Double Thompson Sampling for Dueling Bandits

Huasen Wu University of California, Davis hswu@ucdavis.edu Xin Liu University of California, Davis xinliu@ucdavis.edu

Abstract

In this paper, we propose a Double Thompson Sampling (D-TS) algorithm for dueling bandit problems. As its name suggests, D-TS selects both the first and the second candidates according to Thompson Sampling. Specifically, D-TS maintains a posterior distribution for the preference matrix, and chooses the pair of arms for comparison according to two sets of samples independently drawn from the posterior distribution. This simple algorithm applies to general *Copeland dueling bandits*, including *Condorcet dueling bandits* as a special case. For general Copeland dueling bandits, we show that D-TS achieves $O(K^2 \log T)$ regret. Moreover, using a back substitution argument, we refine the regret to $O(K \log T + K^2 \log \log T)$ in Condorcet dueling bandits and most practical Copeland dueling bandits. In addition, we propose an enhancement of D-TS, referred to as D-TS⁺, to reduce the regret in practice by carefully breaking ties. Experiments based on both synthetic and real-world data demonstrate that D-TS and D-TS⁺ significantly improve the overall performance, in terms of regret and robustness.

1 Introduction

The dueling bandit problem [1] is a variant of the classical multi-armed bandit (MAB) problem, where the feedback comes in the form of pairwise comparison. This model has attracted much attention as it can be applied in many systems such as information retrieval (IR) [2, 3], where user preferences are easier to obtain and typically more stable. Most earlier work [1, 4, 5] focuses on Condorcet dueling bandits, where there exists an arm, referred to as the Condorcet winner, that beats all other arms. Recent work [6, 7] turns to a more general and practical case of a Copeland winner(s), which is the arm (or arms) that beats the most other arms. Existing algorithms are mainly generalized from traditional MAB algorithms along two lines: 1) UCB (Upper Confidence Bound)-type algorithms, such as RUCB [4] and CCB [6]; and, 2) MED (Minimum Empirical Divergence)-type algorithms, such as RMED [5] and CW-RMED/ECW-RMED [7].

In traditional MAB, an alternative effective solution is Thompson Sampling (TS) [8]. Its principle is to choose the optimal action that maximizes the expected reward according to the randomly drawn belief. TS has been successfully applied in traditional MAB [9, 10, 11, 12] and other online learning problems [13, 14]. In particular, empirical studies in [9] show that TS not only achieves lower regret than other algorithms in practice, but is also more robust as a randomized algorithm.

In the wake of the success of TS in these online learning problems, a natural question is whether and how TS can be applied to dueling bandits to further improve the performance. However, it is challenging to apply the standard TS framework to dueling bandits, because not all comparisons provide information about the system statistics. Specifically, a good learning algorithm for dueling bandits will eventually compare the winner against itself. However, comparing one arm against itself does not provide any statistical information, which is critical in TS to update the posterior distribution. Thus, TS needs to be adjusted so that 1) comparing the winners against themselves is allowed, but, 2) trapping in comparing a non-winner arm against itself is avoided.

In this paper, we propose a Double Thompson Sampling (D-TS) algorithm for dueling bandits, including both Condorcet dueling bandits and general Copeland dueling bandits. As its name suggests, D-TS typically selects both the first and the second candidates according to samples independently drawn from the posterior distribution. D-TS also utilizes the idea of confidence bounds to eliminate the likely non-winner arms, and thus avoids trapping in suboptimal comparisons. Compared to prior studies on dueling bandits, D-TS has both practical and theoretical advantages.

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

First, the double sampling structure of D-TS better suits the nature of dueling bandits. Launching two independent rounds of sampling provides us the opportunity to select the same arm in both rounds and thus to compare the winners against themselves. This double sampling structure also leads to more extensive utilization of TS (e.g., compared to RCS [3]), and significantly reduces the regret. In addition, this simple framework applies to general Copeland dueling bandits and achieves lower regret than existing algorithms such as CCB [6]. Moreover, as a randomized algorithm, D-TS is more robust in practice.

Second, this double sampling structure enables us to obtain theoretical bounds for the regret of D-TS. As noted in traditional MAB literature [10, 15], theoretical analysis of TS is usually more difficult than UCB-type algorithms. The analysis in dueling bandits is even more challenging because the selection of arms involves more factors and the two selected arms may be correlated. To address this issue, our D-TS algorithm draws the two sets of samples independently. Because their distributions are fully captured by historic comparison results, when the first candidate is fixed, the comparison between it and all other arms is similar to traditional MAB and thus we can borrow ideas from traditional MAB. Using the properties of TS and confidence bounds, we show that D-TS achieves $O(K^2 \log T)$ regret for a general K-armed Copeland dueling bandit. More interestingly, the property that the sample distribution only depends on historic comparing results (but not t) enables us to refine the regret using a *back substitution* argument, where we show that D-TS achieves $O(K \log T + K^2 \log \log T)$ in Condorcet dueling bandits and many practical Copeland dueling bandits.

Based on the analysis, we further refine the tie-breaking criterion in D-TS and propose its enhancement called D-TS⁺. D-TS⁺ achieves the same theoretical regret bound as D-TS, but performs better in practice especially when there are multiple winners.

In summary, the main contributions of this paper are as follows:

- We propose a D-TS algorithm and its enhancement D-TS⁺ for general Copeland dueling bandits. The double sampling structure suits the nature of dueling bandits and leads to more extensive usage of TS, which significantly reduces the regret.
- We obtain theoretical regret bounds for D-TS and D-TS⁺. For general Copeland dueling bandits, we show that D-TS and D-TS⁺ achieve $O(K^2 \log T)$ regret. In Condorcet dueling bandits and most practical Copeland dueling bandits, we further refine the regret bound to $O(K \log T + K^2 \log \log T)$ using a back substitution argument.
- We evaluate the D-TS and D-TS⁺ algorithms through experiments based on both synthetic and real-world data. The results show that D-TS and D-TS⁺ significantly improve the overall performance, in terms of regret and robustness, compared to existing algorithms.

2 Related Work

Early dueling bandit algorithms study finite-horizon settings, using the "explore-then-exploit" approaches, such as IF [1], BTM [16], and SAVAGE [17]. For infinite horizon settings, recent work has generalized the traditional MAB algorithms to dueling bandits along two lines. First, RUCB [4] and CCB [6] are generalizations of UCB for Condorcet and general Copeland dueling bandits, respectively. In addition, [18] reduces dueling bandits to traditional MAB, which is then solved by UCB-type algorithms, called MutiSBM and Sparring. Second, [5] and [7] extend the MED algorithm to dueling bandits, where they present the lower bound on the regret and propose the corresponding optimal algorithms, including RMED for Condorcet dueling bandits [5], CW-RMED and its computationally efficient version ECW-RMED for general Copeland dueling bandits [7]. Different from such existing work, we study algorithms for dueling bandits from the perspective of TS, which typically achieves lower regret and is more robust in practice.

Dated back to 1933, TS [8] is one of the earliest algorithms for exploration/exploitation tradeoff. Nowadays, it has been applied in many variants of MAB [11, 12, 13] and other more complex problems, e.g., [14], due to its simplicity, good performance, and robustness [9]. Theoretical analysis of TS is much more difficult. Only recently, [10] proposes a logarithmic bound for the standard frequentist expected regret, whose constant factor is further improved in [15]. Moreover [19, 20] derive the bounds for its Bayesian expected regret through information-theoretic analysis.

TS has been preliminarily considered for dueling bandits [3, 21]. In particular, recent work [3] proposes a Relative Confidence Sampling (RCS) algorithm that combines TS with RUCB [4] for Condorcet dueling bandits. Under RCS, the first arm is selected by TS while the second arm is selected according to their RUCB. Empirical studies demonstrate the performance improvement of using RCS in practice, but no theoretical bounds on the regret are provided.

3 System Model

We consider a dueling bandit problem with K ($K \ge 2$) arms, denoted by $\mathcal{A} = \{1, 2, \ldots, K\}$. At each time-slot t > 0, a pair of arms $(a_t^{(1)}, a_t^{(2)})$ is displayed to a user and a noisy comparison outcome w_t is obtained, where $w_t = 1$ if the user prefers $a_t^{(1)}$ to $a_t^{(2)}$, and $w_t = 2$ otherwise. We assume the user preference is stationary over time and the distribution of comparison outcomes is characterized by the preference matrix $\mathbf{P} = [p_{ij}]_{K \times K}$, where p_{ij} is the probability that the user prefers arm i to arm j, i.e., $p_{ij} = \mathbb{P}\{i \succ j\}$, $i, j = 1, 2, \ldots, K$. We assume that the displaying order does not affect the preference, and hence, $p_{ij} + p_{ji} = 1$ and $p_{ii} = 1/2$. We say that arm i beats arm j if $p_{ij} > 1/2$.

We study the general Copeland dueling bandits, where the Copeland winner is defined as the arm (or arms) that maximizes the number of other arms it beats [6, 7]. Specifically, the Copeland score is defined as $\sum_{j\neq i} \mathbb{1}(p_{ij} > 1/2)$, and the normalized Copeland score is defined as $\zeta_i = \frac{1}{K-1} \sum_{j\neq i} \mathbb{1}(p_{ij} > 1/2)$, where $\mathbb{1}(\cdot)$ is the indicator function. Let ζ^* be the highest normalized Copeland score, i.e., $\zeta^* = \max_{1 \le i \le K} \zeta_i$. Then the Copeland winner is defined as the arm (or arms) with the highest normalized Copeland score, i.e., $\mathcal{C}^* = \{i : 1 \le i \le K, \zeta_i = \zeta^*\}$. Note that the Condorcet winner is a special case of Copeland winner with $\zeta^* = 1$.

A dueling bandit algorithm Γ decides which pair of arms to compare depending on the historic observations. Specifically, define a filtration \mathcal{H}_{t-1} as the history before t, i.e., $\mathcal{H}_{t-1} = \{a_{\tau}^{(1)}, a_{\tau}^{(2)}, w_{\tau}, \tau = 1, 2, \ldots, t-1\}$. Then a dueling bandit algorithm Γ is a function that maps \mathcal{H}_{t-1} to $(a_t^{(1)}, a_t^{(2)})$, i.e., $(a_t^{(1)}, a_t^{(2)}) = \Gamma(\mathcal{H}_{t-1})$. The performance of a dueling bandit algorithm Γ is measured by its expected cumulative regret, which is defined as

$$R_{\Gamma}(T) = \zeta^* T - \frac{1}{2} \sum_{t=1}^T \mathbb{E} \big[\zeta_{a_t^{(1)}} + \zeta_{a_t^{(2)}} \big].$$
(1)

The objective of Γ is then to minimize $R_{\Gamma}(T)$. As pointed out in [6], the results can be adapted to other regret definitions because the above definition bounds the number of suboptimal comparisons.

4 Double Thompson Sampling

4.1 D-TS Algorithm

We present the D-TS algorithm for Copeland dueling bandits, as described in Algorithm 1 (time index t is omitted in pseudo codes for brevity). As its name suggests, the basic idea of D-TS is to select both the first and the second candidates by TS. For each pair (i, j) with $i \neq j$, we assume a beta prior distribution for its preference probability p_{ij} . These distributions are updated according to the comparison results $B_{ij}(t-1)$ and $B_{ji}(t-1)$, where $B_{ij}(t-1)$ (resp. $B_{ji}(t-1)$) is the number of time-slots when arm i (resp. j) beats arm j (resp. i) before t. D-TS selects the two candidates by sampling from the posterior distributions.

Specifically, at each time-slot t, the D-TS algorithm consists of two phases that select the first and the second candidates, respectively. When choosing the first candidate $a_t^{(1)}$, we first use the RUCB [4] of p_{ij} to eliminate the arms that are unlikely to be the Copeland winner, resulting in a candidate set C_t (Lines 4 to 6). The algorithm then samples $\theta_{ij}^{(1)}(t)$ from the posterior beta distribution, and the first candidate $a_t^{(1)}$ is chosen by "majority voting", i.e., the arm within C_t that beats the most arms according to $\theta_{ij}^{(1)}(t)$ will be selected (Lines 7 to 11). The ties are broken randomly here for simplicity and will be refined later in Section 4.3. A similar idea is applied to select the second candidate $a_t^{(2)}$, where new samples $\theta_{ia_t^{(1)}}^{(2)}(t)$ are generated and the arm with the largest $\theta_{ia_t^{(1)}}^{(2)}(t)$ among all arms with $l_{ia_t^{(1)}} \leq 1/2$ is selected as the second candidate (Lines 13 to 14).

The double sampling structure of D-TS is designed based on the nature of dueling bandits, i.e., at each time-slot, two arms are needed for comparison. Unlike RCS [3], D-TS selects both candidates using TS. This leads to more extensive utilization of TS and thus achieves much lower regret. Moreover, the two sets of samples are independently distributed, following the same posterior that is only determined by the comparison statistics $B_{ij}(t-1)$ and $B_{ji}(t-1)$. This property enables us to obtain an $O(K^2 \log T)$ regret bound and further refine it by a back substitution argument, as discussed later.

