

Statistical Consequences of Dueling Bandits

Nayan Saxena
Department of Statistical
Sciences
University of Toronto
nayan.saxena@mail.utoronto.ca

Pan Chen
Department of Computer
Science
University of Toronto
pan.chen@mail.utoronto.ca

Emmy Liu
Department of Computer
Science, Cognitive Science
Program
University of Toronto
me.liu@utoronto.ca

ABSTRACT

Multi-Armed-Bandit frameworks have often been used by researchers to assess educational interventions, however, recent work has shown that it is more beneficial for a student to provide qualitative feedback through preference elicitation between different alternatives, making a dueling bandits framework more appropriate. In this paper, we explore the statistical quality of data under this framework by comparing traditional uniform sampling to a dueling bandit algorithm and find that dueling bandit algorithms perform well at cumulative regret minimisation, but lead to inflated Type-I error rates and reduced power under certain circumstances. Through these results we provide insight into the challenges and opportunities in using dueling bandit algorithms to run adaptive experiments.

Keywords

randomized experiments; adaptive experiments; dueling bandit algorithms; double Thompson sampling; statistical inference; hypothesis testing; reinforcement learning

1. INTRODUCTION

Rapid growth in online learning has provided scientists and researchers a new digital platform to conduct randomised experiments with students in real-world settings. This data can be leveraged by education researchers to adaptively assign students carefully tailored educational material, and further improve the quality of their content through careful exploratory assessment of the assigned conditions [19, 23]. To achieve these goals, one of the most popular problem frameworks is the Multi-Armed Bandit (MAB) problem which focuses on optimal assignment of conditions based on numerical reward signals from the participant (student) whose goal is to trade-off between exploration and exploitation of conditions [3, 20]. Historically, MAB algorithms have been used in industrial settings to leverage user feedback and adaptively display more popular advertisements, websites and produce content [8, 13], whereas in research settings these algorithms have been used for clinical trials [14, 21, 4], robot control [15, 16], and by behavioral and social scientists for crowd-sourcing experiments [1, 5, 2].

While MAB algorithms focus entirely on quantitative feedback (scalar reward), it has been shown that human motor learning is maximised when subjected to both qualitative and quantitative feedback [12] and recent work suggests that quantitative metrics might not be the best indicators of true human preferences [18]. Therefore, a more qualitative ap-

proach for preference elicitation from human participants should be more ideal to leverage student feedback, which can in turn be achieved through preference based MAB (dueling bandit) algorithms like Double Thompson Sampling (DTS) [7, 24], which focuses on participants simply choosing between two presented alternatives.

Through this paper, we take a closer look at leveraging these dueling bandit algorithms for adaptive experimentation by empirically investigating the statistical properties of the DTS algorithm against uniform random experimental condition selection. Our study is primarily motivated by prior work assessing the statistical properties of MAB algorithms in an educational setting [19], and builds upon recent work exploring the challenges that come with conducting hypothesis tests to analyze data from adaptive experiments using bandit algorithms [22], such as proposing strategies for modifying MAB algorithms to trade-off reward and power [25] and improving coverage of confidence intervals [10].

In summary, our main contributions are:

- (i) We emphasise the conceptual significance of quantifying the differences in statistical power, regret and false positive rate between uniform sampling and Double Thompson Sampling
- (ii) Through simulation experiments we show that for a reasonable decrease in power, it may be advantageous to run experiments with dueling bandits when the number of arms is small and the expected effect size is large
- (iii) By applying the same analysis on the real-world Microsoft Learning to Rank (LTR) dataset we confirm that using DTS results in lowered power, regret and a higher proportion of people being assigned to the better alternative.

