et al. [2014a], Komiyama et al. [2015], Wu and Liu [2016] for stochastic preferences in presence of CW, Saha et al. [2021], Gajane et al. [2015] for adversarial sequence of preferences, etc. However, it might not always be realistic to assume complete knowledge of the properties underlying preference matrices. Thus the daunting question to ask in this regard is

Is it possible to design an order optimal ‘best-of-both-world’ algorithm for dueling bandits?

That is, an algorithm that can adapt itself to the underlying structures of the preference environments and give optimal regret for both stochastic and adversarial settings? There has been a series of work on this line for the MAB framework [e.g., Bubeck and Slivkins, 2012, Auer and Chiang, 2016, Zimmert and Seldin, 2021], but unfortunately there has not been any existing ‘best-of-both-world’ attempt for general dueling bandits.

Unresolved Question #3 In any real world situation, the true feedback are often corrupted with some form of system noise. Undoubtedly, the binary 0/1 bit dueling preferences are extremely prone to such noises when the learner might get to observe a flipped bit (adversarially corrupted) instead of the true dueling feedback.

Can we design an efficient dueling bandit algorithm which is robust to adversarial corruptions and provably optimal?

Our Contributions In this paper, we answer all of the above three questions affirmatively. The list of our specific contributions can be summarized as follows:

1. **A novel insight on the reduction from MAB to DB.** Ailon et al. [2014] proposed a reduction from MAB to utility-based dueling bandits. We show that the latter can easily be extended to more general dueling bandit problems (including CW), with significant consequences (see below) on the state of the art in dueling bandits theory. We believe that the reduction will find wider application in solving a diverse class of dueling bandit settings, using analyses of their MAB counterparts, which are otherwise studied separately from MAB with often more complex solutions and worse guarantees.
2. **First Best-of-Both World regret for DB.** Applying the above reduction to a Best-of-Both world algorithm from MAB, we provide an algorithm that simultaneously guarantees a pseudo-regret bound $O(\sqrt{KT})$ in the adversarial setting and $O(K \log(T)/\Delta_{\min})$ in the stochastic one.
3. **Robustness to adversarial corruptions.** Our algorithm is robust to adversarial corruptions and significantly improves existing results in DB Agarwal et al. [2021].
4. **Optimal stochastic gap-dependent Regret.** Our algorithm also provides the first optimal Condorcet regret, which suffers neither from a suboptimal dependence on Δ_{\min}^{-2} nor from a quadratic dependence on the number of arms.
5. **Another easy algorithm for stochastic DB.** We also propose a new very simple elimination algorithm with $O(\sum_{i=2}^K \frac{K \log T}{\Delta_i})$ Condorcet regret.
6. **Experimental evaluations.** Finally we corroborate our theoretical results with extensive empirical evaluations (Sec. 7).

2 Problem Formulation

Notations. Decision space (or item/arm set) $[K] := \{1, 2, \dots, K\}$. For any matrix $\mathbf{M} \in \mathbb{R}^{K \times K}$, we define $m_{ij} := M(i, j)$, $\forall i, j \in [K]$. $\mathbf{1}(\cdot)$ denotes the indicator random variable

which takes value 1 if the predicate is true and 0 otherwise and \lesssim a rough inequality which holds up to universal constants. For any two items $x, y \in [K]$, we use the symbol $x \succ y$ to denote x is preferred over y . By convention, we set $\frac{0}{0} := 0.5$.

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

Objective. 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 $k^* \in [K]^1$ in hindsight by calculating the regret w.r.t. $k^* \in [K]$

$$R_T(k^*) := \sum_{t=1}^T \frac{1}{2} (P_t(k^*, k_{+1,t}) + P_t(k^*, k_{-1,t}) - 1). \quad (1)$$

For the *stochastic setting* where P_t s are fixed across all time steps $t \in [T]$, we denote $P_t = P \forall t \in [T]$. Further assuming there exists a Condorcet winner for P , i.e. fixed arm $k^{(\text{cw})} \in [K]$ such that $P(k^{(\text{cw})}, k) > 0.5 \forall k \in [K] \setminus \{k\}$, the above notion of regret boils down to the regret with respect to the Condorcet winner for the choice of $k^* = k^{(\text{cw})}$, as widely studied in many dueling bandit literature Zoghi et al. [2014a], Wu and Liu [2016], Komiyama et al. [2015], Bengs et al. [2021], defined as:

$$R_T^{(\text{cw})} := \sum_{t=1}^T \frac{1}{2} (\Delta(k^{(\text{cw})}, k_{+1,t}) + \Delta(k^{(\text{cw})}, k_{-1,t})), \quad (2)$$

where $\Delta(i, j) := P(i, j) - 1/2$ being the suboptimality gap of item i and j in terms of their relative preferences.

3 Warm-Up: Near-Optimal Algorithm

In this section, we propose a new simple UCB based algorithm for stochastic dueling bandit, which is shown to have a nearly optimal gap-dependent Condorcet regret of $R_T^{(\text{cw})} = O\left(\sum_{i=2}^K i \Delta(k^{(\text{cw})}, k)^{-1} \log T\right)$. Note that existing dueling bandit algorithms, that satisfy a non-asymptotic Condorcet regret bound, suffer an additional constant of order Δ_{\min}^{-2} [Bengs et al., 2021], which implies a worst-case regret of order $O(T^{2/3})$ when $\Delta_{\min} \rightarrow 0$. The simple elimination algorithm below solves this drawback and depends only on Δ_{\min}^{-1} but at the cost of a suboptimal quadratic dependence in the number of arms K . Despite our efforts, we could not avoid this suboptimal factor by following the classical stochastic dual bandit analysis. In the following sections, we will show how to easily reach the optimal dependence in both Δ_{\min} and K using a simple reduction from standard MAB.

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

Main Ideas: Algorithm RR-DB The high-level idea of Algorithm 1 is to sequentially compare arms in a round-robin fashion and eliminate arms when they are significantly suboptimal compared to any other arm. Typically, after t rounds, a suboptimal arm k has been compared at least t/K times with the Condorcet winner. Denoting by Δ_k its suboptimality gap, the arm is eliminated after at most t_k rounds, where $(t_k/K)^{-1/2} \approx \Delta_k$. At that time the arm has been played t_k/K times, yielding a regret of order $(t_k/K) \times \Delta_k \approx (K/\Delta_k^2) \times \Delta_k = K/\Delta_k$. Summing over the arms yields a final regret of order $O(K^2/\Delta_{\min})$.

Algorithm 1 RR-DB (Near Optimal DB)

1: **input:** Arm set: $[K]$, Confidence parameter $\delta \in (0, 1)$
2: **init:** Active arms: $\mathcal{A}_1 := [K]$, $n_{ij}(t) \leftarrow 0$, $\forall i, j \in [K]$
3: **for** $t = 1, 2, \dots, T$ **do**
4: Play $(k_{+1,t}, k_{-1,t}) \in \operatorname{argmin}_{i,j \in \mathcal{A}_t} \{n_{ij}(t-1)\}$
5: Observe $o_t(k_{+1,t}, k_{-1,t}) = 1 - o_t(k_{-1,t}, k_{+1,t})$
6: **for** $i, j \in \mathcal{A}_t$ **do**
7: Define $\mathbf{1}_t(i, j) := \mathbf{1}\{\{i, j\} = \{k_{-1,t}, k_{+1,t}\}\}$ and
 $n_{ij}(t) := \sum_{s=1}^t \mathbf{1}_t(i, j)$
 $\hat{p}_{ij}(t) := \frac{1}{n_{ij}(t)} \sum_{s=1}^t o_t(i, j) \mathbf{1}_t(i, j)$
 $u_{ij}(t) := \hat{p}_{ij}(t) + \sqrt{\frac{\log(Kt/\delta)}{n_{ij}(t)}}$
 where we assume $x/0 = +\infty$.
8: **end for**
9: $\mathcal{A}_{t+1} := \mathcal{A}_t \setminus \left\{ i \in \mathcal{A}_t : \exists j \in \mathcal{A}_t \text{ s.t. } u_{ij}(t) < \frac{1}{2} \right\}$.
10: **end for**

Without loss of generality assume the Condorcet winner $k^{(cw)} = 1$, and denote $\Delta_i = \Delta(1, i)$, $\forall i \in [K] \setminus \{1\}$.