We also note that RUCB-based elimination (Lines 4 to 6) and RLCB (Relative Lower Confidence Bound)-based elimination (Line 14) are essential in D-TS. Without these eliminations, the algorithm

Algorithm 1 D-TS for Copeland Dueling Bandits

1: Init: $B \leftarrow \mathbf{0}_{K \times K}$; // B_{ij} is the number of time-slots that the user prefers arm i to j. 2: for t = 1 to T do

- *|| Phase 1: Choose the first candidate* $a^{(1)}$ 3:
- $U := [u_{ij}], L := [l_{ij}], \text{ where } u_{ij} = \frac{B_{ij}}{B_{ij} + B_{ji}} + \sqrt{\frac{\alpha \log t}{B_{ij} + B_{ji}}}, l_{ij} = \frac{B_{ij}}{B_{ij} + B_{ji}} \sqrt{\frac{\alpha \log t}{B_{ij} + B_{ji}}}, \text{ if } i \neq j, \text{ and } u_{ii} = l_{ii} = 1/2, \forall i; // \frac{x}{0} := 1 \text{ for any } x.$ 4:
- $\hat{\zeta}_i \leftarrow \frac{1}{K-1} \sum_{j \neq i} \mathbb{1}(u_{ij} > 1/2); // Upper bound of the normalized Copeland score.$ 5:
- 6:
- $\mathcal{C} \leftarrow \{i : \hat{\zeta}_i = \max_j \hat{\zeta}_j\};$ for $i, j = 1, \dots, K$ with i < j do 7:
- 8:
- 9:
- 10:
- Sample $\theta_{ij}^{(1)} \sim \text{Beta}(B_{ij}+1, B_{ji}+1);$ $\theta_{ji}^{(1)} \leftarrow 1 \theta_{ij}^{(1)};$ end for $a^{(1)} \leftarrow \underset{i \in \mathcal{C}}{\arg \max} \sum_{j \neq i} \mathbb{1}(\theta_{ij}^{(1)} > 1/2); // Choosing from C to eliminate likely non-winner$ 11: arms; Ties are broken randomly.
- // Phase 2: Choose the second candidate $a^{(2)}$ 12:
- Sample $\theta_{ia^{(1)}}^{(2)} \sim \text{Beta}(B_{ia^{(1)}} + 1, B_{a^{(1)}i} + 1)$ for all $i \neq a^{(1)}$, and let $\theta_{a^{(1)}a^{(1)}}^{(2)} = 1/2$; $a^{(2)} \leftarrow \underset{ia^{(1)} = 1/2}{\operatorname{arg\,max}} \theta_{ia^{(1)}}^{(2)}$; // Choosing only from uncertain pairs. 13:
- 14:
- $i:l_{ia^{(1)}} \le 1/2$
- // Compare and Update 15:
- Compare pair $(a^{(1)}, a^{(2)})$ and observe the result w; 16:
- Update **B**: $B_{a^{(1)}a^{(2)}} \leftarrow B_{a^{(1)}a^{(2)}} + 1$ if w = 1, or $B_{a^{(2)}a^{(1)}} \leftarrow B_{a^{(2)}a^{(1)}} + 1$ if w = 2; 17:
- 18: end for

may trap in suboptimal comparisons. Consider one extreme case in Condorcet dueling bandits¹: assume arm 1 is the Condorcet winner with $p_{1j} = 0.501$ for all j > 1, and arm 2 is not the Condorcet winner, but with $p_{2j} = 1$ for all j > 2. Then for a larger K (e.g., K > 4), without RUCB-based elimination, the algorithm may trap in $a_t^{(1)} = 2$ for a long time, because arm 2 is likely to receive higher score than arm 1. This issue can be addressed by RUCB-based elimination as follows: when chosen as the first candidate, arm 2 has a great probability to compare with arm 1; after sufficient comparisons with arm 1, arm 2 will have $u_{21}(t) < 1/2$ with high probability; then arm 2 is likely to be eliminated because arm 1 has $\hat{\zeta}_1(t) = 1 > \hat{\zeta}_2(t)$ with high probability. Similarly, RLCB-based elimination (Line 14, where we restrict to the arms with $l_{ia_t^{(1)}}(t) \le 1/2$) is important especially for non-Condorcet dueling bandits. Specifically, $l_{ia_t^{(1)}}(t) > 1/2$ indicates that arm i beats $a_t^{(1)}$ with high probability. Thus, comparing $a_t^{(1)}$ and arm *i* brings little information gain and thus should be

Regret Analysis 4.2

eliminated to minimize the regret.

Before conducting the regret analysis, we first introduce certain notations that will be used later.

Gap to 1/2: In dueling bandits, an important benchmark for p_{ij} is 1/2, and thus we let Δ_{ij} be the gap between p_{ij} and 1/2, i.e., $\Delta_{ij} = |p_{ij} - 1/2|$.

Number of Comparisons: Under D-TS, (i, j) can be compared in the form of $(a_t^{(1)}, a_t^{(2)}) = (i, j)$ and $(a_t^{(1)}, a_t^{(2)}) = (j, i)$. We consider these two cases separately and define the following counters: $N_{ij}^{(1)}(t) = \sum_{\tau=1}^t \mathbb{1}(a_{\tau}^{(1)} = i, a_{\tau}^{(2)} = j)$ and $N_{ij}^{(2)}(t) = \sum_{\tau=1}^t \mathbb{1}(a_{\tau}^{(1)} = j, a_{\tau}^{(2)} = i)$. Then the total number of comparisons is $N_{ij}(t) = N_{ij}^{(1)}(t) + N_{ij}^{(2)}(t)$ for $i \neq j$, and $N_{ii}(t) = N_{ii}^{(1)}(t) = N_{ii}^{(2)}(t)$ for i = j.

4.2.1 $O(K^2 \log T)$ Regret

To obtain theoretical bounds for the regret of D-TS, we make the following assumption:

¹A Borda winner may be more appropriate in this special case [22], and we mainly use it to illustrate the dilemma.

Assumption 1: The preference probability $p_{ij} \neq 1/2$ for any $i \neq j$.

Under Assumption 1, we present the first result for D-TS in general Copeland dueling bandits:

Proposition 1. When applying D-TS with $\alpha > 0.5$ in a Copeland dueling bandit with a preference matrix $P = [p_{ij}]_{K \times K}$ satisfying Assumption 1, its regret is bounded as:

$$R_{\text{D-TS}}(T) \le \sum_{i \ne j: p_{ij} < 1/2} \left[\frac{4\alpha \log T}{\Delta_{ij}^2} + (1+\epsilon) \frac{\log T}{D(p_{ij}||1/2)} \right] + O(\frac{K^2}{\epsilon^2}),$$
(2)

where $\epsilon > 0$ is an arbitrary constant, and $D(p||q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ is the KL divergence.

The summation operation in Eq. (2) is conducted over all pairs (i, j) with $p_{ij} < 1/2$. Thus, Proposition 1 states that D-TS achieves $O(K^2 \log T)$ regret in Copeland dueling bandits. To the best of our knowledge, this is the first theoretical bound for TS in dueling bandits. The scaling behavior of this bound with respect to T is order optimal, since a lower bound $\Omega(\log T)$ has been shown in [7]. The refinement of the scaling behavior with respect to K will be discussed later.

Proving Proposition 1 needs to bound the number of comparisons for all pairs (i, j) with $i \notin C^*$ or $j \notin C^*$. When fixing the first candidate as $a_t^{(1)} = i$, the selection of the second candidate $a_t^{(2)}$ is similar to a traditional *K*-armed bandit problem with expected utilities p_{ji} (j = 1, 2, ..., K). However, the analysis is more complex here since different arms are eliminated differently depending on the value of p_{ji} . We prove Proposition 1 through Lemmas 1 to 3, which bound the number of comparisons for all suboptimal pairs (i, j) under different scenarios, i.e., $p_{ji} < 1/2$, $p_{ji} > 1/2$, and $p_{ji} = 1/2$ $(j = i \notin C^*)$, respectively.

Lemma 1. Under D-TS, for an arbitrary constant $\epsilon > 0$ and one pair (i, j) with $p_{ji} < 1/2$, we have

$$\mathbb{E}[N_{ij}^{(1)}(T)] \le (1+\epsilon) \frac{\log T}{D(p_{ji}||1/2)} + O(\frac{1}{\epsilon^2}).$$
(3)

Proof. We can prove this lemma by viewing the comparison between the first candidate arm i and its inferiors as a traditional MAB. In fact, it may be even simpler than that in [15] because under D-TS, arm j with $p_{ji} < 1/2$ is competing with arm i with $p_{ii} = 1/2$, which is known and fixed. Then we can bound $\mathbb{E}[N_{ij}^{(1)}(T)]$ using the techniques in [15]. Details can be found in Appendix B.1.

Lemma 2. Under D-TS with $\alpha > 0.5$, for one pair (i, j) with $p_{ji} > 1/2$, we have

$$\mathbb{E}[N_{ij}^{(1)}(T)] \le \frac{4\alpha \log T}{\Delta_{ji}^2} + O(1).$$

$$\tag{4}$$

Proof. We note that when $a_t^{(1)} = i$, arm j can be selected as $a_t^{(2)}$ only when its RLCB $l_{ji}(t) \le 1/2$. Then we can bound $\mathbb{E}[N_{ij}^{(1)}(T)]$ by $O(\frac{4\alpha \log T}{\Delta_{ji}^2})$ similarly to the analysis of traditional UCB algorithms [23]. Details can be found in Appendix B.2.

Lemma 3. Under D-TS, for any arm $i \notin C^*$, we have

$$\mathbb{E}[N_{ii}(T)] \le O(K) + \sum_{k:p_{ki} > 1/2} \Theta\left(\frac{1}{\Delta_{ki}^2} + \frac{1}{\Delta_{ki}^2 D(1/2||p_{ki})} + \frac{1}{\Delta_{ki}^4}\right) = O(K).$$
(5)

Before proving Lemma 3, we present an important property for $\hat{\zeta}^*(t) := \max_{1 \le i \le K} \hat{\zeta}_i(t)$. Recall that ζ^* is the maximum normalized Copeland score. Using the concentration property of RUCB (Lemma 6 in Appendix A), the following lemma shows that $\hat{\zeta}^*(t)$ is indeed a UCB of ζ^* .

Lemma 4. For any
$$\alpha > 0.5$$
 and $t > 0$, $\mathbb{P}\{\hat{\zeta}^*(t) \ge \zeta^*\} \ge 1 - K\left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right]t^{-\frac{2\alpha}{\alpha + 1/2}}$

Return to the proof of Lemma 3. To prove Lemma 3, we consider the cases of $\hat{\zeta}^*(t) < \zeta^*$ and $\hat{\zeta}^*(t) \geq \zeta^*$. The former case $\hat{\zeta}^*(t) < \zeta^*$ can be bounded by Lemma 4. For the latter case, we note that when $\hat{\zeta}^*(t) \geq \zeta^*$, the event $(a_t^{(1)}, a_t^{(2)}) = (i, i)$ occurs only if: a) there exists at least one $k \in \mathcal{K}$ with $p_{ki} > 1/2$, such that $l_{ki}(t) \leq 1/2$; and b) $\theta_{ki}^{(2)}(t) \leq 1/2$ for all k with $l_{ki}(t) \leq 1/2$. In this case, we can bound the probability of $(a_t^{(1)}, a_t^{(2)}) = (i, i)$ by that of $(a_t^{(1)}, a_t^{(2)}) = (i, k)$, for k with $p_{ki} > 1/2$ but $l_{ki}(t) \leq 1/2$, where the coefficient decays exponentially. Then we can bound $\mathbb{E}[N_{ii}(T)]$ by O(1) similar to [15]. Details of proof can be found in Appendix B.4.

The conclusion of Proposition 1 then follows by combining Lemmas 1 to 3.

4.2.2 Regret Bound Refinement

In this section, we refine the regret bound for D-TS and reduce its scaling factor with respect to the number of arms K.

We sort the arms for each $i \notin C^*$ in the descending order of p_{ji} , and let $(\sigma_{i(1)}, \sigma_{i(2)}, \ldots, \sigma_{i(K)})$ be a permutation of $(1, 2, \ldots, K)$, such that $p_{\sigma_{i(1)}, i} \ge p_{\sigma_{i(2)}, i} \ge \ldots \ge p_{\sigma_{i(K)}, i}$. In addition, for a Copeland winner $i^* \in C$, let $L_C = \sum_{j=1}^K \mathbb{1}(p_{ji^*} > 1/2)$ be the number of arms that beat arm i^* . To refine the regret, we introduce an additional no-tie assumption:

Assumption 2: For each arm $i \notin C^*$, $p_{\sigma_{i(L_C+1)},i} > p_{\sigma_{i(j)},i}$ for all $j > L_C + 1$.