2. PRELIMINARIES

2.1 Multi-armed Bandits

As a precursor to dueling bandits, we consider a central problem in sequential decision making and adaptive experiment design known as the multi-armed bandit (MAB) problem [20]. The problem setup consists of k arms (actions) such that for every action $a_i \in \mathcal{A}$ the corresponding probability of yielding a success is $p_i \in [0, 1]$. With no a priori information about these success probabilities, the central goal for an agent is to maximise the overall number of successes by pulling these arms, or performing these actions, and ultimately settling on one over a fixed period of time denoted

by t . Assume that under this framework the expected reward of arm a_i is given by $\mu(a_i)$ and $\mu^* = \arg \max \mu(a_i)$, for $i \in \{1 \dots k\}$. The objective of an agent is to then minimise the cumulative regret, $\mathcal{R}_t^{\text{MAB}} = \sum_{i=1}^t [\mu^* - \mu_{a(i)}]$.

Here the agent is allowed to only choose one action, $a_{(i)}$ at a given time step. However, in many cases during adaptive experimentation, it may be difficult to frame the result of an action as a success or failure. In particular, in many cases what experimenters are interested in is simply which arm is preferred to the others, in which case pairwise preferences may be more appropriate for the subject to make. To accommodate pairwise preferences we now consider a generalised form of the MAB problem also known as the dueling bandit problem.

2.2 Dueling Bandits

The dueling bandit problem can be characterised as a special case of the popular Multi-armed Bandit (MAB) problem that focuses on pairwise comparisons between actions at every iteration [26]. The problem setup is similar to the MAB problem, except at every iteration the agent chooses two actions $a_m, a_n \in \mathcal{A}$ and performs a comparison before choosing one action that they prefer¹. Throughout this paper we specifically utilise the Double Thompson Sampling (DTS) algorithm, which is an adaptation of the popular Thompson Sampling algorithm in the regular MAB context [24]. A comprehensive outline of the DTS procedure can be found in Appendix A.

Formally, the probability of one arm winning over another is given by $P(a_m \succ a_n)$ which we abbreviate as $P_{mn} \equiv P(a_m \succ a_n)$. For each trial i the outcome of each comparison is binary $x_i \sim \text{Bernoulli}(P_{mn})$ where the probability of one arm winning is formally given by, $P(a_m \succ a_n) = \Delta(a_m, a_n) + 0.5$. Here, $\Delta(a_m, a_n) \in [-0.5, 0.5]$ is the difference measure between the two actions, abbreviated as $\Delta_{mn} \equiv \Delta(a_m, a_n)$, such that $a_m \succ a_n \iff \Delta_{mn} > 0$. It is further assumed that the probabilities representing user preferences are stationary over time and modelled by the unknown preference matrix denoted by $\mathbf{P} = [P_{mn}]_{k \times k}$ where the entries satisfy $P_{mn} + P_{nm} = 1$.

The goal of an agent operating in this environment is to choose the optimal *winning* arm when conducting pairwise comparisons and ultimately minimise regret. Motivated by classical voting theory, there are different notions of a *winning* arm like the Borda winner [11] and Copeland winner [27, 9]. In this paper we focus on using the notion of a Condorcet winner [6] and strong regret [24] as described below.

Condorcet Winner: The Condorcet winner is a single action a^* that beats all other actions in a pairwise comparison, such that $\Delta_{a^*a} > 0$ for all $a \neq a^*$. The Copeland winner, found by maximising the normalised Copeland score $\frac{1}{k-1} \sum_{m \neq n} \mathbb{1}[P_{mn} > 0.5]$ is a Condorcet criterion, which means that it always finds the Condorcet winner if one exists. It should be noted that it is possible for no Condorcet

¹The traditional dueling bandit framework allows comparisons between the same actions (self-dueling) where $m = n$, but throughout this paper we assume $m \neq n$ as it is unlikely for users to be presented with comparisons between the same actions during adaptive experimentation.

winner to exist, in which case the same definition can be extended to obtain a set of Copeland winners which always exist. [27, 24].

Strong Regret: Suppose $a_m^{(i)}$ and $a_n^{(i)}$ are actions chosen at a given timestep i and a_o is the optimal action. Then under this framework the overall objective for an agent is to minimise the cumulative strong regret, $\mathcal{R}_t^{\text{DUBEL}} = \sum_{i=1}^t [\Delta_{om} + \Delta_{on}]$.