Theorem 1. *Let $\delta \in (0, 1/2)$, for any $T \geq 1$, the regret of Algorithm 1 is upper-bounded with probability at least $1 - \delta$ as*

$$R_T^{(cw)} \leq \frac{K^2}{2} + 4 \sum_{i=2}^K \frac{(i-1) \log(KT/\delta)}{\Delta_i}.$$

Further, when $T \geq K^2$, in the worst case (over the problem instance, i.e. $\Delta_2, \dots, \Delta_K$), the regret of Algorithm 1 can be upper bounded as:

$$R_T^{(cw)} \leq 2K \sqrt{T \log(KT/\delta)}.$$

Remark 1. *In particular, our regret analysis shows that, except a logarithmic factor, the regret bound of RR-DB (Alg. 1) is off only by a multiplicative factor of K , as follows from the known $\Omega(\sum_{k=1}^K \frac{\log T}{\Delta_k})$ Condorcet winner regret lower bound Yue et al. [2012], Komiyama et al. [2015].*

The proof is postponed to Appendix A.

4 Reduction from MAB to DB

We now present a simple reduction from a multi-armed bandit algorithm to a dueling bandit one. The latter was already proposed by Gupta and Saha [2021] to show worst-case guarantees and by Ailon et al. [2014] for utility based dueling bandits only. We recall it here since it is central to our analysis and we believe that it is of important interest for the dueling bandit community that usually uses significantly different algorithms and analysis than the ones from standard multi-armed bandits.

The main idea is to apply the multi-armed bandit algorithm independently to two players $i \in \{-1, +1\}$ respectively with losses defined for any $k \in [K]$ and $t \in [T]$, by

$$\ell_{i,t}(k) := o_t(k_{-i,t}, k),$$

where $o_t(k, k') = 1 - o_t(k', k)$ for $1 < k' \leq K$ follows a Bernoulli with parameter $P_t(k, k')$ (and we assume $o_t(k, k) = 1/2$).

Algorithm 2 Reduction from MAB to DB

- 1: **input:** Arm set: $[K]$, two instances \mathcal{A}_i of an algorithm for MAB, $i \in \{-1, +1\}$.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: **for** $i \in \{+1, -1\}$ **do**
 - 4: choose $\mathbf{p}_{i,t}$ from \mathcal{A}_i
 - 5: sample $k_{i,t}$ from the distribution $p_{i,t}$
 - 6: **end for**
 - 7: Observe preference feedback $o_t(k_{+1,t}, k_{-1,t})$ and set $o_t(k_{-1,t}, k_{+1,t}) = 1 - o_t(k_{+1,t}, k_{-1,t})$.
 - 8: Feed \mathcal{A}_i with loss $\ell_{i,t}(k) := o_t(k_{-i,t}, k_{i,t})$ for $i \in \{-1, +1\}$.
 - 9: **end for**
-

We show below that any MAB regret upper-bound satisfied by \mathcal{A}_i can be transformed into a DB regret upper-bound.

Theorem 2. *Define for $i \in \{-1, 1\}$ and $k \in [K]$ by*

$$R_{i,T}(k) := \sum_{t=1}^T \ell_{i,t}(k_{i,t}) - \ell_{i,t}(k)$$

the regret achieved by algorithm \mathcal{A}_i . Then, the expected regret (1) of Algorithm 2 for dueling bandits can be decomposed as

$$\mathbf{E}[R_T(k)] = \frac{1}{2} \mathbf{E}[R_{-1,T}(k) + R_{+1,T}(k)].$$

Proof. The proof follows from

$$\begin{aligned} \mathbf{E}[\ell_{-1,t}(k) + \ell_{+1,t}(k)] &= \mathbf{E}[o_t(k_{+1,t}, k) + o_t(k_{-1,t}, k)] \\ &= \mathbf{E}[P_t(k_{+1,t}, k) + P_t(k_{-1,t}, k)] \\ &= 2 - \mathbf{E}[P_t(k, k_{+1,t}) + P_t(k, k_{-1,t})] \end{aligned}$$

and

$$\begin{aligned} & \mathbf{E}[\ell_{-1,t}(k_{+1,t}) + \ell_{+1,t}(k_{-1,t})] \\ &= \mathbf{E}[o_t(k_{+1,t}, k_{-1,t}) + o_t(k_{-1,t}, k_{+1,t})] = 1. \end{aligned}$$

We conclude by summing over $t = 1, \dots, T$ both equations and by substituting them into the definition of the regret $R_T(k)$ in (1). \square

Note that such a reduction can also be used to bound $R_T(k)$ directly rather than its expectation. Gupta and Saha [2021] indeed use this reduction to show a $O(\sqrt{KT})$ high-probability regret upper-bound for adversarial dueling bandit. They also obtain non-stationary regret bounds.

The main message of this paper is that this reduction can be used to transpose many results from standard multi-armed bandit to general dueling bandits. For instance, applying a subroutine \mathcal{A}_i which is robust to delays [Thune et al., 2019, Zimmert and Seldin, 2020, e.g.], one directly obtains a dueling bandit with the same robustness guarantees.

As we said, this reduction is not new. However, to date, it has only been considered in two specific contexts: the adversarial setting with worst-case regret bounds of order $O(\sqrt{KT})$ and the utility-based setting. Since the losses $\ell_{i,t} = o_t(k_{-i,t}, k)$ are not i.i.d. but depend on an adaptive adversary which chooses $k_{-i,t}$, one cannot use stochastic bandit algorithms. And the dueling bandit community usually needs to resort to more sophisticated algorithms and arguments to obtain logarithmic regret bounds for Condorcet stochastic dueling bandits. The only stochastic dueling bandit for which such a reduction was considered [Ailon et al., 2014, Zimmert and Seldin, 2021, e.g.] was the utility based-dueling bandits, which is overly restrictive in practice. That is, when the preference matrix is of the form $P_t(k, k') : (1 + u_t(k) + u_t(k'))/2$ for some sequence of utility vectors $(u_t)_{t \geq 1}$. Utility based dueling bandit are known to be easily reduced to two independent multi-armed bandit problems [Ailon et al., 2014].

Our main contribution is to show that this reduction can in fact be easily extended to the much weaker Condorcet winner hypothesis. To do this, as we show in the next sections, we simply apply the reduction with a best-of-both-worlds multi-armed bandit algorithm. As we will see, this recovers and improves the best existing upper bounds on the Condorcet regret for dueling bandits.