We present a refined regret bound for D-TS as follows:

Theorem 1. When applying D-TS with $\alpha > 0.5$ in a Copeland dueling bandit with a preference matrix $P = [p_{ij}]_{K \times K}$ satisfying Assumptions 1 and 2, its regret is bounded as:

$$R_{\text{D-TS}}(T) \leq \sum_{i \in \mathcal{C}^*} \left[\sum_{j: p_{ji} > 1/2} \frac{4\alpha \log T}{\Delta_{ji}^2} + \sum_{j: p_{ji} < 1/2} (1+\epsilon) \frac{\log T}{D(p_{ji}||1/2)} \right] + \sum_{i \notin \mathcal{C}^*} \sum_{j=1}^{L_C+1} \frac{4\alpha \log T}{\Delta_{\sigma_i(j),i}^2} + \beta (1+\epsilon)^2 \sum_{i \notin \mathcal{C}^*} \sum_{j=L_C+2}^K \frac{\log \log T}{D(p_{\sigma_i(j),i}||p_{\sigma_i(L_C+1),i})} + O(K^3) + O(\frac{K^2}{\epsilon^2}),$$
(6)

where $\beta > 2$ and $\epsilon > 0$ are constants, and $D(\cdot || \cdot)$ is the KL-divergence.

In (6), the first term corresponds to the regret when the first candidate $a_t^{(1)}$ is a winner, and is $O(K|\mathcal{C}^*|\log T)$. The second term corresponds to the comparisons between a non-winner arm and its first $L_C + 1$ superiors, which is bounded by $O(K(L_C + 1)\log T)$. The remaining terms correspond to the comparisons between a non-winner arm and the remaining arms, and is bounded by $O(K^2 \log \log T)$. As demonstrated in [6], L_C is relatively small compared to K, and can be viewed as a constant. Thus, the total regret $R_{\text{D-TS}}(T)$ is bounded as $R_{\text{D-TS}}(T) = O(K \log T + K^2 \log \log T)$. In particular, this asymptotic trend can be easily seen for Condorcet dueling bandits where $L_C = 0$.

Comparing Eq. (6) with Eq. (2), we can see the difference is the third and fourth terms in (6), which refine the regret of comparing a suboptimal arm and its last $(K - L_C - 1)$ inferiors into $O(\log \log T)$. Thus, to prove Theorem 1, it suffices to show the following additional lemma:

Lemma 5. Under Assumptions 1 and 2, for any suboptimal arm $i \notin C^*$ and $j > L_C + 1$, we have

$$\mathbb{E}[N_{i\sigma_{i(j)}}^{(1)}(T)] \le \frac{\beta(1+\epsilon)^2 \log\log T}{D(p_{\sigma_{i(j)},i}||p_{\sigma_{i(L_C+1)},i})} + O(K) + O(\frac{1}{\epsilon^2}),\tag{7}$$

where $\beta > 2$ and $\epsilon > 0$ are constants.

Proof. We prove this lemma using a *back substitution* argument. The intuition is that when fixing the first candidate as $a_t^{(1)} = i$, the comparison between $a_t^{(1)}$ and the other arms is similar to a traditional MAB with expected utilities p_{ji} $(1 \le j \le K)$. Let $N_i^{(1)}(T) = \sum_{t=1}^T \mathbb{1}(a_t^{(1)} = i)$ be the number of time-slots when this type of MAB is played. Using the fact that the distribution of the samples only depends on the historic comparison results (but not t), we can show $\mathbb{E}[N_{i,\sigma_{i(j)}}^{(1)}(T)|N_i^{(1)}(T)] = O(\log N_i^{(1)}(T))$, which holds for any $N_i^{(1)}(T)$. We have shown that $\mathbb{E}[N_i^{(1)}(T)] = O(K \log T)$ for any $i \ne C^*$ when proving Proposition 1. Then, substituting the bound of $\mathbb{E}[N_i^{(1)}(T)]$ back and using the concavity of the $\log(\cdot)$ function, we have $\mathbb{E}[N_{i,\sigma_{i(j)}}^{(1)}(T)] = \mathbb{E}[\mathbb{E}[N_{i,\sigma_{i(j)}}^{(1)}(T)|N_i^{(1)}(T)]] \le O(\log \mathbb{E}[N_i^{(1)}(T)]) = O(\log \log T + \log K)$. Details can be found in Appendix C.1

4.3 Further Improvement: D-TS⁺

D-TS is a TS framework for dueling bandits, and its performance can be improved by refining certain components of it. In this section, we propose an enhanced version of D-TS, referred to as D-TS⁺, that carefully breaks the ties to reduce the regret.

Note that by randomly breaking the ties (Line 11 in Algorithm 1), D-TS tends to explore all potential winners. This may be desirable in certain applications such as restaurant recommendation, where

users may not want to stick to a single winner. However, because of this, the regret of D-TS scales with the number of winners $|C^*|$ as shown in Theorem 1. To further reduce the regret, we can break the ties according to estimated regret.

Specifically, with samples $\theta_{ij}^{(1)}(t)$, the normalized Copeland score for each arm i can be estimated as $\tilde{\zeta}_i(t) = \frac{1}{K-1} \sum_{j \neq i} \mathbb{1}(\theta_{ij}^{(1)}(t) > 1/2)$. Then the maximum normalized Copeland score is $\tilde{\zeta}^*(t) = \max_i \tilde{\zeta}_i(t)$, and the loss of comparing arm i and arm j is $\tilde{r}_{ij}(t) = \tilde{\zeta}^*(t) - \frac{1}{2} [\tilde{\zeta}_i(t) + \tilde{\zeta}_j(t)]$. For $p_{ij} \neq 1/2$, we need about $\Theta(\frac{\log T}{D(p_{ij})|1/2})$ time-slots to distinguish it from 1/2 [5]. Thus, when choosing i as the first candidate, the regret of comparing it with all other arms can be estimated by $\tilde{R}_i^{(1)}(t) = \sum_{j:\theta_{ij}^{(1)}(t)\neq 1/2} \tilde{r}_{ij}(t)/D(\theta_{ij}^{(1)}(t))|1/2)$. We propose the following D-TS⁺ algorithm that breaks the ties to minimize $\tilde{R}_i^{(1)}(t)$.

D-TS⁺: Implement the same operations as D-TS, except for the selection of the first candidate (Line 11 in Algorithm 1) is replaced by the following two steps:

$$\mathcal{A}^{(1)} \leftarrow \{i \in \mathcal{C} : \zeta_i = \max_{i \in \mathcal{C}} \sum_{j \neq i} \mathbb{1}(\theta_{ij}^{(1)} > 1/2)\};$$
$$a^{(1)} \leftarrow \operatorname*{arg\,min}_{i \in \mathcal{A}^{(1)}} \tilde{R}_i^{(1)};$$

D-TS⁺ only changes the tie-breaking criterion in selecting the first candidate. Thus, the regret bound of D-TS directly applies to D-TS⁺:

Corollary 1. The regret of D-TS⁺, $R_{D-TS^+}(T)$, satisfies inequality (6) under Assumptions 1 and 2.

Corollary 1 provides an upper bound for the regret of D-TS⁺. In practice, however, D-TS⁺ performs better than D-TS in the scenarios with multiple winners, as we can see in Section 5 and Appendix D. Our conjecture is that with this regret-minimization criterion, the D-TS⁺ algorithm tends to focus on one of the winners (if there is no tie in terms of expected regret), and thus reduces the first term in (6) from $O(K|\mathcal{C}^*|\log T)$ to $O(K\log T)$. The proof of this conjecture requires properties for the evolution of the statistics for all arms and the majority voting results based on the Thompson samples, and is complex. This is left as part of our future work.

In the above D-TS⁺ algorithm, we only consider the regret of choosing *i* as the first candidate. From Theorem 1, we know that comparing other arms with their superiors will also result in $\Theta(\log T)$ regret. Thus, although the current D-TS⁺ algorithm performs well in most practical scenarios, one may further improve its performance by taking these additional comparisons into account in $\tilde{R}_i^{(1)}(t)$.

5 Experiments

To evaluate the proposed D-TS and D-TS⁺ algorithms, we run experiments based on synthetic and real-world data. Here we present the results for experiments based on the Microsoft Learning to Rank (MSLR) dataset [24], which provides the relevance for queries and ranked documents. Based on this dataset, [6] derives a preference matrix for 136 rankers, where each ranker is a function that maps a user's query to a document ranking and can be viewed as one arm in dueling bandits. We use the two 5-armed submatrices in [6], one for Condorcet dueling bandit and the other for non-Condorcet dueling bandit. More experiments and discussions can be found in Appendix D².

We compare D-TS and D-TS⁺ with the following algorithms: BTM [16], SAVAGE [17], Sparring [18], RUCB [4], RCS [3], CCB [6], SCB [6], RMED1 [5], and ECW-RMED [7]. For BTM, we set the relaxed factor $\gamma = 1.3$ as [16]. For algorithms using RUCB and RLCB, including D-TS and D-TS⁺, we set the scale factor $\alpha = 0.51$. For RMED1, we use the same settings as [5], and for ECW-RMED, we use the same setting as [7]. For the "explore-then-exploit" algorithms, BTM and SAVAGE, each point is obtained by resetting the time horizon as the corresponding value. The results are averaged over 500 independent experiments, where in each experiment, the arms are randomly shuffled to prevent algorithms from exploiting special structures of the preference matrix.

In Condorcet dueling bandits, our D-TS and D-TS⁺ algorithms achieve almost the same performance and both perform much better than existing algorithms, as shown in Fig. 1(a). In particular, compared with RCS, we can see that the full utilization of TS in D-TS and D-TS⁺ significantly reduces the

²Source codes are available at https://github.com/HuasenWu/DuelingBandits.

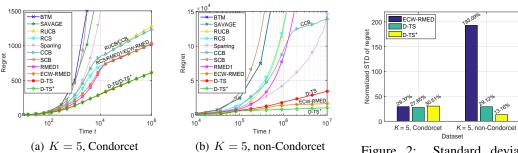


Figure 1: Regret in MSLR dataset. In (b), there are 3 Copeland winners with normalized Copeland score $\zeta^* = 3/4$.

Figure 2: Standard deviation (STD) of regret for $T = 10^6$ (normalized by $R_{\text{ECW}-\text{RMED}}(T)$).

regret. Compared with RMED1 and ECW-RMED, our D-TS and D-TS⁺ algorithms also perform better. [5] has shown that RMED1 is optimal in Condorcet dueling bandits, not only in the sense of asymptotic order, but also the coefficients in the regret bound. The simulation results show that D-TS and D-TS⁺ not only achieve the similar slope as RMED1/ECW-RMED, but also converge faster to the asymptotic regime and thus achieve much lower regret. This inspires us to further refine the regret bounds for D-TS and D-TS⁺ in the future.

In non-Condorcet dueling bandits, as shown in Fig. 1(b), D-TS and D-TS⁺ significantly reduce the regret compared to the UCB-type algorithm, CCB (e.g., the regret of D-TS⁺ is less than 10% of that of CCB). Compared with ECW-RMED, D-TS achieves higher regret, mainly because it randomly explores all Copeland winners due to the random tie-breaking rule. With a regret-minimization tie-breaking rule, D-TS⁺ further reduces the regret, and outperforms ECW-RMED in this dataset. Moreover, as randomized algorithms, D-TS and D-TS⁺ are more robust to the preference probabilities. As shown in Fig. 2, D-TS and D-TS⁺ have much smaller regret STD than that of ECW-RMED in the non-Condorcet dataset, where certain preference probabilities (for different arms) are close to 1/2. In particular, the STD of regret for ECW-RMED is almost 200% of its mean value, while it is only 13.16% for D-TS⁺. In addition, as shown in Appendix D.2.3, D-TS and D-TS⁺ are also robust to delayed feedback, which is typically batched and provided periodically in practice.

Overall, D-TS and D-TS⁺ significantly outperform all existing algorithms, with the exception of ECW-RMED. Compared to ECW-RMED, D-TS⁺ achieves much lower regret in the Condorcet case, lower or comparable regret in the non-Condorcet case, and much more robustness in terms of regret STD and delayed feedback. Thus, the simplicity, good performance, and robustness of D-TS and D-TS⁺ make them good algorithms in practice.

6 Conclusions and Future Work

In this paper, we study TS algorithms for dueling bandits. We propose a D-TS algorithm and its enhanced version D-TS⁺ for general Copeland dueling bandits, including Condorcet dueling bandits as a special case. Our study reveals desirable properties of D-TS and D-TS⁺ from both theoretical and practical perspectives. Theoretically, we show that the regret of D-TS and D-TS⁺ is bounded by $O(K^2 \log T)$ in general Copeland dueling bandits, and can be refined to $O(K \log T + K^2 \log \log T)$ in Condorcet dueling bandits and most practical Copeland dueling bandits. Practically, experimental results demonstrate that these simple algorithms achieve significantly better overall-performance than existing algorithms, i.e., D-TS and D-TS⁺ typically achieve much lower regret in practice and are robust to many practical factors, such as preference matrix and feedback delay.