3. METHODS

3.1 Experimental Setup

In order to examine the impact of dueling bandit algorithms as compared to uniform random assignment during adaptive experimentation, we looked at several conditions relevant to experimenters by varying the number of arms and the effect sizes between pairs of arms. For each set of conditions 5000 simulations were carried out, each corresponding to a real-world study, where data at each timestep would correspond to data collected from a participant exposed to a treatment arm. Furthermore, to assess the statistical quality of each algorithm, we considered the statistical power in comparing pairs of arms, the regret accumulated, and the false positive rate. During its operation, each simulation assigned hypothetical participants using either the DTS algorithm, or uniformly at random and was initialized using a preference matrix $P \in \mathbb{R}^{n \times n}$, where n is the number of arms. For each given pair of arms (i, j) , the effect size was calculated (Cohen’s w of 0.1, 0.3 or 0.5) of the comparison based on the difference (0.05, 0.15, or 0.25) between P_{ij} and P_{ji} , noting that $P_{ij} + P_{ji} = 1$. Finally, in order to compute the false positive rate, preference matrices with zero effect size for all comparisons between pairs of arms were also used. These experimental conditions are also summarized in Table 1. Data and code are publicly available at <https://bit.ly/35J2xH0>.

3.2 Experimental Dataset

In addition to simulated data, we also present experiments using data collected from users in a real-world context. In particular, we perform simulations using the Microsoft Learning to Rank (LTR) dataset which presents pairs of search queries and documents, along with the relevance ratings of the document [17]. As A/B testing is commonly used to test ranking and recommendation algorithms, this is an appropriate context in which we benchmark outcomes of such experiments. We used an implicit preference matrix over all 136 rankers derived from this dataset similar to [27], and then randomly sampled Condorcet and non-Condorcet submatrices from this larger matrix in order to simulate experiments comparing different rankers. The sample size is set as $\binom{n}{2} \times m$, where arm i and arm j are the pair of arms with the smallest effect size, and m is the number of participants needed to achieve expected statistical power of 0.8 for effect size 0.1. This simulates the real-world scenario where the number of required participants is unknown. To further evaluate the long-term performance of the two assignment methods, we ran the simulations with sample sizes up to 10 times more than the initial sample size and tracked the same set of metrics as in the main analysis.

3.3 Analysis

Condition		Description
Sampling type		Double Thompson Sampling Uniform Random Sampling
Number of arms		$n \in \{3, 5\}$
Sample sizes	Uniform effect sizes	Non-zero effect sizes: $s = 10mn$ where m is the number of participants needed to achieve expected statistical power of 0.8 and n is the number of paired comparisons in that simulation
	Zero effect size	Same sample sizes as in the non-zero effect case calculated using the number of participants needed to achieve expected statistical power of 0.8 for effect size 0.3.
	Learning to Rank dataset	Different effect sizes: $s = 10mn$, where m is the number of participants needed to achieve expected statistical power of 0.8 for 0.1 effect size, and n is the number of paired comparisons in that simulation.
Effect sizes	Simulated dataset	None (0), Small (0.1), Medium (0.3), Large (0.5). Same effect size between each pair of arms. The winning arm in each pair is randomly assigned in non-zero cases.
	Learning to Rank dataset	For each pair of arms i and j , the effect size is in $[0, 1)$.

Table 1: Summary of conditions varied across simulations

After each simulation, we shift our focus to analyzing several outcomes relevant to experimenters: (a) statistical power to detect an effect between each pair of arms, (b) false positive rate in detecting effects when none exist, (c) regret over time, and (d) percentage of participants assigned to the Condorcet winner when one exists. For simulations where there was some difference between pairs of arms, a Chi-squared contingency test was performed for each pair of arms, with significance level $\alpha = 0.05$, which is the standard across multiple domains including educational experiments. For each simulation, regret over time, reward over time, and percentage of participants assigned to the Condorcet winner was recorded and aggregated. Finally, to calculate the false positive rate, we considered a pair of arms as a false positive if the comparison between them reached significance, given there should be no effect between these two arms.