5 Best-of-Both Dueling: Optimal Algorithm for Stochastic and Adversarial DB

This section contains our main result, which is a simple reduction of the best-of-both-worlds result from Zimmert and Seldin [2021] to dueling bandits. In particular, it allows us to improve the best existing upper bound on the regret for stochastic and corrupted dueling bandits. The main idea is to apply the reduction of Algorithm 2 with the multi-armed bandit algorithm (Online Mirror Descent with Tsallis regularizer) of Zimmert and Seldin [2021]. Of

course, as for the classical adversarial multi-armed bandits, the losses $\ell_{i,t}(k) = o_t(k_{-i,t}, k)$ cannot be observed for all $k \in [K]$. Therefore, they are estimated in the algorithm with the importance weight estimators

$$\widehat{\ell}_{i,t}(k) = \begin{cases} \ell_{i,t}(k)/p_{i,t}(k) & \text{if } k = k_{i,t} \\ 0 & \text{otherwise} \end{cases}. \quad (3)$$

The resulted algorithm is described in Algorithm 3.

Algorithm 3 Versatile-DB (Best-of-Both DB)

- 1: **input:** Arm set: $[K]$, $(\Psi_t)_{t=1,2,\dots}$
 - 2: **init:** $G_{i,0} \leftarrow \mathbf{0}_K$ for $i \in \{+1, -1\}$
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: choose $\mathbf{p}_{i,t} = \nabla(\Psi_t + \mathcal{I}_\Delta)^*(-\widehat{L}_{i,t-1})$
 - 5: For $i \in \{+1, -1\}$, sample $k_{i,t}$ from the distribution $(p_{i,t}(1), \dots, p_{i,t}(K))$
 - 6: Observe preference feedback $o_t(k_{+1,t}, k_{-1,t})$
 - 7: Compute $\widehat{\ell}_{i,t}(k)$ for $i \in \{+1, -1\}$ and $k \in [K]$ using (3)
 - 8: update $\widehat{L}_{i,t} = \widehat{L}_{i,t-1} + \widehat{\ell}_{i,t}$
 - 9: **end for**
-

Theorem 3. *For any sequence of preference matrices P_t , the pseudo-regret of Algorithm 3 with $\Psi_t(w) = \sqrt{t} \sum_{k=1}^K (\sqrt{w_k} - w_k/2)/8$ satisfies for any $T \geq 1$*

$$\overline{R}_T := \max_{k \in [K]} \mathbf{E}[R_T(k)] \leq 4\sqrt{KT} + 1.$$

Furthermore, if there exists a gap vector $\Delta \in [0, 1]^K$ with a unique zero coordinate $k^* \in [K]$ and $C \geq 0$ such that

$$\overline{R}_T \geq \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T \sum_{k \neq k^*} (p_{+1,t}(k) + p_{-1,t}(k)) \Delta_k \right] - C, \quad (4)$$

the pseudo regret also satisfies

$$\overline{R}_T \leq \sum_{k \neq k^*} \frac{4 \log T + 12}{\Delta_k} + 4 \log T + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 8 + C,$$

where $\Delta_{\min} = \min_{k \neq k^*} \Delta_k$.

The proof is postponed to Appendix B. Note that the theorem largely follows from (and is itself highly similar to) the best-of-both worlds regret-bound of [Zimmert and Seldin, 2021, Theorem 1] for MAB. We insist on the fact that our contribution should not be seen as technical. Indeed, the proof is just a clever combination of their MAB analysis with our black box reduction (Theorem 2). But we believe that the simplicity of our approach is its strength

that can benefit the community of dueling bandits. As we will see, several state-of-the-art results of dueling bandits can be simultaneously improved as a direct consequence of this theorem.

Note that for simplicity, we restricted ourselves to importance weighted estimators (3). By using more sophisticated variance reduced estimators, as in Zimmert and Seldin [2021], the multiplicative constants can be reduced. Furthermore, similar to Zimmert and Seldin [2021], the result holds only for the pseudo-regret and not for the true regret. Auer and Chiang [2016] have indeed proven that no optimal adversarial and stochastic high probability regret bounds can be obtained simultaneously for standard stochastic bandits. The learner must pay suboptimal logarithmic factors. The result can be extended to dueling bandits.

It is worth to emphasize that this is the first best-of-both worlds regret bound for general dueling bandits (the stochastic bound follows from the choice $C = 0$, see Sec 6.1). Zimmert and Seldin [2021, Cor. 10] obtain a similar result for the very same algorithm but for utility based dueling bandits only.

Remark 2. *Note that a single sub-routine of OMD to optimize the weights is actually enough to get the same regret guarantee. To do so, one samples both $k_{-1,t}$ and $k_{+1,t}$ independently from the same distribution $p_t = \nabla(\psi_t + \mathcal{I}_\Delta)^*(-\hat{L}_{t-1})$. Here, $\hat{L}_t = \sum_{s=1}^t \hat{\ell}_s \in \mathbb{R}_+^K$ and the importance weight estimator are defined for all $k \in [K]$ by*

$$\hat{\ell}_t(k) = \frac{1}{2}(\hat{\ell}_{-1,t}(k) + \hat{\ell}_{+1,t}(k)).$$

Noting that $\mathbf{E}[\hat{\ell}_t(k)] = \mathbf{E}[P_t(k_{-1,t}, k) + P_t(k_{+1,t}, k)]/2$ and $\mathbf{E}[\sum_{k=1}^K p_t(k)\hat{\ell}_t(k)] = \mathbf{E}[P_t(k_{-1,t}, k_{+1,t}) + P_t(k_{+1,t}, k_{-1,t})]/2 = 1/2$, the proof follows similarly to the other one. Though the regret upper-bound is exactly the same, we believe that this version might lead to better performance because the two players share information.

6 Improvements Over Existing Dueling Bandit Results

Also our approach and analysis is rather simple, it allows to outperform the best existing regret-upper bounds for stochastic dueling bandits with or without corruption. We believe that the dueling bandit community will benefit from this reduction and that it may be applied to a wider scope such as to deal with non-stationarity, delays, or non-standard feedbacks (graphs) for which many results already exist in standard multi-armed bandits.

6.1 Stochastic Dueling Bandits with Condorcet Winner

In stochastic dueling bandits, the preference matrices P_t are fixed over time $P_t = P$ for all $t \geq 1$. Under the Condorcet winner assumption there exists $k^{(cw)} \in [K]$ such that $P(k^{(cw)}, k) > \frac{1}{2}$ for all $k \neq k^{(cw)}$. Then, the suboptimality gaps of all actions $k \in [K]$ are defined as $\Delta_k := P(k^{(cw)}, k) - \frac{1}{2}$. Remarking that in this case the self-bounding assumption (4)

is satisfied with $C = 0$, since

$$\begin{aligned}\bar{R}_T &= \frac{1}{2} \sum_{t=1}^T \mathbf{E} \left[P_t(k^{(\text{cw})}, k_{+1,t}) + P_t(k^{(\text{cw})}, k_{-1,t}) - 1 \right] \\ &= \frac{1}{2} \sum_{t=1}^T \mathbf{E} \left[\Delta_{k_{+1,t}} + \Delta_{k_{-1,t}} \right] \\ &= \frac{1}{2} \sum_{t=1}^T \mathbf{E} \left[\sum_{k \neq k^{(\text{cw})}} (p_{-1,t} + p_{+1,t}) \Delta_k \right],\end{aligned}$$

we get the following corollary from Theorem 3.