Although logarithmic regret bounds have been obtained for D-TS and D-TS⁺, our analysis relies heavily on the properties of RUCB/RLCB and the regret bounds are likely loose. In fact, we see from experiments that RUCB-based elimination seldom occurs under most practical settings. We will further refine the regret bounds by investigating the properties of TS-based majority-voting. Moreover, results from recent work such as [7] may be leveraged to improve TS algorithms. Last, it is also an interesting future direction to study D-TS type algorithms for dueling bandits with other definition of winners.

Acknowledgements: This research was supported in part by NSF Grants CCF-1423542, CNS-1457060, and CNS-1547461. The authors would like to thank Prof. R. Srikant (UIUC), Prof. Shipra Agrawal (Columbia University), Masrour Zoghi (University of Amsterdam), and Dr. Junpei Komiyama (University of Tokyo) for their helpful discussions and suggestions.

References

- [1] Y. Yue, J. Broder, R. Kleinberg, and T. Joachims. The *k*-armed dueling bandits problem. *Journal* of Computer and System Sciences, 78(5):1538–1556, 2012.
- [2] Y. Yue and T. Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *International Conference on Machine Learning (ICML)*, pages 1201–1208, 2009.
- [3] M. Zoghi, S. A. Whiteson, M. De Rijke, and R. Munos. Relative confidence sampling for efficient on-line ranker evaluation. In ACM International Conference on Web Search and Data Mining, pages 73–82, 2014.
- [4] M. Zoghi, S. Whiteson, R. Munos, and M. D. Rijke. Relative upper confidence bound for the k-armed dueling bandit problem. In *International Conference on Machine Learning (ICML)*, pages 10–18, 2014.
- [5] J. Komiyama, J. Honda, H. Kashima, and H. Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *Proceedings of Conference on Learning Theory*, 2015.
- [6] M. Zoghi, Z. S. Karnin, S. Whiteson, and M. de Rijke. Copeland dueling bandits. In *Advances in Neural Information Processing Systems*, pages 307–315, 2015.
- [7] J. Komiyama, J. Honda, and H. Nakagawa. Copeland dueling bandit problem: Regret lower bound, optimal algorithm, and computationally efficient algorithm. In *International Conference on Machine Learning (ICML)*, 2016.
- [8] W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, pages 285–294, 1933.
- [9] O. Chapelle and L. Li. An empirical evaluation of Thompson Sampling. In Advances in Neural Information Processing Systems, pages 2249–2257, 2011.
- [10] S. Agrawal and N. Goyal. Analysis of Thompson Sampling for the multi-armed bandit problem. In *Conference on Learning Theory (COLT)*, 2012.
- [11] J. Komiyama, J. Honda, and H. Nakagawa. Optimal regret analysis of Thompson Sampling in stochastic multi-armed bandit problem with multiple plays. In *International Conference on Machine Learning (ICML)*, 2015.
- [12] Y. Xia, H. Li, T. Qin, N. Yu, and T.-Y. Liu. Thompson sampling for budgeted multi-armed bandits. In *International Joint Conference on Artificial Intelligence*, 2015.
- [13] A. Gopalan, S. Mannor, and Y. Mansour. Thompson sampling for complex online problems. In International Conference on Machine Learning (ICML), pages 100–108, 2014.
- [14] A. Gopalan and S. Mannor. Thompson sampling for learning parameterized Markov decision processes. In *Proceedings of Conference on Learning Theory*, pages 861–898, 2015.
- [15] S. Agrawal and N. Goyal. Further optimal regret bounds for Thompson Sampling. In International Conference on Artificial Intelligence and Statistics, pages 99–107, 2013.
- [16] Y. Yue and T. Joachims. Beat the mean bandit. In International Conference on Machine Learning (ICML), pages 241–248, 2011.
- [17] T. Urvoy, F. Clerot, R. Féraud, and S. Naamane. Generic exploration and k-armed voting bandits. In *International Conference on Machine Learning (ICML)*, pages 91–99, 2013.
- [18] N. Ailon, Z. Karnin, and T. Joachims. Reducing dueling bandits to cardinal bandits. In Proceedings of The 31st International Conference on Machine Learning, pages 856–864, 2014.
- [19] D. Russo and B. Van Roy. An information-theoretic analysis of Thompson Sampling. *arXiv* preprint arXiv:1403.5341, 2014.
- [20] D. Russo and B. Van Roy. Learning to optimize via posterior sampling. Mathematics of Operations Research, 39(4):1221–1243, 2014.
- [21] N. Welsh. Thompson sampling for the dueling bandits problem. In *Large-Scale Online Learning and Decision Making (LSOLDM) Workshop*, 2012. available at http://videolectures.net/lsoldm2012_welsh_bandits_problem/.
- [22] K. Jamieson, S. Katariya, A. Deshpande, and R. Nowak. Sparse dueling bandits. In *Conference on Learning Theory (COLT)*, 2015.
- [23] S. Bubeck. *Bandits games and clustering foundations*. PhD thesis, Université des Sciences et Technologie de Lille-Lille I, 2010.
- [24] Microsoft Research, Microsoft Learning to Rank Datasets. http://research.microsoft.com/enus/projects/mslr/, 2010.

Appendices

A Preliminary: Concentration of RUCB/RLCB

We first present the concentration properties of RUCB/RLCB. By relating RUCB/RLCB to UCB/LCB in traditional MAB, we can adjust the results in [23] for RUCB/RLCB as follows. **Lemma 6.** 1) When $\alpha > 0.5$, for any (i, j) and t > 0,

$$\mathbb{P}\{p_{ij} \ge u_{ij}(t)\} \le \left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right] t^{-\frac{2\alpha}{\alpha + 1/2}},\tag{8}$$

$$\mathbb{P}\{p_{ij} \le l_{ij}(t)\} \le \left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right] t^{-\frac{2\alpha}{\alpha + 1/2}}.$$
(9)

2) For any $\alpha > 1/2$,

$$\sum_{t=1}^{T} \mathbb{P}\{p_{ij} \ge u_{ij}(t)\} \le \frac{2}{\log(\alpha + 1/2)[2\alpha/(\alpha + 1/2) - 1]^2} = O(1),$$
(10)

$$\sum_{t=1}^{T} \mathbb{P}\{p_{ij} \le l_{ij}(t)\} \le \frac{2}{\log(\alpha + 1/2)[2\alpha/(\alpha + 1/2) - 1]^2} = O(1).$$
(11)

Proof. We prove this lemma using the techniques in the proof of Theorem 2.2 in [23].

In fact, RUCB (resp., RLCB) in dueling bandits are essentially the same as UCB (resp., LCB) in traditional MAB. Thus, Part 1) of this lemma can be proved using the peeling argument in [23].

For Part 2), the sum can be bounded by the integration $\int_1^\infty \left[\frac{\log t}{\log(\alpha+1/2)}+1\right]t^{-\frac{2\alpha}{\alpha+1/2}} dt$ as in [23]. \Box

B Regret Analysis: $O(K^2 \log T)$ Regret

B.1 Proof of Lemma 1

For a pair (i, j) with $p_{ji} < 1/2$, let x_{ji} be a number satisfying $p_{ji} < x_{ji} < 1/2$. Let $\bar{p}_{ji}(t) = \frac{B_{ji}(t-1)}{B_{ji}(t-1)+B_{ij}(t-1)}$ be the empirical estimation for the probability that arm j beats arm i. Define the following events:

$$\mathcal{E}_{ji}^{p}(t) = \{\bar{p}_{ji}(t) < x_{ji}\},\$$
$$\mathcal{E}_{ji}^{\theta}(t) = \{\theta_{ji}^{(2)}(t) < 1/2\}.$$

For an event \mathcal{E} , we let $\neg \mathcal{E}$ be the event of "not \mathcal{E} ". Then

$$\begin{split} \mathbb{E}[N_{ij}^{(1)}(T)] &= \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j)\} \\ &= \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), \mathcal{E}_{ji}^p(t), \mathcal{E}_{ji}^\theta(t)\} \\ &+ \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), \mathcal{E}_{ji}^p(t), \neg \mathcal{E}_{ji}^\theta(t)\} \\ &+ \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), \neg \mathcal{E}_{ji}^p(t)\}. \end{split}$$

The first term is zero, because $\mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), \mathcal{E}_{ji}^p(t), \mathcal{E}_{ji}^\theta(t)\} = 0$ for all t, due to the fact that $a_t^{(2)} \neq j$ when $\theta_{ji}^{(2)}(t) < 1/2 = \theta_{ii}^{(2)}(t)$.

The second and third terms can be bounded similarly to the analysis of TS in traditional MABs [15]. To see this, we note that when fixing the first candidate as $a_t^{(1)} = i$, the comparison between i and other arms is similar to a traditional MAB problem with expected reward p_{ji} $(1 \le j \le K)$. For the case of $p_{ji} < 1/2$, we only need to care about two differences: first, $p_{ii} = 1/2$ is fixed and known; second, in addition to $(a_t^{(1)}, a_t^{(2)}) = (i, j)$, arm i and arm j could also be compared when $(a_t^{(1)}, a_t^{(2)}) = (j, i)$. By capturing the second difference with $N_{ij}(t-1) = N_{ij}^{(1)}(t-1) + N_{ij}^{(2)}(t-1)$, we can leverage the techniques in [15] to prove our results.

Specifically, the second term can be bounded by using the concentration property of the Thompson samples. Letting $L_{ji}(T) = \frac{\log T}{D(x_{ji})|1/2}$, similar to the proof of Lemma 4 in [15], we have

$$\sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, j), \mathcal{E}_{ji}^{p}(t), \neg \mathcal{E}_{ji}^{\theta}(t)\}$$

$$= \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, j), \mathcal{E}_{ji}^{p}(t), \neg \mathcal{E}_{ji}^{\theta}(t), N_{ij}(t-1) \leq L_{ji}(T)\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, j), \mathcal{E}_{ji}^{p}(t), \neg \mathcal{E}_{ji}^{\theta}(t), N_{ij}(t-1) > L_{ji}(T)\}$$

$$\leq L_{ji}(T) + \sum_{t=1}^{T} \frac{1}{T}$$

$$= \frac{\log T}{D(x_{ji}||1/2)} + 1.$$
(12)

The third term can be bounded similarly to Lemma 3 in [15]. Specifically, let τ_n be the slot index when i and j are compared for the n-th time, including both cases $(a_t^{(1)}, a_t^{(2)}) = (i, j)$ and $(a_t^{(1)}, a_t^{(2)}) = (j, i)$. Let $\tau_0 = 0$. Then, $\bar{p}_{ji}(t)$ is fixed between $\tau_n + 1$ and τ_{n+1} , and $\sum_{t=\tau_n+1}^{\tau_{n+1}} \mathbb{1}((a_t^{(1)}, a_t^{(2)}) = (i, j)) \leq 1$ (it is 0 if the (n + 1)-th comparison is implemented in the form of $(a_t^{(1)}, a_t^{(2)}) = (j, i)$). Then

$$\sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, j), \forall \mathcal{E}_{ji}^{p}(t)\}$$

$$\leq \sum_{n=0}^{T-1} \mathbb{E}\left[\sum_{t=\tau_{n}+1}^{\tau_{n+1}} \mathbb{1}((a_{t}^{(1)}, a_{t}^{(2)}) = (i, j)) \cdot \mathbb{1}(\forall \mathcal{E}_{ji}^{p}(t))\right]$$

$$\leq \sum_{n=0}^{T-1} \mathbb{E}\left[\mathbb{1}(\forall \mathcal{E}_{ji}^{p}(\tau_{n}+1)) \sum_{t=\tau_{n}+1}^{\tau_{n+1}} \mathbb{1}((a_{t}^{(1)}, a_{t}^{(2)}) = (i, j))\right]$$

$$\leq \sum_{n=0}^{T-1} \mathbb{P}(\forall \mathcal{E}_{ji}^{p}(\tau_{n}+1))$$

$$\leq 1 + \sum_{n=1}^{T-1} e^{-nD(x_{ji}||p_{ji})}$$

$$\leq 1 + \frac{1}{D(x_{ji}||p_{ji})}.$$
(13)

For any $\epsilon \in (0, 1]$, we choose $x_{ji} \in (p_{ji}, 1/2)$ such that $D(x_{ji}||1/2) = D(p_{ji}||1/2)/(1+\epsilon)$, which also implies $\frac{1}{D(x_{ji}||p_{ji})} = O(\frac{1}{\epsilon^2})$ as shown in [15]. The conclusion then follows by combining the bounds for all the above three terms.