4. RESULTS

4.1 Synthetic Data

4.1.1 Conditions that differ in terms of effectiveness

Average Power Over Time: In educational settings, different conditions, such as sample solutions, may have a different impact on students’ learning efficiency and engagement with the educational resources. However, one condition that outperforms the others is often not statistically significant until testing is conducted with many students involved. Yet, different ways of conducting the experiments may require a different number of participants to reach significant effects amongst conditions.

We modeled the means of the power of different pairs of conditions, as this measures how well each assignment did in finding statistically significant effects among all the conditions. It can be observed in Figure 4.1 that uniform sampling consistently reached the 0.8 power threshold with fewer participants as compared to DTS, in both 3 and 5-condition settings. This pattern was also observed when final power was recorded for both uniform sampling and DTS, as seen in Figure B.1 (Appendix B.1), thus indicating that DTS is more susceptible to inflated Type-II error rates (reduced power). This shows that when there are insufficient students to test among multiple variants, uniform sampling is more often a better way of assignment given the effects among different conditions are of the highest priority.

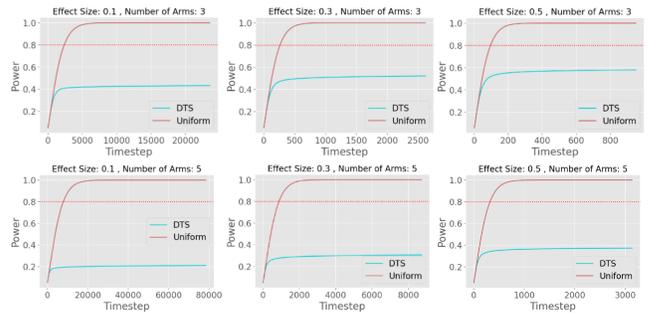


Figure 4.1: Average power over time for DTS algorithm (blue) and Uniform assignment (red) with synthetic data.

Proportion of Condorcet Winners: In educational experiments, given that there is one unknown best condition, we examined how uniform sampling and double Thompson sampling fulfill the goal of providing students with the best condition by calculating the proportion of trials in which a student would be assigned to the Condorcet winner.

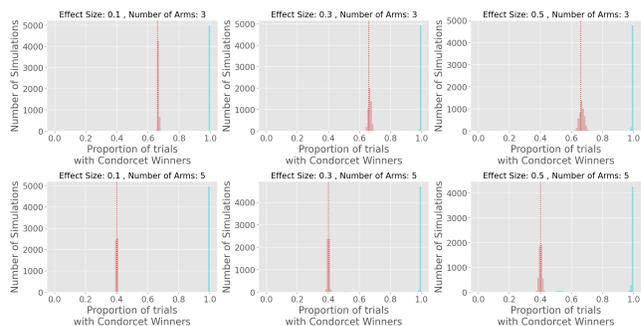


Figure 4.2: Proportion of trials with Condorcet winners for DTS algorithm (blue) and Uniform assignment (red) with synthetic data. Here, the dotted line indicates expected assignment.

From Figure 4.2, it can be observed that, on average, DTS assigned a higher number of Condorcet winners to students who participated in the experiments. Across both 3 and

5-condition simulations, we observe that DTS assigned the Condorcet winner to almost every student, whereas with uniform sampling the number of students receiving the Condorcet winner is in line with the statistical expectation.

Cumulative Strong Regret: Apart from the proportion of Condorcet winners, we used cumulative strong regret to measure the different experiences uniform sampling and DTS give to students. As seen in Figure B.2 (Appendix B.1), a noticeable difference in cumulative strong regret between uniform sampling and DTS can be seen, which grew more evident with increase in effect size from 0.5 to 0.1, and number of arms from 3 to 5. Across all simulations, DTS performed better by accumulating less regret as compared to uniform sampling.

4.1.2 Conditions that are equally-effective

False Positive Rate: To simulate scientific settings like clinical trials, along with behavioral and social sciences where little to no difference between arms is commonly observed, we measured the false positive rate when conditions are equally effective. From Figure 4.3, we observe that the false positive rates for uniform sampling remain consistent across 3-condition and 5-condition scenarios and are relatively lower when compared to DTS algorithm.