Corollary 1. *For stochastic dueling bandits with Condorcet winner, the pseudo-regret of Algorithm 3 with well-chosen parameters satisfies*

$$\bar{R}_T \leq \sum_{k \neq k^{(\text{cw})}} \frac{4 \log T + 12}{\Delta_k} + 4 \log T + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 8.$$

Note that the above bound is the first pseudo-regret upper-bound for stochastic dueling bandit that does not suffer from a Δ_{\min}^{-2} dependence under the Condorcet winner assumption, without a quadratic dependence on the number of arms, K , which is a concern when it comes to dealing with large-scale problems. For instance, RUCB Zoghi et al. [2014a] satisfies a regret bound of order $O(K \log(T) \Delta_{\min}^{-2} + K^2)$, MergeRUCB Zoghi et al. [2015b] has linear dependence on K but suffers $O(K \log(T) \Delta_{\min}^{-2})$. Finally, RMED from Komiyama et al. [2015] is asymptotically optimal when $T \rightarrow \infty$ but also suffers from large constant terms (K^2 and Δ_{\min}^{-2}) and is only valid for $K \rightarrow \infty$. We refer the reader to Bengs et al. [2021] for existing results on stochastic dueling bandits.

6.2 Corrupted Regime

Here, we consider the stochastic dueling bandit problem in the presence of adversarial corruptions. The robustness to adversarial corruption has known recent progress in the MAB setting [Gupta et al., 2019, Lykouris et al., 2018, Zimmert and Seldin, 2021] and has recently been extended to the DB framework [Agarwal et al., 2021]. The preference matrices are fixed $P_t = P$ for all $t \geq 1$ and we assume the existence of a Condorcet winner $k^{(\text{cw})}$. Furthermore, an adversary may corrupt the outcomes of some duels by replacing the results of the duels $o_t(k, k')$ with corrupted ones $\tilde{o}_t(k, k')$. At the end of each round, the player only observes $\tilde{o}_t(k_{+1,t}, k_{-1,t})$. The objective of the player is to minimize the pseudo-regret \bar{R}_T under a bounded total amount of corruption

$$C := \sum_{t=1}^T \sum_{k \neq k^{(\text{cw})}} \left| o_t(k^{(\text{cw})}, k) - \tilde{o}_t(k^{(\text{cw})}, k) \right|.$$

We show here that similarly to what happens for standard Multi-armed bandits in Zimmert and Seldin [2021], this corrupted setting is a special case of the self-bounding assumption (4).

Indeed, defining \tilde{P}_t the corrupted preference matrices by $\tilde{P}_t(k, k') = \mathbf{E}[\tilde{o}_t(k, k')]$ and the corrupted pseudo-regret

$$\tilde{R}_T(k) = \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T \tilde{P}_t(k, k_{+1,t}) + P_t(k, k_{-1,t}) - 1 \right],$$

we have

$$\begin{aligned} & \tilde{R}_T(k^{(\text{cw})}) \\ &= \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T \tilde{o}_t(k^{(\text{cw})}, k_{+1,t}) + \tilde{o}_t(k^{(\text{cw})}, k_{-1,t}) - 1 \right] \\ &\geq \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T o_t(k^{(\text{cw})}, k_{+1,t}) + o_t(k^{(\text{cw})}, k_{-1,t}) - 1 \right] - C \\ &= \bar{R}_T - C \\ &= \frac{1}{2} \sum_{t=1}^T \mathbf{E} \left[\sum_{k \neq k^{(\text{cw})}} (p_{-1,t} + p_{+1,t}) \Delta_k \right] - C. \end{aligned}$$

Therefore, the corrupted regime satisfies the self-bounding assumption (4). Applying Theorem 3 on the corrupted regime and using that we also have $\bar{R}_T(k^{(\text{cw})}) \leq \tilde{R}_T(k^{(\text{cw})}) + C$, we get the following corollary.

Corollary 2. *For stochastic dueling bandits with Condorcet winner and corruptions, whose total amount is bounded by C , the pseudo-regret \bar{R}_T Algorithm 3 is upper-bounded as*

$$\bar{R}_T \leq \sum_{k \neq k^{(\text{cw})}} \frac{4 \log T + 12}{\Delta_k} + 4 \log T + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 8 + 2C.$$

Although Corollary 2 easily follows from Theorem 3 which itself easily follows from Theorem 1 of Zimmert and Seldin [2021], the latter result significantly improves upon the recent results on dueling bandit with corruptions obtained by Agarwal et al. [2021]. Indeed, the latter provide for the same setting and a significantly more sophisticated procedure a high-probability regret bound of order

$$O \left(\frac{K^2 C}{\Delta_{\min}} + \sum_{k \neq k^{(\text{cw})}} \frac{K^2}{\Delta_k^2} \log \left(\frac{K}{\Delta_k} \right) + \sum_{k \neq k^{(\text{cw})}} \frac{\log T}{\Delta_k} \right).$$

As often in the dueling bandit literature, it suffers from both a quadratic dependence on the number of actions and Δ_{\min}^{-1} . Furthermore, our regret bound is sublinear in T as soon as the corruption level is $o(T)$ while Agarwal et al. [2021] can only afford $o(\Delta_{\min} T / K^2)$.

Moreover, Zimmert and Seldin [2021] also provide an upper-bound for stochastic bandits with adversarial corruption. The latter is of order $O \left(\sum_{k \neq k^{(\text{cw})}} \frac{\log T}{\Delta_k} + \sqrt{C \sum_{k \neq k^{(\text{cw})}} \frac{\log T}{\Delta_k}} \right)$, which seems to outperform our bound when C is large. The difference is since they upper-bound the corrupted regret $\tilde{R}_T(k^{(\text{cw})})$ and not \bar{R}_T .

7 Experiments

Algorithms. We compared the performances of the following algorithms: 1. VDB: Our proposed **Versatile-DB** (Alg. 3) algorithm. 2. RUCB: The algorithm proposed in Zoghi et al. [2014a] for K-armed stochastic dueling bandits for CW regret. (We set the algorithm parameter $\alpha = 0.6$). 3. RMED: Another algorithm for CW regret as proposed in Komiyama et al. [2015]. (We set the algorithm parameter $f(K) = 0.3K^{1.01}$ as suggested in their experimental evaluation). 4. DTS: The double thompson sampling algorithm of Wu and Liu [2016]. (Here again we set the similar algorithm parameter $\alpha = 0.6$). 5. REX3: As introduced in Gajane et al. [2015]. Note that their suggested optimal tuning parameters, i.e. the uniform exploration rate γ as well as the learning rate η requires the knowledge of problem dependent parameters τ (see Thm. 1 of Gajane et al. [2015]) which are unknown to the learner. We used T in place of τ henceforth.

Performance Measures. We report the average cumulative regret (Eqn. (1)) of the algorithms averaged over 20 runs.

7.1 Stochastic Preferences

We compared their regret performance across the following stochastic environments:

Constructing Preference Matrices (\mathbf{P}). We use four different utility parameter $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)$ based preference models where the underlying preference model is defined as $P(i, j) := \frac{\theta_i}{\theta_i + \theta_j} \forall i, j \in [K]$. The model is famously studied as BTL model or model generally Plackett-Luce model Saha and Gopalan [2019b], Chen et al. [2018], Negahban et al. [2017]. Note this ensures P to have *total-ordering* Yue and Joachims [2011], Falahatgar et al. [2017].