B.2 Proof of Lemma 2

We prove Lemma 2 by using the concentration property of RLCB $l_{ji}(t)$. According to the definition of $N_{ii}^{(1)}(T)$, we have

$$\mathbb{E}[N_{ij}^{(1)}(T)] = \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j)\}$$

$$= \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) \ge \frac{4\alpha \log T}{\Delta_{ji}^2}\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) < \frac{4\alpha \log T}{\Delta_{ji}^2}\}$$

For the first term, we note that when $a_t^{(1)} = i$, arm j can be selected as $a_t^{(2)}$ only when the $l_{ji}(t) \le 1/2$. When $N_{ij}(t-1) \ge \frac{4\alpha \log T}{\Delta_{ji}^2}$, we have $\Delta_{ji} \ge 2\sqrt{\frac{\alpha \log t}{N_{ij}(t-1)}}$. Thus, $l_{ji}(t) + \Delta_{ji} \ge u_{ji}(t)$. Because $p_{ji} > 1/2$, its RLCB satisfies

$$\begin{aligned} \mathbb{P}\big\{l_{ji}(t) \leq 1/2, N_{ij}(t-1) \geq \frac{4\alpha \log T}{\Delta_{ji}^2}\big\} &\leq \mathbb{P}\big\{l_{ji}(t) &\leq p_{ji} - \Delta_{ji}, N_{ij}(t-1) \geq \frac{4\alpha \log T}{\Delta_{ji}^2}\big\} \\ &\leq \mathbb{P}\big\{u_{ji}(t) \leq p_{ji}\big\}. \end{aligned}$$

Using Lemma 6, we have

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) \ge \frac{4\alpha \log T}{\Delta_{ji}^2} \right\} \le \sum_{t=1}^{T} \mathbb{P}\left\{ u_{ji}(t) \le p_{ji} \right\} = O(1).$$

For the second term, we can bound it as follows:

$$\sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) < \frac{4\alpha \log T}{\Delta_{ji}^2}\}$$

= $\mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}((a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) < \frac{4\alpha \log T}{\Delta_{ji}^2})\right] \le \frac{4\alpha \log T}{\Delta_{ji}^2},$ (14)

because $\sum_{t=1}^{T} \mathbb{1}\left((a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) < \frac{4\alpha \log T}{\Delta_{ji}^2}\right) \leq \frac{4\alpha \log T}{\Delta_{ji}^2}$ due to the fact that: when $(a_t^{(1)}, a_t^{(2)}) = (i, j)$ at $t, N_{ij}(t-1)$ will be increased by 1, but $\mathbb{1}\left((a_t^{(1)}, a_t^{(2)}) = (i, j), N_{ij}(t-1) < \frac{4\alpha \log T}{\Delta_{ji}^2}\right) = 0$ as long as $N_{ij}(t-1) \geq \frac{4\alpha \log T}{\Delta_{ji}^2}$.

The conclusion then follows by combining the bounds for the above two terms.

B.3 Proof of Lemma 4

Let i^* be the Copeland winner (or any one of them if there are multiple Copeland winners) in the dueling bandit. We prove Lemma 4 by analyzing the RUCB $u_{i^*j}(t)$ at t. According to Lemma 6, we have that for any $j \neq i^*$,

$$\mathbb{P}\{u_{i^*j}(t) < p_{i^*j}\} \le \left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right] t^{-\frac{2\alpha}{\alpha + 1/2}}.$$
(15)

Note that $\zeta^* = \frac{1}{K-1} \sum_{j \neq i} \mathbb{1}(p_{i^*j} > 1/2)$. Let $\mathcal{L}_{i^*} = \{j : 1 \leq j \leq K, p_{i^*j} > 1/2\}$ be the set of arms that lose to i^* . Thus,

$$\hat{\zeta}^*(t) < \zeta^* \Rightarrow \exists j \in \mathcal{L}_{i^*}, \text{ such that } u_{i^*j}(t) < p_{i^*j}.$$
(16)

Consider all elements in \mathcal{L}_{i^*} , we have

$$\mathbb{P}\{\hat{\zeta}^{*}(t) \geq \zeta^{*}\} = 1 - \mathbb{P}\{\hat{\zeta}^{*}(t) < \zeta^{*}\} \\
\geq 1 - \mathbb{P}\{\exists j \in \mathcal{L}_{i^{*}}, \text{ s.t. } u_{i^{*}j}(t) < p_{i^{*}j}\} \\
\geq 1 - |\mathcal{L}_{i^{*}}| \left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right] t^{-\frac{2\alpha}{\alpha + 1/2}} \\
\geq 1 - K \left[\frac{\log t}{\log(\alpha + 1/2)} + 1\right] t^{-\frac{2\alpha}{\alpha + 1/2}}.$$
(17)

B.4 Proof of Lemma 3

To bound the number of time-slots when we compare one non-winner arm against itself, we need to investigate the necessary conditions for this event.

Specifically, when the upper bound of the Copeland score $\hat{\zeta}^*(t) \ge \zeta^*$, the event $(a_t^{(1)}, a_t^{(2)}) = (i, i)$ for $i \notin \mathcal{C}^*$ occurs only if: a) there exists at least one $k \in \mathcal{K}$ with $p_{ki} > 1/2$, such that $l_{ki}(t) \le 1/2$; and b) $\theta_{ki}^{(2)}(t) \le 1/2$ for all k with $l_{ki}(t) \le 1/2$. Now we bound $\mathbb{E}[N_{ii}(T)]$ by bounding the probability of these two conditions.

For k with $p_{ki} > 1/2$, we define the following probability

$$q_{ki}(t) = \mathbb{P}\{\theta_{ki}^{(2)}(t) > 1/2 | \mathcal{H}_{t-1}\}.$$
(18)

Note that the value of $\hat{\zeta}^*(t)$, $l_{ki}(t)$, and $q_{ki}(t)$ depends on the history, and thus is determined by \mathcal{H}_{t-1} . Similar to Lemma 1 in [15], we bound the probability of comparing *i* against itself (accompanied by $\hat{\zeta}^*(t) \geq \zeta^*$ and $l_{ki}(t) \leq 1/2$, which is different from TS for traditional MABs) by that of comparing *i* with *k*.

Lemma 7. Given (i, k) with $p_{ki} > 1/2$, we have

$$\mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, i), \hat{\zeta}^*(t) \ge \zeta^*, l_{ki}(t) \le 1/2 |\mathcal{H}_{t-1} \right\} \\
\le \frac{1 - q_{ki}(t)}{q_{ki}(t)} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, k) |\mathcal{H}_{t-1} \right\}.$$
(19)

Proof. First of all, the value of $\hat{\zeta}^*(t)$ and $l_{ki}(t)$ depends on \mathcal{H}_{t-1} . Thus, if \mathcal{H}_{t-1} satisfies that $\hat{\zeta}^*(t) < \zeta^*$ or $l_{ki}(t) > 1/2$, then (19) holds because the left hand side of is zero.

Now we consider \mathcal{H}_{t-1} satisfying $\hat{\zeta}^*(t) \geq \zeta^*$ and $l_{ki}(t) \leq 1/2$. For the left hand side, we have

$$\mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, l_{ki}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\} \\
= \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i) | \mathcal{H}_{t-1}\right\} \\
\leq \mathbb{P}\left\{\theta_{k'i}^{(2)}(t) \leq 1/2, \forall k', \text{ s.t. } l_{k'i}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\} \\
= \mathbb{P}\left\{\theta_{ki}^{(2)}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\} \cdot \mathbb{P}\left\{\theta_{k'i}^{(2)}(t) \leq 1/2, \forall k' \neq k, \text{ s.t. } l_{k'i}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\} \\
= [1 - q_{ki}(t)] \mathbb{P}\left\{\theta_{k'i}^{(2)}(t) \leq 1/2, \forall k' \neq k, \text{ s.t. } l_{k'i}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\}.$$
(20)

For the right hand side, we have

$$\mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, k) | \mathcal{H}_{t-1} \right\} \\
\geq \mathbb{P}\left\{ \theta_{ki}^{(2)}(t) > 1/2 \ge \theta_{k'i}^{(2)}(t), \forall k' \neq k, \text{ s.t. } l_{k'i}(t) \le 1/2 | \mathcal{H}_{t-1} \right\} \\
= \mathbb{P}\left\{ \theta_{ki}^{(2)}(t) > 1/2 | \mathcal{H}_{t-1} \right\} \cdot \mathbb{P}\left\{ \theta_{k'i}^{(2)}(t) \le 1/2, \forall k' \neq k, \text{ s.t. } l_{k'i}(t) \le 1/2 | \mathcal{H}_{t-1} \right\} \\
= q_{ki}(t) \mathbb{P}\left\{ \theta_{k'i}^{(2)}(t) \le 1/2, \forall k' \neq i \text{ or } k, \text{ s.t. } l_{k'i}(t) \le 1/2 | \mathcal{H}_{t-1} \right\}.$$
(21)

The conclusion then follows by combining (20) and (21).

Now we return to the proof of Lemma 3. We divide the probability of $(a_t^{(1)}, a_t^{(2)}) = (i, i)$ into two terms according to the value of $\hat{\zeta}^*(t)$.

$$\mathbb{E}[N_{ii}(T)] = \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i)\} \\
\leq \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, \exists k, p_{ki} > 1/2, l_{ki}(t) \leq 1/2\} + \sum_{t=1}^{T} \mathbb{P}\{\hat{\zeta}^{*}(t) < \zeta^{*}\} \\
\leq \sum_{k:p_{ki} > 1/2} \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, l_{ki}(t) \leq 1/2\} + \sum_{t=1}^{T} K[\frac{\log t}{\log(\alpha + 1/2)} + 1]t^{-\frac{2\alpha}{\alpha + 1/2}} \\
\leq \sum_{k:p_{ki} > 1/2} \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, l_{ki}(t) \leq 1/2\} + O(K).$$
(22)

In the above equation, we have already bounded the second term by using Lemmas 4 and 6.

Next, we bound the first term by analyzing the bound for each k with $p_{ki} > 1/2$. Let τ_n be the timeslot index where k and i are compared for the n-th time, including both cases $(a_t^{(1)}, a_t^{(2)}) = (i, k)$ and $(a_t^{(1)}, a_t^{(2)}) = (k, i)$, and let $\tau_0 = 0$. Then by Lemma 7, we have that for each k with $p_{ki} > 1/2$,

$$\begin{split} &\sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, l_{ki}(t) \leq 1/2\right\} \\ &= \sum_{t=1}^{T} \mathbb{E}\left[\mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, i), \hat{\zeta}^{*}(t) \geq \zeta^{*}, l_{ki}(t) \leq 1/2 | \mathcal{H}_{t-1}\right\}\right] \\ &\leq \sum_{t=1}^{T} \mathbb{E}\left[\frac{1 - q_{ki}(t)}{q_{ki}(t)} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, k) | \mathcal{H}_{t-1}\right\}\right] \\ &\leq \sum_{t=1}^{T} \mathbb{E}\left[\mathbb{E}\left[\frac{1 - q_{ki}(t)}{q_{ki}(t)} \mathbb{1}\left((a_{t}^{(1)}, a_{t}^{(2)}) = (i, k)\right) | \mathcal{H}_{t-1}\right]\right] \\ &\stackrel{(a)}{=} \sum_{n=0}^{T-1} \mathbb{E}\left[\frac{1 - q_{ki}(\tau_{n} + 1)}{q_{ki}(\tau_{n} + 1)} \sum_{t=\tau_{n}+1}^{\tau_{n+1}} \mathbb{1}\left((a_{t}^{(1)}, a_{t}^{(2)}) = (i, k)\right)\right] \\ &\leq \sum_{n=0}^{T-1} \mathbb{E}\left[\frac{1}{q_{ki}(\tau_{n} + 1)} - 1\right]. \end{split}$$

The equality (a) follows from the fact that the distribution of $\theta_{ki}^{(2)}(t)$ only changes after k and i are compared. According to Lemma 2 in [15], $\mathbb{E}\left[\frac{1}{q_{ki}(\tau_n+1)}\right]$ is bounded as follows:

$$\mathbb{E}\Big[\frac{1}{q_{ki}(\tau_n+1)}\Big] \\ \leq \begin{cases} 1+\frac{3}{\Delta_{ki}}, & \text{for } n < \frac{8}{\Delta_{ki}}; \\ 1+\Theta\Big(e^{-n\Delta_{ki}^2/2} + \frac{1}{(n+1)\Delta_{ki}^2}e^{-nD(1/2||p_{ki})} + \frac{1}{e^{n\Delta_{ki}^2/4}-1}\Big), & \text{for } n \ge \frac{8}{\Delta_{ki}}. \end{cases}$$

Thus,

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, i), \hat{\zeta}^*(t) \ge \zeta^*, l_{ki}(t) \le 1/2 \right\}$$
$$\le \quad \frac{24}{\Delta_{ki}^2} + \sum_{n=0}^{T-1} \Theta\left(e^{-n\Delta_{ki}^2/2} + \frac{1}{(n+1)\Delta_{ki}^2} e^{-nD(1/2||p_{ki}|)} + \frac{1}{e^{n\Delta_{ki}^2/4} - 1} \right)$$

$$\leq \frac{24}{\Delta_{ki}^{2}} + \Theta\left(\frac{1}{\Delta_{ki}^{2}} + \frac{1}{\Delta_{ki}^{2}D(1/2||p_{ki})} + \frac{1}{\Delta_{ki}^{4}} + \frac{1}{\Delta_{ki}^{2}}\right)$$
$$= \Theta\left(\frac{1}{\Delta_{ki}^{2}} + \frac{1}{\Delta_{ki}^{2}D(1/2||p_{ki})} + \frac{1}{\Delta_{ki}^{4}}\right).$$

The conclusion then follows by summing over all k with $p_{ki} > 1/2$.