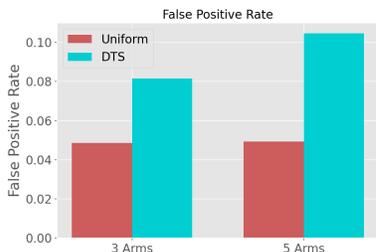


Figure 4.3: False positive rate for DTS algorithm (blue) and Uniform assignment (red) with simulated data.

This indicates that when there is little to no difference between conditions the DTS algorithm will result in a higher Type-I error as compared to uniform sampling. These issues are especially problematic for scientific research, since erroneously believing an intervention is better than a control may result in lack of reproducibility.

4.2 Learning to Rank Dataset

In real settings the effect sizes between any two conditions are not limited to the set $\{0.1, 0.3, 0.5\}$, therefore we ran the same experiments on the LTR dataset. Overall, our results on the real-world dataset were consistent with our simulation studies. We noticed that DTS provided the best condition to students more often and accrued relatively lower average power as seen in Figure 4.4. It should further be noted that in a few trials DTS did not assign the Condorcet winner to the majority of students possibly because the differences between the best condition and other arms is miniscule, thus requiring more participants for the algorithm to learn the best condition. Consistent with our results on synthetic data, it can further be observed in Figure 4.4 that uniform sampling reached the 0.8 power threshold with fewer participants as compared to DTS across both 3 and 5-condition

settings. This pattern was also observed when final power was recorded for both uniform sampling and DTS, as seen in Figure B.3 (Appendix B.2), thus reinforcing the finding that DTS is more susceptible to reduced power even in real world settings. Finally, we also observed that DTS does a much better job at minimising regret than uniform sampling on the LTR data set as seen in Figure B.4 (Appendix B.2). Overall, we note that both uniform sampling and DTS are useful algorithms in educational contexts depending on the available resources and potential intervention goals. Uniform sampling is more appropriate in cases where finding significant effects is of higher priority and the number of students is small, whereas DTS is more advantageous when better student experience is more important through optimal allocation.

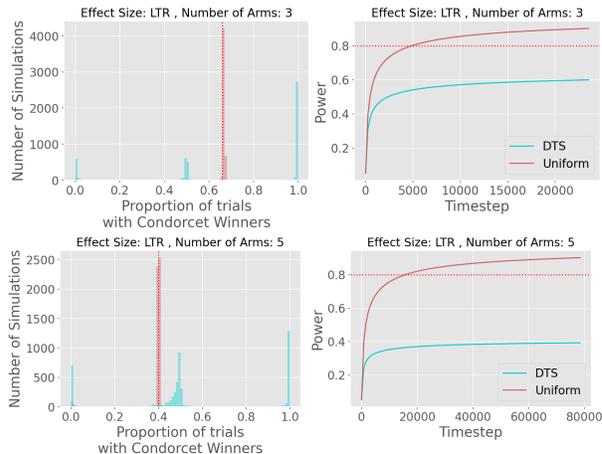


Figure 4.4: Proportion of trials with Condorcet winners alongside average power over time for DTS algorithm (blue) and Uniform assignment (red) with LTR dataset.

5. CONCLUSION

Adaptive experimentation through bandits presents an important opportunity for researchers, as it allows more participants to receive the probable best intervention. To explore the statistical quality of data during randomised experimentation under this framework we conducted an analysis of statistical power, regret, and false positive rates comparing traditional uniform sampling to the DTS algorithm. Our findings indicate that the DTS algorithm shows excellent performance with respect to cumulative regret minimisation but in the case when there is little to no difference between arms, DTS leads to higher Type-I error rates and reduced power. This makes it increasingly difficult to collect data that enables statistical hypothesis tests between conditions, while also balancing this with reward maximization and minimizing the false positive rate. These findings pave the way for researchers to better leverage machine learning for adaptive experimentation, while enabling sound statistical inference. They also present a new opportunity for future work to explore ways of incorporating the DTS algorithm in adaptive experiments where both reward optimisation and statistical inference are important, possibly through algorithm-driven dynamic interpolation between DTS and uniform exploration for better performance.