In particular we consider the following choices of $\boldsymbol{\theta}$: 1. *Trivial* 2. *Easy* 3. *Medium*, and 4. *Hard* with their respective $\boldsymbol{\theta}$ parameters are given by: 1. *Trivial*: $\boldsymbol{\theta}(1) = 1, \boldsymbol{\theta}(2 : K) = 0.5$. 2. *Easy*: $\boldsymbol{\theta}(1 : \lfloor K/2 \rfloor) = 1, \boldsymbol{\theta}(\lfloor K/2 \rfloor + 1 : K) = 0.5$. 3. *Medium*: $\boldsymbol{\theta}(1 : \lfloor K/3 \rfloor) = 1, \boldsymbol{\theta}(\lfloor K/3 \rfloor + 1 : \lfloor 2K/3 \rfloor) = 0.7, \boldsymbol{\theta}(\lfloor 2K/3 \rfloor + 1 : K) = 0.4$. 4. *Hard*: $\boldsymbol{\theta}(i) = 1 - (i-1)/K, \forall i \in [K]$. Note for each $\boldsymbol{\sigma}^* = (1 > 2 > \dots > K)$. For the purpose of our experiments we set $K = 10$. We also evaluated the algorithms on two general 10×10 preference matrices *Car* and *Hurdy* as also used in Niranjana and Rajkumar [2017], Saha and Gopalan [2018a].

Regret vs Time. Fig. 1 shows the relative performances of different algorithms with time. As follows from the plots, in general *VDB* (**Versatile-DB**) outperforms the rest in all instances, with *DTS* being closely competitive in some cases. In terms of the problem hardness, as their names suggest too, the *Trivial* and *Easy* instances are easiest to learn as the best-vs-worst item preferences are well separated in these cases and the diversity of the item preferences across different groups are least. Consequently the algorithms yield slightly more regret on *instance-Medium* due to higher preference diversity, and the hardest instance being *Hard* where the algorithms require maximum time to converge, though *VDB* reaching the convergence fastest still.




Figure 1: Averaged cumulative regret over time

7.2 Corrupted Preferences

We also evaluated the performances of algorithms in presence of corruption (Sec. 6.2). In particular, Fig. 2 and 3 respectively shows the relative performances of the algorithms with 20% and 40% corrupted feedback (at each round, we flip the winner feedback with that certain probability) respectively on *Medium* and *Hard* Plackett-Luce instances. As expected, the performances of all the algorithms decay significantly with increasing degree of feedback-corruption, however as before, *VDB* consistently performed best over all the baselines and tend to converge the fastest among all.




Figure 2: Averaged cumulative regret (20% corrupted feedback)




Figure 3: Averaged cumulative regret (40% corrupted feedback)

8 Discussions

We studied the problem of **Versatile Dueling Bandits**, which gives the first *‘best-of-both’ world* result for the problem of Dueling Bandits. The crux of our analyses relies on a *novel idea of decomposing the dueling bandit regret into multiarmed bandit (MAB) regret* by interpreting the dueling preference feedback as a certain realization of adversarial reward sequence. An important byproduct of our best-of-both dueling analysis is, this gives the first order optimal gap-dependent regret bound for K -armed stochastic dueling bandits, closing the decade-long open problem of tightness of ‘Condorcet dueling bandit regret’. Further we also analyze the robustness of our algorithm under corrupted preference feedback setting, which provably improves over the state-of-the art corrupted dueling bandits algorithms.

Future Works. Proving the first *best-of-both world* result for dueling bandits using our novel reduction idea is just a first step towards exploring the possibility of understanding how far this idea can be extended to apply existing multiarmed bandits results to dueling bandits frameworks, instead of putting individual and isolated efforts in developing dueling bandit algorithms, taking inspirations from existing MAB generalizations. Some such extensions could be to analyze dynamic dueling bandit regret under non-stationary preferences Wei and Luo [2021], Chen et al. [2019], Besbes et al. [2015], item non-availability Neu and Valko [2014], Kanade et al. [2009], delayed feedback Vernade et al. [2018], Pike-Burke et al. [2018], Thune et al. [2019], budget constraints Immorlica et al. [2019], Zhou and Tomlin [2018], Ding et al. [2013], or even the more general reinforcement learning (RL) scenarios Auer et al. [2009], Talebi and Maillard [2018], Ng et al. [2006], for which there are already well established theory of works with MAB framework. Also under what settings of Dueling Bandits, its corresponding MAB counterpart based reductions are bound to fail?

Finally it is worth mentioning that, an ambitious (and broad) objective along these line of thoughts is to understand the connection between different learning scenarios to dueling bandits, e.g. feedback graphs Alon et al. [2015, 2017], partial monitoring problems Lattimore and Szepesvári [2019], Mannor et al. [2014], Lin et al. [2014], markov games Xie et al. [2020], Bai et al. [2020], Bai and Jin [2020], etc. The obvious motivation being to understand how far we can re-engineer the existing results from related learning literature to solve the online preference bandits problems.

Acknowledgment

Thanks to Julian Zimmert and Karan Singh for the useful discussions on the existing best-of-both-world multiarmed bandits results.

References

- Arpit Agarwal, Shivani Agarwal, and Prathamesh Patil. Stochastic dueling bandits with adversarial corruption. In *Algorithmic Learning Theory*, pages 217–248. PMLR, 2021.
- Nir Ailon, Zohar Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *International Conference on Machine Learning*, pages 856–864. PMLR, 2014.
- Noga Alon, Nicolo Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *JMLR Workshop and Conference Proceedings*, volume 40. Microtome Publishing, 2015.
- Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 46(6):1785–1826, 2017.
- Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 116–120. PMLR, 2016.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in neural information processing systems*, pages 89–96, 2009.
- Yu Bai and Chi Jin. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR, 2020.
- Yu Bai, Chi Jin, and Tiancheng Yu. Near-optimal reinforcement learning with self-play. In *Advances in Neural Information Processing Systems*, 2020.
- Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *J. Mach. Learn. Res.*, 22:7–1, 2021.
- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- Brian Brost, Yevgeny Seldin, Ingemar J. Cox, and Christina Lioma. Multi-dueling bandits and their application to online ranker evaluation. *CoRR*, abs/1608.06253, 2016.
- Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 42–1. JMLR Workshop and Conference Proceedings, 2012.

- Róbert Busa-Fekete, Balazs Szorenyi, Weiwei Cheng, Paul Weng, and Eyke Hüllermeier. Top-k selection based on adaptive sampling of noisy preferences. In *International Conference on Machine Learning*, pages 1094–1102, 2013.
- Róbert Busa-Fekete, Balázs Szörényi, and Eyke Hüllermeier. PAC rank elicitation through adaptive sampling of stochastic pairwise preferences. In *AAAI*, pages 1701–1707, 2014.
- Bangrui Chen and Peter I Frazier. Dueling bandits with weak regret. *arXiv preprint arXiv:1706.04304*, 2017.
- Xi Chen, Yuanzhi Li, and Jieming Mao. A nearly instance optimal algorithm for top-k ranking under the multinomial logit model. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2504–2522. SIAM, 2018.
- Yifang Chen, Chung-Wei Lee, Haipeng Luo, and Chen-Yu Wei. 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.
- Wenkui Ding, Tao Qin, Xu-Dong Zhang, and Tie-Yan Liu. Multi-armed bandit with budget constraint and variable costs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- Miroslav Dudík, Katja Hofmann, Robert E Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Conference on Learning Theory*, pages 563–587, 2015.
- Moein Falahatgar, Yi Hao, Alon Orlitsky, Venkatadheeraj Pichapati, and Vaishakh Ravindrakumar. Maxing and ranking with few assumptions. In *Advances in Neural Information Processing Systems*, pages 7063–7073, 2017.
- Moein Falahatgar, Ayush Jain, Alon Orlitsky, Venkatadheeraj Pichapati, and Vaishakh Ravindrakumar. The limits of maxing, ranking, and preference learning. In *International Conference on Machine Learning*, pages 1427–1436. PMLR, 2018.
- Pratik Gajane, Tanguy Urvoy, and Fabrice Clérot. A relative exponential weighing algorithm for adversarial utility-based dueling bandits. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 218–227, 2015.
- Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Conference on Learning Theory*, pages 1562–1578. PMLR, 2019.
- Shubham Gupta and Aadirupa Saha. Optimal and efficient dynamic regret algorithms for non-stationary dueling bandits. *arXiv preprint arXiv:2111.03917*, 2021.
- Nicole Immorlica, Karthik Abinav Sankararaman, Robert Schapire, and Aleksandrs Slivkins. Adversarial bandits with knapsacks. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 202–219. IEEE, 2019.