C Regret Bound Refinement

Theorem 1 can be proved by combining Lemmas 1 to 3 and Lemma 5. In this appendix, we present the proof of Lemma 5.

C.1 Proof of Lemma 5

We first consider the event of $(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)})$ with $j > L_C + 1$ in two cases with different values of $\hat{\zeta}^*(t)$. Recall that $\zeta^* = (K - L_C)/(K - 1)$ is the maximum normalized Copeland score. Then,

$$\mathbb{E}[N_{i\sigma_{i(j)}}^{(1)}(T)] = \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)})\} \\
\leq \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \ge \zeta^{*}\} + \sum_{t=1}^{T} \mathbb{P}\{\hat{\zeta}^{*}(t) < \zeta^{*}\} \\
\leq \sum_{t=1}^{T} \mathbb{P}\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \ge \zeta^{*}\} + O(K),$$
(23)

where the second term is bounded by $\sum_{t=1}^{T} K \left[\frac{\log t}{\log(\alpha+1/2)} + 1 \right] t^{-\frac{2\alpha}{\alpha+1/2}} = O(K)$ according to Lemmas 4 and 6.

To bound the first term in (23), with a slight abuse of notation, we choose two numbers $x_{\sigma_{i(j)},i}$ and $y_{\sigma_{i(j)},i}$ such that $p_{\sigma_{i(j)},i} < x_{\sigma_{i(j)},i} < y_{\sigma_{i(j)},i} < p_{\sigma_{i(L_C+1)},i}$, and define the following events:

$$\mathcal{E}^{p}_{\sigma_{i(j)},i}(t) = \{\bar{p}_{\sigma_{i(j)},i}(t) < x_{\sigma_{i(j)},i}\},\$$
$$\mathcal{E}^{\theta}_{\sigma_{i(j)},i}(t) = \{\theta^{(2)}_{\sigma_{i(j)},i}(t) < y_{\sigma_{i(j)},i}\},\$$

where the existence of $x_{\sigma_{i(j)},i}$ and $y_{\sigma_{i(j)},i}$ is guaranteed under Assumption 2. Then, the first term can be decomposed as

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \geq \zeta^{*} \right\}$$

$$= \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \geq \zeta^{*}, \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t) \right\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \geq \zeta^{*}, \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \neg \mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t) \right\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \geq \zeta^{*}, \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t) \right\}$$

$$\leq \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^{*}(t) \geq \zeta^{*}, \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t) \right\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \neg \mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t) \right\}$$

$$+ \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \neg \mathcal{E}_{\sigma_{i(j)},i}^{p}(t) \right\}.$$
(24)

Now we bound each term in Eq. (24) respectively.

a) First term:
$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t) \right\}$$

For the first term, we note that when $\hat{\zeta}^*(t) \ge \zeta^*$, the first candidate $a_t^{(1)}$ could be i only when there exists a $j' \le L_C + 1$, such that $l_{\sigma_{i(j')},i}(t) \le 1/2$. Thus,

$$\mathbb{P}\left\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t)\right\}$$

$$\leq \sum_{j'=1}^{L_C+1} \mathbb{P}\left\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, l_{\sigma_{i(j')}, i}(t) \le 1/2, \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t)\right\}.$$

For each $j' \leq L_C + 1$, define the following probability:

$$q_{j'j}^{(i)}(t) = \mathbb{P}\{\theta_{\sigma_{i(j')},i}^{(2)}(t) > y_{\sigma_{i(j)},i}|\mathcal{H}_{t-1}\}.$$
(25)

Similar to Lemma 7, we can show that

$$\mathbb{P}\left\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, l_{\sigma_{i(j')}, i}(t) \le 1/2, \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t) | \mathcal{H}_{t-1}\right\} \le \frac{1 - q_{j'j}^{(i)}(t)}{q_{j'j}^{(i)}(t)} \mathbb{P}\left\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j')}), \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t) | \mathcal{H}_{t-1}\right\},$$

and its summation over t can be bounded as

$$\sum_{t=1}^{I} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, l_{\sigma_{i(j')}, i}(t) \le 1/2, \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t) \right\} = O(1).$$

Considering all j' from 1 to $L_C + 1$, we have

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \hat{\zeta}^*(t) \ge \zeta^*, \mathcal{E}^p_{\sigma_{i(j)}, i}(t), \mathcal{E}^\theta_{\sigma_{i(j)}, i}(t) \right\} = O(L_C + 1).$$
(26)

b) Second term: $\sum_{t=1}^{T} \mathbb{P}\left\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \exists \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t)\right\}$

We use the back substitution argument to refine the second term to $O(\log \log T)$. When fixing the first candidate as $a_t^{(1)} = i$, the comparison between $a_t^{(1)}$ and other arms is similar to a traditional MAB. Let $N_i^{(1)}(T) = \sum_{t=1}^T \mathbb{1}(a_t^{(1)} = i)$ be the number of time-slots when this type of MAB is played, and let

$$L_{ji}^{\beta}(n) = \frac{\beta \log n}{D(x_{\sigma_{i(j)},i} || y_{\sigma_{i(j)},i})}$$

Then, considering all possible cases of $N_i^{(1)}(T)$ and $N_{i\sigma_{i(j)}}^{(1)}(t-1)$, we have

$$\sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t)\right\}$$

$$\leq \sum_{n=0}^{T} \sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) \leq L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n\right\}$$

$$+ \sum_{n=0}^{T} \sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n\right\}$$

For the first case, note that

m

$$\sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) \leq L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n\right\}$$

$$\leq \sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), N_{i\sigma_{i(j)}}^{(1)}(t-1) \leq L_{ji}^{\beta}(n) | N_{i}^{(1)}(T) = n\right\} \mathbb{P}\left\{N_{i}^{(1)}(T) = n\right\}$$

$$\leq L_{ji}^{\beta}(n) \mathbb{P}\left\{N_{i}^{(1)}(T) = n\right\}, \qquad (27)$$

similar to the analysis for Eq. (14).

Then, we have

$$\sum_{n=0}^{T} \sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \neg \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t), N_{ij_{ik}}^{(1)}(t-1) \leq L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n\right\}$$

$$\leq \mathbb{E}[L_{ji}^{\beta}(n)] \leq \frac{\beta \log(\mathbb{E}[N_{i}^{(1)}(T)])}{D(x_{\sigma_{i(j)}, i}||y_{\sigma_{i(j)}, i}|)} \leq \frac{\beta \left(\log \log T + O(\log K)\right)}{D(x_{\sigma_{i(j)}, i}||y_{\sigma_{i(j)}, i}|)},$$

where the last inequality follows from the concavity of the $\log(\cdot)$ function and the fact that $\mathbb{E}[N_i^{(1)}(T)] = O(K \log T)$ as shown when proving Proposition 1.

For the second case, let $\tau_m^{(i)}$ be time-slot where $a_t^{(1)} = i$ for the *m*-th time and $\tau_0^{(i)} = 0$. Then

$$\sum_{t=1}^{I} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathbb{V}\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n \right\}$$

$$\leq \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}\left((a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathbb{V}\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n \right) \right]$$

$$\leq \sum_{m=0}^{n} \mathbb{E}\left[\sum_{t=\tau_{m}+1}^{\tau_{m+1}} \mathbb{1}\left((a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathbb{V}\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n) \right]$$

$$\stackrel{(a)}{\leq} \sum_{m=0}^{n} \mathbb{E}\left[\mathbb{1}\left((a_{\tau_{m+1}}^{(1)}, a_{\tau_{m+1}}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)},i}^{p}(\tau_{m+1}), \mathbb{V}\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(\tau_{m+1}), N_{i\sigma_{i(j)}}^{(1)}(\tau_{m+1}-1) > L_{ji}^{\beta}(n) \right]$$

$$\leq \sum_{m=0}^{n} \mathbb{P}\left\{\mathcal{E}_{\sigma_{i(j)},i}^{p}(t), \mathbb{V}\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n) \right\}$$

$$^{(b)}_{\leq} n \cdot \frac{1}{n^{\beta}} = \frac{1}{n^{\beta-1}},$$
(28)

where (a) holds because $a_t^{(1)} = i$ could only happen at $t = \tau_{m+1}^{(i)}$; (b) is true because, the two sets of samples in D-TS are drawn independently and their distributions only depend on the historic comparison results; thus, given $N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\alpha}(n)$, the events $\mathcal{E}_{\sigma_{i(j)},i}^{p}(t)$ and $\mathcal{E}_{\sigma_{i(j)},i}^{\theta}(t)$ are independent of $t = \tau_{m+1}^{(i)}$, and the probability can be bounded according to the concentration property of Thompson samples (in the proof of Lemma 3 in [15]):

$$\mathbb{P}\{\mathcal{E}^{p}_{\sigma_{i(j)},i}(t), \neg \mathcal{E}^{\theta}_{\sigma_{i(j)},i}(t), N^{(1)}_{i\sigma_{i(j)}}(t-1) > L^{\beta}_{ji}(n)\} \le e^{-L^{\beta}_{ji}(n)D(x_{\sigma_{i(j)},i}||y_{\sigma_{i(j)},i}|} = \frac{1}{n^{\beta}}$$

Then for $\beta > 2$,

$$\sum_{n=0}^{T} \sum_{t=1}^{T} \mathbb{P}\left\{(a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t), \neg \mathcal{E}_{\sigma_{i(j)}, i}^{\theta}(t), N_{i\sigma_{i(j)}}^{(1)}(t-1) > L_{ji}^{\beta}(n), N_{i}^{(1)}(T) = n\right\}$$

$$\leq \sum_{n=0}^{T} \frac{1}{n^{\beta-1}} = O(1).$$
(29)

Combining the above two cases, we can bound the second term of Eq. (24) as:

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \mathcal{E}_{\sigma_{i(j)}, i}^p(t), \neg \mathcal{E}_{\sigma_{i(j)}, i}^\theta(t) \right\} \le \frac{\beta \left(\log \log T + O(\log K) \right)}{D(x_{\sigma_{i(j)}, i} || y_{\sigma_{i(j)}, i})} + O(1).$$
(30)

c) Third term: $\sum_{t=1}^{T} \mathbb{P}\{(a_t^{(1)}, a_t^{(2)}) = (i, \sigma_{i(j)}), \exists \mathcal{E}_{\sigma_{i(j)}, i}^p(t)\}$

For the third term, we can bound it as

$$\sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t) \right\}$$

$$\leq \sum_{t=1}^{T} \mathbb{P}\left\{ (a_{t}^{(1)}, a_{t}^{(2)}) = (i, \sigma_{i(j)}), \forall \mathcal{E}_{\sigma_{i(j)}, i}^{p}(t) \right\}$$

$$\leq \frac{1}{D(x_{\sigma_{i(j)}, i} || p_{\sigma_{i(j)}, i})} + 1, \qquad (31)$$

where the last inequality follows from Lemma 3 in [15].

Combining the analysis for all above three terms, we can bound Eq. (24) by choosing appropriate $x_{\sigma_{i(j)},i}$ and $y_{\sigma_{i(j)},i}$. Specifically, for any $\epsilon > 0$, similar to [15], we can choose appropriate $x_{\sigma_{i(j)},i}$ and $y_{\sigma_{i(j)},i}$ such that $D(x_{\sigma_{i(j)},i}||y_{\sigma_{i(j)},i}) = D(p_{\sigma_{i(j)},i}||p_{\sigma_{i(L_{C}+1)},i})/(1+\epsilon)^{2}$ and $\frac{1}{D(x_{\sigma_{i(j)},i}||p_{\sigma_{i(j)},i}||} = O(\frac{1}{\epsilon^{2}})$. The conclusion of Lemma 5 then follows.

D Additional Experimental Results

This appendix presents additional experimental results using both synthetic and real-world data, to further evaluate the performance of the proposed algorithms, in comparison to the state-of-the-art schemes.