6. REFERENCES

- [1] I. Abraham, O. Alonso, V. Kandydas, and A. Slivkins. Adaptive crowdsourcing algorithms for the bandit survey problem. In *Conference on learning theory*, pages 882–910. PMLR, 2013.
- [2] S. Athey, S. Wager, et al. Efficient policy learning. *arXiv preprint arXiv:1702.02896*, 78, 2017.
- [3] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- [4] J. Bather. Randomised allocation of treatments in sequential trials. *Advances in Applied Probability*, pages 174–182, 1980.
- [5] J. P. Berkowitz, C. A. Depken II, and J. M. Gandar. One-armed bandits and behavioral economics. In *Handbook of Behavioral Industrial Organization*. Edward Elgar Publishing, 2018.
- [6] D. Black et al. The theory of committees and elections. 1958.
- [7] R. Busa-Fekete and E. Hüllermeier. A survey of preference-based online learning with bandit algorithms. In *International Conference on Algorithmic Learning Theory*, pages 18–39. Springer, 2014.
- [8] D. Chakrabarti, R. Kumar, F. Radlinski, and E. Upfal. Mortal multi-armed bandits. *Advances in neural information processing systems*, 21:273–280, 2008.
- [9] A. Copeland. A ‘reasonable’ social welfare function, seminar on mathematics in social sciences, university of michigan. *Cited indirectly from its mention by Luce and Raiffa (1957)*, page 358, 1951.
- [10] Y. Deshpande, L. Mackey, V. Syrgkanis, and M. Taddy. Accurate inference for adaptive linear models. In *International Conference on Machine Learning*, pages 1194–1203. PMLR, 2018.
- [11] K. Jamieson, S. Katariya, A. Deshpande, and R. Nowak. Sparse dueling bandits. In *Artificial Intelligence and Statistics*, pages 416–424. PMLR, 2015.
- [12] N. C. Kilduski and M. S. Rice. Qualitative and quantitative knowledge of results: effects on motor learning. *American Journal of Occupational Therapy*, 57(3):329–336, 2003.
- [13] P. Kohli, M. Salek, and G. Stoddard. A fast bandit algorithm for recommendation to users with heterogeneous tastes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 2013.
- [14] V. Kuleshov and D. Precup. Algorithms for multi-armed bandit problems. *arXiv preprint arXiv:1402.6028*, 2014.
- [15] P. Matikainen, P. M. Furlong, R. Sukthankar, and M. Hebert. Multi-armed recommendation bandits for selecting state machine policies for robotic systems. In *2013 IEEE International Conference on Robotics and Automation*, pages 4545–4551. IEEE, 2013.
- [16] S. McGuire, P. M. Furlong, C. Heckman, S. Julier, D. Szafir, and N. Ahmed. Failure is not an option: Policy learning for adaptive recovery in space operations. *IEEE Robotics and Automation Letters*, 3(3):1639–1646, 2018.
- [17] T. Qin and T. Liu. Introducing LETOR 4.0 datasets. *CoRR*, abs/1306.2597, 2013.
- [18] F. Radlinski, M. Kurup, and T. Joachims. How does clickthrough data reflect retrieval quality? In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 43–52, 2008.
- [19] A. Rafferty, H. Ying, and J. Williams. Statistical consequences of using multi-armed bandits to conduct adaptive educational experiments. *JEDM Journal of Educational Data Mining*, 11(1):47–79, 2019.
- [20] H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- [21] Y. Sui and J. Burdick. Clinical online recommendation with subgroup rank feedback. In *Proceedings of the 8th ACM conference on recommender systems*, pages 289–292, 2014.
- [22] S. S. Villar, J. Bowden, and J. Wason. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 30(2):199, 2015.
- [23] J. J. Williams, A. N. Rafferty, D. Tingley, A. Ang, W. S. Lasecki, and J. Kim. Enhancing online problems through instructor-centered tools for randomized experiments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2018.
- [24] H. Wu and X. Liu. Double thompson sampling for dueling bandits. *arXiv preprint arXiv:1604.07101*, 2016.
- [25] J. Yao, E. Brunskill, W. Pan, S. Murphy, and F. Doshi-Velez. Power-constrained bandits. *arXiv preprint arXiv:2004.06230*, 2020.
- [26] 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.
- [27] M. Zoghi, Z. Karnin, S. Whiteson, and M. De Rijke. Copeland dueling bandits. *arXiv preprint arXiv:1506.00312*, 2015.