- Kevin Jamieson, Sumeet Katariya, Atul Deshpande, and Robert Nowak. Sparse dueling bandits. In *Artificial Intelligence and Statistics*, pages 416–424. PMLR, 2015.
- Varun Kanade, H Brendan McMahan, and Brent Bryan. Sleeping experts and bandits with stochastic action availability and adversarial rewards. 2009.
- Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *COLT*, pages 1141–1154, 2015.
- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Copeland dueling bandit problem: Regret lower bound, optimal algorithm, and computationally efficient algorithm. *arXiv preprint arXiv:1605.01677*, 2016.
- Wataru Kumagai. Regret analysis for continuous dueling bandit. In *Advances in Neural Information Processing Systems*, 2017.
- Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Conference on Learning Theory*, pages 2111–2139. PMLR, 2019.
- Tian Lin, Bruno Abrahao, Robert Kleinberg, John Lui, and Wei Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *International Conference on Machine Learning*, pages 901–909. PMLR, 2014.
- Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–122, 2018.
- Shie Mannor, Vianney Perchet, and Gilles Stoltz. Set-valued approachability and online learning with partial monitoring. *The Journal of Machine Learning Research*, 15(1): 3247–3295, 2014.
- Sahand Negahban, Sewoong Oh, and Devavrat Shah. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 65(1):266–287, 2017.
- Gergely Neu and Michal Valko. Online combinatorial optimization with stochastic decision sets and adversarial losses. In *Advances in Neural Information Processing Systems*, pages 2780–2788, 2014.
- Andrew Y Ng, Adam Coates, Mark Diel, Varun Ganapathi, Jamie Schulte, Ben Tse, Eric Berger, and Eric Liang. Autonomous inverted helicopter flight via reinforcement learning. In *Experimental robotics IX*, pages 363–372. Springer, 2006.
- UN Niranjana and Arun Rajkumar. Inductive pairwise ranking: going beyond the $n \log(n)$ barrier. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Ciara Pike-Burke, Shipra Agrawal, Csaba Szepesvari, and Steffen Grunewalder. Bandits with delayed, aggregated anonymous feedback. In *International Conference on Machine Learning*, pages 4105–4113, 2018.

- Wenbo Ren, Jia Liu, and Ness B Shroff. PAC ranking from pairwise and listwise queries: Lower bounds and upper bounds. *arXiv preprint arXiv:1806.02970*, 2018.
- Aadirupa Saha. Optimal algorithms for stochastic contextual dueling bandits. In *Advances in Neural Information Processing Systems*, 2021.
- Aadirupa Saha and Pierre Gaillard. Dueling bandits with adversarial sleeping. *Advances in Neural Information Processing Systems*, 34, 2021.
- Aadirupa Saha and Aditya Gopalan. Battle of bandits. In *Uncertainty in Artificial Intelligence*, 2018a.
- Aadirupa Saha and Aditya Gopalan. Active ranking with subset-wise preferences. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018b.
- Aadirupa Saha and Aditya Gopalan. Combinatorial bandits with relative feedback. In *Advances in Neural Information Processing Systems*, 2019a.
- Aadirupa Saha and Aditya Gopalan. PAC Battling Bandits in the Plackett-Luce Model. In *Algorithmic Learning Theory*, pages 700–737, 2019b.
- Aadirupa Saha and Aditya Gopalan. From pac to instance-optimal sample complexity in the plackett-luce model. In *International Conference on Machine Learning*, pages 8367–8376. PMLR, 2020.
- Aadirupa Saha and Akshay Krishnamurthy. Efficient and optimal algorithms for contextual dueling bandits under realizability. *arXiv preprint arXiv:2111.12306*, 2021.
- Aadirupa Saha, Tomer Koren, and Yishay Mansour. Adversarial dueling bandits. In *International Conference on Machine Learning*, pages 9235–9244. PMLR, 2021.
- Yanan Sui, Masrour Zoghi, Katja Hofmann, and Yisong Yue. Advancements in dueling bandits. In *IJCAI*, pages 5502–5510, 2018.
- Balazs Szorenyi, Robert Busa-Fekete, Adil Paul, and Eyke Hullermeier. Online rank elicitation for plackett-luce: A dueling bandits approach. In *Advances in Neural Information Processing Systems*, pages 604–612, 2015.
- Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. *arXiv preprint arXiv:1803.01626*, 2018.
- Tobias Sommer Thune, Nicolo Cesa-Bianchi, and Yevgeny Seldin. Nonstochastic multiarmed bandits with unrestricted delays. In *Advances in Neural Information Processing Systems*, pages 6541–6550, 2019.
- Claire Vernade, Alexandra Carpentier, Tor Lattimore, Giovanni Zappella, Beyza Ermis, and Michael Brueckner. Linear bandits with stochastic delayed feedback. *arXiv preprint arXiv:1807.02089*, 2018.

- Chen-Yu Wei and Haipeng Luo. Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. In *Proceedings of the 32nd International Conference on Learning Theory*, 2021.
- Huasen Wu and Xin Liu. Double Thompson sampling for dueling bandits. In *Advances in Neural Information Processing Systems*, pages 649–657, 2016.
- Qiaomin Xie, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In *Conference on Learning Theory*, 2020.
- Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208. ACM, 2009.
- Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 241–248, 2011.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k -armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.
- Datong P Zhou and Claire J Tomlin. Budget-constrained multi-armed bandits with multiple plays. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Julian Zimmert and Yevgeny Seldin. An optimal algorithm for adversarial bandits with arbitrary delays. In *International Conference on Artificial Intelligence and Statistics*, pages 3285–3294. PMLR, 2020.
- Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *J. Mach. Learn. Res.*, 22:28–1, 2021.
- Masrour Zoghi, Shimon Whiteson, Remi Munos, Maarten de Rijke, et al. Relative upper confidence bound for the k -armed dueling bandit problem. In *JMLR Workshop and Conference Proceedings*, number 32, pages 10–18. JMLR, 2014a.
- Masrour Zoghi, Shimon A Whiteson, Maarten De Rijke, and Remi Munos. Relative confidence sampling for efficient on-line ranker evaluation. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 73–82. ACM, 2014b.
- Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and Maarten De Rijke. Copeland dueling bandits. In *Advances in Neural Information Processing Systems*, pages 307–315, 2015a.
- Masrour Zoghi, Shimon Whiteson, and Maarten de Rijke. Mergerucb: A method for large-scale online ranker evaluation. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 17–26. ACM, 2015b.