Because the simulation complexity is $O(K^2T)$, we adjust the number of independent experiments to save time, and run 500, 100, and 10 independent experiments for $K < 10, 10 \le K \le 100$, and K > 100, respectively. In each experiment, the arms are randomly shuffled to prevent algorithms from exploiting special structures of the preference matrix, except for ECW-RMED in the "Gap" dataset, where we run experiments for both fixed and shuffled arm orders.

D.1 Datasets

D.1.1 Condorcet Dueling Bandits

Cyclic: A dataset adopted from [5], where the preference matrix is given by Table 1. In this dataset, Arm 1 is the Condorcet winner with $p_{1j} = 0.6$, and the other arms have a cyclic preference relationship with one arm beating another with high probability. Strong transitivity does not hold in this dataset, and the Condorcet winner is not necessary the best arm when comparing all other arms with a fixed arm, i.e., $p_{1i} < \max_i p_{ii}$ for $i \neq 1$.

StrongBorda: A 5-armed dueling bandit with a preference matrix in Table 2. In addition to a Condorcet winner, there is a strong Borda winner, which is not the Condorcet winner, but beats the other arms with high probability. To validate the correctness of algorithms, we still treat this problem as a Condorcet dueling bandit problem and try to find the Condorcet winner, although a Borda winner may be more appropriate in this case [22].

ArXiv: A 6-armed dueling bandits with a preference matrix given in Table 3, which is derived by conducting pairwise interleaving experiments [16] based on the search engine of ArXiv.org.

Sushi: A 16-armed dueling bandits with a preference matrix derived by [5, 7] from a Sushi preference dataset, where the matrix can be found in the appendix of [7].

Table 2: StrongBorda												Tał	ole 3:	ArXiv				
Table 1: Cyclic				1	2. 54	2	1		_	1	1	2	3	4	5	6		
	1	2	3	4	1	0.5	0.55	0.55	0.55	0.55		2	0.30	0.55	0.55	0.0 .	0.61 0.58	0.61
1	0.5	0.6	0.6	0.6	2	0.45	0.5	0.95	0.95	0.95		3	0.45	0.45	0.50	0.54	0.51	0.56
23	0.4	0.5 0.1	0.9 0.5	0.1 0.9	3 4	0.45	$0.05 \\ 0.05$	0.5 0.05	0.95 0.5	0.95 0.95		4 5	0.46	0.45 0.42	0.46 0.49	0.50 0.46	0.54 0.50	0.50 0.51
4	0.4	0.9	0.1	0.5	5	0.45	0.05	0.05	0.05	0.5		6	0.39	0.40	0.44	0.50	0.49	0.50

D.1.2 Non-Condorcet Dueling Bandits

Non-Condorcet Cyclic: A 9-armed dueling bandit with a preference matrix given by Table 6. In this dataset, there is a cyclic preference relationship among arms, and the arms can be divided into 3 groups with Copeland scores 6, 4, and 2, respectively. Due to this cyclic symmetry, there are multiple Copeland winners with exactly the same performance.

Non-Condorcet StrongBorda: Similar to the Condorcet dueling bandits, we consider this case where there is a Borda winner different from the Copeland winner. In this non-Condorcet setting, we even assume that when comparing the Copeland winner and the Borda winner, the user prefers the Borda winner to the Copeland winner with high probability. Again, this is a extreme case used to the validate the correctness of algorithms.

Gap: A dataset adopted from [7], which is a 5-armed dueling bandits with a preference matrix given by Table 5. In this dataset, the ratio between the regret bound for ECW-RMED and the regret bound for the optimal CW-RMED algorithm is very large.

500-Armed Dueling Bandits: The 500-armed dueling bandit constructed in [6], where there are three Copeland winners that form a cycle and each has a Copeland score 498, and the other arms have Copeland scores ranging from 0 to 496. We use this dataset to evaluate the scaling behaviors of the algorithms.

Table 4: Gap

	1	2	3	4	5
	0.5				
2	0.2 0.2	0.5	0.8	0.2	0.8
3	0.2	0.2	0.5	0.8	0.8
4	0.49 0.8	0.8	0.2	0.5	0.2
5	0.8	0.2	0.2	0.8	0.5

Table 5: Non-Condorcet StrongBorda

		2				
1	0.5	0.05	0.55	0.55	0.55	0.55
2	0.95	0.5	0.95	0.95	0.45	0.45
3	0.45	0.05	0.5	0.95	0.95	0.95
4	0.45	0.05	0.05	0.5	0.95	0.95
5	0.45	0.55	0.05	0.05	0.5	0.95
6	0.45	0.55	0.05	0.05	$\begin{array}{c} 0.95\\ 0.45\\ 0.95\\ 0.95\\ 0.5\\ 0.05\end{array}$	0.5

Table 6: Non-Condorcet Cyclic

									9
1	0.5	0.4	0.6	0.1	0.6	0.6	0.6	0.6	0.6 0.6
2	0.6	0.5	0.4	0.6	0.1	0.6	0.6	0.6	0.6
3	0.4	0.6	0.5	0.6	0.6	0.1	0.6	0.6	0.6
4	0.9	0.4	0.4	0.5	0.1	0.9	0.6	0.6	0.4
5	0.4	0.9	0.4	0.9	0.5	0.1	0.4	0.6	0.6
6	0.4	0.4	0.9	0.1	0.9	0.5	0.6	0.4	0.6
7	0.4	0.4	0.4	0.4	0.6	0.4	0.5	0.1	0.9
8	0.4	0.4	0.4	0.4	0.4	0.6	0.9	0.5	0.1
9	0.4	0.4	0.4	0.6	0.4	0.4	0.1	0.9	0.5

D.1.3 MSLR (Condorcet and non-Condorcet)

For the MSLR dataset, we have evaluated the algorithms in the two 5-arm cases in Section 5, where the preference matrices are given by Tables 7 and 8. In this appendix, we also run experiments for larger scale dueling bandits, K = 16 and 32, consisting of arms randomly selected from the 136 rankers (in the *MSLR_Informational_PMat.npz* file on http://bit.ly/nips15data, [6]; to see the asymptotic performance within an acceptable T, we eliminate the arms with $|p_{ij} - 1/2| < 0.003$). The indices of the chosen rankers and their Copeland scores are presented in Table 9, where the arms are indexed from 1. Note that due to the randomness, our chosen rankers are likely different from those in [5] and [7], even for the same K.

Table 7: MSLR (K = 5, Condorcet)

Table 8: MSLR	(K = 5,	non-Condorcet)
---------------	---------	----------------

	1	2	3	4	5
1	0.500	0.535	0.613	0.757	0.765 0.738 0.669
2	0.465	0.500	0.580	0.727	0.738
3	0.387	0.420	0.500	0.659	0.669
4	0 243	0 273	0 341	0 500	0 510
5	0.235	0.262	0.331	0.490	0.500

	1	2	3	4	5
1	0.500	0.484	0.519	0.529 0.530	0.518
2	0.516	0.500	0.481	0.530	0.539
				0.504	
4	0.471	0.470	0.496	0.500	0.503
5	0.482	0.461	0.488	0.497	0.500

Subset name	Chosen rankers (Copeland score)	Winners (Copeland score)
K = 16, Condorcet	10(0), 22(6), 36(12), 58(9), 59(10), 66(5), 67(3), 68(4), 77(2), 98(1), 109(11), 112(7), 115(15), 116(13), 117(8), 125(14)	115(15)
K = 32, Condorcet	7(9), 21(26), 24(12), 28(5), 32(13), 35(24), 37(15), 38(19), 43(6), 44(4), 45(10), 48(21), 52(3), 56(28), 61(29), 68(8), 71(27), 73(20), 82(14), 85(25), 87(17), 91(11), 96(1), 99(2), 100(0), 105(7), 112(16), 115(31), 117(18), 121(30), 123(22), 133(23)	115(31)
K = 16, non-Condorcet	2(6), 21(11), 42(4), 60(14), 67(5), 70(8), 79(3), 90(12), 91(7), 98(0), 99(1), 103(2), 106(14), 109(10), 117(9), 130(14)	60(14), 106(14), 130(14)
K = 32, non-Condorcet	4(8), 8(4), 12(24), 16(1), 20(0), 24(13), 28(7), 32(14), 36(25), 40(26), 44(6), 48(18), 52(5), 56(28), 60(30), 64(22), 68(10), 72(16), 76(19), 80(21), 84(12), 88(17), 92(9), 96(3), 100(2), 104(11), 108(23), 112(15), 116(28), 120(30), 124(20), 128(29)	60(30), 120(30)

Та	ble	9:	MSLR	(K	= 16	and	32,	arms	are	indexed	from	1))
----	-----	----	------	----	------	-----	-----	------	-----	---------	------	----	---

D.2 Performance Comparisons

We first analyze the regret performance of algorithms for Condorcet and non-Condorcet dueling bandits, respectively, and then discuss their robustness with respect to preference matrix and delayed feedback, as well as the impact of RUCB/RLCB elimination.

D.2.1 Condorcet Dueling Bandits

Fig. 3 shows the cumulative regret of all algorithms in Condorcet dueling bandits. From Figs. 3(a) to 3(f), we can see that D-TS and D-TS⁺ achieve similar performance, because there is a unique winner in the system and few ties occur when choosing the first candidate. Compared to existing algorithms, D-TS and D-TS⁺ perform much better, except for the StrongBorda dataset, as discussed later.

Compared with the earlier trial RCS, our D-TS and D-TS⁺ algorithms lead to more extensive utilization of TS and significantly reduce the regret. Although the RCS algorithm achieves better performance than RUCB and CCB in real-world datasets (Figs. 3(c) to 3(f)) by leveraging TS for selecting the first candidate, the improvement is limited. This is because RCS requires a "100%-pass" when using "majority voting" to select the first candidate, i.e., the arm that beats all the other arms with respect to the samples, and randomly picks one if no such an arm exists. Thus, it may miss many opportunities to choose the Condorcet winner as the first candidate. In contrast, using RUCB-based elimination when choosing the first candidate, D-TS/D-TS⁺ only need a simple "majority voting" rule. Thus, under D-TS/D-TS⁺, the Condorcet winner can be chosen as the first candidate even it is beaten by (a few) other arms with respect to the samples. In addition, without requiring "100%-pass", D-TS and D-TS⁺ directly apply to general Copeland dueling bandits. Moreover, by launching another round of sampling for selecting the second candidate, D-TS and D-TS⁺ further reduce the regret.

Compared with RMED1 and ECW-RMED, D-TS and D-TS⁺ still perform better. Note that without knowing the existence of a Condorcet winner, ECW-RMED performs slightly worse than RMED1, while achieving similar asymptotic performance. RMED1 has been shown to achieve the optimal asymptotic performance, in both asymptotic order $O(\log T)$ and *coefficients* [5]³. Interestingly, D-TS and D-TS⁺ not only approach the similar asymptotic performance as RMED1 and ECW-RMED, but also achieve a smaller *constant term*. This is because, at the beginning stage, when the empirical estimates for p_{ij} deviate from the true value, RMED-type algorithms may temporally trap in exploring the non-winner arms and result in a very large constant term. In contrast, by TS, the best arm always has a positive probability to be explored [13], and will be identified as the winner much faster. This property makes TS-type algorithms better for the scenarios where the system statistics are not stationary and slowly varying over time.

For the StrongBorda dataset, Fig. 3(b) shows that D-TS and D-TS⁺ can escape from the suboptimal comparisons and achieve very good performance compared to existing algorithms, except for (Condorcet-)SAVAGE. Somewhat surprisingly, SAVAGE performs best in this dataset with respect to the Copeland-score-based regret. This is because in this dataset, all arms except for the Condorcet winner (arm 1) are beaten by the Borda winner (arm 2) with a preference probability close to 1. Thus, arms 3 to 5 can be eliminated without causing too much loss by comparing with arm 2. This property in fact leads to a very small lower regret bound according to [7]. As an "explore-then-exploit" algorithm and with the awareness of the existence of a Condorcet winner, SAVAGE can eliminate the suboptimal arms quickly and achieve low regret. ECW-RMED does not perform very well in this dataset as its regret bound is much higher than the optimal lower bound. Note that these results do not conflict with the optimality of RMED1 in Condorcet dueling bandits with respect to the regret

³The optimality of RMED1 is shown under another definition of regret [5], which depends on the probabilitygap-to-the-winner (G2W). The trend based on the G2W regret is similar to that in this paper, except for the StrongBorda dataset.

defined in [5]. We can see that SAVAGE approaches the same asymptotic performance to RMED1 if we compare them based probability-gap-to-the-winner regret (for other datasets, the trend and relative relationship is similar for both definitions of regret, and the results are omitted here).