APPENDIX

A. DOUBLE THOMPSON SAMPLING

Algorithm 1 D-TS for Copeland Dueling Bandits [24]

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

B. SUPPLEMENTARY RESULTS

B.1 Synthetic Data

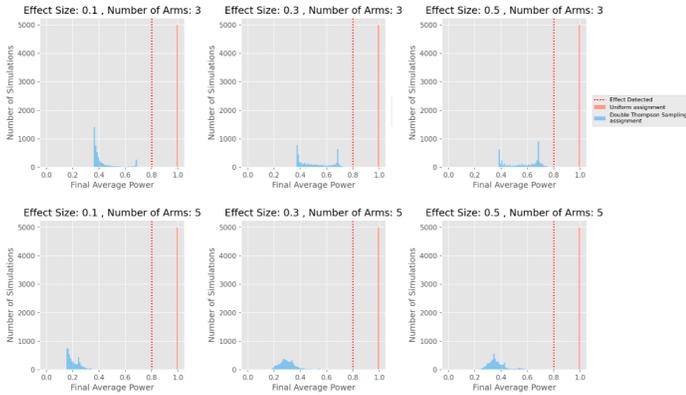


Figure B.1: Comparison of final power between DTS algorithm (blue) and Uniform assignment (red) for simulated dataset.

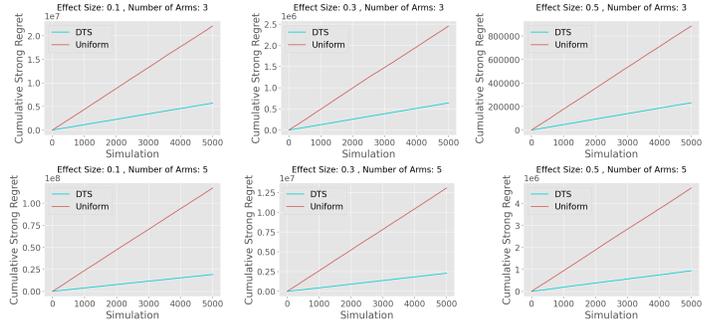


Figure B.2: Comparison of cumulative strong regret over time between DTS algorithm (blue) and Uniform assignment (red) for simulated dataset.

B.2 Real World Data

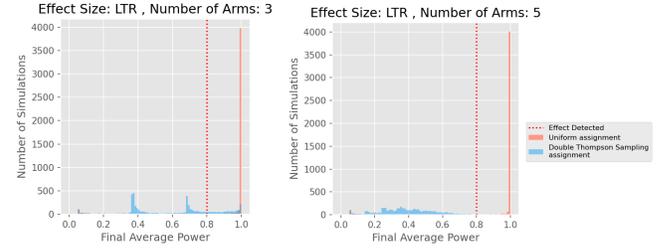


Figure B.3: Comparison of final power between DTS algorithm (blue) and Uniform assignment (red) for LTR dataset.

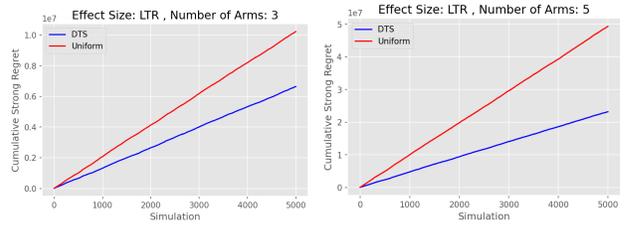


Figure B.4: Comparison of cumulative strong regret over time between DTS algorithm (blue) and Uniform assignment (red) for LTR dataset