Supplementary for Versatile Dueling Bandits: Best-of-both-World Analyses for Online Learning from Preferences

A Regret analysis of Algorithm 1

Theorem 1. *Let $\delta \in (0, 1/2)$, for any $T \geq 1$, the regret of Algorithm 1 is upper-bounded with probability at least $1 - \delta$ as*

$$R_T^{(cw)} \leq \frac{K^2}{2} + 4 \sum_{i=2}^K \frac{(i-1) \log(KT/\delta)}{\Delta_i}.$$

Further, when $T \geq K^2$, in the worst case (over the problem instance, i.e. $\Delta_2, \dots, \Delta_K$), the regret of Algorithm 1 can be upper bounded as:

$$R_T^{(cw)} \leq 2K \sqrt{T \log(KT/\delta)}.$$

Proof of Theorem 1. Let us denote by

$$u_{ij}(t) := \hat{p}_{ij}(t) + c_{ij}(t)$$

for any pair (i, j) and time t , where

$$c_{ij}(t) := \sqrt{\frac{\log(Kt/\delta)}{n_{ij}(t)}},$$

and assume $\Delta_2 \leq \Delta_3 \leq \dots \leq \Delta_K$ without loss of generality. We will also assume the confidence bounds of Lem. 4 holds good for all $t \in [T]$ and all pairs (i, j) , which is shown to hold good with probability at least $1 - \delta$. In particular, this implies that the best arm cannot be eliminated, i.e., $1 \in \mathcal{A}_t$ for all $t \geq 1$.

We start by noting that if the worst arm K (since $\Delta_K = \max_{i=2}^K \Delta_i$, arm- K gets maximally beaten by the CW) is played at time t , it means $u_{K1}(t) \geq \frac{1}{2}$. However we also have,

$$\begin{aligned} u_{K1}(t) &= \hat{p}_{K1}(t) + c_{K1}(t) \\ &\leq p_{K1} + 2c_{K1}(t) = 1/2 - \Delta_K + 2c_{K1}(t), \end{aligned}$$

where the inequality holds due to Lem. 4 and the last equality holds by noting $p_{K1} = 1 - p_{1K} = 1 - (1/2 + \Delta_K) = 1/2 - \Delta_K$.

So $u_{K1}(t) > 1/2$ can only hold good if $c_{K1}(t) > \Delta_K/2$ which implies,

$$n_{K1}(t) \leq \frac{4 \log(Kt/\delta)}{\Delta_K^2}. \tag{5}$$

But by the our algorithm design since all the pairs are drawn in a round robin fashion, at any round $t \in [T]$, for any two distinct pairs (i, j) and (i', j') that are in \mathcal{A}_t note that

$$|n_{ij}(t) - n_{i'j'}(t)| \leq 1. \quad (6)$$

Thus the total regret incurred by Alg. 1 at rounds where $K \in \{k_{+1,t}, k_{-1,t}\}$, can be upper bounded as:

$$\begin{aligned} & \sum_{t=1}^T \sum_{k < K} \mathbf{1}(\{k_{+1,t}, k_{-1,t}\} = \{k, K\}) \frac{\Delta_K + \Delta_k}{2} \\ & \leq \sum_{k=1}^{K-1} n_{kK}(T) \Delta_K \leq (K-1)(1 + n_{K1}(T)) \Delta_K \\ & \leq (K-1) \left(1 + \frac{4 \log(Kt/\delta)}{\Delta_K^2} \right) \Delta_K \\ & = (K-1) \left(\Delta_k + \frac{4 \log(Kt/\delta)}{\Delta_K} \right). \end{aligned}$$

Similarly, note for any $i \in \{2, 3, \dots, K-1\}$, we can upper bound the regret of rounds where i was played in the duel as:

$$\begin{aligned} & \sum_{t=1}^T \sum_{k < i} \mathbf{1}(\{k_{+1,t}, k_{-1,t}\} = \{k, i\}) \frac{\Delta_i + \Delta_k}{2} \\ & \leq \sum_{k=1}^{i-1} n_{ki}(T) \Delta_i \leq (i-1)(1 + n_{1i}(T)) \Delta_i \\ & \leq (i-1) \left(\Delta_i + \frac{4 \log(Kt/\delta)}{\Delta_i} \right). \end{aligned}$$

Thus we can bound the total regret of Algorithm 1 as:

$$\begin{aligned} & \sum_{t=1}^T \sum_{i=2}^K \sum_{k=1}^{i-1} \mathbf{1}(\{k_{+1,t}, k_{-1,t}\} = \{k, i\}) \frac{\Delta_i + \Delta_k}{2} \\ & \leq \sum_{i=2}^K (i-1) \left(\Delta_i + \frac{4 \log(Kt/\delta)}{\Delta_i} \right) \\ & \stackrel{(\Delta_i \leq 1/2)}{\leq} \frac{K^2}{4} + 4 \sum_{i=2}^K (i-1) \frac{\log(Kt/\delta)}{\Delta_i} \end{aligned}$$

which concludes the first half of the proof. Further, to show the second part of the claim (analyzing worst-case gap-independent regret bound of Algorithm 1), note that Eqn. (5) equivalently implies for any $i \in [K] \setminus \{1\}$:

$$\Delta_i \leq \sqrt{\frac{4 \log(Kt/\delta)}{n_{i1}(t)}}.$$

Hence we can alternatively upper bound the regret as:

$$\begin{aligned}
R_T &= \sum_{i=2}^K \sum_{k=1}^{i-1} n_{ik}(T) \frac{\Delta_j + \Delta_k}{2} \leq \sum_{i=2}^K \sum_{k=1}^{i-1} n_{ik}(T) \Delta_i \\
&\leq \sum_{i=2}^K \sum_{k=1}^{i-1} n_{ik}(T) \sqrt{\frac{4 \log(KT/\delta)}{n_{i1}(T)}} \\
&\stackrel{(a)}{\leq} 2 \sum_{i=2}^K \sum_{k=1}^{i-1} \sqrt{2n_{ik}(T) \log(Kt/\delta)} \\
&\stackrel{(b)}{\leq} \sum_{i=2}^K 2 \sqrt{K^2 \sum_{i=2}^K \sum_{k=1}^{i-1} n_{ik}(T) \log(KT/\delta)} \\
&\leq 2K \sqrt{T \log(KT/\delta)},
\end{aligned}$$

where (a) follows from the observation of Eqn. (6) which implies $n_{i1}(T) \geq n_{ik}(T)$ when $T \geq K^2$ and (b) from Jensen's inequality and $\sum_{i=2}^K (i-1) \leq K^2/2$. \square

Lemma 4. For any $\delta \in (0, 1/2)$. Then, with probability at least $1 - \delta$, for any pair $i, j \in [K]$ and any time $t \in [T]$

$$\hat{p}_{ij}(t) - c_{ij}(t) \leq p_{ij} \leq \hat{p}_{ij}(t) + c_{ij}(t), \quad \forall t \in [T],$$

where $c_{ij}(t) := \sqrt{\frac{\log(Kt/\delta)}{n_{ij}(t)}}$.

Proof. Let us denote by $u_{ij}(t) := \hat{p}_{ij}(t) + c_{ij}(t)$ and $\ell_{ij}(t) := \hat{p}_{ij}(t) - c_{ij}(t)$. Note the inequality holds trivially at round t , for any pair (i, j) for which $n_{ij}(t) = 0$ since in these cases $\ell_{ij}(t) \leq 0$ and $u_{ij}(t) \geq 1$.