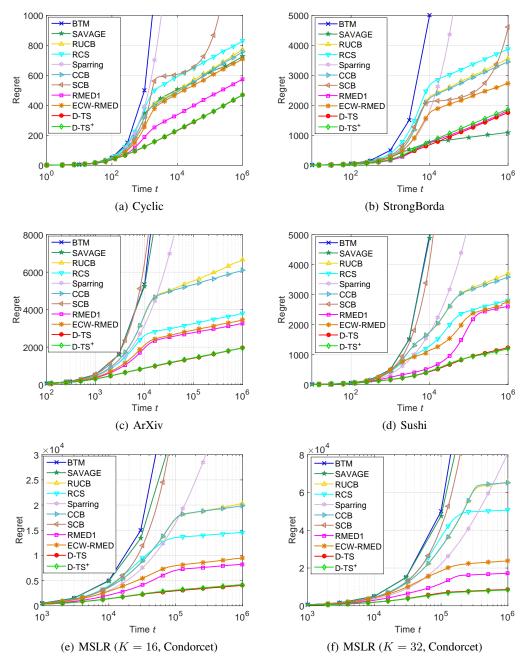


Figure 3: Cumulative regret in Condorcet dueling bandits.

D.2.2 Non-Condorcet Dueling Bandits

Fig. 4 presents the cumulative regret in non-Condorcet dueling bandits. Note that one simulation in the large scale 500-armed dueling bandit takes very long time, and we only compare CCB, SCB, ECW-RMED, and D-TS, for this dataset.

From Fig. 4, we can see that the regret of algorithms dedicated to Condorcet dueling bandits, including BTM, RUCB, RCS, and RMED1, grows rapidly, because these algorithms keep exploring for the

Condorcet winner, which does not exist in these datasets. The Sparring algorithm, whose regret is not theoretically guaranteed, also results in a very large regret in non-Condorcet dueling bandits. The SAVAGE algorithm requires to explore all pairs, and usually leads to large regret.

Compared D-TS and D-TS⁺ with the UCB-type algorithms for general Copeland dueling bandits, CCB and SCB, we can see that D-TS and D-TS⁺ perform much better. In particular, as shown in Fig. 4(d) for the 500-armed scenario, D-TS still achieves much lower regret than SCB, the scalable version of CCB for large scale systems.

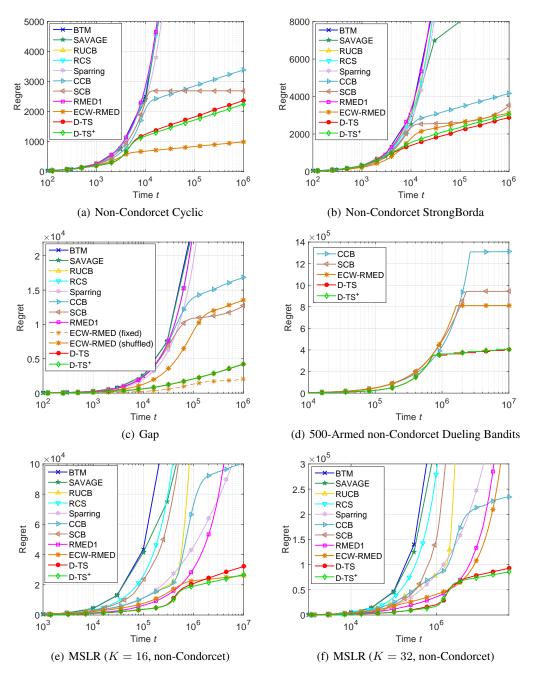


Figure 4: Cumulative regret in non-Condorcet dueling bandits.

Compared with the recently developed algorithm ECW-RMED, D-TS and D-TS+ usually converge faster to the asymptotic regime and thus can achieve smaller regret for small to relatively large T. In theory, the optimal CW-RMED achieves the best asymptotic performance. ECW-RMED

is its efficient but approximate implementation, with a larger coefficient in terms of asymptotic performance. Therefore, in the asymptotic regime, ECW-RMED could outperform D-TS⁺ in certain scenarios. In practice, however, we notice that D-TS⁺ could perform better than ECW-RMED, even for relatively large T, e.g., $T = 5 \times 10^6$. This is because ECW-RMED estimates the required number of comparisons based on the empirical preference probability. Then ECW-RMED may temporally trap in suboptimal comparisons at the beginning stage when the empirical preference probability is likely to deviate from its true value. In contrast, D-TS and D-TS⁺ make decisions in a random manner, and the winner(s) has a positive probability to be explored even when the empirical estimates deviate from the true values. Thus, D-TS and D-TS⁺ usually converge to the asymptotic regime more quickly than ECW-RMED. Because of this, D-TS and D-TS⁺ may outperform ECW-RMED even for a relatively large T, especially when the number of arms is larger. For example, as shown in Fig. 4(e), ECW-RMED performs worse than D-TS⁺ when $t \leq 5 \times 10^6$, although it has better asymptotic performance. This situation becomes more serious when the number of arms increases, as we can see from Fig. 4(f).

Our D-TS algorithm performs worse than ECW-RMED when there are multiple winners with similar performance, e.g., for the non-Condorcet Cyclic dataset shown in Fig. 4(a). One main reason is that with a random tie-breaking rule, D-TS randomly explores all potential winners in each individual sample path ⁴, as shown in Fig. 5. Thus, the regret of D-TS scales with the number of winners $|C^*|$. By carefully breaking the ties, D-TS⁺ can reduce the regret in many practical scenarios, as shown in Fig. 1(b), Fig. 4(e), and Fig. 4(f). However, the improvement is limited especially when the winners have very similar or exactly the same performance, as shown in Fig. 4(a). From the perspective of regret optimization, this may be a disadvantage of TS, where based on randomly sampled belief, all winners have a positive probability to be explored [13]. However, from the diversity perspective, this may be desirable in the application scenarios such as restaurant recommendation, where users may not want to stick to a single winner.

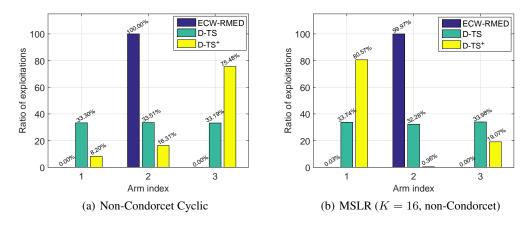


Figure 5: Diversity of exploitations: distribution of exploitations over different Copeland winners in one sample path. This is calculated when an arm is compared against itself.

D.2.3 Robustness

We study the robustness of ECW-RMED, D-TS, and D-TS⁺, with respect to the preference matrix and delayed feedback.

Influence of preference matrix: We have seen from Fig. 2 that, when some preference probabilities for different arms are close to 1/2 (5-armed non-Condorcet MSLR dataset), the regret of ECW-RMED fluctuates significantly and has a very large standard deviation, while D-TS and D-TS⁺ have much smaller regret deviation.

The robustness of D-TS and D-TS⁺ to the preference matrix can also be seen from the results in the Gap dataset. Recall that in this dataset, the regret bound of ECW-RMED is much larger than the optimal lower bound. By experiments, we find that the performance of ECW-RMED is significantly affected by the order of arms (and essentially the preference matrix). For example,

⁴Another reason is that D-TS explores the superiors for each arm sequentially and results in a regret higher than the lower bound in [7].

when the order of arms is fixed as that in Table 5, ECW-RMED performs very well. Specifically, as shown by the dashed line in Fig. 4(c), the regret of ECW-RMED is very small and its asymptotic behavior is even better than its asymptotic bound (similar results can be found in [7]). However, when the arms are randomly shuffled in each experiment, ECW-RMED achieves much larger regret that is consistent with its asymptotic bound. In contrast, D-TS and D-TS⁺ do not depend on the order of arms and perform much better than ECW-RMED on average in this dataset.

Influence of delayed feedback: In practice, it may be difficult and costly to process each individual comparison result immediately. Typically, feedback will be batched and provided periodically, say every d time-slots. We evaluate the algorithm performance with respect to the feedback delay. As we can see from Fig. 6, as the feedback delay increases, the regret of ECW-RMED increases much faster than D-TS and D-TS⁺. In particular, even in the non-Condorcet Cyclic dataset with multiple winners (Fig. 6(b)), the regret of ECW-RMED becomes larger than D-TS and D-TS⁺ when the feedback delay is larger than about 300 time-slots.

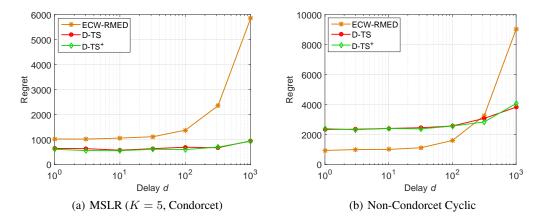


Figure 6: Influence of feedback delay: the regret when the feedback is batched and provided every d time-slots ($T = 10^6$).

D.2.4 Impact of RUCB/RLCB Elimination

In this section, we illustrate the necessity of RUCB/RLCB elimination in D-TS and D-TS⁺. A pure Thompson Sampling algorithms without the auxiliary RUCB/RLCB elimination step seem to be more elegant and may be more efficient (without the limitation of RUCB/RLCB). However, we notify that the RUCB/RLCB elimination is necessary to guarantee sublinear or logarithmic regret in general settings, especially in non-Condorcet dueling bandits.

Specifically, we consider the following "pure D-TS" algorithm, which is similar to D-TS in Algorithm 1, except that the RUCB/RCLB elimination step is ignored, i.e., the candidates $a^{(1)}$ and $a^{(2)}$ are selected from all arms according to the samples $\theta_{ij}^{(1)}$ and $\theta_{ia^{(1)}}^{(2)}$. As shown in Fig 7, pure D-TS may result in large regret in certain scenarios. For example, in the (Condorcet) StrongBorda dataset (Table 2), the Borda winner (Arm 2) beats all the other arms with high probability except for the Condorcet winner (Arm 1), and hence, it is easier for the Borda winner to get more votes at the beginning and to be chosen as the first candidate. Thus, pure D-TS achieves higher regret in this case as shown in Fig 7(a). In non-Condorcet dueling bandits, without RLCB elimination, pure D-TS could achieve linear regret if a Copeland winner is beaten by a non-winner arm. For example, as shown in Fig. 7(b), the algorithm fails to converge to comparing the Copeland winners against themselves, as the non-winner arm will have higher samples with high probability at the second round. By introducing RLCB elimination, D-TS/D-TS⁺ can avoid trapping in theses suboptimal comparisons and achieve much better performance.

We also point out without the limitation of RUCB/RLCB, pure D-TS may achieve in certain practical scenarios, e.g., the 5-arm MSLR data sets as shown in Figs 7(c) and 7(d). In fact, the RLCB elimination can be ignored in Condorcet dueling bandits, because none of the other arms can beat the Condorcet winner. It is an interesting direction to identify the conditions under which pure D-TS can perform better and obtain its theoretical performance under these conditions.

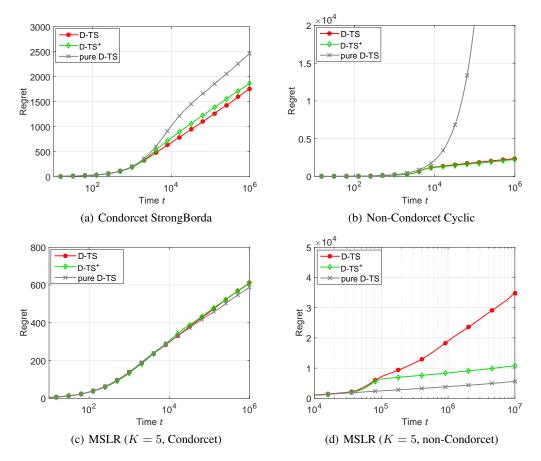


Figure 7: Impact of RUCB/RLCB elimination: "pure D-TS" is similar to D-TS, except that the RUCB/RLCB elimination is ignored and the candidates are selected among all arms according to the Thompson samples only.

D.2.5 Summary of Experimental Results

We summarize the performance evaluation results based on synthetic and real-world data:

- In Condorcet dueling bandits, D-TS and D-TS⁺ achieve similar performance, and both perform much better than existing algorithms that work for unknown/infinite time-horizon settings. This benefits from the double sampling structure that we proposed for D-TS and D-TS⁺.
- In non-Condorcet dueling bandits, D-TS and D-TS⁺ performs much better than UCB-type algorithms, CCB and SCB. Again, this benefits from the double sampling structure of D-TS and D-TS⁺, assisted by RUCB/RLCB elimination that guarantees sublinear or logarithmic regret in general settings.
- Compared with ECW-RMED, D-TS and D-TS⁺ achieve much better performance in Condorcet dueling bandits, better or close-to performance in non-Condorcet dueling bandits, especially when T is small to relatively large. Furthermore, D-TS and D-TS⁺ are also much more robust with respect to the preference matrix and delayed feedback.

In practice, we may not know in advance whether we have a Condorcet and non-Condorcet dueling bandit. We may have in practice a time-varying system and delayed feedback. Overall, good performance, and robustness of D-TS and D-TS⁺ make them strong candidates in practice.