Now consider any pair (i, j) and round $t \in [T]$ such that $n_{ij}(t) > 0$. Note in this case by Hoeffding's Inequality:

$$\begin{aligned}
Pr \left(|p_{ij} - \hat{p}_{ij}(t)| > \sqrt{\frac{\ln(Kt/\delta)}{n_{ij}(t)}} \right) \\
\leq 2e^{-2n_{ij}(t) \frac{\ln(Kt/\delta)}{n_{ij}(t)}} = \frac{2\delta^2}{K^2 t^2} \leq \frac{\delta}{K^2 t^2}.
\end{aligned}$$

Taking union bound over all $\binom{K}{2}$ pairs and time $t \in [T]$ we get:

$$\begin{aligned}
Pr \left(\exists i, j \in [K], t \in [T] \text{ s.t. } |p_{ij} - \hat{p}_{ij}(t)| > \sqrt{\frac{\ln(Kt/\delta)}{n_{ij}(t)}} \right) \\
\leq \sum_{t=1}^T \sum_{i=2}^K \sum_{j=1}^{i-1} \frac{\delta}{K^2 t^2} \leq \sum_{t=1}^{\infty} \frac{\delta}{2t^2} \leq \frac{\delta \pi^2}{12} \leq \delta,
\end{aligned}$$

where in the second last inequality we used $\sum_{t=1}^{\infty} \frac{1}{t^2} < \frac{\pi^2}{6}$. This concludes the claim. \square

B Regret Analysis of Alg. 3

Theorem 3. For any sequence of preference matrices P_t , the pseudo-regret of Algorithm 3 with $\Psi_t(w) = \sqrt{t} \sum_{k=1}^K (\sqrt{w_k} - w_k/2)/8$ satisfies for any $T \geq 1$

$$\bar{R}_T := \max_{k \in [K]} \mathbf{E}[R_T(k)] \leq 4\sqrt{KT} + 1.$$

Furthermore, if there exists a gap vector $\Delta \in [0, 1]^K$ with a unique zero coordinate $k^* \in [K]$ and $C \geq 0$ such that

$$\bar{R}_T \geq \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T \sum_{k \neq k^*} (p_{+1,t}(k) + p_{-1,t}(k)) \Delta_k \right] - C, \quad (4)$$

the pseudo regret also satisfies

$$\bar{R}_T \leq \sum_{k \neq k^*} \frac{4 \log T + 12}{\Delta_k} + 4 \log T + \frac{1}{\Delta_{\min}} + \frac{3}{2} \sqrt{K} + 8 + C,$$

where $\Delta_{\min} = \min_{k \neq k^*} \Delta_k$.

Proof of Theorem 3. The analysis follows from carefully combining our reduction (Theorem 2) with Theorem 1 of Zimmert and Seldin [2021] for MAB to both of the players. Indeed, for each player $i \in \{-1, 1\}$, Algorithm 3 chooses $k_{i,t}$ by following the decisions of Tsallis-INF [Zimmert and Seldin, 2021, Alg. 1] with $\alpha = 1/2$, symmetric regularization, learning rate $\eta_t = 4/\sqrt{t}$ and losses $\ell_{i,t}$ estimated in (3) with standard importance sampling (IW).

Adversarial regime A direct application of Theorem 1 of Zimmert and Seldin [2021], upper-bounds the pseudo-regret for each player $i \in \{-1, 1\}$ as

$$\max_{k \in [K]} \mathbf{E}[R_{i,T}(k)] \leq 4\sqrt{KT} + 1.$$

Combining the about bounds with the reduction from MAB to DB (Theorem 2) yields the adversarial pseudo-regret upper-bound

$$\mathbf{E}[R_T(k)] = \frac{1}{2} \mathbf{E}[R_{-1,T}(k) + R_{+1,T}(k)] \leq 4\sqrt{KT} + 1.$$

Adversarial regime with a self-bounding constraint Our self-bounding constraint is slightly different from that of Zimmert and Seldin [2021], since it involves both players simultaneously. This is necessary so that our gap vector Δ can recover the standard suboptimality gaps used in stochastic dueling bandits. Thus, we cannot directly combine their result with our black-box reduction in this case. However, the proof largely follows their analysis, except that the upper-bounds on the regret of both players must be combined in the middle of their analysis, just before they apply their self-bounding constraint assumption. Thus, we give here only the modification to the proof of their Theorem 1.

Following their proof of Thm. 1 until their pseudo-regret bound at the top of p. 23, we get for each player $i \in \mathcal{I} := \{-1, +1\}$:

$$\begin{aligned} & \mathbf{E}[R_{i,T}(k)] \\ & \leq \sum_{k \neq k^*} \left(\sum_{t=1}^T \frac{\sqrt{\mathbf{E}[p_{i,t}(k)]}}{\sqrt{t}} + \sum_{t=T_0+1}^T \frac{\mathbf{E}[p_{i,t}(k)]}{4\sqrt{t}} \right) + M, \end{aligned}$$

where $M := \sqrt{T_0} + \frac{3}{4}\sqrt{K} + 15 + 14K \log(T)$ and $T_0 := \lceil \Delta_{\min}^{-2}/4 \rceil$. Together with Theorem 2 and taking the max over k , \bar{R}_T is thus upper-bounded by

$$\frac{1}{2} \sum_{i \in \mathcal{I}} \sum_{k \neq k^*} \left(\sum_{t=1}^T \frac{\sqrt{\mathbf{E}[p_{i,t}(k)]}}{\sqrt{t}} + \sum_{t=T_0+1}^T \frac{\mathbf{E}[p_{i,t}(k)]}{4\sqrt{t}} \right) + M.$$

Now, applying the self-bounding property (4) we get for any $\lambda \in [0, 1]$

$$\bar{R}_T \leq \bar{R}_T + \lambda \left(\bar{R}_T - \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^T \sum_{k \neq k^*} (p_{+1,t}(k) + p_{-1,t}(k)) \Delta_k \right] + C \right)$$

Thus, combined with the previous bound using $1 + \lambda \leq 2$

$$\begin{aligned} \bar{R}_T & \leq \frac{1}{2} \sum_{i \in \mathcal{I}} \sum_{k \neq k^*} \left(\sum_{t=1}^T \left(\frac{2\sqrt{\mathbf{E}[p_{i,t}(k)]}}{\sqrt{t}} - \lambda \Delta_k \mathbf{E}[p_{i,t}] \right) \right. \\ & \quad \left. + \sum_{t=T_0+1}^T \frac{\mathbf{E}[p_{i,t}(k)]}{2\sqrt{t}} \right) + 2M + \lambda C. \\ & \leq \sum_{k \neq k^*} \left(\sum_{t=1}^{T_0} \max_{z \geq 0} \left\{ \frac{2\sqrt{z}}{\sqrt{t}} - \lambda \Delta_k z \right\} \right. \\ & \quad \left. + \sum_{t=T_0+1}^T \max_{z \geq 0} \left\{ \frac{2\sqrt{z} + \frac{1}{2}z}{\sqrt{t}} - \lambda \Delta_k z \right\} \right) + 2M + \lambda C \end{aligned}$$

Now, we are back with the same upper-bound Zimmert and Seldin [2021] have in the middle of their page 23. Following their analysis by solving the optimization problems, summing over t , and optimizing λ concludes. \square